Accounting for Education in Earnings Inequality: New Evidence from Survey-Linked Administrative Data

Michael D. Carr* Emily E. Wiemers†

August 27, 2018

Abstract

This paper uses survey-linked administrative earnings data to document trends since 1980 in short-run and long-run within-group, or residual, earnings inequality and within group intragenerational earnings mobility. Our administrative data are unique in allowing the estimation of earnings inequality overall and earnings inequality within education and gender and education groups, allowing for the first analysis of residual inequality and mobility with administrative earnings data. Consistent with other work, we find large increases in top-end residual inequality in both short- and long-run earnings. However, while there is a small decline in short-run residual inequality at the bottom of the distribution, consistent with polarizing earnings, there is a substantial increaes in bottom-end residual inequality in long-run earnings. This result points towards rising permanent earnings inequality and declining mobility. Indeed, we find notable declines in within group intragenerational mobility over the same time period. However, given that the skills and gender composition of the labor force has changed considerably, and the levels and trends in long-run earnings inequality could reflect this compositional change, we follow the wage inequality literature and employ a kernel reweighting method to assess the impact of the skills by gender composition on this trend. Preliminary results from this counterfactual exercise demonstrate that the bulk of the rise in cross-sectional earnings inequality is due to a changing distribution of within group prices, not changes in the characteristics of the labor force. This finding is consistent with the wage inequality literature. We also find that declining within group mobility is due to rising inequality, and not to changes in the distribution of characteristics of the labor force. A new finding which sheds important light on the changing role of education in generating upward mobility.

*Department of Economics, University of Massachusetts-Boston, michael.carr@umb.edu

†Department of Economics, University of Massachusetts-Boston, emily.wiemers@umb.edu

The authors gratefully acknowledge funding for this project from the Russell Sage Foundation (#83-15-09) and the Washington Center for Equitable Growth.

1 Introduction

Over the last forty years, the United States has experienced a very well-documented increase in cross-sectional wage inequality, with an analogous increase in annual income inequality (Autor, Katz, and Kearney, 2008; Gottschalk and Danziger, 2005; Lemieux, 2006a; Piketty and Saez, 2003). Despite the focus on measuring short-term inequality, inequality of longer-run economic outcomes is perhaps more relevant to individual decisions about labor supply, wealth building, and investment in human capital. Moreover, short-run earnings inequality does not translate directly into long-run inequality because increases in mobility in economic outcomes over the life-cycle could offset the consequences of short-run inequality for the distribution longer-run economic outcomes.

The increase in cross-sectional wage inequality was driven by increases in inequality both between education, age, and occupation subgroups and within these groups.[1] One of the most important insights from this literature is that rising residual inequality is due both to skill and routine biased technological change that altered the relative return to particular skills, and the increase in the relative shares of groups in the labor force who tend to have higher within group inequality (Autor, Katz, and Kearney, 2008; Lemieux, 2006a,b). Similarly, rising between group inequality is also partly the result of changes in relative supply and demand of educational attainment groups, and demographic and institutional factors that differentially affect groups.

Wage inequality is critical for understanding changes in how the labor market values workers with different characteristics, but it is not a complete description of inequality in labor market outcomes more broadly. This paper builds on the literature on short-run wage inequality to ask whether the increase in short-run earnings inequality over the last 40 years

[1]This literature is quite large. See Berman, Bound, and Griliches (1994); Bound and Johnson (1992); Katz and Murphy (1992) for key analyses documenting the initial rise in wage inequality, and Autor, Katz, and Kearney (2008); DiNardo, Fortin, and Lemieux (1996); Gottschalk and Danziger (2005); Lemieux (2006a) for literatures documenting the continued rise in inequality and resulting polarization.

has translated into long-run earnings inequality. We use administrative-linked survey data from the Survey of Income and Program Participation Gold Standard File (SIPP GSF) to document trends in short-run earnings inequality, long-run earnings inequality, and lifetime earnings mobility since 1980. Drawing on the evidence of the importance of both returns to skills and labor force composition on wage inequality, we consider the evolution of inequality and mobility both between and within education and gender subgroups. Our use of survey-linked administrative data is crucial for these analyses because the SIPP GSF is, to our knowledge, the only source of data with administrative earnings histories that also includes self-reported measures of educational attainment. Finally, we follow the wage inequality literature and use a kernel reweighting procedure, first developed in DiNardo, Fortin, and Lemieux (1996), to examine the relative importance of the returns to skills and the changes in the relative shares of groups in the labor force to the overall trends in long-run inequality.

This paper is the first to use administrative data to examine long-run earnings inequality both between and within education groups and to consider the role of labor force composition in shaping these trends in long-run earnings inequality, though it is not the first paper to consider the link between short-run and long-run earnings inequality. There is a substantial literature that decomposes the increase in annual earnings inequality into its permanent and transitory components (Baker and Solon, 2003; Carr and Wiemers, 2016b; Debacker et al., 2013; Gottschalk and Moffitt, 2009; Gottschalk et al., 1994; Haider, 2001; Moffitt and Gottschalk, 2002, 2012). The broad consensus in this literature is that there has been both a widening of the permanent component of earnings and an increase in transitory earnings instability since the 1970s. This literature informs our understanding of whether increasing annual inequality accumulates over a working lifetime or whether short-term variability mitigates increases in inequality, though it does not address this issue directly. Kopczuk, Saez, and Song (2010) tackle the question of an explicit link between short-run inequality, long-run inequality, and mobility directly by examining trends in these three concepts over

the period between 1950 and 1980. They find that increases in short-run inequality translated into increases in long-run inequality though levels of long-run inequality remain lower than short-run inequality. Moreover, they find that lifetime earnings mobility was relatively stable in the period through 1980 which suggests that changes in mobility were not crucial for understanding long-run economic inequality.[2]

We build on Kopczuk, Saez, and Song (2010) who, because of data constraints, are unable to consider the differences in trends in short-run and long-run earnings inequality within education subgroups. As the literature on wage inequality would suggest, we show that understanding trends in long-run inequality within education subgroups is crucially important. We show that short- and long-run increases in earnings inequality are driven largely by substantial increases in earnings inequality for college-educated workers, particularly college-educated men. In contrast to recent narratives about rising wage inequality which focuses on the polarization of job growth, we show that both top-end and bottom-end earnings inequality rose for college-educated men, though there is some evidence of polarization for college-educated women. Moreover, the rise in long-run inequality is much larger than the rise in short-run inequality. While short-run inequality for college graduates increased by 36 percent between 1980 and 2007, long-run inequality increased by a staggering 70 percent. The implication of this pattern, as we show, is that, in contrast to an earlier period of stable lifetime earnings mobility, declining lifetime earnings mobility reinforced the effects of increases in short-run inequality. Finally, we show that the increase in long-run inequality is driven largely by changes in relative returns rather than in the composition of the labor force though these results are still preliminary.

[2]There is a small literature that looks only at lifetime earnings mobility, with no directly link to inequality. This work uses either administrative earnings records or the Panel Study of Income Dynamics (PSID) and suggests that overall, mobility has been stable or slightly increasing over time (Acs and Zimmerman, 2008; Auten and Gee, 2009; Auten, Gee, and Turner, 2013), though Bradbury and Katz (2009) find evidence of declining mobility.

2 Data and Samples

2.1 Data

The data for this project come from the Survey of Income and Program Participation Gold Standard File (SIPP GSF), which takes all members of families who participate in the SIPP survey and links them to administrative earnings histories. The SIPP survey is a nationally representative sample of the civilian noninstitutionalized population of the U.S. that began in 1984. There have been 14 SIPP panels since 1984, with each panel comprised of a new nationally representative sample of 14,000 to 52,000 households. The SIPP GSF links *each individual* in a SIPP household in the 1984, and 1990 – 2008 SIPP panels to their IRS and SSA earnings and benefits records through 2011.[3] The administrative data provides non-topcoded total earnings from 1978 to 2011, including deferred and non-deferred earnings from all jobs and from self-employment but do not include under the table earnings not reported to the IRS. Prior to 1978, the dataset includes FICA taxable earnings back to 1951. If all earnings values are zero or missing, then the individual had zero earnings for that year.[4] In

[3]This analysis was first performed using the SIPP Synthetic Beta (SSB) on the Synthetic Data Server housed at Cornell University which is funded by NSF Grant #SES-1042181. These data are public use and may be accessed by researchers outside secure Census facilities. For more information, visit https://www.census.gov/programs-surveys/sipp/methodology/sipp-synthetic-beta-data-product.html. Final results for this paper were obtained from a validation analysis conducted by Census Bureau staff using the SIPP Completed Gold Standard Files and the programs written by this author and originally run on the SSB. The validation analysis does not imply endorsement by the Census Bureau of any methods, results, opinions, or views presented in this paper.

