Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?

Thomas Lemieux
University of British Columbia and NBER

Revised, May 2005

Abstract

Using data from the May and Outgoing Rotation Group (ORG) supplements of the CPS, this paper shows that a large fraction of the growth in residual wage inequality between 1973 and 2003 is due to spurious composition effects. These composition effects are linked to the secular increase in the level of experience and education of the workforce, two factors associated with higher within-group wage dispersion. I also show that both the level and growth in residual wage inequality are overstated in March CPS data that have been used in most previous studies. Measured wages are noisier in the March than in the May/ORG CPS because the March CPS does not measure directly the hourly wages of workers paid by the hour. The extent of measurement error in CPS wages also increases over time.

Once these factors are corrected for, I find that residual wage inequality only accounts for a small share of the overall growth in wage inequality. Furthermore, *all* of the growth in residual wage inequality occurs during the 1980s. This closely mirrors the pattern of change in "between-group" wage differentials like the college-high school wage premium. Overall, the magnitude and timing of the growth in residual wage inequality provides little evidence of a pervasive increase in the demand for skill due, for instance, to skill-biased technological change.

1. Introduction

The growth in wage inequality over the last three decades is one of the most extensively researched topic in labor economics. An important part of the change in wage inequality has been linked to the growth in the college-high school wage premium since the late 1970s (Bound and Johnson, 1992, Katz and Murphy, 1992). However, explanations for the growth in wage inequality linked to standard human capital variables like experience and education are limited by the fact that these variables only explain about a third of the variance of wages.[1] Perhaps not surprisingly, residual or within-group wage inequality – i.e. wage dispersion among workers with the same education and experience–, is generally believed to account for most of the growth in overall wage inequality (Juhn, Murphy and Pierce, 1993, JMP hereinafter). According to JMP, residual wage inequality increased steadily throughout the 1970s and 1980s. More recently, Acemoglu (2002) and Katz and Autor (1999) argue that residual wage inequality kept increasing steadily in the 1990s.

Understanding the sources of growth in residual wage inequality is problematic as there are many reasons why workers with the same level of experience and education may report different wages. Perhaps these workers have different levels of valuable but unobserved skills linked to school quality, intrinsic ability, effort, etc. Or perhaps the reported wage differences are simply due to measurement error. In this simple setting, there are already three possible reasons why residual wage inequality may be increasing over time. First, the "price" or return to unobserved skills may be increasing because of an increase in the demand for skill. JMP argue that this is in fact the main factor behind the growth in residual wage inequality during the 1970s and 1980s. Second, the dispersion in unobserved skills may be growing over time. For example, if unobserved skills are more dispersed among older and more educated workers, dispersion in unobserved skills could increase because of composition effects linked to the aging and increasing educational achievement of the workforce. Third, the extent of measurement error may be increasing over time.

---

[1] For example, the R-square of the regression models estimated later in the paper range from 0.19 to 0.38.

In this paper, I show that all three factors played an important role in the increase in residual wage inequality over the last three decades. In other words, the growth in residual inequality cannot simply be equated to a rise in the demand for skill. In fact, I show that increases in the return to unobserved skills account for no more than 25 percent of the overall increase in wage inequality over the last three decades. Moreover, I show that the all of the increase in the return to unobserved skills is concentrated in the 1980s.

These findings have important implications for understanding the sources of change in wage inequality. While JMP did not elaborate on the underlying source of growth in the demand for skill, most subsequent studies have argued that skill-biased technological change (SBTC) was the main factor responsible for the steady growth in the demand for skill. In particular, the computer and information technology revolution has emerged as the leading hypothesis for explaining the growth in the relative demand for skills since the early 1970s (Berman, Bound and Griliches, 1994, Autor, Katz and Krueger, 1998). As pointed out by Card and DiNardo (2002), however, SBTC should have resulted in an increase in the demand for skill in both the 1980s and 1990s since computer technologies kept advancing rapidly in the 1990s. The fact that the return to unobserved skill only grew in the 1980s is a major challenge for the SBTC explanation.

These findings are also at odds with most of the previous literature that generally suggests that residual wage inequality increased steadily over time and accounts for most of the increase in overall wage inequality. I show that the difference between my findings and those of earlier studies is due to a combination of several factors. First, I use data on hourly wages from the May and Outgoing Rotation Group (ORG) supplements of the Current Population Survey (CPS), while earlier studies have typically used the March Supplement of the CPS. In Section 6, I explain why the May/ORG CPS is better suited than the March CPS for studying the evolution of residual, or within-group, wage dispersion. The main problem with the March CPS is that it poorly measures the wages of workers paid by the hour (the majority of the workforce). The fraction of workers paid by the hour has grown substantially over time which results in spurious growth in residual wage inequality in the March CPS. I also show that the variance of measurement error in wages has increased over time. This has also resulted in spurious growth in residual wage inequality, especially in the March CPS.

2

A second important difference is that, unlike most other studies, I control for composition effects. [2] Wage dispersion among narrowly defined groups of workers is substantially larger for older and more educated workers than for younger and less-educated works. As result, I show that a large fraction part of the increase in residual wage inequality is a spurious consequence of the fact that the workforce has grown older and more educated since the early 1980s. A final difference with earlier studies is that much more data are now available for studying secular changes in residual wage inequality. For example, I use wage data for up to 2003 while the last year of wage data available to JMP was 1989. Composition effects play a much bigger role in changes in residual inequality in the 1990s and early 2000s than in the 1970s and 1980s. This may explain why composition effects remained relatively unnoticed in the earlier literature.

The paper is organized as follows. In Section 2, I discuss in more detail the link between residual wage inequality, unobserved skill prices, and composition effects. I explain how to account for composition effects in Section 3. Section 4 presents the May/ORG CPS data and shows basic trends in within-group wage inequality for twenty experience-education groups. The main results on the evolution of residual wage inequality once composition effects are adjusted for are presented in Section 5. Section 6 shows why both the level and growth in residual wage inequality are overstated in the March CPS. Section 7 concludes by suggesting possible explanations for the trends in residual wage inequality since the early 1970s.

2. Residual Wage Inequality: Skill Prices, Composition Effects, and Measurement Error

a. Determinants of Residual Wage Inequality

As mentioned in the Introduction, changes in residual wage inequality can only be interpreted as evidence of changing skill prices when both the distribution of unobserved skills and the variance of measurement error are constant over time. To see this, first consider a standard Mincer-type wage equation:

---

[2] One earlier study that controls for compositions effects is DiNardo, Fortin, and Lemieux (1996). They show that a third of the growth in residual inequality between 1979 and 1988 is due to composition effects.

(1) $w_{it} = x_{it}b_t + \varepsilon_{it}$,

where w_{it} is the natural logarithm of the hourly wage rate of individual i at time t; x_{it} is a
vector of observed skills (education and labor market experience); b_t is the return (or
price) to observed skills; ε_{it} is the standard regression residual. JMP assume that the
residual is the product of some unobserved skills, e_{it}, with the return to unobserved skills,
p_t. Allowing for a measurement error v_{it} yields an error component model for the residual
similar to the one considered by Chay and Lee (2000):

(2) $\varepsilon_{it} = p_t\, e_{it} + v_{it}$.

What I call "residual wage inequality" is simply the measured inequality in the residual,
ε_{it}. The main inequality measure used in the paper is the variance because, unlike other
popular measures like the difference between the 90[th] and the 10[th] percentile of log wages
(the "90-10 gap"), it is a decomposable measure of inequality.[3] Using equation (2), the
residual variance can be written as:

(3) $\mathrm{Var}(\varepsilon_{it}) = p_t^2\, \mathrm{Var}(e_{it}) + \mathrm{Var}(v_{it})$.

Equation (3) shows that changes in the residual variance can only be interpreted as
evidence of changing skills prices, p_t, when both the variance of unobserved skills,
$\mathrm{Var}(e_{it})$, and the variance of measurement error, $\mathrm{Var}(v_{it})$, remain constant over time.
However, most of the existing literature simply interprets growing residual wage
inequality as evidence of rising unobserved skill prices without controlling for changes in
the dispersion unobserved skills or measurement error.

---

[3] The total variance of wages can be written as $\mathrm{Var}(w_{it}) = \mathrm{Var}(x_{it}b_t) + \mathrm{Var}(\varepsilon_{it})$ and the residual variance is
the fraction $(1-R^2)$ of the overall variance of wages. By contrast, the 90-10 gap in w_{it} is not generally equal
to the sum of the 90-10 gap in $x_{it}b_t$ and in ε_{it}. As a result, arbitrary choices have to be made when trying to
quantify the contribution of residual inequality to overall inequality, which complicates the economic
interpretation of the results.

In particular, JMP use a residual imputation procedure to compute the contribution of changes in unobserved skill prices to the growth in wage inequality. Consider computing, for example, the contribution of changes in unobserved skill prices to the growth in wage inequality between period s and t. JMP's procedure consists of replacing each period t residual by a period s residual at the same position in the residual wage distribution.[4] For instance, if the residual ε_{it} turns out to be at the 92nd centile of the wage distribution in period t, it simply gets replaced by the 92nd centile of the residual wage distribution in period s. In the case of the variance, replacing the period t residuals by the period s residuals amounts to simply replacing the period t residual variance by the period s residual variance.[5] JMP's procedure thus imposes, *by assumption*, that the growth in the residual variance is solely due to changes in skill prices.

b. Changes in Observed and Unobserved Skill Prices: Is There Really a Puzzle?
When combined with the steady growth in residual wage inequality in the March CPS in the 1970s, 1980s, and 1990s (JMP, Acemoglu, 2002, Katz and Autor, 1999), the assumption that the residual variance and skill prices are the two sides of the same coin means that skill prices have been steadily rising over the last 30 years. While this is generally interpreted as support for SBTC, the difference in the timing of changes in residual wage inequality and other wage differentials like the college-high school wage premium was initially viewed as a puzzle in the literature (Levy and Murnane, 1992, Mincer, 1997). After all, if these various dimensions of wage dispersion were all linked to similar factors like technological change, they should more or less vary in a similar way over time.

JMP attempt to reconcile this initial puzzle using a two-factor model for observed (education and experience) and unobserved skills. Like Bound and Johnson (1992) and Katz and Murphy (1992), JMP view the increase in the relative demand for skills as the driving force behind the expansion of the college-high school wage premium during the

---

[4] The procedure described here is based on how JMP explain in words how their procedure works. The equations in their paper describe the conditional distribution of the residuals for given values of the regressors, suggesting that the procedure could, in principle, control for composition effects. See Lemieux (2002) for more discussion of this issue.
[5] One main advantage of JMP procedure is that, like the re-weighting procedure presented below, it can be used to decompose measures of inequality other than the variance.

1980s. They argue that during the 1970s, however, the growth in the relative demand for skill was offset by an even stronger growth in the relative supply of college-educated labor. By contrast, residual wage inequality grew in the 1970s because the relative supply of unobserved skills did not increase during this period (or other periods). The role of growing relative demand on residual wage inequality was not "masked" by a large increase supply in the 1970s as in the case of the college-high school premium,.

While JMP's synthesis of the causes of growing wage inequality had a major impact in labor economics and other fields, there are a growing number of problems with the story.[6] First, it only works for a very specific production function, namely a CES function in education (or experience) and unobserved skills (Acemoglu, 2002).[7] By contrast, when unobserved skills (e.g. school quality or cognitive skills) are close substitutes for education, increases in the supply of education should depress both the college-high school premium and the return to unobserved skills.[8]

Second, both Card and DiNardo (2002) and Beaudry and Green (2004) point out that the pattern of change in wage inequality in the 1990s is hard to reconcile with a traditional supply and demand explanation. In particular, the college-high school premium increased much less in the 1990s than in the 1980s despite the fact that relative supply kept increasing at the same rate. The "supply" explanation for why residual wage inequality grew in the 1970s while the college-high school wage premium did not grow does not work for the 1990s.

Third, if the increase in the price of unobserved skills was the most important source of growing wage inequality, we should have seen a large increase in the return to various measures of "ability" and in the male-female, or black-white, wage gap (to the extent that part of these gaps are due to differences in unobserved skills). The fact that

---

[6] Beyond labor economics, JMP's interpretation of growing residual inequality as an increase in the skill premium has laid the foundations for a large and influential literature on economic growth, technical change, and inequality (see Acemoglu, 2002 and Aghion, 2002 for recent surveys of this literature).
[7] Acemoglu (2002) illustrates this point using a "two-by-two" factor model. The first factor is education (college and high school) and the second factor is unobserved skills (low and high). The CES assumption implies that the substitutability between college workers with low and high (unobserved) skills is the same as the substitutability between high skill college workers and low skill high school workers, which is not a very appealing property of the production function.
[8] Years of schooling and school quality are perfect substitutes in an "efficiency units" model of schooling where schooling (in efficiency units) is the product of years of schooling and school quality. In this model, residual inequality and the college-high school premium should move exactly together over time.

none of those wage differentials expanded over the last three decades is a major challenge to the view that the return to unobserved skills grew substantially during this period.[9]

Figure 1 illustrates how the results of the paper suggest a surprising "explanation" for these various puzzles in the inequality literature. The explanation is that there was simply not a puzzle in the first place because the return to unobserved skills 1) only increased in the 1980s, and 2) does not account for much of the overall increase in wage inequality over the last three decades. This is illustrated in Figure 1 that plots the between-group variance and the composition–adjusted residual variance (based on 1973 characteristics) for men using the May/ORG CPS data for 1973 to 2003. I explain in detail later how these two series are computed. The important point is that, unlike the unadjusted residual variance computed from the March CPS, the composition–adjusted residual variance reported in Figure 1 can be interpreted as reflecting underlying changes in unobserved skill prices.

The figure clearly shows that unobserved skills prices only increases in the 1980s, just like education and experience differentials, which are summarized by the between-group variance in Figure 1, mostly increased in the 1980s. There is thus no difference between the timing of changes in the prices of observable and unobservable skills, which was the source of the initial puzzle in the inequality literature. More importantly, the between-group variance increases much more between 1973 and 2003 than the composition-adjusted residual variance, suggesting a modest role for unobserved skill prices in the overall growth in wage inequality.

Even if the puzzles listed above are no longer so puzzling, after all, the obvious question that emerges from Figure 1 is why so much of the growth in inequality concentrated in the 1980s? I return to this question in Section 7.

3. Accounting for Composition Effects

Leaving aside measurement error, equation (3) shows that the residual variance depends both on the price of unobserved skills, p_t, and on the variance of unobserved skills,

---

[9] See Heckman and Vytlacil (2001) who show that the return to cognitive ability has not changed much during the 1980s and 1990s. See Card and DiNardo (2002) for evidence that the black-white wage gap was relatively stable during the 1980s and 1990s, while the male-female wage gap *declined* substantially during this period.

Var(e_{it}). In this Section, I argue that there are strong empirical and theoretical reasons to believe that the variance of unobserved skills has increased over the last 30 years because of composition effects. I then propose a simple method for controlling for these effects.

The role of composition effects is easily illustrated using a standard variance decomposition formula. Consider the case where observed skills, x_{it}, are divided into a finite number of groups (or cells) j. The unconditional variance of unobserved skills, Var(e_{it}), is linked to the conditional variance, σ_{jt}, by the formula

(4) $Var(e_{it}) = \sum_j \theta_{jt}\, \sigma_{jt}^2$,

where $\sigma_{jt}^2 = Var(e_{it} \mid x_{it} \in j)$, and θ_{jt} is the share of workers in experience-education group j at time t. Unless wages are homoskedastic ($\sigma_{jt}^2 = \sigma_{kt}^2$ for all j, k), changes in the shares θ_{jt} will result in changes in the unconditional variance Var(e_{it}) even if the conditional variances σ_{jt}^2 are constant over time.

There is pervasive evidence of heteroskedasticity in wages, however. For example, Mincer (1974) and more recently Chay and Lee (2000) show that the variance of wages generally grows with both education and labor market experience. Since the conditional variance in wages, V_{jt}, is linked to the conditional variance of unobserved skills by the equation

(5) $V_{jt} = p_t^2 \sigma_{jt}^2$,

this suggest that σ_{jt}^2 also increases as a function of experience and education. There are a number of possible explanations for this link. In particular, Mincer (1974) argues that wage dispersion increases as a function of experience (past the overtaking point) because of differential investments in on-the-job training (OJT). In other words, inequality in the distribution of unobserved skills (OJT) increases with experience. Similarly, Farber and Gibbons (1996) show that inequality in wages and unobserved skills (as valued by the market) also increases as a function of experience in a simple learning model.[10] In both

---

[10] In Farber and Gibbons (1996), wages are equal to the expected value of productivity given the available information about the past productivity of workers. There is little wage inequality among inexperienced

of these models, the aging of the workforce results in a more dispersed unconditional distribution of unobserved skills as increasingly more weight is put on older workers with more unequally distributed skills. This can result in significant composition effects in the 1980s, 1990s, and 2000s because of the aging of the baby-boom generation. Similar arguments can be made in the case of education. For example, Mincer (1997) shows that the within-group variance of wages increases as a function of education in a standard Becker (1967) human capital model with heterogeneity in the marginal costs and benefits of investments in education. Alternatively, there may be more heterogeneity in school quality at the college than high school level.