[4]Missing data can arise either because the SIPP survey participant refused to answer a specific demographic question or because the SIPP respondent could not be matched to administrative earnings or benefits data. The match rate for most panels is quite high. In the 1980's and 1990's panels, the match rate hovers around 80%. In 2001, the match rate dropped to 47% because many SIPP participants refused to provide social security numbers. Beginning with the 2004 panel, the match rate increased to around 90% because the Census Bureau changed its matching procedures removing the necessity to explicitly ask for social security numbers. While the public use SIPP has missing observations that are imputed using a hot-deck method, the Gold Standard File uses a substantially more sophisticated multiple imputation method to replace missing observations (see Abowd and Stinson (2013) for details). The Census Bureau advises against excluding imputed observations and we have thus included these observations. It is important to note that the low match rate in 2001 only affects individuals interviewed in the 2001 SIPP panel, it has no implications for the ability to follow individuals interviewed in other panels through the 2000s.

addition to the administrative earnings records, the SIPP GSF has basic demographic and human capital variables, though variables collected in the SIPP panels that are not linked to administrative data cover only the years of the individual's SIPP panel.

The SIPP GSF has some important advantages for understanding the link between short-run earnings inequality, long-run earnings inequality, and lifetime earnings mobility because it contains both long earnings panels with no attrition and large cross-sections. Critically, it is also the only data on the U.S. we are aware of that combines administrative earnings with data on educational attainment. The earnings data in the SIPP GSF are similar to the administrative data used in Kopczuk, Saez, and Song (2010) and Auten and Gee (2009), however their data do not include human capital. In addition, the fact that the data is closer to nationally representative of both workers and non-workers creates an advantage over Kopczuk, Saez, and Song (2010), whose long time series of mobility from the early 1950s necessitated the use of workers in a subset of industries. The disadvantage of the SIPP GSF compared with the SSA data in Kopczuk, Saez, and Song (2010) is that the SIPP GSF does not contain the quarter in which the Social Security earnings cap was reached and so does not allow us to impute non-topcoded earnings prior to 1978. The cross-sectional sample size is also smaller, though still sizeable. For this reason, we focus on the more recent period for which the data are not topcoded. The PSID is the obvious alternative data source. However, samples of long panels in the PSID typically do not have a large cross-section, and may not be population representative. Moreover, the sample is not large enough to examine inequality within subgroups.

2.2 Samples

We use a sample of individuals 25 to 59 who meet a minimum earnings threshold based on a multi-year average of earnings and have positive earnings in year t. We use annual earnings to measure short-run earnings inequality and seven-year average earnings to measure long-

run earnings inequality. To reduce the impact of individuals with very marginal labor force attachment, average earnings for individual i over a seven-year period $(t-3, t+3)$ must be above a minimum threshold of one-fourth of a full-year full-time minimum wage in 2013 ($3770) indexed to inflation. For both short-run and long-run earnings we measure inequality using the residuals from a regression of earnings on a quadratic in age, estimated separately by year. To calculate seven-year average earnings, we average the age-adjusted residuals over a seven-year period. Short-run earnings inequality and long-run earnings inequality are calculated using the same sample for each year t to preserve the intuition that short-run inequality should be at least as large as long-run inequality because of the smoothing of annual earnings fluctuations. Described in more detail below, we measure mobility in long-run earnings over a 15-year window and so our mobility sample has the additional restriction that we observe individual i in year t and $t+15$.

3 Methods

3.1 Measuring Mobility

The measure of mobility used here is a rank-based measure, meaning it represents positional (or relative) mobility. Specifically, we measure the change in rank based on seven-year average age-adjusted earnings over a 15 year period. Individuals are ranked relative to other individuals in their cohort, rather than using the entire sample of 25 to 59 year olds in a given calendar year. That is, rank in years t and $t+15$, respectively, is determined relative to individuals present in both years t and $t+15$. Mobility is thus measured relative to an individual's cohort, and does not include the impact of new entrants into or exit from the labor force during the 15-year window.

As a simple summary of trends in mobility through time, we use a rank-rank regression

similar to that used in Chetty et al. (2014). Given in equation 1,

$$rank_{i,t+15} = \beta_{0t} + \beta_{1t} rank_{it} + \varepsilon_{it} \tag{1}$$

the rank-rank regression regresses ending rank $(t+15)$ on starting rank (t) for all individuals present in both t and $t+15$. The coefficient β_{1t} represents the correlation in rank between long-run earnings measured 15 years apart, or the amount of relative mobility in the cohort, while the coefficient β_{0t} represents the average change in the number of ranks within the cohort.

What makes the rank-rank regression appealing as a summary measure of mobility is that the combination of β_{0t} and β_{1t} is all one needs to describe expected mobility for any point in the earnings distribution. Of course, as with any linear regression, this is only true if the relationship between starting and ending rank is linear. If the relationship is not linear, then expected mobility depends on starting rank, and the rank-rank regression no longer summarizes mobility for the entire distribution. A second limitation is that it is also possible for different patterns of mobility at different points in the earnings distribution to result in the same average relationship between starting and ending rank. In fact, in a cohort based model of mobility, we would expect this to be true as individuals who start at the top of the earnings distribution clearly can only stay at the top or move down, while the opposite is true for individuals who start at the bottom. Though not reported here, the relationship between starting and ending rank is approximately linear, with the exception of the expected non-linearity at the top and bottom of the earnings distribution. Further, for the present purposes, it is sufficient to know that broad qualitative trends in the rank-rank regressions are born out in transition matrices (not reported below), though as Carr and Wiemers (2016a) it is possible for rank-rank regressions to show qualitatively different patterns in mobility.

4 Inequality

Overall inequality increased about 21%. The cyclicality is notable, with spikes in inequality during recessions, followed by periods of decline. The increase is due almost entirely to changing prices, not composition. As with hourly wages, changes at the top are different than at the bottom. Top end inequality increased steadily, with a substantial surge in the early 1990s. In total, top end inequality increased 46%. This increase was due almost entirely to prices, with no evident cyclicality. Bottom end inequality is high cylical. Given recessions in in the early 1980s and 2007, it is difficult to tell whether there there is any trend in the 50/10, though it ends the period somewhat higher than it began. Here, there is a small role for compositional changes: applying the 1996 skills distribution to the 1981 earnings distribution results in an increase in inequality that is about half as large as the overall increase, while inequality would have fallen slightly if the 1981 skills distribution had held through the entire period.

In sum, changes in earnings inequality are less about the type of polarization seen in wages and more about the top pulling away from everyone else. This represents a widening distribution of prices and is not about changes in the age by education distribution.

Overall inequality for women has increased somewhat, though not nearly as much as for men. This overall increase hides a notable increase in top-end inequality and a notable decrease in bottom end inequality. Thus, female annual earnings are becoming strongly polarized. The role of changing characteristics for women is more complicated, in part possibly because the changes in the level of inequality are smaller. Most notably, while top-end inequality would have been higher for all years had the 2008 skills distribution held througout the time period, bottom-end inequality would have been lower had the 2008 skills distribution held throughout.

Residual, or within group, earnings inequality has increased about 17%, implying that

Figure 1: Annual Earnings Inequality

(a) 90/10: Men

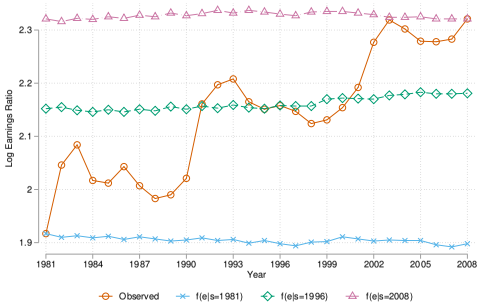

(b) 90/50: Men

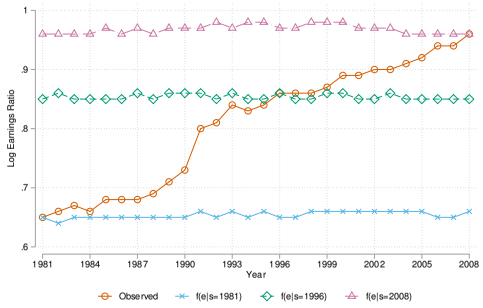

(c) 50/10: Men

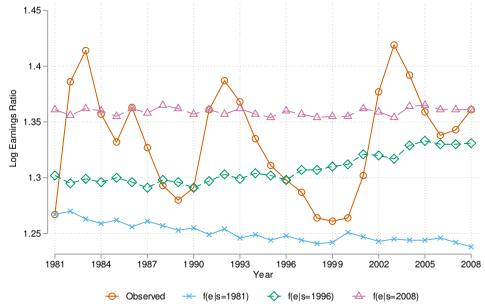

(d) 90/10: Women

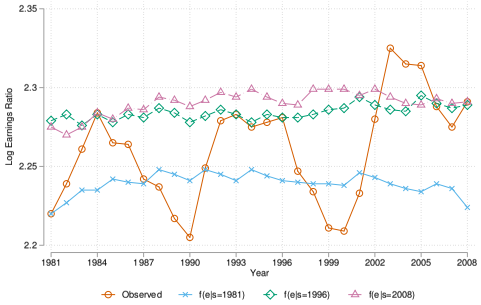

(e) 90/50: Women

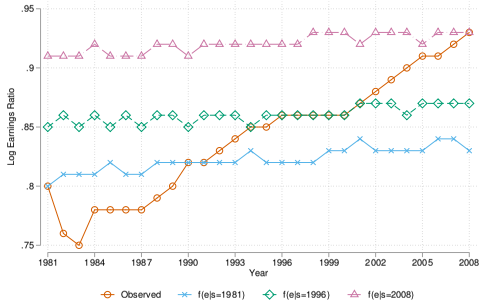

(f) 50/10: Women

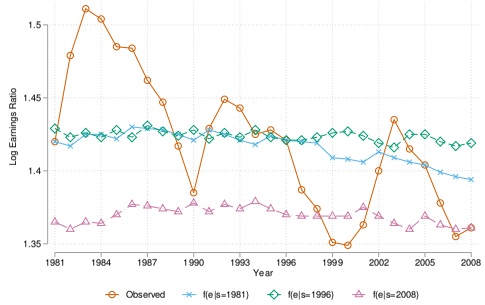

Figure 2: Annual Residual Earnings Inequality

(a) 90/10: Men

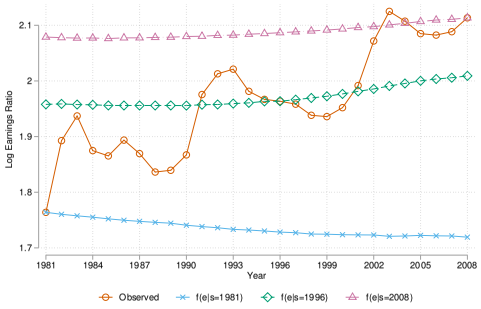

(b) 90/50: Men

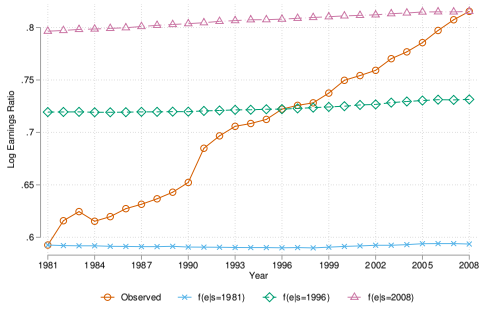

(c) 50/10: Men

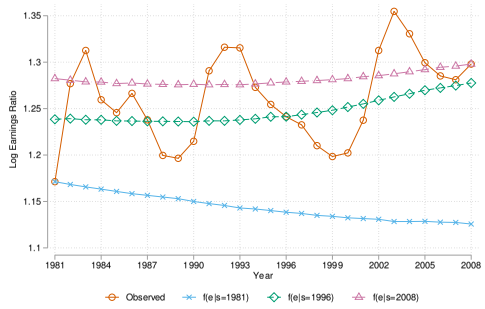

(d) 90/10: Women

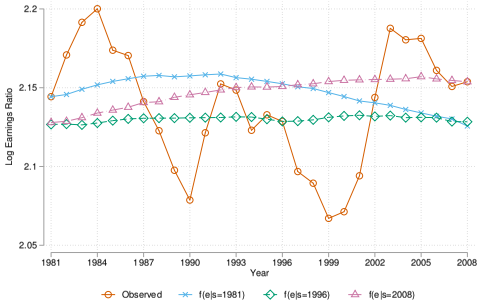

(e) 90/50: Women

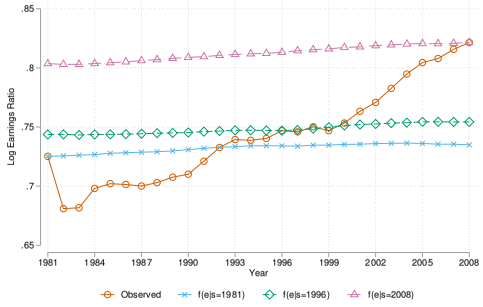

(f) 50/10: Women

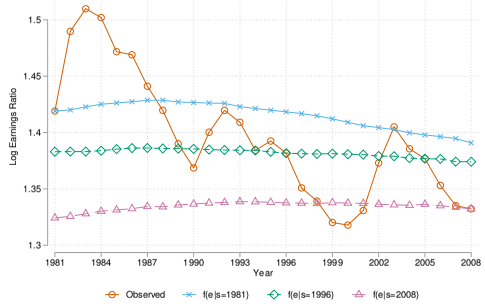

about 80% of the total increase in inequality is happening within age by education groups. Perhaps not surprisingly, given the fact that the distribution of skills has changed considerably, there is a somewhat larger role for compositional changes in explaining the observed change, though it is again small. If the 1981 skills distribution had prevailed for the entire period, then inequality would have fallen slightly. Both the 1996 and 2008 skills distribution contribute a small amount to rising inequality. However, these changes are dwarfed by the effect of changing prices over ths ttime period. This same basic pattern holds for both top end and bottom end inequality as well. Top end inequality shows a substantial and steady increase, with changes in the skills distribution playing a small role. Bottom end inequality also increases somewhat, with a slightly larger role for compositional changes.

Overall, residual earnings show an increase in inequality rather than polarization as both the 90/50 and 50/10 are increasing. However, the 90/50 increases much more, again demonstrating that the top is pulling away from everyone else. This is about the distribution of prices, not skills.

As is the norm, long-run inequality is lower than short-run inequality, however the magnitude of the increase is similar with a 26% increase in the 90/10. Again there is little role for compositional changes in explaining the overall increase in inequality. Top-end inequality increases 43%, also with no role for compositional changes. Unlike bottom-end inequality in annual earnings, bottom-end long-run inequality shows a steady increase, with a small cyclical decline in the 1990s. Overall, bottom-end inequalit increases 15%. Here we see a small role for compositional changes, as inequality grows with the 1996 skills distribution but falls with the 1981 skills distribution, highlighting the surge in educational attainment during this period. The fact that annual bottom end inequality is highly cyclical and shows little trend, while long-run inequality increases steadily highlights the rise in permanent inequality and provides suggestive evidence of an increase in the number of people with persistently low (though cyclical) earnings.

Figure 3: Long-Run Earnings Inequality for Men

(a) 90/10: Men

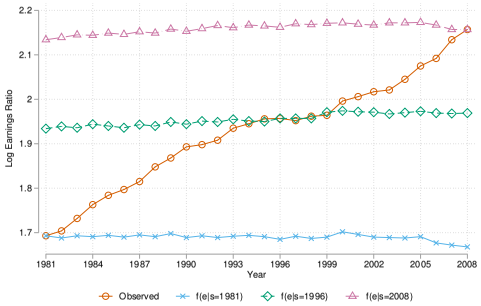

(b) 90/50: Men

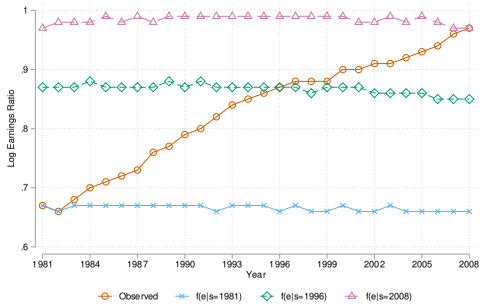

(c) 50/10: Men

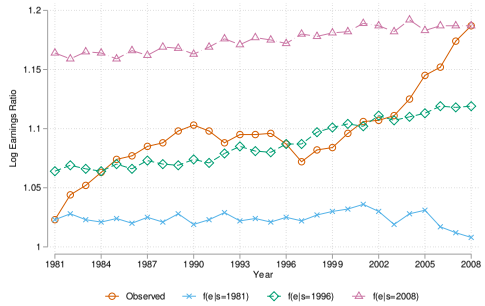

(d) 90/10: Women

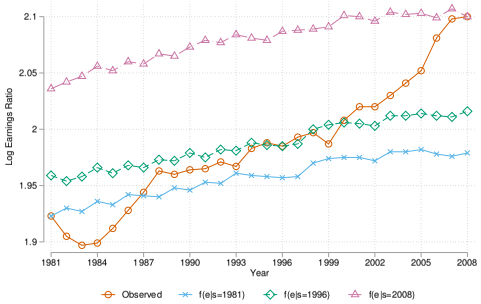

(e) 90/50: Women

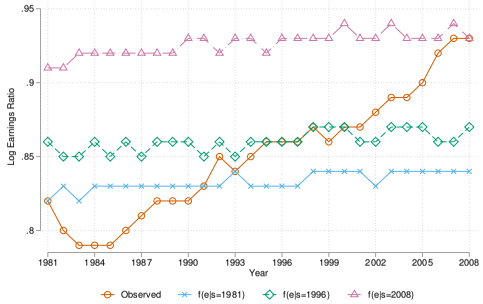

(f) 50/10: Women

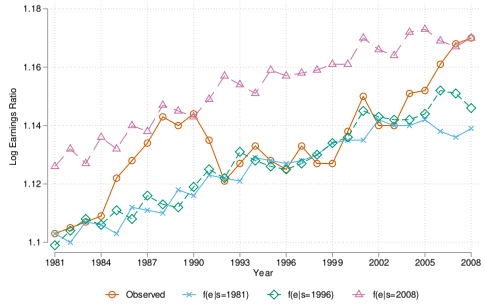

Similar to men, long-run earnings inequality for women increases overall, in the top-end, and in the bottom-end. While the increase in top-end inequality is explained almost entirely by changing prices, the increase in bottom-end inequality is about 50% due to changing skill distribution. That said, the overall increases in inequality remain considerably smaller for women than for men.

Long-run residual earnings looks basically the same, though the level of inequality is lower: steady increases in overall inequality, driven both by rising top-end and bottom-end inequality.