While it is important to allow for heteroskedasticity in wages, some restrictions nonetheless need to be imposed to identify the effects of changes in skill prices as residual wage inequality. Following Chay and Lee (2000), I assume that the distribution of unobserved skills among workers with the same level of experience and education is stable over time.[11] In terms of variances, this amounts to assuming that: [12]

(6) $\sigma_{jt}^2 = \sigma_j^2$ for all time periods t.

In the absence of measurement error (I return to this issue in Section 6), the residual variance of wages, Var (ε_{it}), is then obtained by substituting equations (6) and (4) into equation (3)

(7) $\text{Var}(\varepsilon_{it}) = p_t^2 \sum_j \theta_{jt} \sigma_j^2.$

workers since the market does not yet know who is productive and who is not. Inequality increases as a function of experience as the market learns who is productive (skilled) and who is not. From the point of view of the econometrician, inequality in unobserved skills (what is valued by the market) thus grows as a function of labor market experience.

[11] Whether or not this assumption is reasonable is discussed in more detail by JMP and Chay and Lee (2000). The problem is that younger cohorts of workers may have different distributions of unobserved skills because, for example, of inter-cohort changes in the distribution of school quality. JMP convincingly argue that the steep growth in residual inequality in the late 1970s and early 1980s cannot be due to cohort effects because inequality growth accelerated uniformly for all cohorts. Unfortunately, this argument cannot be used to rule out smooth long-run trends in inequality linked to changing cohort composition because of the well-known problem that linear cohort, age, and time effects cannot be separately identified.

[12] More generally, the assumption means that $F_t(e_{it}|x_{it}) = F(e_{it}|x_{it})$ for all time periods t.

In this model, an increase in the residual variance can now be interpreted as an increase in skill prices, p_t, when the skill composition of the workforce (the θ_{jt}'s) is held constant.

Equation (7) suggests a straightforward way of holding the skill composition of the workforce constant. The residual variance just has to be recomputed at some counterfactual values of the shares, θ_j^*, that remain constant over time. To see this, rewrite the residual variance as a function of V_{jt}, the variance of wages within each skill group j

$$(8) \qquad \text{Var}\,(\varepsilon_{it}) = \sum_j \theta_{jt} V_{jt},$$

where we now have $V_{jt} = p_t^2 \sigma_j^2$. Assuming that changing the skill composition of the workforces has no general equilibrium effects on skill prices, the counterfactual residual variance, V_t^*, is[13]

$$(9) \qquad V_t^* = \sum_j \theta_j^* \, V_{jt}.$$

When the number of skill groups is small relative to sample sizes, the within-group variance V_{jt} can be computed for each skill group j. It is then straightforward to estimate the counterfactual variance by replacing the year-specific shares, θ_{jt}, by some average or base year shares, θ_j^*.[14]

In Section 3, I present some basic trends in residual and within-group inequality by dividing data in a limited number of experience-education cells (twenty). Working with these coarse cells helps illustrate which factors are driving the overall changes in residual inequality. To see this, consider the following decomposition of the change in the residual variance between a base period s and an end period t:

---

[13] Increasing the share of more educated and experienced workers depresses the return to these observed skills in a standard supply and demand model. The effect on unobserved skill prices depends, however, on the substitutability between observed and unobserved skills (see Section 2).

[14] Mincer (1974) computed such counterfactual variances. After dividing the data in about one hundred experience-education cells, he shows that the variance of wages would have been substantially larger in 1959 if older workers had been as highly educated as younger workers, which is basically was happened in the U.S. labor market over the last 40 years (Card and Lemieux, 2001a, 2001b). Mincer shows that the variance of log annual earnings in 1959 would have increased from 0.668 to 0.721 if workers at all experience levels had had the same level of education as younger workers (7-9 years of experience). This suggests that compositions effects can be quite important empirically.

(10) $V_t - V_s = \Sigma_j \, (\theta_{jt} V_{jt} - \theta_{js} V_{js})$

$= \Sigma_j \, \theta_{js}(V_{jt} - V_{js}) + \Sigma_j \, (\theta_{jt} - \theta_{js})V_{jt}.$

Equation (10) shows that the overall change in the residual variance can be decomposed into two terms. The first term on the right hand side of equation (10), $\Sigma_j \, \theta_{js}(V_{jt} - V_{js})$, is a weighted average of changes in the within-group variance. In terms of equation (8), this represents the change in the counterfactual variance, V_t^*, when the counterfactual weights, θ_j^*, are set at the base period level ($\theta_j^* = \theta_{js}$).

The second term on the right hand side of equation (10), $\Sigma_j \, (\theta_{jt} - \theta_{js})V_{jt}$, captures composition effects. Composition effects result in a spurious growth in the residual variance when changes in the weights, $\theta_{jt} - \theta_{js}$, are positively correlated with the within-group variances, V_{jt}.

When the number of cells is small enough, equation (10) suggests a simple approach for separating the role of rising skill prices from composition effects. Since $V_{jt} = p_t^2 \sigma_j^2$, the most direct evidence on rising skill prices is that the within-group variances, V_{jt}, are also growing over time. This can be readily checked by comparing these variances in a base and end period. Equation (10) then shows how these changes can be aggregated into a single factor, $\Sigma_j \, \theta_{js}(V_{jt} - V_{js})$.

From an estimation point of view, however, dividing the data in a limited number of coarse experience-education cells may be too restrictive. One alternative is to construct finer cells based on single years of education and experience. Unfortunately, cell sizes based on single years of age and education are often "too small" (and sometimes empty) in most CPS samples. Following Lemieux (2002) and DiNardo, Fortin and Lemieux (1996), I address this problem by estimating a flexible logit model to re-weight the data in a way that keeps the distribution of skills constant over time.

To see how this procedure works, note that residual variance can be computed directly from the individual level data as

(11) $V_t = \Sigma_i \, \omega_{it} \, r_{it}^2,$

where r_{it} and ω_{it} are the estimated wage residual and sample weight, respectively, for worker i at time t. In pure random samples, ω_{it} is simply defined as the inverse of the number of observations. In the CPS, however, ω_{it} differs across observations to correct for non-random features of the sample.

Equation (11) can be thought as the weighted sum of the contribution r_{it}^2 of each worker to the overall variance. By analogy, with grouped data the variance is also the weighted sum $V_t = \Sigma_j \, \theta_j \, V_{jt}$ of the contribution (V_{jt}) of each skill group to the overall variance. By analogy with equation (9), the micro-data based counterfactual variance is

$$V_t^* = \Sigma_i \, \omega_{it}^* \, r_{it}^2.$$

The estimation problem simply consists of finding the counterfactual weights ω_{it}^* that make the (counterfactual) distribution of skills at time t the same as in an appropriate base year (for example 1973). These weights are obtained by multiplying the sample weights ω_{it} by a re-weighting factor. Intuitively, to transform the skill distribution of 2003 back to its 1973 level we need to put less weight on more educated and experienced workers since the share of these workers has increased over time. In practice, the re-weighting factor is computed using the estimates from a logit model for the probability of being in year t relative to the base year. For example, the counterfactual weights for 2003 when 1973 is used as base year are computed by estimating a logit model on data for years 1973 and 2003 pooled together. The dependent variable is a dummy variable for year 2003, while the explanatory variables are the age and education variables.[15] The predicted probability that worker i is in year 2003, P_i, is then used to compute the counterfactual weight as

$$\omega_{it}^* = [(1 - P_i)/P_i] \, \omega_{it}.$$

---

[15] I use the same set of explanatory variables in the logit as in the wage regressions (full set of indicators for age and education plus interactions between education and a quartic in age). Note also this reweighting procedure is similar to the propensity score reweighting method used in the program evaluation literature. One can think of the period t sample as the "treatment group", the base period sample as the "control group", and P_i as the (estimated) propensity score.

Older and more educated workers are relatively more likely to be observed in 2003 than in 1973, suggesting a larger value of P_i and lower value of $(1-P_i)/P_i$. These workers are thus "downweighted" when ω_{it} is replaced by $\omega_{it}*$.

Once the counterfactual weights have been computed, it is straightforward to compute alternative measures of residual wage dispersion beside the variance. For example, the actual 90-10 residual gap is defined as the difference between the 90[th] and the 10[th] centile of the wage residuals when the usual sample weights ω_{it} are used. The counterfactual 90-10 residual gap is readily obtained by recomputing the 90[th] and the 10[th] percentiles using the counterfactual weights, $\omega_{it}*$, instead of the regular weights, ω_{it}.

4. Data and Trends in Within-Group Inequality by Skill Groups.

In this Section, I briefly present the May/ORG CPS data and show the basic trends in within-group wage dispersion for twenty experience-education groups. I use equation (10) to illustrate which factors —rising skill prices or composition effects— are driving the growth in the residual variance.

a. May/ORG Data

Data issues are discussed in detail in Appendix A that compares the hourly wage measure constructed from the May/ORG and March CPS Supplements. I only briefly discuss how the May and ORG supplements of the CPS are processed here. Following most of the literature, the wage measure I use is the hourly wage rate. The main advantage of this measure is that theories of wage determination typically pertain to the hourly wage rate. For example, the interplay of demand and supply considerations has direct implications for the hourly price of labor. By contrast, the impact of these factors on weekly or annual earnings also depends on the responsiveness of labor supply to changes in the hourly wage rate.

The Dual Jobs Supplement of the May CPS for 1973 to 1978 asks questions about wages on the main job held during the survey week to all wage and salary workers. For workers paid by the hour, the May CPS asks workers directly about their hourly rate of pay. This is the hourly wage measure that I use for this group of workers (about sixty percent of the workforce). For the other workers, I compute an hourly wage rate by

dividing usual weekly earnings by usual weekly hours of work. I use the same procedure for the 1979 to 1993 ORG supplements that ask the same wage questions as the May CPS. The wage questions in the 1994 to 2003 ORG supplements are similar except that workers not paid by the hour can choose the periodicity at which they report earnings. I compute their hourly wage rate by dividing earnings by hours over the corresponding time period. The merged outgoing rotation group (MORG) files combine this information for all 12 months of the year. One important advantage of the MORG supplement is that it roughly three times as large as the May of March supplements of the CPS.[16]

Unlike in the ORG and March supplements of the CPS, in the May CPS wages were not allocated for workers who refused to answer the wage questions. To be consistent, I only keep workers with non-allocated wages in the 1979-2003 ORG supplement. As a consequence, I have to drop observations for 1994 and the first eight months of 1995 in which the CPS did not flag workers with missing wages. Following most of the literature, I trim extreme values of wages (less than $1 and more than $100 in 1979 $), adjust top-coded earnings by a factor of 1.4, and weight wage observations by hours of work (in addition to the usual CPS weights). I also keep workers age 16 to 64 with positive potential experience.

All the measures of residual wage inequality are computed from the residuals of a regression of log wages on an unrestricted set of dummies for age, years of schooling, as well as interactions between nine schooling dummies and a quartic in age.[17][18] Separate regressions are estimated for both men and women in each year.

---

[16] The May 1973-78 and March supplements are administered to all (eight) rotation groups of the CPS during these months. By contrast, only one quarter of respondents (in rotation groups 4 and 8) are asked the questions from the ORG supplement each month. Combining the 12 months of data into a single MORG file yields wage data for 24 rotation groups compared to 8 in the May or March supplements (plus the hispanic and Medicare (post-2000) over-samples in the March CPS).

[17] One well-known problem with using schooling as a regressor in wage equations is that schooling is not measured in a consistent fashion over time in the CPS. Prior to 1992, the CPS asked about the highest grade attended, and whether the highest grade was completed. Starting in 1992, however, the CPS switches to a question about the highest grade or diploma completed. It is nonetheless possible to construct a relatively consistent variable for years of schooling completed over the whole sample period. The nine categories I use for years of schooling completed are 0-4, 5-8, 9, 10, 11, 12, 13-15, 16, and 17+.

[18] While it would be ideal to use an unrestricted set of age-education dummies in the wage regressions, in practice many age-education cells are quite small in the March and May supplements of the CPS. The flexible specification I use fits the data quite well. In the larger ORG samples, using a full set of age-education dummies only raises the R-square by about half a percentage point relative to the specification

b. Basic trends in within-group variances

I first divide workers into twenty skill groups based on five education categories (high school dropouts, high school graduates, some college, college graduates, and college post-graduates) and four experience categories (1-10, 11-20, 21-30, and 31 years or more of potential experience). Table 1 shows the within-group variances for each experience-education group at the beginning and end of the sample period. Since the residuals are computed from a very flexible regression, the within-group variance (variance of residuals) for a given group is smaller than the variance of unadjusted wages for the same group. To improve the precision of the estimates, I pool years 1973 to 1975 for the base period, and years 2000 to 2002 for the end period.

Tables 1a (men) and 1b (women) show the within-group variances for each of the 20 groups in 1973-75 (column 1) and 2000-02 (column 2). The change in the within-group variance is reported in column 3. Table 1a shows that, for men, changes in the within-group variance are not uniformly large and positive across skill groups. For four of the twenty groups (college graduates with 1-10, 11-20, and 21-30 year of experience, and college post-graduates with 1-10 years of experience), the changes are large and positive, and exceed the overall change in the residual variance (0.041). These groups are highlighted (in bold) in column 3. For the sixteen other education-experience groups, however, there is no systematic pattern of increase in the within-group variance. The variance grows for most groups, but declines for all high-school dropouts and for the two older groups of college post-graduates. Changes are positive and significant for four groups, but negative and significant for four other groups.

Several other clear patterns also emerge from Table 1a. In particular, the within-group variance grows as a function of both experience and education. For example, in both 1973-75 and 2000-02, high school dropouts with 1 to 10 years of experience have the lowest variance (around 0.10) while college post-graduates with 31 years and more of

---

experience have the largest variance (around 0.40). This suggests that composition effects may be quite important since both experience and education increase over time.

The results for women in Table 1b are qualitatively similar to those for men with the exception of women with some college education. For this education group, the within-group variance systematically increases between 1973-75 and 2000-02 (as in the case of men, groups for which the variance grows by more than 0.04 are highlighted in column 3). As in the case of men, the within-group variance increases for college graduates, decreases for high school dropouts, and remains relatively unchanged for high school graduates and college post-graduates.

Columns 4 and 5 show the share of each skill group in the workforce in 1973-75 and 2000-02, respectively, while column 6 shows the change in the shares over time. For both men and women, there is a large and systematic decline in the share of workers in groups with low within-group wage dispersion. This is most obvious when looking at education. For women, column 6 of Table 1b shows that, for all experience groups, the share of women with a high school degree or less has declined over time. By contrast, the share of women with some college education or more has increased for each and every experience group. With two small exceptions, the same pattern holds for men in Table 1a. The other clear pattern is that the share of more experienced workers relative to young workers systematically increases for all education groups (except high school dropouts). This reflects the aging of the baby boom generation.

Overall, Table 1a and 1b clearly show a systematic growth in the share of experience-education groups that exhibit large within-group variances. The correlation coefficient between the within-group variance in 2000-02 (V_{jt}) and the change in share ($\theta_{jt} - \theta_{js}$) is 0.55 for men and 0.68 for women. This suggests a large and positive composition effect term in equation (10).

Panel B of Table 1a and 1b shows more explicitly the magnitude of composition effects. The first row of the panel shows the weighted average of the within-group variances when the weights used are the actual shares in the corresponding year. The 1973-75 shares are used to weight the 1973-75 variances, and the 2000-02 shares are used to weight the 2000-02. The weighted averages correspond to the unadjusted

residual variances for 1973-75 and 2000-02, respectively. Table 1a shows that the residual variance for men increases by 0.041 between these two time periods.

The second row of Panel B shows that the change in the residual variance is much smaller (0.012) when the shares are held at their 1973-75 level. As a result, only about a quarter of the 0.041 change in the residual variance is due to the increase in the within-group variances. The remaining change in the residual variance, 0.029 (0.041 minus 0.012) is due to composition effects. Note that education, as opposed to experience, accounts for the bulk of the composition effects.[19]

The results for women reported in Table 1b are quite similar.   Composition effects account for 0.035 of the 0.047 increase in the residual variance. Only a quarter of the total increase (0.012) is due to the changes in within-group group variances, holding the shares constant at their 1973-75 level.

Interestingly, the last row of Panel B shows that the residual variance increases more when shares are held at their 2000-02 levels instead. The intuition for this result is that using the 2000-02 shares instead of the 1973-75 shares puts more weight on college graduates who experience a sharp increase in their within-group variances, and less weight on high school dropouts who experience a decline in their within-group variances. In other words, the base period matters in the decompositions.

Figure 2 provides some information on the detailed year-by-year evolution of the within-group variance for each of the five education groups. To control for changes in the experience distribution of the workforce, the variance for each education group is defined as the simple average of the within-group variances over the four experience groups. For example, the within-group variance for college graduates in Figure 2 is the arithmetic average of the within-group variances for college graduates with 1-10, 11-20, 21-30 and 31 or more years of experience.