Overall residual inequality increases slowly until 1999, and the surges between 1999 and 2008. Prior to 1999, the increase in inequality is due entirely to changing characteristics, after which the surge is due almsot entirely to prices. After a small decrease in top-end inequality in the early 1980s, top-end residual inequality increases throughout the period, with the increase due almost entirely to prices. Bottom-end inequality falls to 1999, then increases quickly from 1999 to 2008. The total increase in bottom-end residual inequality is due to characteristics, as inequality would have been only slightly higher with 2008 characteristics in 1981, and would have been somewhat smaller in 1981 with 1996 characteristics. The change through time to the 2008 distribution of characteristics (f(e—s=2008)) shows the same total increase in inequality as is observed.

Figure 4: Long-Run Residual Earnings Inequality for Men

(a) 90/10: Men

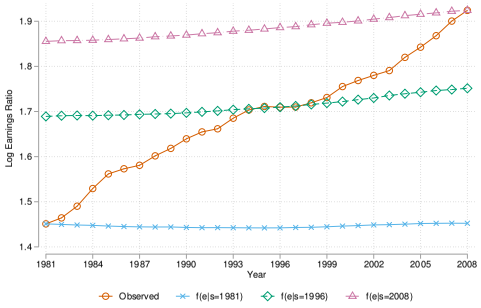

(b) 90/50: Men

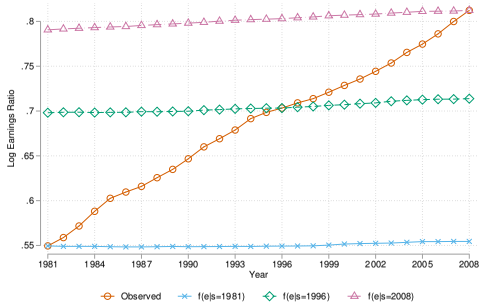

(c) 50/10: Men

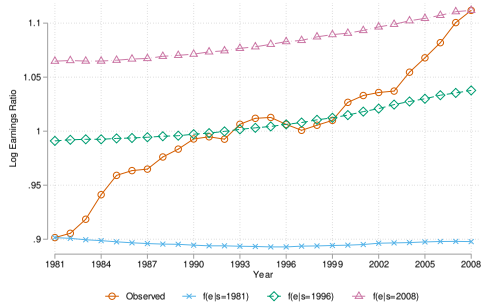

(d) 90/10: Women

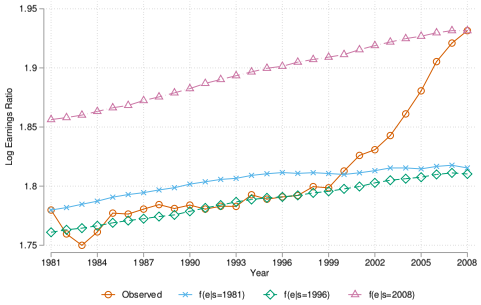

(e) 90/50: Women

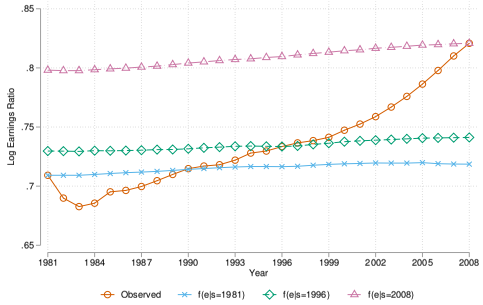

(f) 50/10: Women

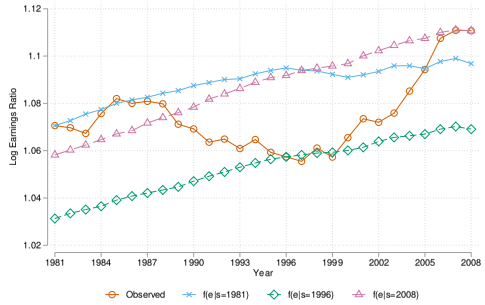

5 Mobility

5.1 Large Jumps

Figure 5: Large Upward Movement, Earnings

(a) Below 40 to Above 80: Men

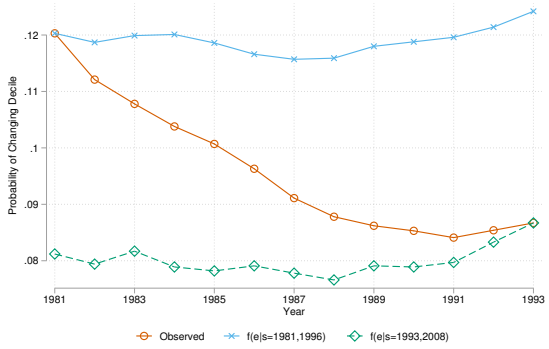

(b) Below 20 to Above 60: Men

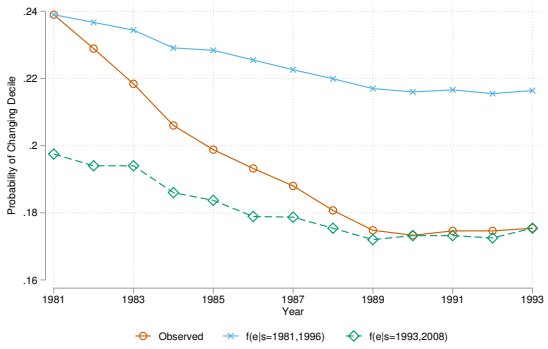

(c) Below 40 to Above 80: Women

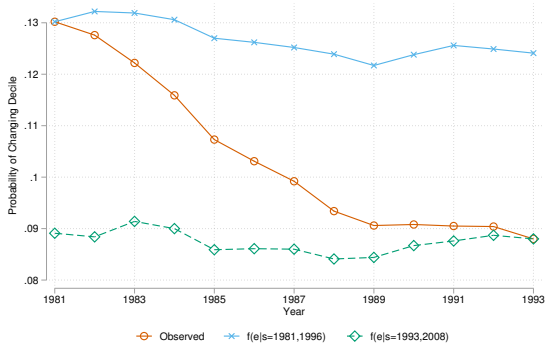

(d) Below 20 to Above 60: Women

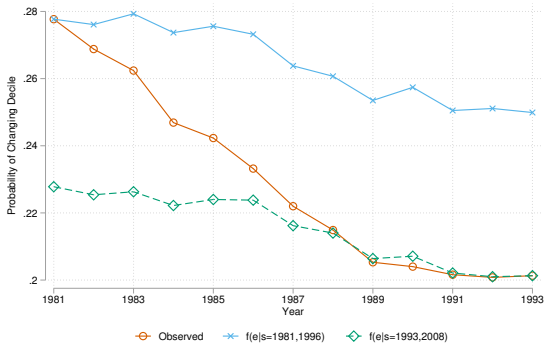

Overall the probability of moving from below 40 to about 80 fell about 27% from men, from 0.12 to 0.87. Between 1981 and 1989 this fall is due entirely to changing prices, as the composition adjusted trends are essentially flat. After 1989, if were not for increasing inequality, the changing composition would have resulted in an increase in upward mobility. This pattern is evident for both the 1981 and 2008 distributions. Observed mobility increases slightly between 1991 and 1993. However, holding prices fixed at 1981 levels, mobility begins to increase in 1987, suggesting that changing composition counteracts rising inequality. The same is true with the 2008 distribution: mobility shows a notable increase

due to compositional changes between 1989 and 1993. In general, mobility is higher in the 1981 distribution than in the 2008 distribution, showing the considerable effects of inequality on mobility. Using single-year weights, which truly reflect the effect of changes in inequality, mobility would have been 35% lower in 1981 if the earnings distribution had looked like it did in 2008. Between 1993 and 2008, mobility would have been 27% higher if the price distribution had looked like it did in 1981.

Upward mobility out of the bottom of the distribution declines steadly between 1981 and 1989, and is flat afterward. It shows no increase in later years, suggesting that the uptick in upward mobility seen above is due to a small increase in the probability of moving from between 20 and 40 to above 80. Here, composition and prices play a roughly equal role: overall observed mobility declines 27%, while the composition adjusted 1981 distribution declines about 11%. Mobility using the 2008 composition adjusted distribution declines steadily to 1989, but largely recovers by 1993, for an overall trend that is relatively flat. As with before, initial compsitional changes tended to result in declining mobility while later changes tended to increase mobility, regardless of which earnings distribution is adjusted.

Upward mobility out of the bottom of residual earnings looks much more like 40/80 mobility. Observed residual mobility declines 16% from the highest level (1981) to the lowest (1991), but regains some of the loss between 1991 and 1993, for an overall decline of 7.5%. How the observed decline is distributed between composition and prices is a little hard to say as it depends on how mobility is weighted. With paired year weights, the entirety of the decline in mobiltiy between 1981 and 1989 is due to composition, while only about 50% of the decline is due to composition with single year weights. Both weighting schemes agree that the uptick in mobility would have increased from 1989 to 1993 if the distribution of prices had not changed.

Updward mobility from below 40 to above 80 declined about 30% for women from a 13% chance to a 9% chance. Composition adjusted mobility is largely flat, demonstrating that

the decline in upward mobility is due to rising inequality in prices. There is a slight uptick in mobility after 1989 using the 2008 distribution of prices, again showing that compositional change tends to counteract the effect of rising inequality somewhat is recent years. This effect is dwarfed by the effect of inequality, however.

Within group mobility shows a much bigger role for compositional change. While the level of mobility is certainly lower with the 2008 earnings distribution than the 1981 distribution, the downward trend in observed mobility is almost identical to the trend in composition adjusted mobility. Again, we see that while observed mobility continues to decline between 1989 and 1993, composition adjusted mobility increases.

Figure 6: Large Upward Movement, Residual Earnings

(a) Below 40 to Above 80: Men

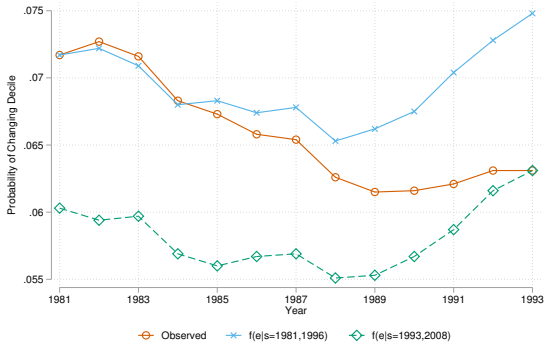

(b) Below 20 to Above 60: Men

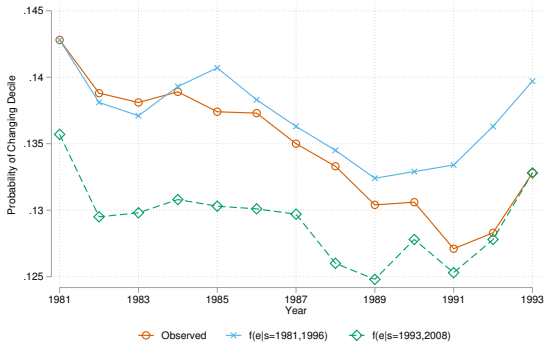

(c) Below 40 to Above 80: Women

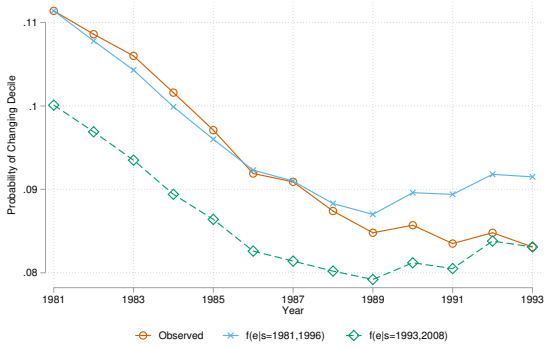

(d) Below 20 to Above 60: Women

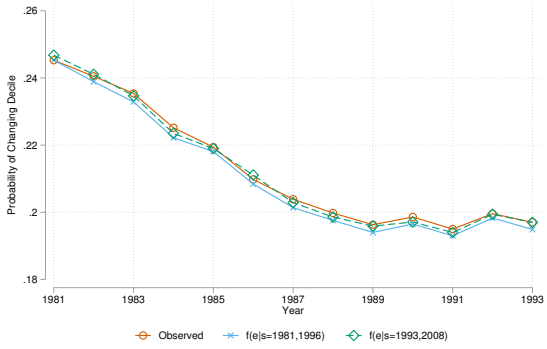

Patterns in residual mobility are similar, though more pronounced. Between 1981 and 1987 declining mobility is due almost entirely to compositional change, as demonstrated by

the fact that the trend in observed mobility is parallel to trends in composition adjusted mobilty regardless of which earnings distribution is adjusted. However, while there is only a small uptick in observed mobility after 1989, there is a substantial uptick in composition adjusted mobility, again regardles of which residual earnings distribution is used. Throughout this period, college educated workers have the highest level of mobility, but their mobility declines the most. Mobility declines for other education groups as well, but not as much. What these results suggest is that, as college graduates and older workers became an increasingly large shar of the labor force, this initiall caused mobility to fall because of steeply declining mobility for this group, but then resulted in mobility rising because of their higher overall level of mobility, holding inequality constant.

Upward mobility out of the bottom of residual earnings looks much more like 40/80 mobility. Observed residual mobility declines 16% from the highest level (1981) to the lowest (1991), but regains some of the loss between 1991 and 1993, for an overall decline of 7.5%. How the observed decline is distributed between composition and prices is a little hard to say as it depends on how mobility is weighted. With paired year weights, the entirety of the decline in mobiltiy between 1981 and 1989 is due to composition, while only about 50% of the decline is due to composition with single year weights. Both weighting schemes agree that the uptick in mobility would have increased from 1989 to 1993 if the distribution of prices had not changed.

Within group mobility shows a much bigger role for compositional change. While the level of mobility is certainly lower with the 2008 earnings distribution than the 1981 distribution, the downward trend in observed mobility is almost identical to the trend in composition adjusted mobility. Again, we see that while observed mobility continues to decline between 1989 and 1993, composition adjusted mobility increases.

Because we rank in the full distribution, trends in upward and downward mobility need not offset each other. In other words, because mobility is estimated for a subset of individuals

in the full earnings distribution, mobility is not zero sum. Thus, we need to look at upward and downward mobility separately.

Figure 7: Large Downward Movement, Earnings

(a) Above 80 to Below 40: Men

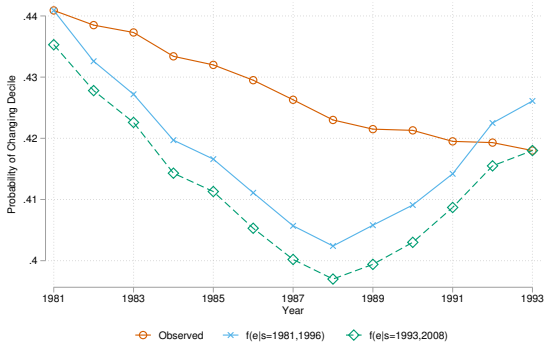

(b) Above 60 to Below 20: Men

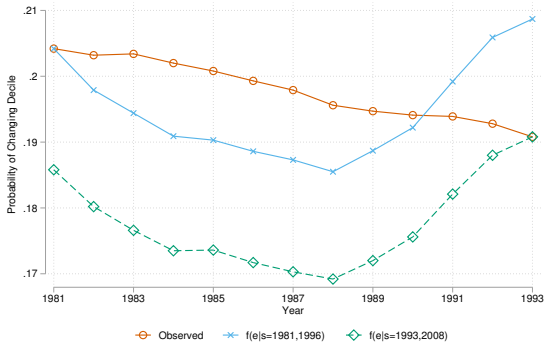

(c) Above 80 to Below 40: Women

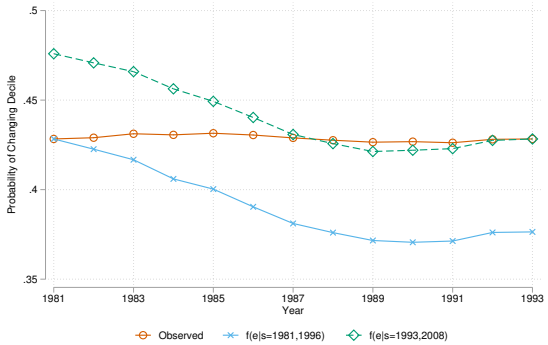

(d) Above 60 to Below 20: Women

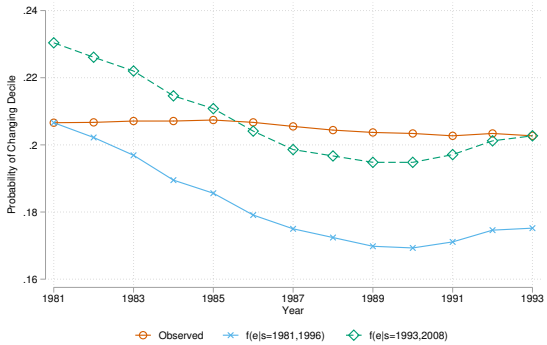

I'm not sure what to make of this one. The rate of moving from about 80 to below 40 is substantially higher than moving from below 40 to above 80, perhaps because a man who starts his career about the 80th percentile is more likely to have outlier earnings or a flatter earnings trajectory? However, downward mobility declines about 2 percentage points, which on a base of 44% is not substantial. Downward mobility would have increased slightly after 1989 had prices not changed, and downward mobility would have been higher with 1981 prices than 2008 prices. However, downward mobiltiy is higher with either 1981 or 2008 prices than what is actually observed. Combine this with upward mobility, and what rising inequality does is reduce churning as a whole, but by reducing upward movements more than

downward movements.

Movement from about 60 to below 20 is qualitatively similar to movement from above 80 to below 40. Also, overall and residual mobility show a similar story. The only real difference is that the declines in downward mobility are much smaller. This makes sense as there are fewer people below 20 so it should be harder to move into the group. Similarly, the increases in mobility that would have occurred had prices not changed are bigger, showing a larger relative role for compositional changes. This is perhaps surprising as the group that is growing (college educated) should be less likely to move down the earnings distribution.