I only show the detailed evolution in the within-group variance by education groups for two reasons. First, it would not be practical to show the detailed evolution in the within-group variance for each of the twenty experience-education groups. Second,

---

[19] The 0.029 composition effect (men) can be decomposed into three subcomponents. Changes in the distribution of education holding experience constant (0.025), changes in the distribution of experience holding education constant (0.008), and the interaction term (-0.004). For women, the overall composition effect (0.035) is the sum of 0.030 (education), 0.011 (experience) and –0.007 (interaction term).

Table 1 suggests that, conditional on education, the change in the within-group variance is relatively similar across experience groups. In other words, education (as opposed to experience) accounts for most of the variation in the growth in the within-group variance across the twenty experience-education groups.

The results for both men (Figure 2a) and women (Figure 2b) are different for different time periods. For both men and women, the within-group variances are either stable or declining during the 1970s. The within-group variances then grow substantially for each and every group during the 1980s. In the 1990s, however, there is a divergence in the trends by education groups. For college graduates and post-graduates, the within-group variance increases slightly or remains constant between 1990 and 2000. For all other education groups, however, the within-group variance declines during the 1990s. The decline is particularly pronounced for high-school dropouts. Finally, the within-group variances grow mildly for most groups during the early 2000s.

Taken together, the results in Table 1 and Figure 2 indicate that, for most groups, there is relatively little change in the within-group variance between 1973 and 2003. The only exception is college graduates and women with some college education. For these particular groups, however, most of the growth in the within-group variance is concentrated in the 1980s.

5. Changes in Residual Inequality: Re-weighting Results

Having established the basic patterns of changes in the within-group variance for twenty coarse experience-education cells, I now turn to a re-weighting approach to analyze in more detail the role of composition effects in changes in residual wage inequality. As discussed in Section 3, one advantage of the re-weighting approach is that it is easily implemented even when the data cannot be divided into fine experience-education cells. Another advantage of the approach is that it can be used to compute counterfactual measures of residual wage dispersion other than the variance.

Figures 3a (men) and 3b (women) compare the actual residual variance from 1973 and 2003 to the counterfactual variances that would have prevailed if the distribution of skills (experience and education) had remained at its 1973 (or 2003) level. The composition effects are the difference between the actual and counterfactual variances.

Figure 3a shows that the residual variance grows by about 0.04 over the whole sample period. Consistent with Figure 2a, most of the growth is concentrated in the first part of the 1980s. The residual variance remains essentially unchanged in the 1970s and 1990s, but grows between 1999 and 2003.

By contrast, the counterfactual variance in the late 1990s / early 2000s is only about 0.01 higher than in the mid-1970s when the distribution of skills is held constant at its 1973 level. Consistent with Table 1a, Figure 3a suggests that about three quarters of the growth in the residual variance is a spurious consequence of composition effects (when the distribution of skills is held at its 1973 level).

In terms of timing, Figure 3a shows that composition effects play a negligible role during the 1970s but become very important during the 1980s and 1990s. It is clear from Appendix Table 1 why composition effects are not important during the 1970s. The table shows that while the workforce became more educated between 1973 and 1980, it also became less experienced with the entry in the labor market of the largest baby boom cohorts (born in the late 1950s). The positive impact of growing educational achievement on the residual variance is thus offset by the fact that the workforce became younger (lower within-group variance) during this period. By contrast, Appendix Table 1 shows that both experience and educational achievement increased in the 1980s and 1990s, leading to an unambiguous positive composition bias in the growth of the residual variance.

A closer examination of Figure 3a also shows evidence of a cyclical effect in the composition effects. During the recessions of 1981-83, 1990-92, and 2000-2002, the actual variance grows faster that the counterfactual variance. This is consistent with less-skilled workers —who tend to have a lower within-group variance— being more adversely affected in terms of their employment during recessions. It is well known that composition effects tend to hide the pro-cyclicality of the level of real wages (Barsky et al, 1994). By analogy, Figure 3a suggests that composition effects tend to over-state the counter-cyclical pattern in wage inequality over the business cycle (inequality grows during recessions).

Figure 3a also shows the counterfactual variance when the distribution of characteristics is held constant at its 2003 level. The results are qualitatively similar,

though not as dramatic, as those obtained by holding characteristics at their 1973 level. The main difference is that the counterfactual variance declines less dramatically in the 1990s when characteristics are held at their 2003 instead of 1973 level.[20]

The results for women in Figure 3b are qualitatively similar to those for men. Composition effects explain most of the growth in the within-group variance between 1973 and 2003 when characteristics are held at their 1973 level. Composition effects also play a qualitatively similar, though less dramatic, role when characteristics are held at their 2003 level instead.

The results for both men and women are summarized in Table 2. The table confirms that composition effects account for most of the growth in the residual variance between 1973 and 2003 when the distribution of experience and education is held at its 1973 level. Once again, the results are less dramatic when the distribution of experience and education is held at its 2003 level instead. Even in this case, however, composition effects still account for about half of the growth in the residual variance for men, and for a third of the growth in the residual variance for women.

Table 2 also compares the growth in the residual variance to the growth in the total variance of wages (both within- and between-group components) over the same periods. Interestingly, over the whole 1973-2003 period, the residual component of the variance accounts for less than half of the growth in the total variance (43 percent for men, 46 percent for women). This finding is at odds with several previous studies that tend to find that most of the growth in wage inequality is due to the residual component. I explain in Section 6 that this earlier finding appears to be an artifact of measurement problems in the March CPS.

When the distribution of experience and education is held at its 1973 level, the remaining growth in the residual variance only accounts for 14 percent of the 1973-2003 growth in the total variance of wages for men, and 5 percent for women. These percentages increase to 23 and 31 percent, respectively, when the distribution of experience and education is held at its 2003 level instead. Table 2 also shows that when

---

[20] The difference stems from the fact that holding characteristics at their 1973 level puts relatively more weight on high-school dropouts who experience a clear decline in their within-group variance (Figure 2). By contrast, holding characteristics at their 2003 level puts relatively more weight on college graduates who experience a clear increase in their within-group variance.

20

the skill distribution is held constant, there is more growth in the residual variance between 1979 and 1989 than for the whole 1973-2003 period. This result holds for both men and women when skills are either held at their 1973 or 2003 levels. For example, the residual variance for men increases by 0.034 between 1979 and 1989 when the distribution of skills is held at its 2003 level. This is larger than the 0.025 change over the whole period. This means that for the other sample periods (1973 to 1979 and 1989 to 2003) pooled together, the residual variance *declined* by 0.009.

In light of the discussion in Section 2, these findings suggest that changes in the prices of unobserved skills only play a modest role in the overall growth in wage inequality between 1973 and 2003. For men, changes in the prices of unobserved skills account for no more than a quarter of the growth in overall inequality. For women, changes in the price of unobserved skills account for between 5 and 31 percent of the overall growth inequality.

The results also imply that all of the growth in the price of unobserved skills is concentrated in the 1980s. This finding is difficult to reconcile with the skill-biased technological change (SBTC) hypothesis that typically states that the relative demand for skills also increased during the 1970s and the 1990s. I return to the question of what else may explain the pattern of growth in residual wage inequality in Section 6.

Finally, the main findings are robust to the choice of alternative measures of wage dispersion. Figure 4 reproduces the results of Figure 3 using the 90-10 residual gap instead of the residual variance. The results are very similar to those for the residual variance. As in the case of the residual variance, almost all the growth in the 90-10 residual gap is concentrated in the 1980s (first half of the 1980s for men in Figure 4a). Furthermore, most of the growth in the 90-10 residual gap appears to be a spurious consequence of composition effects. When the distribution of experience and education is held at its 1973 level (dotted line in the figures), the 90-10 residual gap in the early 2000s is barely higher than in 1973.

Interestingly, all of the remaining growth in the 90-10 is driven by inequality growth in the upper end (90-50) of the distribution. In fact, Appendix Figure 1 shows that, for both men and women, the 50-10 gap declined while the 90-50 gap increased over time once composition effects are controlled for. This mirrors the earlier finding that

21

residual inequality increased in the upper part of the wage distribution (college educated workers) but decreased or remained stable in the lower part of the distribution (high school graduates and dropouts). For the sake of completeness, I also present the re-weighting results using the March CPS in Appendix B.

6. What is wrong with the March CPS?

As mentioned in the Introduction, the findings of Section 4 and 5 are at odds with most of the previous literature on residual wage inequality. In addition to composition effects, one potential explanation for this difference is that I use data on hourly wages from the May and ORG supplements of the CPS, while earlier studies typically use the March Supplement of the CPS.

In this Section, I argue that a key problem with the March CPS is that it poorly measures the wages of workers paid by the hour (the majority of the workforce).   I present to several pieces of evidence to show that both the level and trends in residual inequality are systematically biased in the March CPS because of the mismeasurement of the wages of workers paid by the hour.

I explain in detail in Appendix A how I compute hourly wage rates in the March CPS. Unlike the May/ORG CPS that measures wages at a point-in-time, the March CPS provides a retrospective measure of annual earnings over the previous year. From 1975 on, an hourly wage rate can be computed by dividing annual earnings by annual hours of work (annual hours of work were not collected prior to March 1976).   A number of adjustments are performed to make the hourly wage rates computed in the May/Org and March CPS as comparable as possible (see Appendix A for more detail).

a. Mismeasurement of the Wages of Workers Paid by the Hour in the March CPS
Wages computed using the March and May/ORG CPS could differ for a variety of reasons including the treatment of self-employment earnings, topcoding, etc. Instead of looking systematically at all possible sources of differences between the two data sources, I focus on the fact that earnings are collected on a yearly basis in the March CPS, while workers can report their earnings at different periodicities in the May/ORG CPS.

In particular, around 60 percent of workers in the May/ORG CPS are paid by the hour (see Figure 8). These workers report a direct measure of their hourly wage rate in the May/ORG CPS. In the March CPS, however, they have to report their total annual earnings and hours of work that are then used to compute an average hourly wage rate.

In the absence of measurement error, it should not matter whether hourly wages are computed directly from questions about hourly wage rates, or indirectly by dividing earnings by hours of work. Several validation studies show, however, that there is substantial measurement error in the earnings reported in the CPS or similar surveys.[21]

It is plausible to think that asking directly hourly-rated workers about their hourly wage rates provides a more accurate wage measure than dividing earnings by hours. If it is easier for workers paid by the hour to report directly their hourly wage rate, this direct measure will likely be less affected by measurement error than the indirect wage measure based on average hourly earnings. For example, a minimum wage worker will likely know and correctly report the exact value of the hourly wage at which he or she is paid. The same worker would probably have more difficulty reporting total hours and earnings during the year. In fact, the U.S. Census Bureau and other national statistical offices often mention the case of the minimum wage as one reason for asking directly workers paid by the hour about their hourly wage rate.

My basic hypothesis is that for hourly-rated workers, the hourly wage rate indirectly computed from the March CPS is a more noisy measure of the true hourly rate of pay than the hourly wage rate collected in the May/ORG CPS. For workers not paid by the hour, the hourly wage rate has to be indirectly computed by dividing earnings by hours in both the May/ORG and the March CPS. Therefore, I do not expect the hourly wage from the March CPS to be a more noisy measure for these workers.

Under the additional assumption of classical measurement error, this hypothesis yields several clear empirical predictions.[22] The most direct prediction is that the variance of March CPS wages should be larger than the variance of May/ORG CPS

---

[21] Mellow and Sider (1983) compare employee and employer responses in the January 1977 Validation Study of the CPS. Bound and Krueger (1991) compare employee responses from the March 1977 and 1978 CPS to employer reported Social Security Earnings.
[22] Under classical measurement error, the measurement error in wages is assumed to have mean zero and be independent of all observable variables.

wages among workers paid by the hour. I test this prediction by comparing the variance of the two wage measures for workers paid by the hour and workers not paid by the hour.

One problem with implementing this test is that the March CPS does not ask individuals whether they are paid by the hour or not. Fortunately, this problem can be resolved by exploiting the rotation group feature of the CPS. Among individuals sampled in the March CPS, roughly one quarter of individuals rotate out of the CPS in each of the next four months, including March. This means that from 1979 on, all individuals in the March CPS should eventually be part of the outgoing rotation group in March, April, May or June. In principle, their responses to the ORG supplement questions can thus be matched to their March CPS records. As discussed below, however, not all March respondents can be matched because of attrition and other data problems.

Prior to 1979, it is still possible to match the May CPS responses to the March responses for the March respondents who are still in the CPS in May (half of the March respondents when there is no attrition). My empirical strategy is thus to match the March CPS respondents to either their ORG or May CPS records. From this matched sample, I can then use the information from the ORG or May CPS questions to divide workers into those paid and not paid by the hour.

Working with these matched samples involves a number of data issues that are beyond the scope of this paper. In particular, between five and ten percent of respondents cannot be matched because of attrition and other data problems.[23] Also, while the March and May/ORG wage records are for the same respondent, they are not for the same period (March wage is for last year, May/ORG wage is for the current survey month). This means that some workers coded as "paid by the hour" may not have been paid by the hour in the previous year. Focusing on workers who both report a wage in the month of the survey (the May/ORG wage) and in the previous year also results in a more "stable" sample of workers. Fortunately, Appendix Figures 2a and 2b show that while the level of wage inequality is lower for this matched sample than for the full sample, the trends in inequality are very similar for the two samples.

---

[23] Since I am only matching months close by, the matching rates are much higher than in most applications where records in one year are matched to the record for the same respondent one year later.

Despite these data limitations, a striking pattern of results emerges from Figures 5a (men) and 5b (women). These figures contrast the variance of the March and May/ORG wages for the two groups of workers (paid by the hour or not). For both men and women, the variance of March wages is systematically larger than the variance of May/ORG wages for workers paid by the hour. By contrast, there is no systematic difference in the variance of March and May/ORG wages for workers not paid by the hour. Figures 5a and 5b provide clear evidence that the key difference between the March and May/ORG wages is that the wages of workers paid by the hour are more noisily measured in the March CPS.

The extent of measurement error in March CPS for workers paid by the hour is both quantitatively and statistically significant. For men (Figure 5a), the average difference in the variance is 0.064, which represents about a third of the average variance in the May/ORG CPS (0.198). The results are similar for women. The average difference in variances (0.055) also represents a third of the average variance of wages in the May/ORG CPS (0.167).

Despite these large differences in levels, the pattern of change in the variances over time is relatively similar in the March and May/ORG CPS. For both wage measures, the variance of wages for hourly workers is flat in the 1970s, grows sharply in the 1980s, and remains relatively constant thereafter. For workers not paid by the hour, however, the variance of wages keeps increasing steadily during the 1990s. This is consistent with workers not paid by the hour being much more educated than workers paid by the hour (see below), and within-group inequality increasing for college educated workers in the 1990s.

There is a second empirical prediction about measurement error in the March CPS that can be tested without resorting to the matched sample. Under the assumption of classical measurement error, the additional noise in the March CPS measure of wages (for hourly workers) should not affect estimates of the conditional means of wage (by education, age, etc.).[24] This means that measurement error should have no effect on the between-group variance of wages (i.e. the dispersion in conditional means) when samples

---

[24] The assumption is reasonable since both Mellow and Sider (1983) and Bound and Krueger (1991) find that measurement error in the CPS earnings in the late 1970s is uncorrelated with typical regressors like experience and education.

are large enough. If hourly wages from the March CPS are simply a noisier measure of hourly wages than wages in the May/ORG CPS (for hourly workers), then the two wage measures should yield similar between-group variances of wages. The measurement error should just increase the within-group, or residual, variance of wages in the March relative to the May/ORG CPS.

Figures 6 and 7 confirm this empirical prediction. Figure 6a shows the evolution of the between-group variance for men over the 1975-2003 period for the two measures (March and May/ORG) of hourly wages.[25] In the case of hourly wages computed from the March CPS, I report the between-group variance with and without observations with allocated earnings. The figure shows that including observations with allocated earnings has essentially no impact on the between-group variance. This suggests that the mean of allocated wages by age and education categories are similar to the mean for observation with valid (non-missing) wages.

More importantly, the two wage measures yield very similar between-group variances of log wages. Both the levels and the trends in the two series are very similar. Almost all the growth in the between-group variance is concentrated during the first half of the 1980s. The between-group variance is more or less constant between 1975 and 1980, and after 1985. This finding is very robust to the choice of hourly wage measure.