Figure 8: Large Downward Movement, Residual Earnings

(a) Above 80 to Below 40: Men

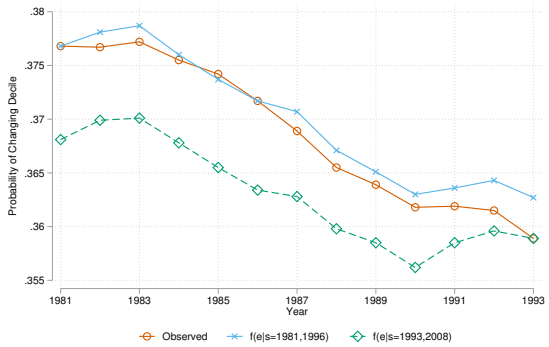

(b) Above 60 to Below 20: Men

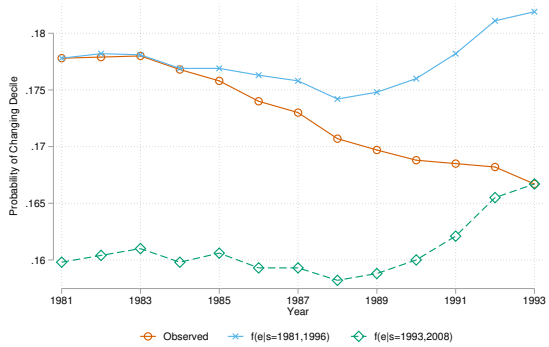

(c) Above 80 to Below 40: Women

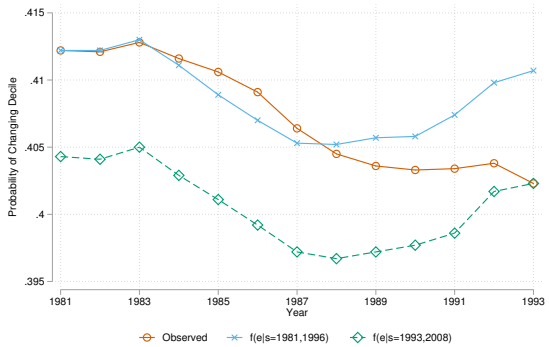

(d) Above 60 to Below 20: Women

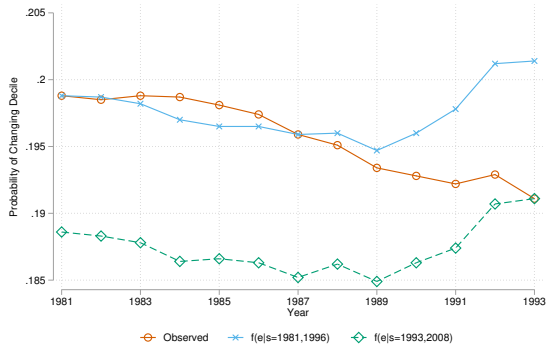

Conclusions on downward residual mobility from above 80 to below 40 are somewhat different depending on weighting method. With single-year weights, the decline in mobility is all about compositional change, though the level of mobility throughout the period would

have been higher under 1981 prices. With paired-year weights, the decline in mobility is about 80% composition and 20% prices.

5.2 Churning

One simple way to summarize mobility across a transition matrix is to look at the probability of making a given size relative jump, regardless of starting decile. Here we look at the probability of moving up two deciles or more. Observed mobility declines 30%. If the 1981 distribution of prices were held fixed, mobility would have been flat. If the 2008 price distribution is held fixed, mobility would have increased about 30%. In 1981, mobility would have been about 30% lower if the disrtribution of 2008 prices held. By 1993, the difference is only about 6%. What this shows is that compositional change is an important counterbalance to the effect of rising inequality: rising inequality depresses mobility, and it would have depresed mobility even more had the skills distribution not changed in the way it did.

The declines in downward movement of 2 or more deciles are substantial in absolute magniuted (10 percentage points) but not as large in relative terms (18%). Overall and within group show almost identical levels of downward mobility. Regardless of weighting method of how earnings are measured, the decline in mobility is almost entirely about compositional change. Downward mobility is rising across the earnings distribution because of a rising share of college educated and older workers. We take this to mean that better educated new entrants to the labor market are entering above existing workers in the earnings distribution, pushing existing workers down the earnings distribution over their working life.

Once again, residual mobility is lower than overall mobility. Observed residual mobility declines about 13%, while composition adjusted mobility is flat, implying that the entirety of the decline in residual mobility is due to rising inequality.

The rates of large upward movements – increases of 4 deciles or more – are lower, but

Figure 9: Moved 2 or More Deciles, Earnings

(a) Moved Up: Men

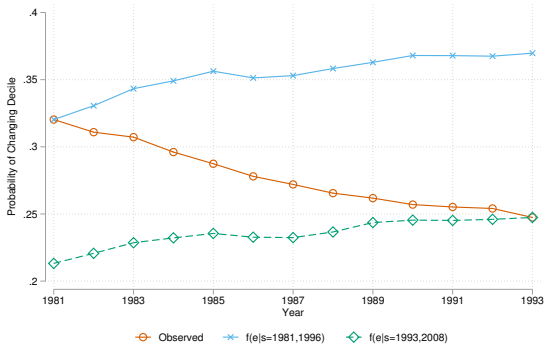

(b) Moved Down: Men

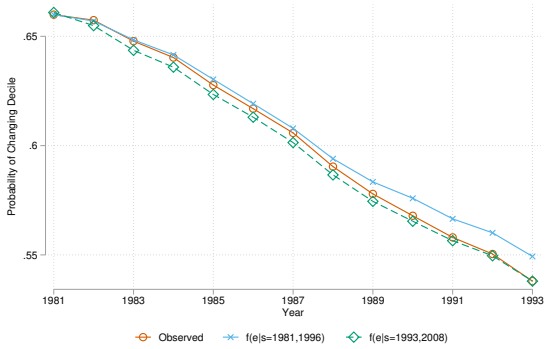

(c) Moved Up: Women

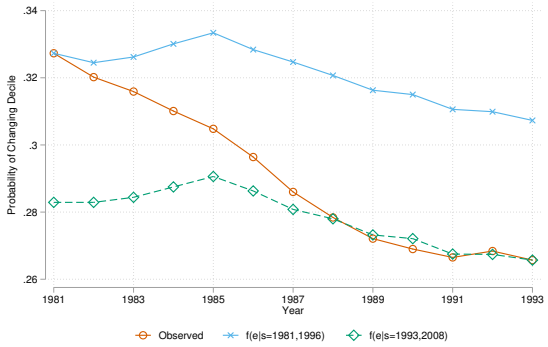

(d) Moved Down: Women

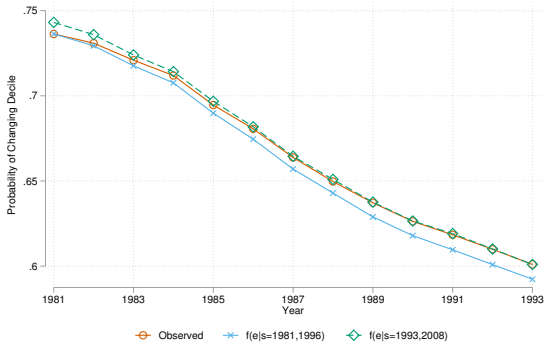

the pattern of change is similar. Overall observed mobility declines 42%. Mobility in the composition adjusted 1981 distribution is relatively flat, while mobility is the composition adjusted 2008 distribution increases 20%, again highlighting the fact that compositional change tends to increase mobility while inequality tends to decrease mobility.

A quantitative and quantitatively similar story holds for larger downward movements, though here we see a somewhat larger role for changes in inequality: observed mobility declines about 17% in both earnings distributions with both weighting schemes, while composition adjusted mobility declines about 12% across weighting methods and earnings distributions.

A similar pattern holds with residual mobility: observed residual mobility declines, while composition adjusted mobility increases slightly, showing the effect of rising within group