The results for women in Figure 6b are also robust to the choice of wage measure. The between-group variance obtained from the May/ORG and the March CPS (with and without allocators included) all show the same basic pattern. The between-group variance declines in the 1970s, grows sharply in the first half of the 1980s, and grows more slowly thereafter. One natural explanation for the continuing growth in the between-group variance throughout the 1980s and 1990s is that age-earnings profiles are getting steeper during this period because of the increased attachment of women to the labor force.[26]

Turning to residual wage dispersion, Figure 7a shows that, for men, the residual variance of March CPS wages (without allocated earnings) is systematically larger than

---

[25] Figures 6 and 7 report the variance of wages by earnings year (year of the survey in the May/ORG CPS, previous year in the March CPS). I report data for 1975 to 2003 that correspond to the 1976 to 2004 survey years in the March CPS.
[26] See Blau and Kahn (1997) and Fortin and Lemieux (1998).

the residual variance of May/ORG wages. The results in Figure 7b for women are very similar.   A set of strong conclusions thus emerges from Figures 5, 6 and 7. First, Figure 5 clearly shows that wages are more noisily measured in the March CPS. Consistent with classical measurement error, Figures 6 and 7 show that these measurement problems do not affect between-group wage dispersion but spuriously inflate residual wage dispersion in the March CPS. These findings strongly support the view that, relative to the May/ORG CPS, residual wage inequality is biased up in the March CPS because this wage measure poorly captures the hourly wage rate for workers paid by the hour.

c. Spurious Trends in Residual Wage Inequality in the March CPS?
If the bias in residual wage inequality in the March CPS were constant over time, using the May/ORG or the March CPS would have little consequence for the interpretation of the sources of *change* in residual wage inequality. Figure 7 shows, however, that both the level and growth in residual wage inequality are larger in the March than in the May/ORG CPS. In the case of men (Figure 7a), the residual variance of wages in the May/ORG CPS is stable during the 1970s, grows rapidly in the early 1980s, and remains fairly constant from the mid-1980s to the late 1990s. By contrast, the residual variance grows steadily from 1975 to 2003 when hourly wages are computed using the March CPS. Among women, there is also more growth in the residual variance of March relative to May/ORG wages, though the difference is not as marked as in the case of men.

One simple explanation for this difference is that measurement problems in the March CPS have been magnified by the growth over time in the fraction of workers paid by the hour. Consistent with Hamermesh (2002), Figure 8 indeed shows that the fraction of workers paid by the hour has increased over time. Since education is by far the most important factor explaining the propensity to be paid by the hour, I only report the fraction of workers paid by the hour by education group in Figure 8.[27] Figure 8a (men) and 8b (women) show that the fraction of workers paid by the hour has increased by up to 15-20 percentage points (depending on the education group) between the mid-1970s and the late 1980s.

---

[27] The fraction of workers paid by the hour declines as a function of experience. Relative to education, however, experience has a smaller impact (in absolute value) on the probability of being paid by the hour.

Recall from Figure 5 that, for hourly-rated workers, the variance of wages in the March CPS exceeds the variance in the May/ORG by 0.05 to 0.07. This provides an estimate of the variance component due to the fact that wages are more poorly measured in the March than in the May/ORG CPS. Combining this spurious variance component with the 10-15 percentage point growth in the fraction of workers paid by the hour (Figure 8) yields a spurious growth of up to 0.01 in the variance of March wages. This is substantial relative to the 0.04 to 0.05 growth in the residual variance in the May/ORG CPS during the same period.

Autor, Katz and Kearney (2004) argue, however, that it is instead the May/ORG CPS that yields downward biased growth in residual wage inequality. Their point is based on a very different assumption about the nature of measurement error in the March and May/ORG CPS. They assume that hourly wages obtained by dividing earning by hours are more noisily measured than direct measures of the hourly wage (as for hourly workers in the May/ORG CPS).[28] They further assume that the variance of measurement error is the same for hourly workers in the March CPS, non-hourly workers in the March CPS, and non-hourly workers in the May/ORG CPS.

In the absence of further information about measurement error in the different wage measures, it is not possible to say whether the growth in residual wage inequality is biased up in the March CPS (as argued above) or biased down in the May/ORG CPS (as argued by Autor, Katz and Kearney, 2004). Fortunately, the matched March-May/ORG samples can be used to probe these assumptions in more detail. Under the assumption that the measurement error in the March and May/ORG wages is uncorrelated, it is possible to estimate the "true" variance in wages along with the variance of measurement error in the March and May/ORG CPS. More formally, assume that

$$w^M = w^* + v^M \text{ , and}$$
$$w^O = w^* + v^O,$$

where w^* is the "true" (log) wage, w^M and w^O are the hourly wages as measured in the March and May/ORG CPS, respectively, and v^M and v^O are the corresponding measurement errors, where $cov(v^M, v^O) = 0$. In this simple model, the true variance of

---

[28] A simple rationale for this assumption is that hours are also measured with error and introduce an additional error component in measures of hourly wages obtained by dividing earnings by hours.

wages, var(w^*), and the measurement error variances, var(v^M) and var(v^O), can be identified from the empirical covariance matrix between w^M and w^O.[29]

Appendix Figure 3a shows the estimated measurement error variances for men paid by the hour and not paid by the hour. Appendix Figure 3b reports the same estimates for women. As expected from Figure 5, the measurement error variances for non-hourly workers are comparable in the March and May/ORG CPS. Table 3 shows that measurement error accounts for about 20 percent of the variance of wages for these workers. Also as expected, the measurement error variance for hourly workers is much larger in the March than in the May/ORG CPS. Measurement error represents about a third of total variance of wages in the March CPS compared to only about 10% of the total variance of wages in the May/ORG CPS. Interestingly, the measurement error variance for non-hourly workers lies more or less in-between the measurement error variance for hourly workers in the May/ORG and March CPS. This is inconsistent with Autor, Katz and Kearney (2004)'s assumption that the variance of measurement error is the same for hourly and non-hourly workers in the March CPS.

The estimates suggest that the growth in the fraction of workers paid by the hour both biases up the growth in inequality in the March CPS, and biases down the growth in inequality in the May/ORG CPS. The magnitude of these biases are shown in row 3a Table 3 under the assumption that the fraction of hourly workers increased by 10% for men and 15% for women (see Figure 8).

Interestingly, Appendix Figure 3 and Table 3 (rows 2a and 2b) also indicate that the measurement error variance has been generally growing over time. This suggests that part of the increase in residual wage inequality is simply a consequence of the fact that wages are increasingly badly measured in both the March and the May/ORG CPS. Row 3b of Table 3 shows that, for men, the variance of measurement error increased by 0.018 in the March CPS compared to 0.004 in the May/ORG CPS. For women, the corresponding measurement error variances grew by 0.014 (March CPS) and 0.013 (May/ORG CPS). Since the two sources of measurement error go in opposite directions for men in the May/ORG CPS, the adjusted change in the residual variance (row 5) is the

---

[29] The three elements of the covariance matrix are var(w^M)=var(w^*)+var(v^M), var(w^O)=var(w^*)+ var(v^O) and cov(w^M, w^O)=var(w^*). The three unknows var(w^*), var(v^M), and var(v^O) can be directly solved from these three equations.

same as the change unadjusted for measurement error (row 4). By contrast, the adjusted change is systematically smaller than the unadjusted change in the March CPS as both sources of measurement error tend to inflate the growth in the residual variance.

I conclude from this detailed examination of the measurement of hourly wages in the CPS that the May/ORG CPS provides a more accurate measure of both the level and the growth in residual wage inequality than the March CPS. For men, the growth in the residual wage inequality is unaffected by measurement error corrections in the May/ORG CPS. For women, adjusting for measurement error reduces the growth in residual inequality in May/ORG CPS from 0.057 to 0.047. Measurement error adjustments are even larger in the March CPS, suggesting that the growth in residual wage inequality as measured in this data set is substantially biased up.

6. Concluding Comments: What Explains the (Modest) Growth in Residual Wage Inequality?

The "common wisdom" about residual wage inequality is that it grows steadily over time and accounts for most of the overall growth in wage inequality. A very different picture emerges when the composition of the workforce is held constant over time and wages are measured using the "better" May/ORG CPS instead of the March CPS. In particular, I find that residual wage inequality only accounts for a modest share of the growth in overall inequality between 1973 and 2003 after these adjustments are made. I also find that residual wage inequality generally moves in tandem with other "between-group" wage differentials. From a time-series point of view, the growth in both residual and between-group wage inequality is all concentrated in the 1980s (Figure 6 and 7). From a cross-sectional point of view, the group of workers for which residual inequality grows the most (college-educated workers) also happens to be a group for which relative wages expanded the most dramatically over the last 30 years.[30]

---

[30] Because of space limitations, I do not present detailed information on other wage differentials. Both Mincer (1997) and Deschênes (2001) show, however, that (log) wages became an increasingly convex function of years of schooling since the 1970s. In other words, the gap between high school and college educated workers expanded dramatically, while the gap between high school graduates and dropouts remained more or less constant (or even declined in some cases).

As discussed in Section 2, these findings have important implications for the interpretation of the role of unobserved skill prices in the overall growth in wage inequality. First, they help resolve several puzzles linked to the timing and extent of changes in observable and unobservable skill prices. Second, the results generally reinforces the conclusion of Card and DiNardo (2002) that the timing of the growth in wage inequality is difficult to reconcile with the SBTC hypothesis. This paper poses a further challenge to the SBTC hypothesis since I find that residual inequality actually *declined* in periods other than the 1980s. Technological change can only explain these changes under the implausible assumption that while technological change was biased in favor of skilled workers during the 1980s, it was biased in favor of unskilled workers during the other periods.

This suggests looking at other possible explanations for changes in residual wage inequality. For example, DiNardo, Fortin, and Lemieux (1996) find that the decline in the real value of the minimum wage during the 1980s accounts for about a third of the increase in residual wage inequality. Lee (1999) finds an even larger effect by allowing for spillover effects of the minimum wage. Interestingly, the basic trends in the real value of the minimum wage are closely related to the trends in residual wage inequality documented above. For example, Row C of Table 2 shows that the real value of the minimum wage declined in the 1980s and early 2000s, while residual inequality increased during those two periods. By contrast, residual inequality declined when the real value of the minimum wage increased during the 1970s and 1990s.

Figure 9 explores in more detail the connection between the evolution of the minimum wage and the residual variance between 1973 and 2003. The figure compares the residual variance when characteristics are held constant at their 1973 level to the predicted variance from a regression that includes a linear trend and the log real minimum wage as regressors.[31] This simple regression model explains the data quite

---

[31] The estimated effect of the minimum wage is very similar when the regressions are fit to the residual variance that holds the distribution of characteristics at its 2003 (instead of 1973) level. In this case, however, the underlying time trend is small and positive, while it is negative and significant when characteristics are held at their 1973 level. These results suggest that there is essentially no growth in residual inequality left once composition effects and the impact of the minimum wage have been accounted for.

31

well. The R-square is 0.81 and 0.88 for men and women, respectively. This is a very good fit since there is almost no time trend in the residual variance (the dependent variable).

For both men (Figure 9a) and women (Figure 9b), the minimum wage has a strong impact on the residual variance. The regression models are reported in the figures and show large t-statistics for the effect of the minimum wage (t-statistic of 9 for men, and 12 for women). Consistent with DiNardo, Fortin and Lemieux (1996), the effect of the minimum wage is also larger for women then men. The "visual fit" of the model is most impressive for women. The large increases in the minimum wage in 1973-74, 1989-91, and 1995-97 all closely match corresponding declines in the residual variance. By contrast, the three periods where the minimum wage declined in real terms for failing to be indexed (1981-1989, 1992-1995, and 1998-2003) all correspond to clear increases in residual wage inequality.[32]

While the minimum wage explains well the time series pattern of the residual variance, it is not a very credible explanation for the substantial growth in within-group inequality for the most highly educated workers, or for the related expansion in wage inequality in the upper part of the residual distribution (residual 90-50 gap).   Clearly, something else needs to be brought in to explain the growth in inequality in the top end of the wage distribution.

DiNardo, Lemieux and Fortin (1996) also show that, for men, a substantial fraction (40 percent) of the increase in the 90-50 gap in the 1980s can be linked to the fall of unionization. Intuitively, the decline in unions mostly affect workers around the middle of the skill distribution who were traditionally more likely to belong to unions, thereby expanding the gap between the median and higher wage quantiles. Other studies like Freeman (1993) and Card (1992) have also shown that de-unionization explains around 20 percent of the increase in inequality for men during the 1980s, though most of the effect appears to be concentrated on the between-group instead of the residual variance (DiNardo, Lemieux and Fortin, 1996). Interestingly unionization did not decline

---

[32] The three most important increases in the minimum wage are: from $1.60 to $2.00 in May 1974, from $3.35 in March 1990 to $4.25 in April 1990, and from $4.25 in September 1996 to $5.15 in September 1997. The real value of the minimum wage was substantially eroded by inflation as the minimum wage remained fixed at $3.35 from January 1981 to March 1990, at $4.25 from April 1991 to September 1996, and at $5.15 from September 1997 on.

nearly as much in the 1990s as in the 1980s. In fact, Card, Lemieux, and Riddell (2004) find that unions had very little impact on changes in wage inequality during the 1990s. Like the minimum wage explanation, de-unionization works quite well in terms of timing as it helps explain the growth in inequality in the period during which most of the changes were concentrates (the 1980s).

Other possible explanations are discussed by Piketty and Saez (2003) who document a dramatic increase in inequality in the top end of the earnings distribution (using tax data) since the 1970s. They argue that both the timing (over the long run) and the extent of the growth in inequality at the top end are hard to reconcile with SBTC. They rather favor an alternative explanation based on social norms. It is nonetheless possible for more nuanced forms of technological change to disproportionally benefit some workers at the top end of the wage distribution. For example, Goos and Manning (2003) show that the model of technological change of Autor, Levy, and Murnane (2003) where computerization only replaces "routine", as opposed to "unskilled", tasks can lead to an expansion of wage inequality in the upper part of the wage distribution, but to a reduction in inequality in the lower part of the distribution.[33]

A more traditional but related "human capital" explanation is that when the return to college education increases, we also expect the return to a "good" college education to increase even more. This idea can be captured in a "single index" model where total educational input is simply the product of years of schooling with school quality. This implies that school quality and years of schooling are perfect substitutes and that the within-group variance for college graduates (variance of school quality times the return to college) is proportional to the return to college. The important point here is that a standard human capital approach could help explain the cross-sectional pattern of changes in both residual wage inequality and between-group wage differentials that are observed in the data.

In summary, the fact that residual wage inequality growth is not that important quantitatively and moves in tandem with between-group inequality generally "simplifies the story" on changes in wage inequality. It leaves only two key questions to be

---

[33] The idea is that workers in the middle of the distribution are those who perform "skilled but routine" task. They are thus the most adversely affected by technological change.

answered. The first question is why the overall growth in wage inequality is so concentrated in the 1980s? The second question is why wage inequality has mostly expanded in the upper end of the wage distribution? I have suggested possible answers to these questions but much nonetheless remain to be done in future research.

Finally, an important by-product of the paper is to show that the March CPS does not provide very accurate measures of wages for the majority of workers who are paid by the hour. The ORG supplement of the CPS provides more accurate measures of hourly wages for much larger samples of workers than the March CPS. Over thirty years of data are now available when the ORG CPS is combined with the 1973-78 May CPS. There is thus a strong case for using the May/ORG CPS, instead of the March CPS, for studying the determinants and the evolution of the structure of wages in the United States since the early 1970s.

References

Acemoglu, Daron, "Technical Change, Inequality, and the Labor Market," *Journal of Economic Literature* 40, March 2002, 7-72

Aghion, Philippe, "Schumpeterian Growth Theory and the Dynamics of Income Inequality," *Econometrica* 70(3), May 2002, 855-82

Autor, David H., Lawrence F. Katz, and Alan B. Krueger, "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113(4), November 1998, 1169-1213

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney, "Trends in U.S. Wage Inequality: Re-assessing the Revisionists," mimeo, August 2004

Autor, David, Frank Levy, and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Investigation," *Quarterly Journal of Economics* 118(4), November 2003, 1279-1333

Barsky, Robert, Parker, Jonathan a, and Gary Solon, "Measuring the Cyclicality of Real Wages: How Important Is Composition Bias?" *Quarterly Journal of Economics* 109(1), February 1994, 1-25

Beaudry, Paul, and David A. Green, "Changes in US Wages, 1976-2000: Ongoing Skill Bias or Major Technological Change?" forthcoming, *Journal of Labor Economics*, 2004

Becker, Gary S. *Woytinski Lecture*, University of Michigan, 1967. Reproduced in *Human Capital (2nd Edition),* University of Chicago Press, 1975.

Berman, Eli, John Bound, and Zvi Griliches, "Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures," *Quarterly Journal of Economics*, 109(2), May 1994, 367-97

Blau, Francine D., and Lawrence M. Kahn, "Swimming Upstream: Trends in the Gender Wage Differential in the 1980s," *Journal of Labor Economics* 14(4) October 1996, 1-42

Bound, John, and George Johnson, "Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations," *American Economic Review* 82, June 1992, 371-92

Bound, John, and Alan B. Krueger, "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?" *Journal of Labor Economics 9*, January 1991, 1-24

Card, David, "The Effects of Unions on the Distribution of Wages: Redistribution or Relabelling?", NBER Working Paper 4195, Cambridge: Mass.: National Bureau of Economic Research, October 1992.

Card, David, and John DiNardo, "Skill Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles," *Journal of Labor Economics* 20, October 2002, 733-83

Card, David, and Thomas Lemieux, "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis," *Quarterly Journal of Economics 116*, May 2001a, 705-46

— , "Dropout and Enrollment Trends in the Post War Period: What Went Wrong in the 1970s?" in Jonathan Gruber (ed.) *An Economic Analysis of Risky Behavior among Youth*, 2001b, Chicago: University of Chicago Press for NBER

Card, David, Thomas Lemieux, and W. Craig Riddell, "Unions and Wage Inequality" *Journal of Labor Research 25*, Fall 2004, 519-62.