Figure 10: Moved 2 or More Deciles, Residual Earnings

(a) Moved Up: Men

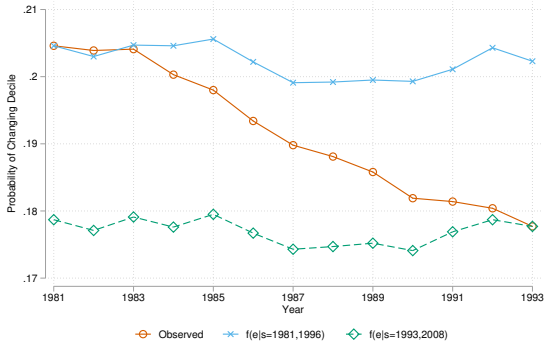

(b) Moved Down: Men

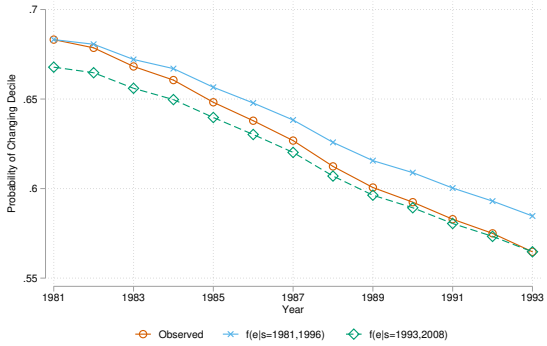

(c) Moved Up: Women

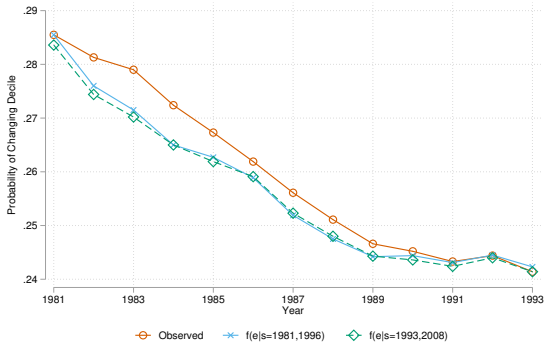

(d) Moved Down: Women

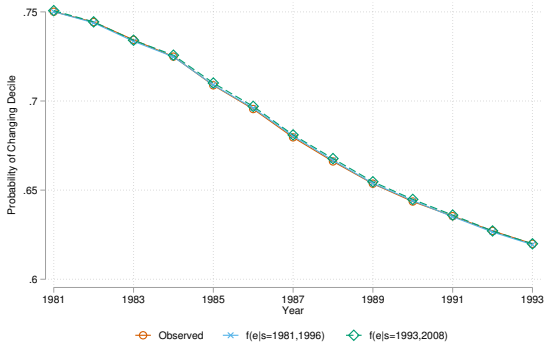

inequality on mobility.

6 Labor Supply

The problem we face when comparing our decomposition to the inequality literature is that we're not actually holding prices fixed, we're holding prices times work hours fixed. So, switching from 1981 to 2008 results in a change in both the distribution of wages and the distribution of hours. If work hours and/or labor force participation change differently at different points of the earnings distribution, than part of what we're implicitly attributing to changes in inequality in prices is actually changes in inequality in work hours.

These changes could occur at both "extensive margin" or the "intensive margin." By

Figure 11: Moved 4 or More Deciles, Earnings

(a) Moved Up: Men

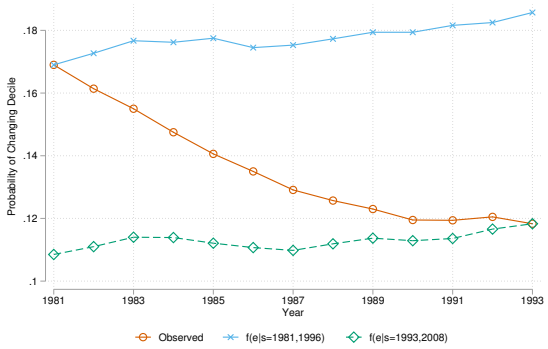

(b) Moved Down: Men

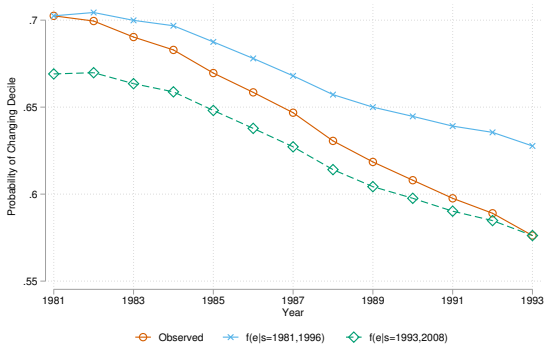

(c) Moved Up: Women

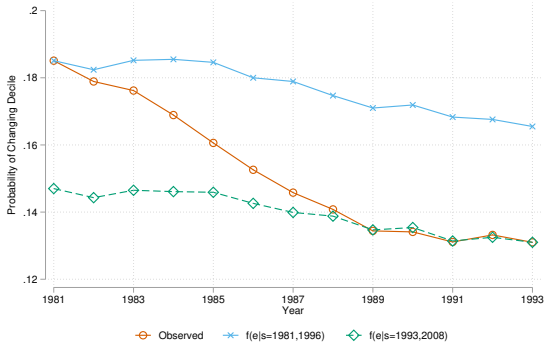

(d) Moved Down: Women

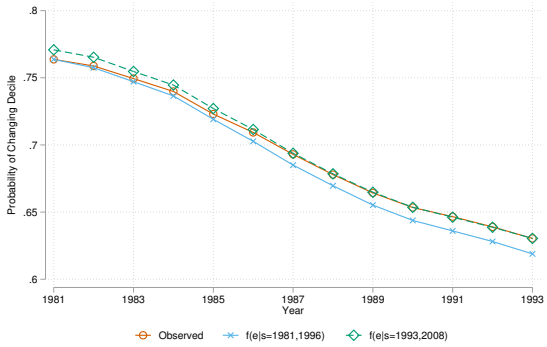

extensive, we mean differential changes in number of days employed, measured by weeks worked and covered quarters. By intensive margin, we mean number of hours worked in a year.

What we see in Figure **??** is that hours worked per week have declines somewhat. Figure **??** shows average work hours of individuals in our mobility sample taken from their SIPP survey responses. Clearly something is wrong with 1984, and there are a meaningful number of people who have positive earnings in the GSF (otherwise they wouldn't be in our sample) but have zero annual work hours. Without having access to both usual weekly work hours last month and weeks worked last month for each month of the SIPP survey, it's impossible to know what's going on here. If we restrict to positive work hours (Figure **??**) and ignore 1984, we get something more reasonable, but this excludes people in the mobility sample

Figure 12: Moved 2 or More Deciles, Residual Earnings

(a) Moved Up: Men

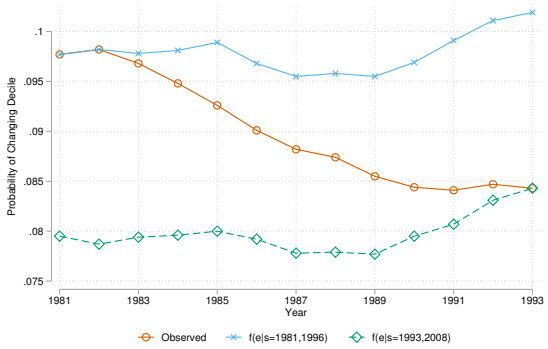

(b) Moved Down: Men

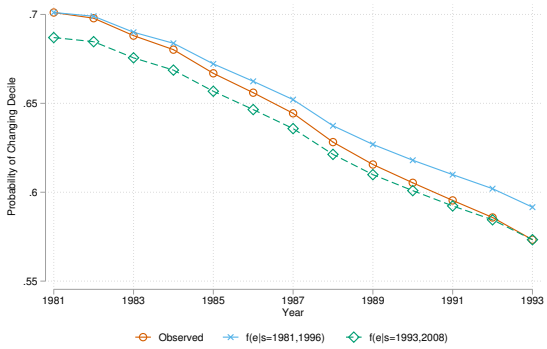

(c) Moved Up: Women

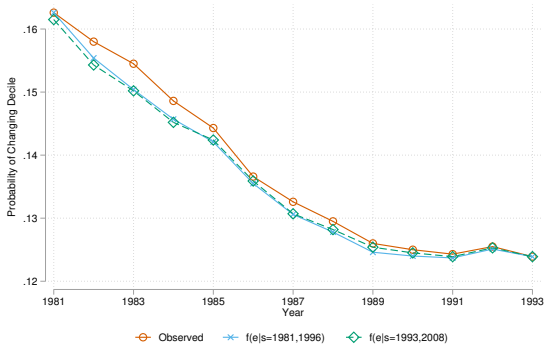

(d) Moved Down: Women

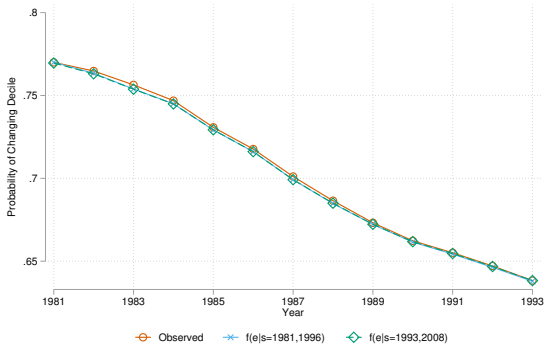

who should have positive hours. We do see a slight decline into 2008.

On the other hand, covered quarters (Figures 13b and 13c) have increased steadily for both men and women. Putting the two together suggests that there are more people working full-year at slightly less than full-time hours per week.

Figure **??** shows that the trend in annual work hours through time is similar for men and women both in the 5th decile (D5) and the 1st decile (D1), and thus similar within genders as well. This shows that, although part of the reason an individual might have low earnings is because of work hours, the trends in declining mobility through time are not driven by differential trends in annual work hours. This is also (mostly) true when measuring annual labor force participation with covered quarters. Men generally have more covered quarters than women, but the trends through time are not quite the same. Covered quarters for the

Figure 13: Labor Supply Overall and by Earnings Decile

(a) Weeks Worked

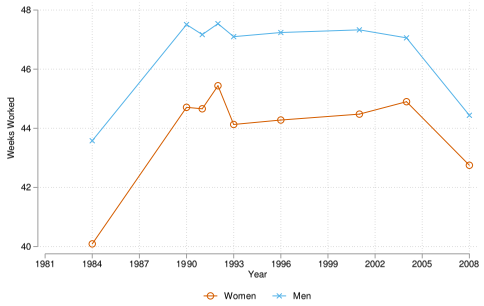

(b) Covered Quarters

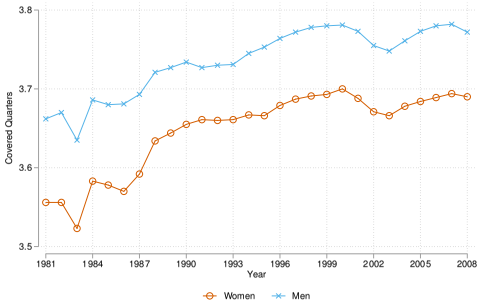

(c) Share 4 Covered Quarters

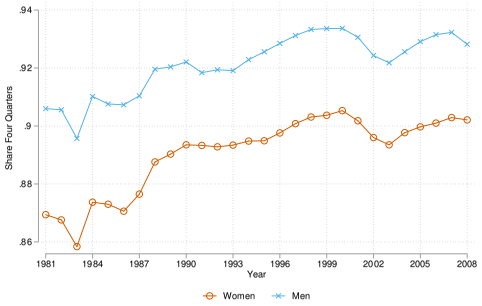

(d) Weeks Worked by Decile

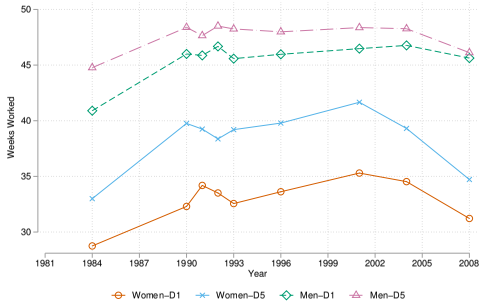

(e) Covered Quarters by Decile

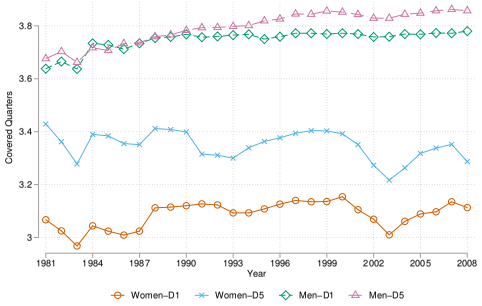

(f) Share 4 Covered Quarters by Decile

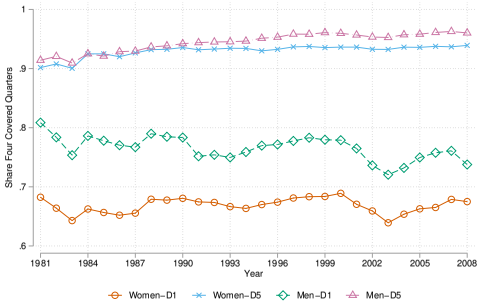

men in the 5th decile rise steadily from 3.68 to 3.82, with about 95% of men covered all four quarters. Covered quarters for 1st decile men are essentially flat from 1990 to 2008, with XX% of men covered for 4 quarters.

Covered quarters for women are lower, more cyclical, and follow slightly different trends through time for the 1st and 5th deciles, respectively. Covered quarters for the 5 decile seem to decline slightly, though it is impossible to rule out the possibility that dips in the 2000s are cyclical rather than part of a secular trend. Covered quarters for the 1st appear to rise slightly. None-the-less, while it is certainly the case that labor supply differences play a role in the level of starting earnings, there is little evidence to suggest that changes mobility through time are the result of changing patterns in labor supply as opposed to changes in hourly wages.

7 Conclusion

Our paper uses administrative-linked survey data to document trends in earnings inequality since the 1980s. We show tends in both long-run and short-run inequality and link these trends through an examination of lifetime earnings mobility. Because we use a unique source of administrative earnings records that also contain data on educational attainment, we are able to follow the wage inequality literature and estimate trends in earnings inequality both overall and within education and gender subgroups and to use a kernel reweighting approach to consider the relative importance of changes in the gender by age by education composition of the labor force to the trend in earnings inequality. Our results show a staggeringly large increase in upper tail earnings inequality in long-run earnings inequality for college-educated workers of over 40%, with an increase of over 70% for men – a trend that is driven both by large increases in upper-tail earnings inequality in the short run coupled with declines in lifetime earnings mobility. Finally, we provide evidence that this trend in rapidly increasing

earnings inequality is not being driven by changes in labor market composition.

Our results show a rapid increase in within-group earnings inequality for college educated workers, despite this, trends in overall and within gender earnings inequality are not not explained by changes over time in the composition of the labor market. However, there are two avenues that we have not yet explored, but that our data support. The first is to build on Lemieux (2006b) work on wage inequality to consider the role of earnings of workers with post-secondary degrees to the overall trends in within-group inequality. Because of changes in labor supply across the distribution of educational attainment over time, the trends in earnings inequality may not match wage inequality trends. Second, because our results on the role of labor market composition on trends in earnings inequality contrast with the wage inequality literature, we plan to further test our reweighting results to alternative methodological approaches, particularly to population subgroups are weighted together or separately and to approaches which include the consideration of post-secondary schooling.

References

Abowd, John M. and Martha H. Stinson. 2013. "Estimating Measurement Error in Annual Job Earnings: A Comparison of Survey and Administrative Data." *Review of Economics and Statistics* 95 (5):1451–1467.

Acs, Gregory and Seth Zimmerman. 2008. *US Intragenerational Economic Mobility from 1984-2004: Trends and Implications*. Washington, D.C.: The Urban Institute.

Auten, Gerald and Geoffrey Gee. 2009. "Income Mobility in the United States: New Evidence from Income Tax Data." *National Tax Journal* 62 (2):301–328.

Auten, Gerald, Geoffrey Gee, and Nicholas Turner. 2013. "Income Inequality, Mobility, and Turnover at the Top in the US, 1987-2010." *American Economic Review: Papers and Proceedings* 103 (3):168–172.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics* 90 (2):300–323.

Baker, Michael and Gary Solon. 2003. "Earnings Dynamics and Inequality Among Canadian Men, 1976-1992." *Journal of Labor Economics* 21 (2):289–321.

Berman, Eli, John Bound, and Zvi Griliches. 1994. "Changes in the Demand for Skilled Labor Within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures." *Quarterly Journal of Economics* 109 (2):367–397.

Bound, John and George Johnson. 1992. "Changes in the Structure of Wages in the 1980s: an Evaluation of Alternative Explanations." *American Economic Review* 82 (3):371–392.

Bradbury, Katharine and Jane Katz. 2009. "Trends in U.S. Family Income Mobility, 1967-2004." *Federal Reserve Bank of Boston Working Paper* 09-7:1–39.

Carr, Michael D. and Emily Wiemers. 2016a. "The Decline in Lifetime Earnings Mobility in the U.S.: Evidence from Survey Linked Administrative Data." *Washington Center for Equitable Growth Working Paper* .

Carr, Michael D. and Emily E. Wiemers. 2016b. "Recent Trends in the Variability of Men's Earnings: Evidence from Administrative and Survey Data." *University of Massachusetts Boston Working Paper* .

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014. "Where is the Land of Opportunity: The Geography of Intergenerational Mobility in the United States." *Quarterly Journal of Economics* 129 (4):1553–1623.

Debacker, Jason, Bradley Heim, Vasia Panousi, Shanthi Ramnath, and Ivan Vidangos. 2013. "Rising Inequality: Transitory or Persistent? New Evidence from a Panel of U.S. Tax Returns." *Brookings Papers on Economic Activity* Spring.

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Disbrtribution of Wages, 1973-1992: A Semiparametric Approach." *Econometrica* 64 (5):1001–1044.

Gottschalk, Peter and Sheldon Danziger. 2005. "Inequality of wage rates, earnings and family income in the United States, 1975–2002." *Review of Income and Wealth* 51 (2):231–254.

Gottschalk, Peter and Robert Moffitt. 2009. "The Rising Instability of U.S. Earnings." *Journal of Economic Perspectives* 23 (4):3–24.

Gottschalk, Peter, Robert Moffitt, Lawrence F Katz, and William T Dickens. 1994. "The Growth of Earnings Instability in the US Labor Market." *Brookings Papers on Economic Activity* 1994 (2):217–272.

Haider, Steven J. 2001. "Earnings instability and earnings inequality of males in the United States: 1967–1991." *Journal of Labor Economics* 19 (4):799–836.

Katz, Lawrence F. and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107 (1):35–78.

Kopczuk, Wojciech, Emmanuel Saez, and Jae Song. 2010. "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data Since 1937." *Quarterly Journal of Economics* 125 (1):91–128.

Lemieux, Thomas. 2006a. "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" *American Economic Review* 96 (3):461–498.

———. 2006b. "Post-Secondary Education and Increasing Wage Inequality." *American Economic Review* 96 (2):195–199.

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

———. 2012. "Trends in the Transitory Variance of Male Earnings Methods and Evidence." *Journal of Human Resources* 47 (1):204–236.

Piketty, Thomas and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913-1998 (Updated in 2016)." *Quarterly Journal of Economics* 118 (1):1–39.