Chay, Kenneth and David Lee, "Changes in Relative Wages in the 1980s: Returns to Observed and Unobserved Skills and Black-White Wage Differentials," *Journal of Econometrics* 99, 2000, 1-38

Deschênes, Olivier, "Unobserved Ability, Comparative Advantage and the Rising Return to Education in the United States: A Cohort-Based Approach," Princeton University Industrial Relations Section Working Paper No. 465, August 2001.

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux (1996) "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica 65*, September 1996, 1001-46

Farber, H. and R. Gibbons, "Learning and Wage Dynamics," *Quarterly Journal of Economics* 111, 1996, 1007-47

Fortin, Nicole, and Thomas Lemieux, "Rank Regressions, Wage Distributions, and the Gender Gap" *Journal of Human Resources* 33, Summer 1998, 610-643

Freeman, Richard B. "How Much has Deunionization Contributed to the Rise of Male Earnings Inequality?", in Sheldon Danziger and Peter Gottschalk (eds), *Uneven Tides: Rising Income Inequality in America*, New York: Russell Sage Foundation.

Goos, Maarten, and Alan Manning, "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," Centre for Economics Performance Discussion Paper No. 604, December 2003.

Hamermesh, Daniel, "12 Million Salaried Workers Are Missing" Industrial and Labor Relations Review 55, No. 4 (July 2002), p. 649.

Heckman, James J., and Edward Vytlacil, "Identifying the Role of Cognitive Ability in Explaining the Level of and Change in the Return to Schooling," *Review of Economics and Statistics* 83(1), February 2001, 1-12

Hirsch, Barry, and Edward Schumacher, "Match Bias in Wage Gap Estimates Due to Earnings Imputation," forthcoming, *Journal of Labor Economics*, 2004

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

Katz, Lawrence, and David Autor, "Changes in the Wage Structure and Earnings Inequality," in O. Ashenfelter and D. Card (eds.) *Handbook of Labor Economics*, volume 3A, 1999, Amsterdam: Elsevier Science

Katz, Lawrence, and Kevin Murphy, "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics 107,* January 1992, 35-78

Lee, David, "Wage Inequality in the United States during the 1980s: Rising Dispersion or Falling Minimum Wage?" *Quarterly Journal of Economics* 114, August 1999, 977-1023

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

Lemieux, Thomas, "Decomposing Wage Distributions: a Unified Approach" *Canadian Journal of Economics 35*, November 2002, 646-88.

Lillard, Lee, James P. Smith, and Finis Welch, "We Really Know about Wages? The Importance of Nonreporting and Census Imputation" *Journal of Political Economy* 94, Part 1, June 1986; 94(3): 489-506.

Mellow, Wesley, and Hal Sider, "Accuracy of Response in Labor Market Surveys: Evidence and Implications" *Journal of Labor Economics 1*, October 1983, 331-44

Mincer, Jacob, *Schooling, Experience, and Earnings*, New York: NBER, 1974

Mincer, Jacob "Changes in Wage Inequality, 1970-1990," *Research in Labor Economics 16*, 1997, 1-18.

Piketty, Thomas, and Emmanuel Saez, "Income Inequality in the United States, 1913-98," *Quarterly Journal of Economics* 118(1), February 2003, 1-39

APPENDIX A: May/ORG and March CPS Data

This appendix explains in more detail how the March and May/ORG CPS are processed to make the wage samples as comparable as possible. Both the May/ORG and the March CPS can be used to compute hourly wage rates. The March Supplement of the CPS asks about total earnings during the previous year. An hourly wage rate can then be computed by dividing last year's earnings by total hours worked last year. The latter variable is computed by multiplying two other variables available in the March CPS, usual weekly hours of work last year and weeks worked last year.

For historical reasons, however, many studies based on March CPS data proxy for hourly wage rates by focusing only on the earnings of full-time (and sometimes full-year) workers. The reason is that prior to 1976, the March CPS only asked about full-time/part-time status last year (instead of usual hours of work last year). Furthermore, the information about weeks worked last year was limited to few intervals (0, 1-13, 14-26, 27-39, 40-47, 48-49, 50-52) in the pre-1976 March CPS. One important drawback of this alternative wage measure, however, is that it is limited to the subset of the workforce that works full-time (and sometimes full-year). This is particularly problematic for women. It also fails to control for the dispersion in hours of work among workers who work full-time (35 hours and more a week).

Since we now have almost 30 years of data for which hourly wages rates can be directly computed for all workers, I limit the analysis of wages in the March CPS to the period starting with the earnings year 1975 (March 1976 survey). Another reason for starting with the wage data for 1975 is that the other wage measure available in the May/ORG CPS is only available from May 1973 on. Since one contribution of the paper is to compare the two data sources, the gain of using a more precise and comparable measure of hourly wages from the March CPS clearly outweighs the cost of losing two years of data for 1973 and 1974.[34]

There are important differences between the way wages are measured in the March and May/ORG CPS. First, while the March CPS asks about retrospective measures of wages and earnings (last year), the May/ORG supplement asks about wages at the time of the survey. Second, the May/ORG wage questions are only asked to wage and salary workers. By contrast, the March CPS asks separate questions about wage and salary earnings and self-employment earnings. To get comparable wage samples, I limit my analysis of the March data to wage and salary earnings. One problem is that when workers have both wage and salary and self-employment earnings, we do not know how many hours of work pertain to wage and salary jobs vs. self-employment. To minimize the impact of these considerations, I limit my analysis to wage and salary workers with very limited self-employment earnings (less than ten percent of wage and salary earnings).

Another difference is that the ORG supplement only asks questions about the worker's main job (at a point in time) while the March CPS includes earnings from all

---

[34] Another problem is that since missing wages were not allocated in the May 1973-78 CPS, allocated wages and earnings should be excluded from the March CPS for the sake of comparability. Unfortunately, individual earnings allocation flags are not available in the March CPS prior to the 1976 survey (Lillard, Smith, and Welch, 1986). Though family earnings allocation flags can be used instead (see JMP), this is one more reason for focusing on the March CPS data starting with the earnings year 1975.

jobs, including second jobs for dual job holders. Fortunately, only a small fraction of workers (around 5 percent typically) hold more than one job at the same time. Furthermore, these secondary jobs represent an even smaller fraction of hours worked.

Finally, since the May/ORG CPS is a "point-in-time" survey, the probability that an individual's wage is collected depends on the number of weeks worked during a year. By contrast, a wage rate can be constructed from the March wage information irrespective of how many weeks (provided that it is not zero) are worked during the year. This means that the May/ORG wage observations are implicitly weighted by the number of weeks worked, while the March wage observations are not.

One related issue is that several papers like DiNardo, Fortin and Lemieux (1996) also weight the observations by weekly hours of work to get a wage distribution representative over the total number of hours worked in the economy. Weighting by weekly hours can also be viewed as a reasonable compromise between looking at full-time workers only (weight of 1 for full-time workers, zero for part-time workers) and looking at all workers as "equal" observations irrespective of the number of hours worked. Throughout the paper, I thus weight the March CPS observations by annual hours of work, and weight the May/ORG observations by weekly hours of work.

In both the March and ORG supplements of the CPS, a growing fraction of workers do not answer questions about wages and earnings. The Census Bureau allocates a wage or earnings item for these workers using the famous "hot deck" procedure. The CPS also provides flags and related sources of information that can be used to identify workers with allocated wages in all years except in the January 1994 to August 1995 ORG supplements.[35] By contrast, in the May 1973-78 CPS, wages were *not* allocated for workers who failed to answer wage and earnings questions.[36] For the sake of consistency across data sources, all results presented in the paper only rely on observations with non-allocated wages, unless otherwise indicated.

Wages and earnings measures are topcoded in both the March and May/ORG CPS. Topcoding is not much of an issue for workers paid by the hour in the May/ORG CPS. Throughout the sample period, the topcode remains constant at $99.99 and only a handful of workers have their wage censored at this value. By contrast, a substantial number of workers in the March CPS, and non-hourly workers in the May/ORG CPS, have topcoded wages. When translated on a weekly basis for full-year workers, the value of the topcode for annual wages in the March CPS tends to be comparable to the value of the topcode for weekly wages in the May/ORG CPS. For instance, in the first sample years (1975 to 1980) the weekly topcode in the May/ORG CPS is $999 compared to $962

---

[35] Allocation flags are incorrect in the 1989-93 ORG CPS and fail to identify most workers with missing wages. Fortunately, the BLS files report both edited (allocated) and unedited (unallocated) measures of wages and earnings. I use this alternative source of information to identify workers with allocated wages in these samples.

[36] There has been some confusion in the literature because of the lack of good documentation on the allocation of missing wages in the 1973-78 CPS. Several papers assume that, like in the March CPS prior to 1976, wages were allocated but not flagged in the May 1973-78 CPS. For example, Katz and Autor (1999) compare a (May CPS) sample without allocated wages in 1973 to a sample with allocated wages in 1979. This likely overstates the growth in residual wage inequality during the 1970s since residual wage dispersion is generally higher when allocated wages are included than when they are not (see Figure 7). See Hirsch and Schumacher (2004) for a detailed discussion of how wages are allocated (or not allocated) in the May/ORG CPS.

for full-year workers in the March CPS (annual topcode of $50,000). In the last sample years (1998 to 2003), the weekly topcode in the ORG CPS is $2884, which is identical to the implied weekly topcode for full-year workers in the March CPS (annual topcode of $150,000 divided by 52). Following most of the literature, I adjust for topcoding in both the May/ORG and the March CPS by multiplying topcoded wages by a factor 1.4.

For a variety of reasons, several data adjustments are performed before applying the 1.4 factor to topcoded wages. In the May/ORG CPS, the topcode on the edited weekly earnings variable for workers not paid by the hour goes from $999 in 1973-1988 to $1923 in 1989-1997, and $2884 in 1998-2002. Between 1986 and 1988, however, it is possible to use the unedited weekly earnings variable which is topcoded at $1999 instead of $999. Though the unedited variable is not computed for workers who fail to respond to the earnings question, this does not matter here since I only use data for workers with unallocated wages and earnings. I thus use the unedited earnings variable for the 1986-88 period.

Several adjustments also have to be performed before applying the 1.4 factor to the March CPS data. Until March 1989, wages and salaries were collected in a single variable pertaining to all jobs, with a topcode at $50,000 until 1981 (survey year), $75,000 from 1982 to 1984, and $99,999 from 1985 to 1988. Beginning in 1989, the March CPS started collecting wage and salary information separately for main jobs and other jobs, with topcodes at $99,999 for each of these two variables. The topcodes were later revised to $150,000 for the main job and $25,000 for other jobs in March 1996.

Prior to March 1996, the earnings variable of workers who are topcoded simply takes the value of the actual topcode. Starting in March 1996, however, the value of earnings for topcoded workers is replaced by the mean earnings among all topcoded workers. Mean earnings are separately computed for different demographic groups. For example, in the March 2001 CPS, the mean for topcoded main job earnings ranges from $195,699 for white females not working full-time full-year, to $335,115 for full-time full-year white males. The corresponding means for these two groups are $39,320 and $56,879 for wage and salary earnings on other jobs.

To maintain consistency over time, I first construct a topcoded variable for total wage and salary earnings from March 1989 on. For 1989-1995, I simply keep the pre-1989 $99,999 topcode. Since both main job and other job earnings are separately topcoded at $99,999, I simply add these two earnings variables and topcode the sum at $99,999. After various experiments, I decided to use a topcode of $150,000 for total wage and salary earnings from 1996 on. Unfortunately, it is not possible to topcode total wage and salary earnings in a way that is completely consistent with the pre-1996 situation. The problem is with workers who earn less that $125,000 on their main job but have earnings from other jobs topcoded at $25,000. It is not possible to know whether total earnings of these workers are above or below $150,000. After some experiments, I decided to compute total earnings as the sum of main job earnings (censored at $150,000) and earnings on other jobs where I use the actual earnings provided in the CPS (where topcoded observations are imputed the actual mean earnings among topcoded workers).

For example, consider a full-time full-year white male who earns $90,000 on his main job but has his earnings topcoded at $25,000 for other jobs in the March 2001 CPS. I compute total earnings as the sum of $90,000 and $56,879 (see above), which yields $146,876. Since this is below the $150,000 topcode, I do not compute further

adjustments for this worker. By contrast, I would censor at $150,000 the total earnings of the same worker if he earned $100,000 instead of $90,000 on his main job (total of $156,876).

These adjustments likely have little impact since, in the March 1996-2003 CPS, less than one percent of workers have main job earnings below $125,000 and are topcoded on their other jobs earnings. Finally, once total wage and salary earnings have been censored in a consistent fashion, I multiply the earnings of workers at this consistent topcode by the standard 1.4 factor.

In both the May/ORG and March CPS, I also follow the existing literature by trimming very small and very large value of wages to remove potential outliers. Following Card and DiNardo (2002), I remove observations with an hourly wage of less than $1 or more than $100 in 1979 dollars. I also limit the analysis to workers age 16 to 64 with positive potential experience (age-education-6).

One last point about the ORG CPS is that, starting in 1994, workers are first asked what is the earnings periodicity (hourly, weekly, bi-weekly, annual, etc.) that they prefer to use in reporting their earnings on their current job. But as before, all workers paid by the hour are asked for their hourly wage rate. Hourly rated workers are asked this question even is "hourly" is not their preferred periodicity in the first question. Workers not paid by the hour are then asked to report their earnings for the periodicity of their choice. An hourly wage rate can again be computed by dividing earnings by usual hours of work over the relevant period. In 1994, The CPS also introduced "variables hours" as a possible answer for usual hours of work. I impute hours of work for these workers using a procedure suggested by Anne Polivka of the BLS.

APPENDIX B: Accounting for Composition Effects in the March CPS

Appendix Figures 4a (men) and 4b (women) compare the actual residual variance using the March CPS hourly wage rate to the residual variance that would have prevailed if the distribution of age and education had remained at its 1975 level. The re-weighting methodology used to compute the counterfactual variance is the same as for the May/ORG CPS (Figure 3). The figures show that the impact holding the distribution of characteristics constant is less dramatic in the March CPS than in the May/ORG CPS data.

Despite these differences, adjusting for composition effects still has a significant impact on the economic interpretation of the trends in the residual variance in the March CPS. In particular, the figures show essentially no growth in the residual variance after 1987-88 when the distribution of experience and education is held at its 1975 level. For women, the pattern of growth in residual inequality in the March CPS is similar to the one in the May/ORG CPS (with or without adjustments for composition effects). All the growth in residual inequality is concentrated in the first half of the 1980s. For men, the post-1980 growth in residual inequality also becomes qualitatively similar to the one in the May/ORG CPS. The only major discrepancy is that residual inequality grows rapidly in the March CPS during the 1970s, while it remains stable in the May/ORG data.

For reasons discussed in detail in Section 6, trends in residual wage inequality appear to be substantially biased up in the March CPS. One problem is that the wages of

hourly rated workers are particularly badly measured in the March CPS, and that the fraction of hourly rated workers has grown over time. The importance of composition effects can also be understated in the March CPS for the same reason. Remember that composition effects represent the difference between the actual residual variance and the counterfactual variance obtained by replacing the skill composition in the end period (say 2003) by the skill composition in the base period (say 1973). The counterfactual puts much more weight on less educated workers and less weight on more educated workers. This results in large composition effects in the May/ORG CPS because the within-group variance among highly-educated workers is much larger than among less-educated workers.

The difference in the within-group variance across education should be lower in the March than in the May/ORG CPS because the variance among less-educated workers is inflated by the larger fraction of those workers being paid by the hour. Consider, for example, the case of college post-graduates relative to high school dropout. In the early 2000s, between 80 and 90 percent of high-school dropouts are paid by the hour compared to just more than 10 percent among college postgraduates, a difference of about 70 percentage points (Figure 8). In light of the evidence in Figure 5, this suggests that the within-group variance of high-school dropouts in the March CPS is inflated by about 0.05 relative to college postgraduates.[37] This represents about a third of the 0.15-0.20 difference in the within-group variance between college post-graduates and high-school dropouts during the same period. Consistent with this prediction, a closer examination of the March data indeed indicates that the difference between the variance of these two groups is about 0.05 lower in the March than in the May/ORG CPS data.[38]

Because of this problem, the reweighting procedure should yield smaller composition effects when applied to the March CPS instead of the May/ORG CPS. Appendix Figure 4 indeed shows that composition effects in the March CPS are about a third smaller than in the May/ORG CPS.

---

[37] In Figure 5, the difference between the March and the May/ORG variance among workers paid by the hour is over 0.07 in the early 2000s. Multiplying this difference by the difference in the fraction of workers paid by the hour (about 0.70) yields about 0.05.
[38] For example, among men in 2000-01, the ORG CPS variance is 0.319 and 0.132 for college postgraduates and high school dropouts, respectively (difference of 0.187). The March CPS variance for the same groups in 2000-01 is 0.374 and 0.226 (difference of 0.146).

Table 1a: Within-group variance of wages by experience-education cell for men, 1973-75 and 2000-02

	Within-group variance			Workforce share		
	1973-75	2000-02	Change	1973-75	2000-02	Change
	(1)	(2)	(3)	(4)	(5)	(6)
A. By education and experience						
Dropout:						
1-10	0.118	0.083	-0.035*	0.065	0.035	-0.030
11-20	0.169	0.130	-0.038*	0.052	0.026	-0.026
21-30	0.170	0.154	-0.017*	0.055	0.025	-0.029
31+	0.180	0.162	-0.019*	0.123	0.028	-0.095
High school graduates:						
1-10	0.130	0.130	0.000	0.137	0.082	-0.055
11-20	0.145	0.181	0.035*	0.094	0.085	-0.009
21-30	0.162	0.196	0.034*	0.069	0.086	0.017
31+	0.188	0.217	0.029*	0.074	0.058	-0.016
Some college:						
1-10	0.143	0.152	0.008	0.076	0.077	0.001
11-20	0.173	0.204	0.031*	0.036	0.075	0.039
21-30	0.216	0.227	0.012	0.025	0.072	0.048
31+	0.245	0.256	0.011	0.020	0.046	0.026
College graduates:						
1-10	0.161	0.224	**0.064***	0.048	0.061	0.014
11-20	0.204	0.276	**0.072***	0.022	0.063	0.041
21-30	0.220	0.310	**0.091***	0.017	0.051	0.034
31+	0.299	0.332	0.033	0.009	0.024	0.015
Post-graduates:						
1-10	0.217	0.316	**0.099***	0.034	0.023	-0.010
11-20	0.324	0.324	0.000	0.023	0.033	0.009
21-30	0.327	0.302	-0.025	0.015	0.033	0.018
31+	0.420	0.369	-0.051	0.006	0.016	0.010
B. Weighted Average (using alternative shares)						
Actual shares	0.173	0.214	0.041			
1973-75 shares	0.173	0.185	0.012			
2000-02 shares	0.191	0.214	0.023			

Notes: "*" indicates that the change in the variance is significantly different from zero at the 95 percent confidence level. Changes that exceed 0.04 are highlighted (bold). The standard errors of the estimated variances are 0.011, on average, in column 1, and 0.005, on average, in column 2. The standard errors range from 0.0025 (high school graduates with 1-10 years of experience) to 0.035 (college post-graduates with 31-40 years of experience) in column 1, and from 0.0023 (high school graduates with 1-10 years of experience) to 0.011 (college post-graduates with 31-40 years of experience) in column 1

Table 1b: Within-group variance of wages by experience-education cell for women, 1973-75 and 2000-02

	Within-group variance			Workforce share		
	1973-75	2000-02	Change	1973-75	2000-02	Change
	(1)	(2)	(3)	(4)	(5)	(6)
A. By education and experience						
Dropout:						
1-10	0.099	0.056	-0.043*	0.057	0.026	-0.031
11-20	0.130	0.090	-0.040*	0.039	0.015	-0.024
21-30	0.125	0.106	-0.019*	0.050	0.018	-0.032
31+	0.139	0.123	-0.017*	0.103	0.023	-0.080
High school graduates:						
1-10	0.106	0.108	0.002	0.179	0.070	-0.109
11-20	0.145	0.157	0.011*	0.095	0.072	-0.023
21-30	0.144	0.172	0.028*	0.092	0.086	-0.006
31+	0.162	0.178	0.016*	0.097	0.074	-0.023
Some college:						
1-10	0.118	0.137	0.019*	0.077	0.091	0.014
11-20	0.134	0.198	**0.065***	0.025	0.081	0.057
21-30	0.152	0.209	**0.057***	0.020	0.084	0.064
31+	0.160	0.220	**0.060***	0.020	0.054	0.034
College graduates:						
1-10	0.134	0.179	**0.045***	0.055	0.076	0.020
11-20	0.170	0.260	**0.090***	0.015	0.058	0.043
21-30	0.173	0.262	**0.088***	0.014	0.052	0.038
31+	0.195	0.254	**0.059***	0.010	0.021	0.010
College post-graduates						
1-10	0.154	0.239	**0.085***	0.022	0.026	0.004
11-20	0.238	0.259	0.021	0.012	0.027	0.015
21-30	0.204	0.217	0.013	0.011	0.034	0.023
31+	0.280	0.234	-0.046	0.006	0.013	0.007
B. Weighted Average (using alternative shares)						
Actual shares	0.136	0.183	0.047			
1973-75 shares	0.136	0.148	0.012			
2000-02 shares	0.149	0.183	0.034			

Notes: "*" indicates that the change in the variance is significantly different from zero at the 95 percent

confidence level. Changes that exceed 0.04 are highlighted (bold). The standard errors of the estimated variances are 0.010, on average, in column 1, and 0.004, on average, in column 2. The standard errors range from 0.0025 (high school graduates with 1-10 years of experience) to 0.035 (college post-graduates with 31-40 years of experience) in column 1, and from 0.0018 (high school dropouts with 1-10 years of experience) to 0.009 (college post-graduates with 31-40 years of experience) in column 1.

Table 2: Composition Effects and Changes in the Residual
Variance of Log Hourly Wages, May/ORG CPS

	1973-1979	1979-1989	1989-1999	1999-2003	1973-2003
A. Men					
Residual variance:					
Actual change	-0.003	0.036	-0.003	0.017	0.047 [43]
1973 skills distribution	-0.003	0.027	-0.019	0.011	0.015 [14]
2003 skills distribution	-0.008	0.034	-0.013	0.012	0.025 [23]
Total variance:	-0.002	0.080	0.007	0.024	0.109 [100]
B. Women					
Residual variance:					
Actual change	-0.014	0.047	-0.001	0.013	0.045 [46]
1973 skills distribution	-0.017	0.036	-0.019	0.005	0.005 [5]
2003 skills distribution	-0.012	0.040	-0.006	0.008	0.030 [31]
Total variance:	-0.026	0.092	0.017	0.015	0.098 [100]
C. Real value of the minimum wage (logs)					
	0.103	-0.391	0.135	-0.099	-0.252

Note: Numbers in square brackets represents the percentage of the
1973-2003 change in the total variance of wages (both within- and
between-group components) that is attributable to this variance
component.

Table 3: Estimates of measurement error in the May/ORG and March CPS

	Men		Women	
	May/ORG	March	May/ORG	March
1. Average measurement error variance (1976-2003)[1]				
a. Paid by the hour	0.017 [8.4%]	0.087 [33.2%]	0.024 [14.3%]	0.077 [35.0%]
b. Not paid by the hour	0.052 [16.8%]	0.065 [20.4%]	0.045 [19.4%]	0.054 [23.2%]
2. 1976-2003 Change in measurement error variance				
a. Paid by the hour	0.006	0.020	0.011	0.016
b. Not paid by the hour	0.000	0.017	0.015	0.011
3. Spurious Change in variance due to:				
a. Growth in fraction of hourly workers[2]	-0.004	0.002	-0.003	0.003
b. Growth in measurement error variance[3]	0.004	0.018	0.013	0.014
c. Total (3a.+3b.)	0.000	0.020	0.010	0.017
4. 1976-2003 change in residual variance	0.046	0.079	0.057	0.074
5. Change adjusted for measurement error (4.-3c.)	0.046	0.059	0.047	0.057

Notes: Measurement error estimated using the matched March-May/ORG sample. See text for details.
1: Numbers in square brackets represents the percentage of the overall variance of wages due to measurement error.
2: Based on Figure 8, it is assumed that the growth in the fraction of workers paid by the hour is 10% for men and 15% for women. These proportions are then multiplied by the difference in the estimated measurement error variances for hourly (row 1a) and non-hourly (row 1b) workers.
3: Change in the weighted average of the measurement error variances for hourly and non-hourly workers.

Appendix Table 1: Percentage distribution of workers by education and experience groups, May/ORG CPS

	Men				Women			
	1973-74	1980	1990	2003	1973-74	1980	1990	2003
A. Education categories								
High School Dropout	30.4	23.0	15.9	11.4	25.7	17.5	11.4	7.4
High School Graduate	37.4	37.9	38.1	31.0	46.3	46.0	41.5	29.3
Some college	15.3	18.1	20.4	26.6	13.7	18.7	23.2	30.8
Bachelors' Degree	9.1	11.6	14.8	20.4	9.3	11.0	14.8	21.5
Post-graduate Degree	7.7	9.4	10.9	10.6	5.0	6.9	9.2	10.9
B. Years of Experience								
0-10	35.8	39.4	31.9	26.6	38.5	41.4	33.8	27.7
11-20	22.7	24.5	32.8	28.2	18.5	22.8	29.5	24.4
21-30	18.2	16.4	19.5	26.3	19.1	16.6	21.0	27.3
31+	23.3	19.7	15.8	19.0	23.9	19.3	15.7	20.7

Figure 1: Between-group variance and composition-adjusted (using 1973 characteristics) residual variance for men, May/ORG CPS
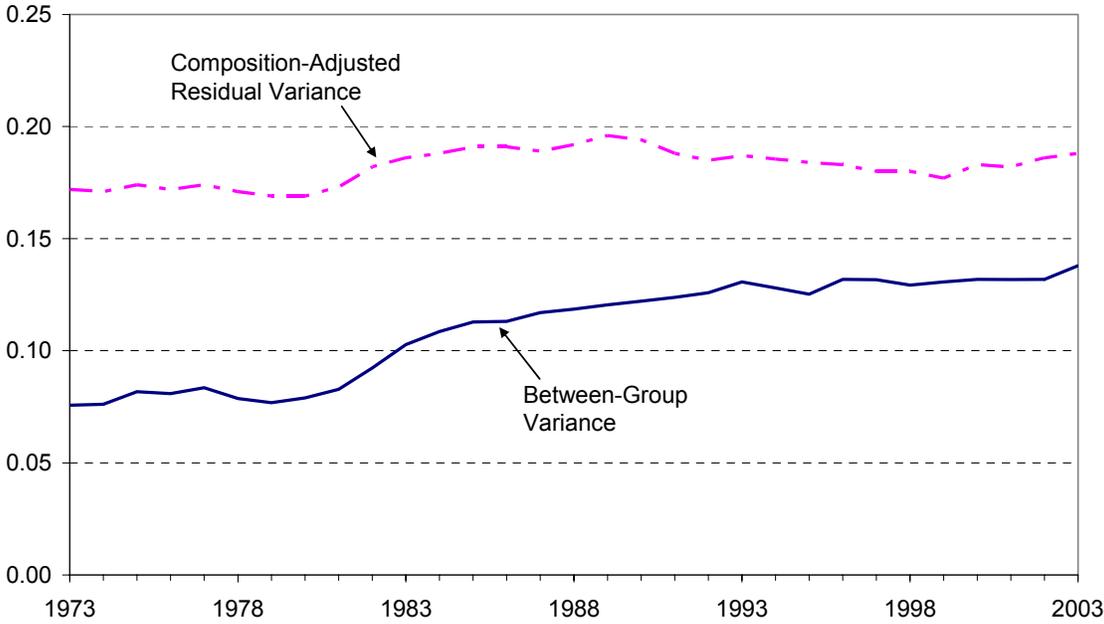
Figure 2a: Within-group variance by education group for men,
(average of the four experience groups)

HS Dropout — HS graduate — Some college — College — Post-graduate
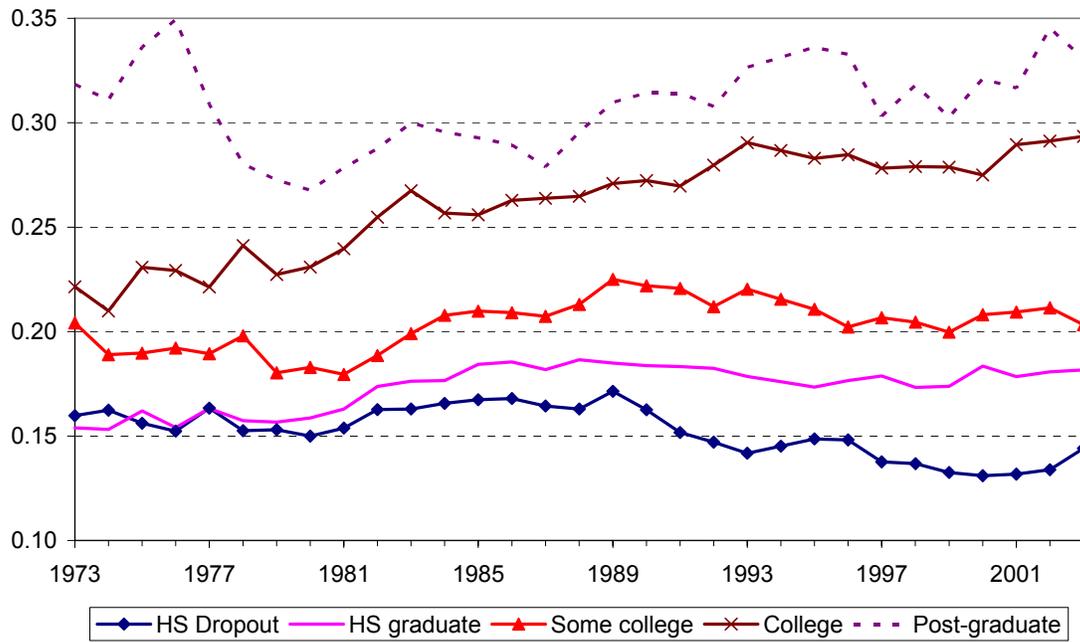
Figure 2b: Within-group variance by education group for women
(average of the four experience groups)

HS Dropout — HS graduate — Some college — College — Post-graduate
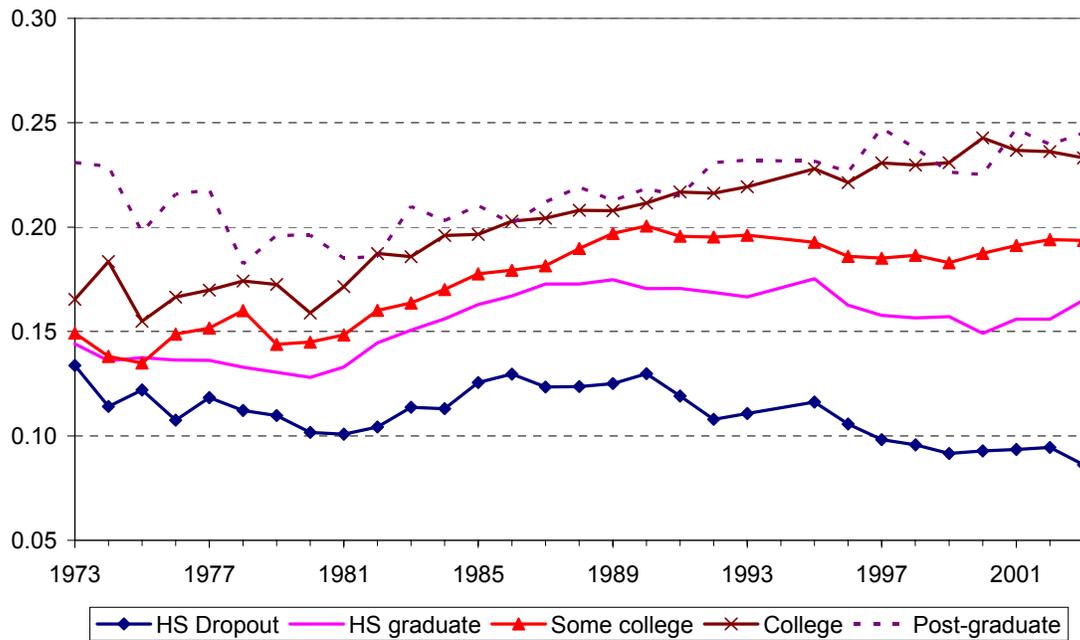
Figure 3a: Actual and counterfactual residual variance of wages for men, 1973 to 2003
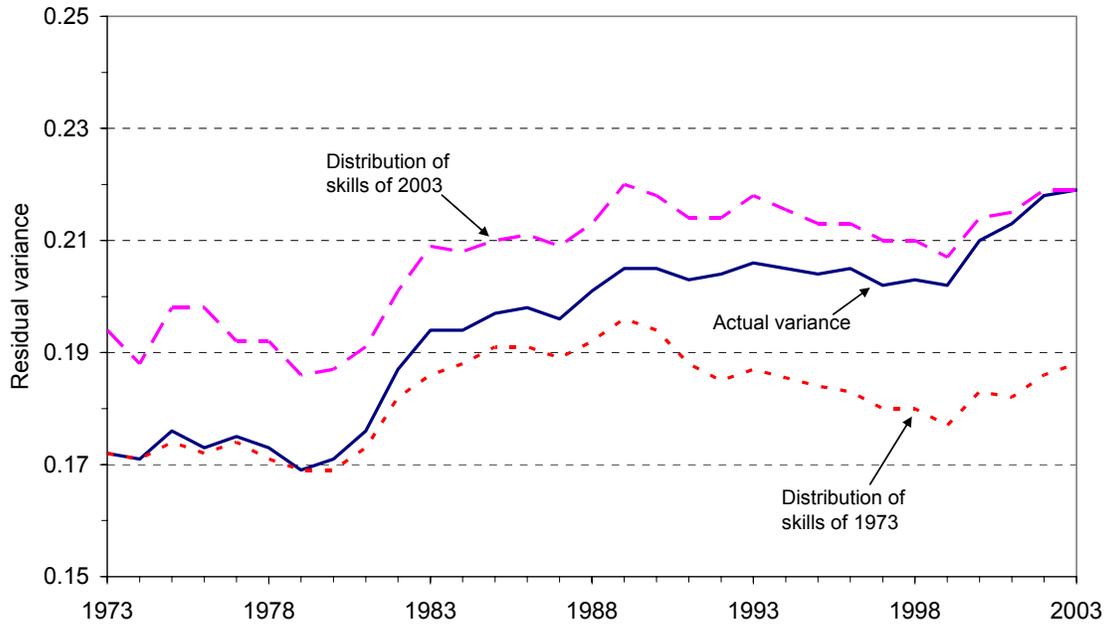
Figure 3b: Actual and counterfactual residual variance of wages for women, 1973 to 2003
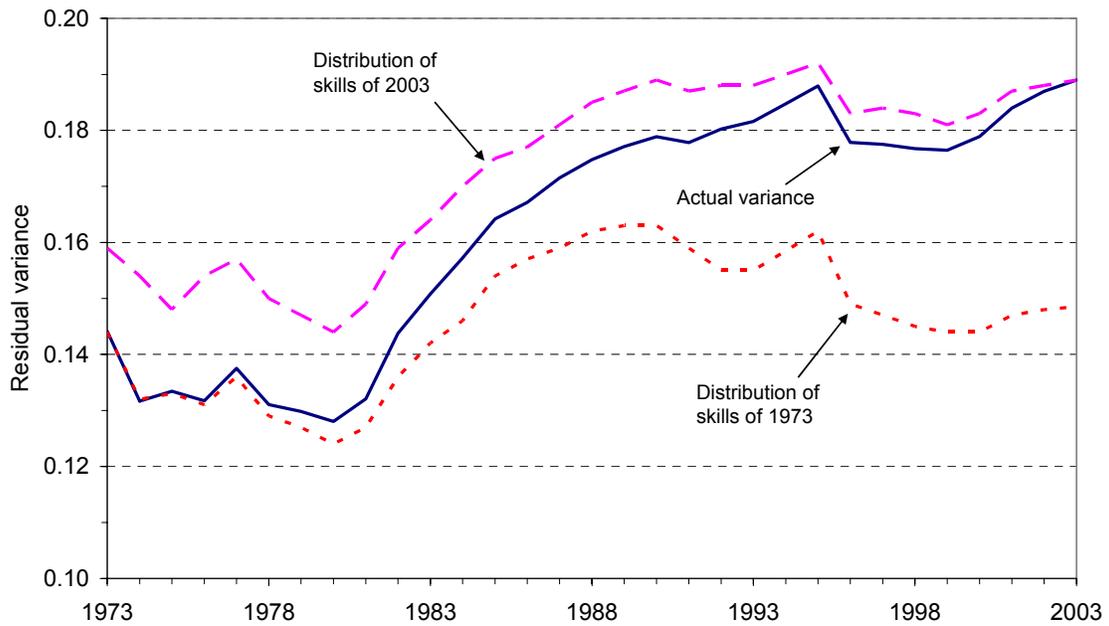
Figure 4a: Residual 90-10 wage gap for men, holding distribution of skills at their 1973 level
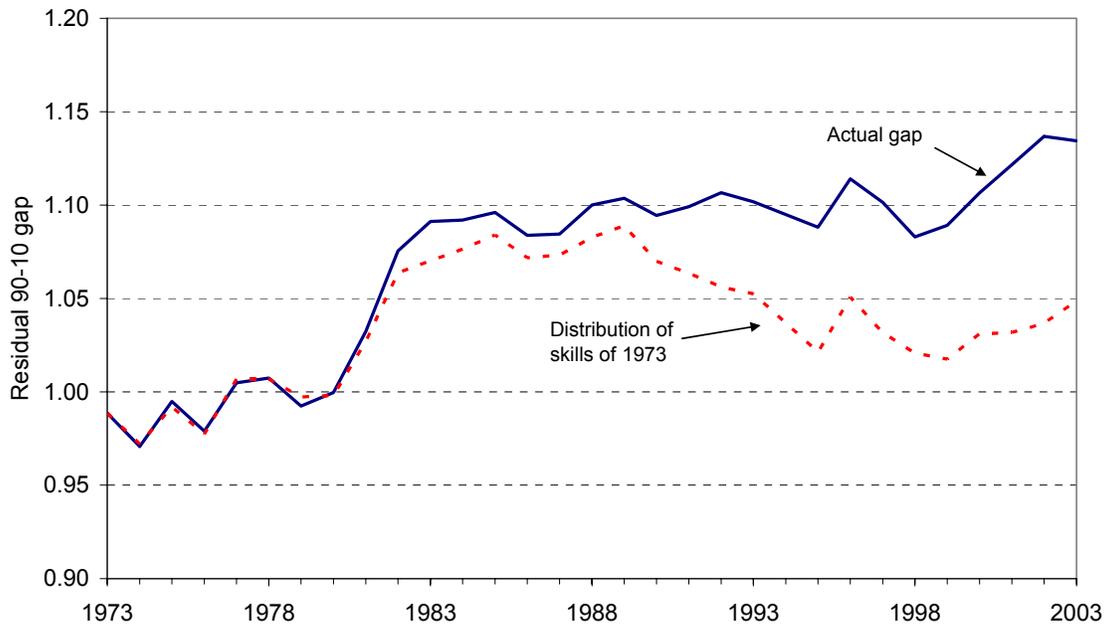
Figure 4b: Residual 90-10 gap for women, holding distribution of skills at their 1973 level
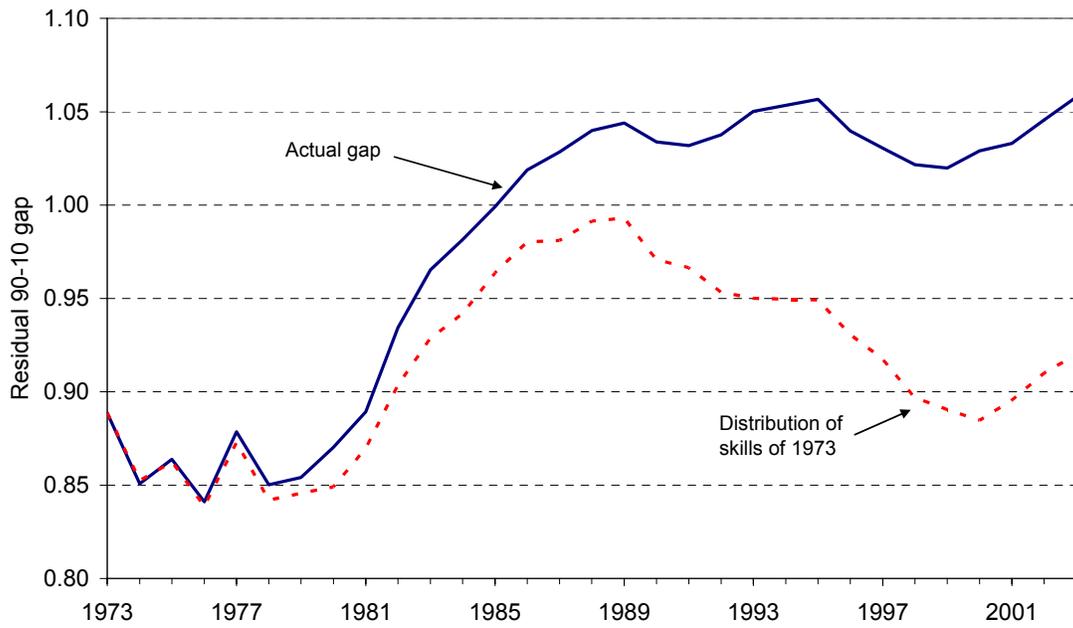
Figure 5a: Variance of log hourly wages of men with both May/ORG and March wages (matched sample)

Non hourly, March

Non hourly, May/ORG

Hourly, March

Hourly, May/ORG
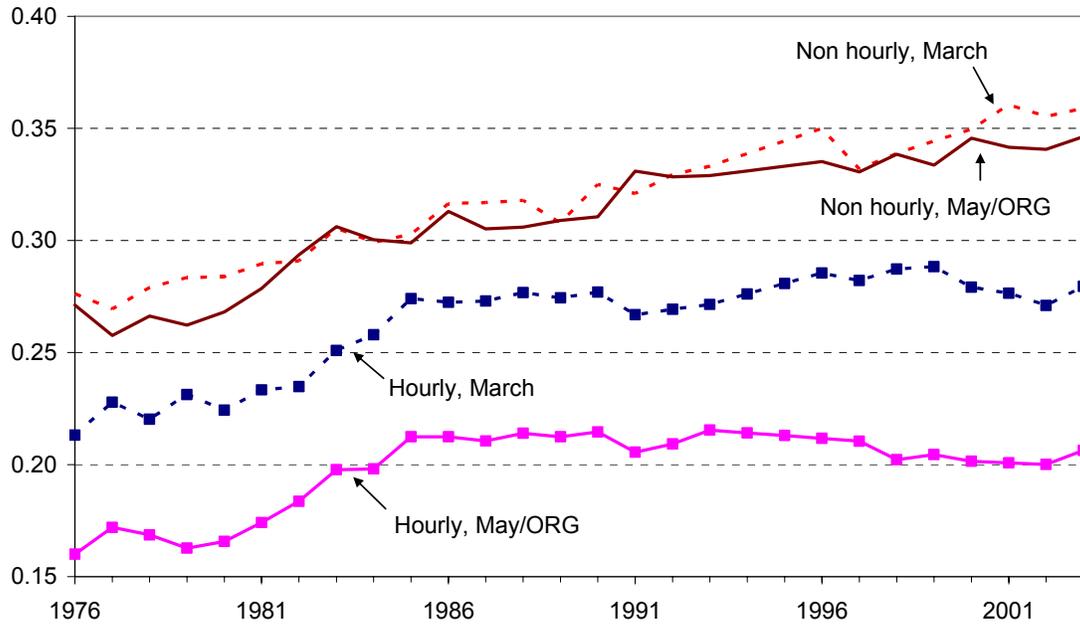
Figure 5b: Variance of log wages of women with both May/ORG and March wages (matched sample)

Non hourly, May/ORG

Non hourly, March

Hourly, March

Hourly, May/ORG
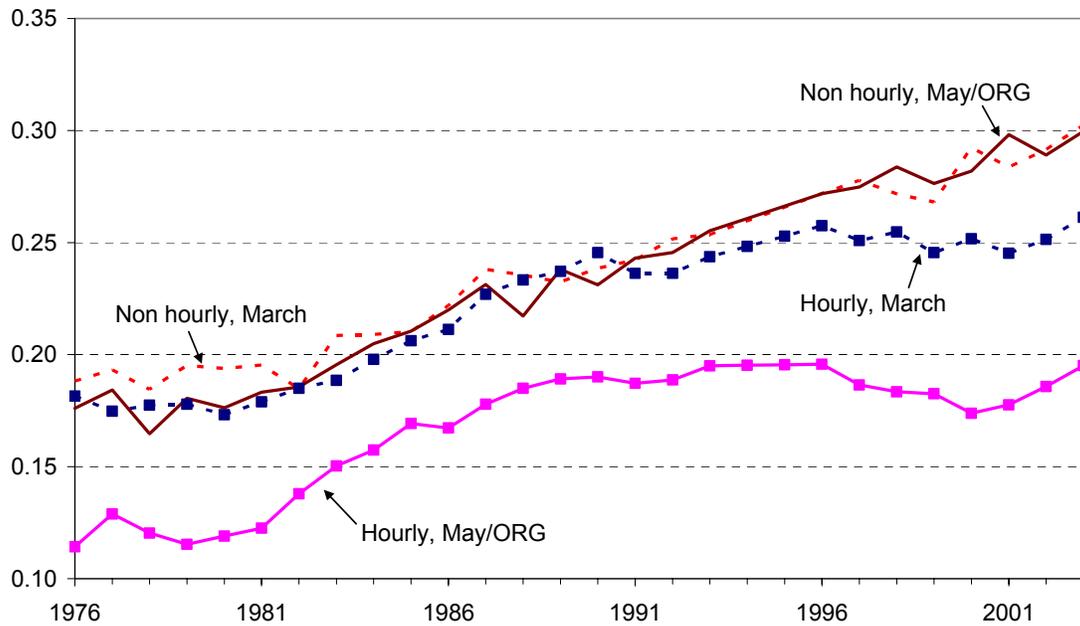
Figure 6a: Between-group variance, men
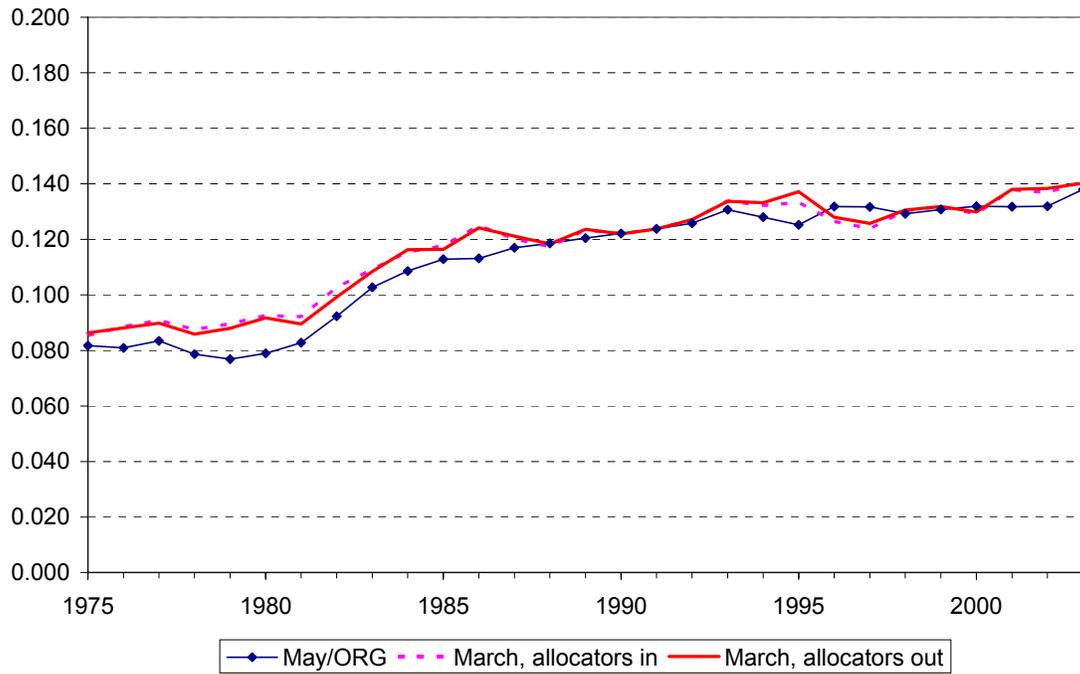
Figure 6b: Between-group variance, women
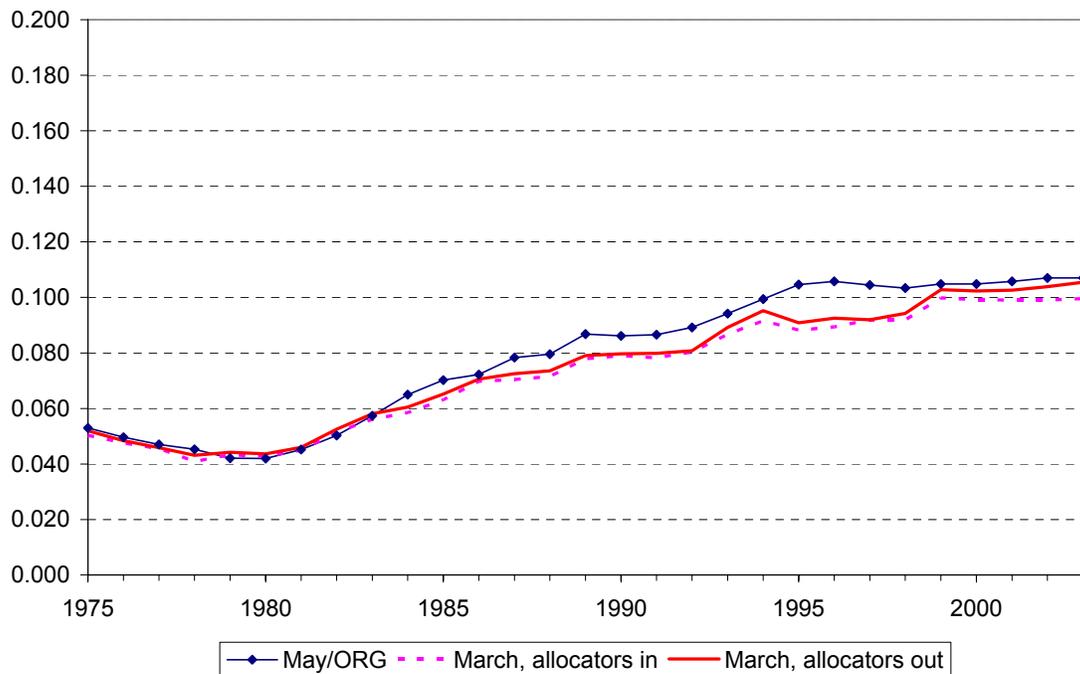
Figure 7a: Residual variance of wages, men

May/ORG ----- March, allocators in ——— March, allocators out
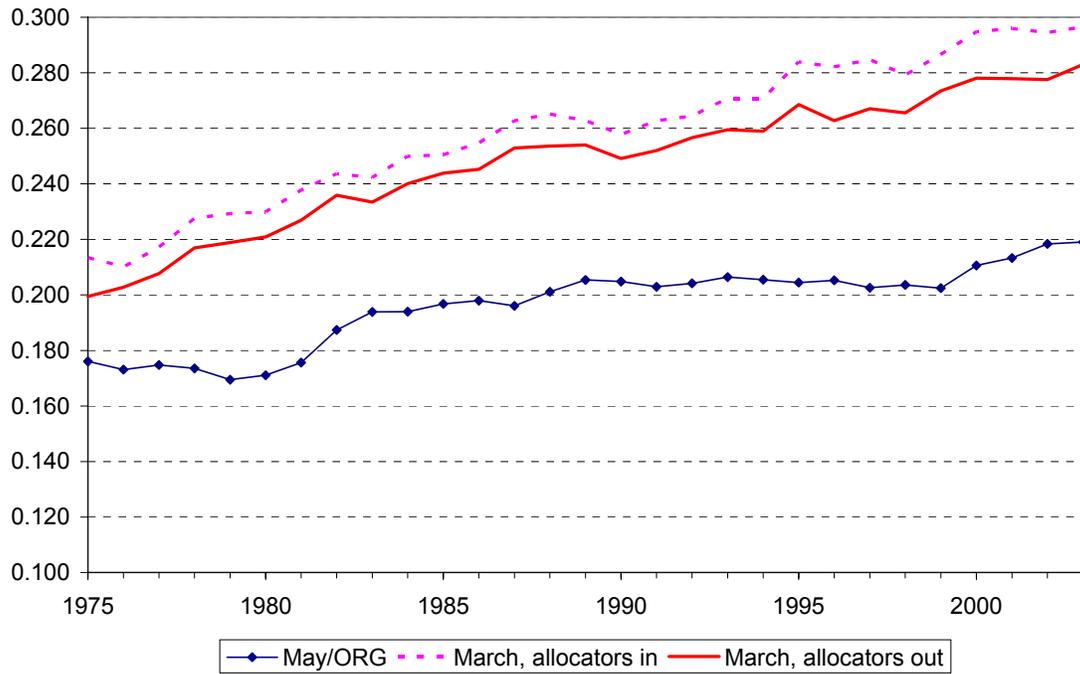
Figure 7b: Residual variance of wages, women

May/ORG ----- March, allocators in ——— March, allocators out
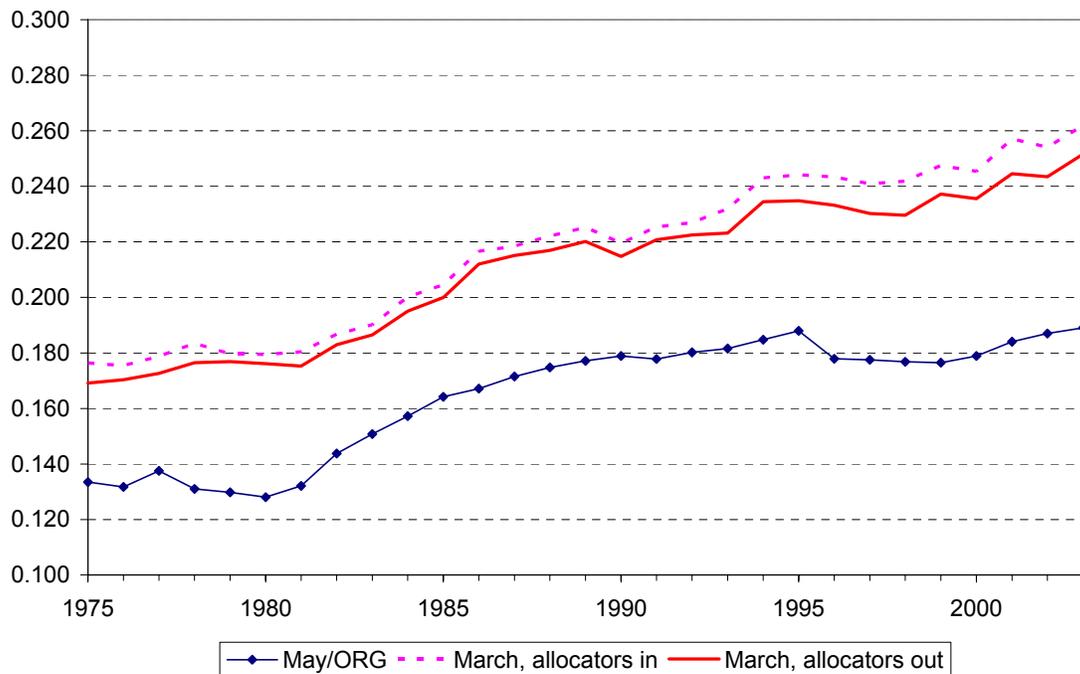
Figure 8a: Fraction of men paid by the hour, by education category

Legend: Dropout · High School · Some college · College · Post-graduate
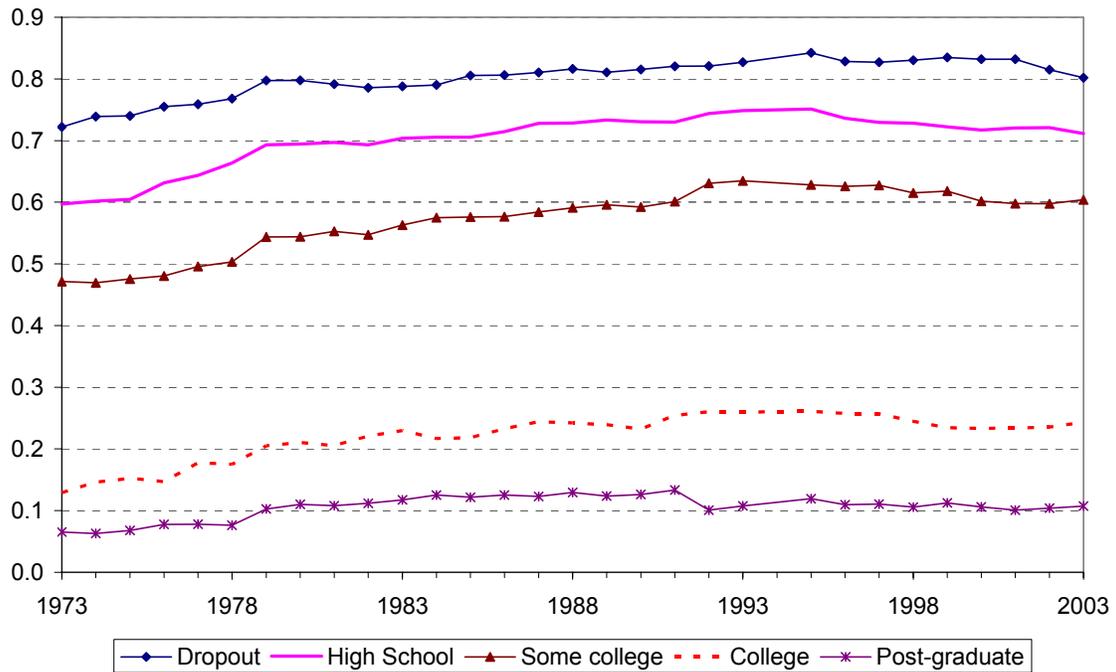
Figure 8b: Fraction of women paid by the hour, by education category

Legend: Dropout · High School · Some college · College · Post-graduate
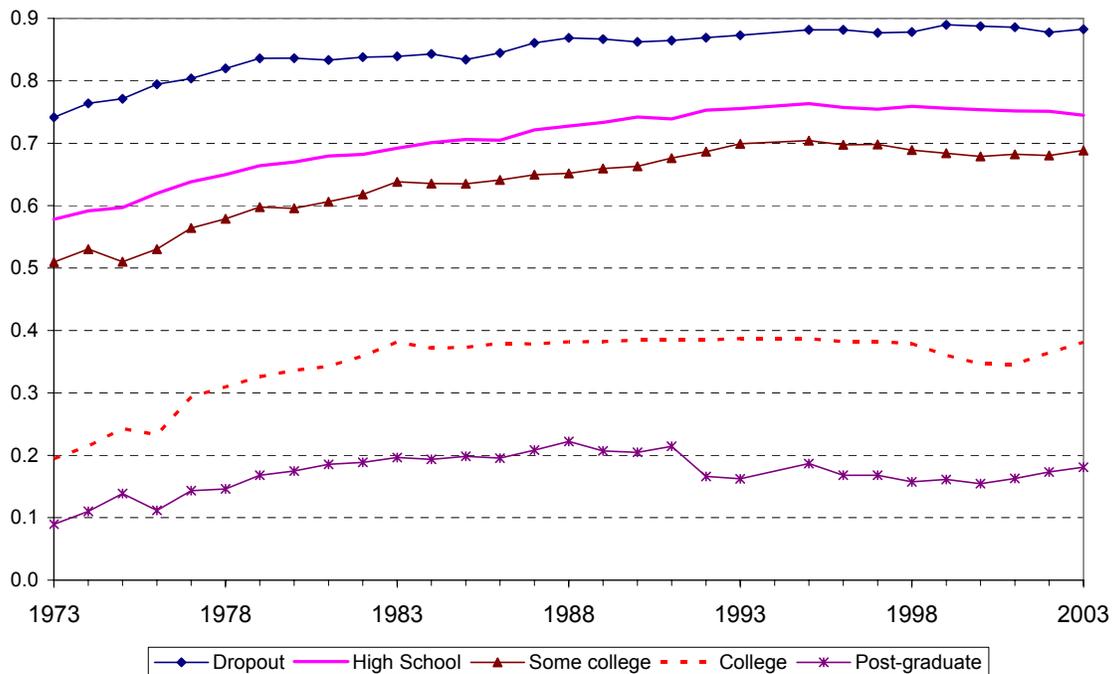
**Figure 9a: Male residual variance predicted using the minimum wage
(holding characteristics at their 1973 level)**
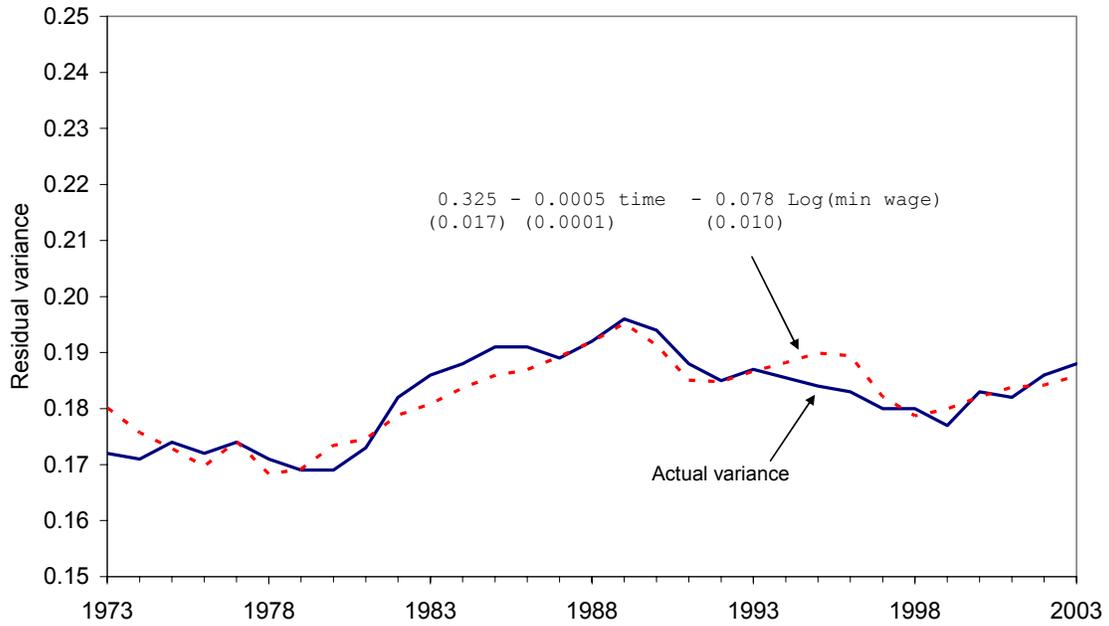
**Figure 9b: Female residual variance predicted using the minimum
wage (holding characteristics at their 1973 level)**
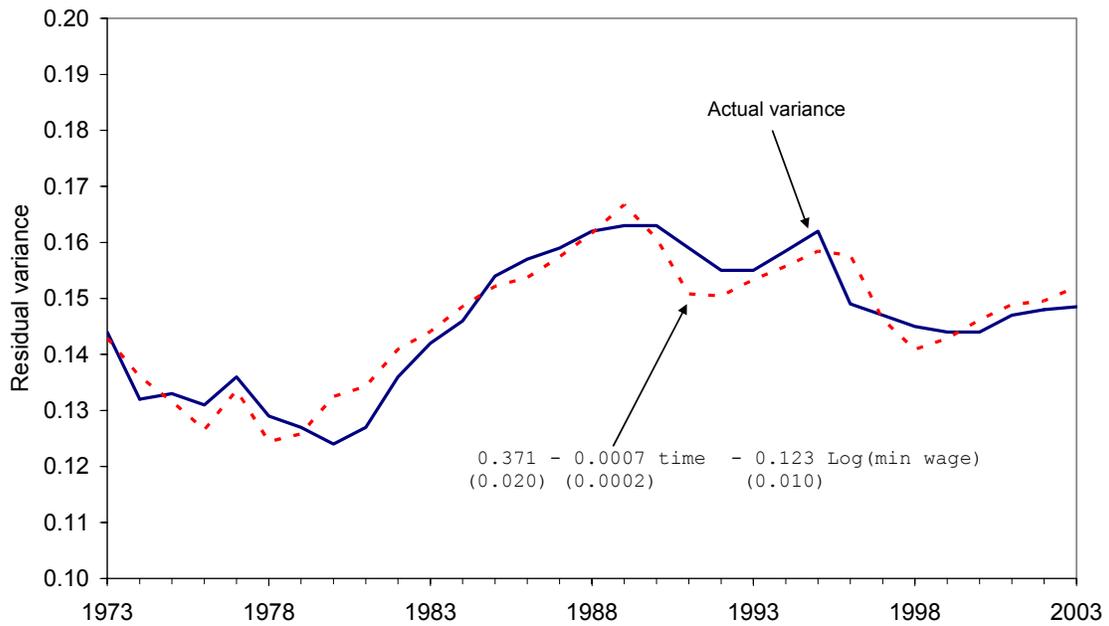
Appendix Figure 1a: Residual 50-10 wage gap for men, holding distribution of skills at their 1973 level
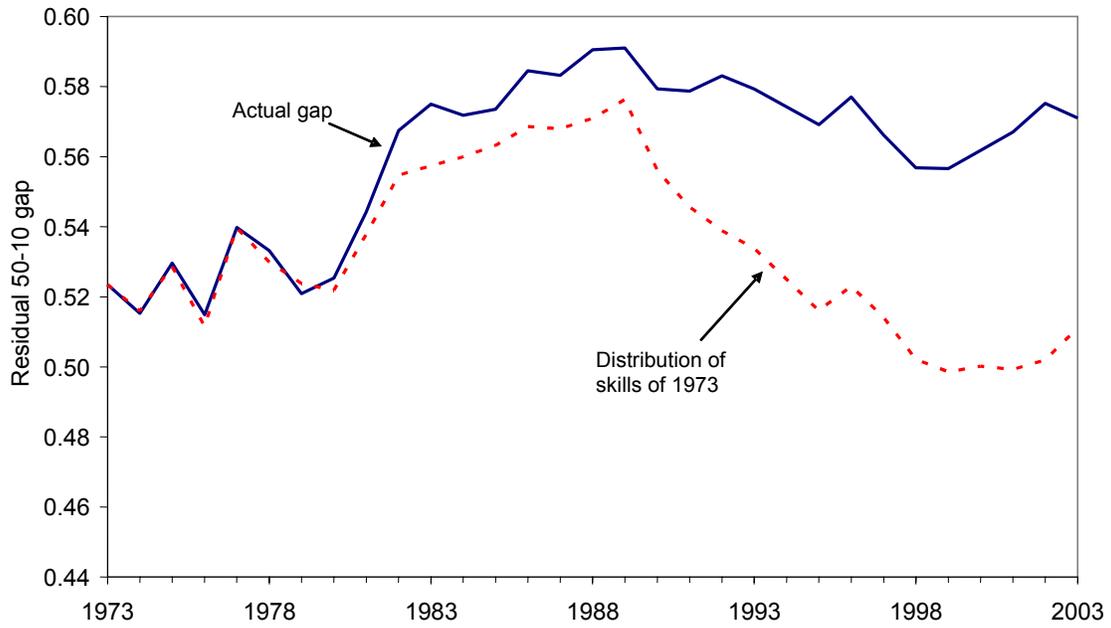

Appendix Figure 1b: Residual 90-50 wage gap for men, holding distribution of skills at their 1973 level
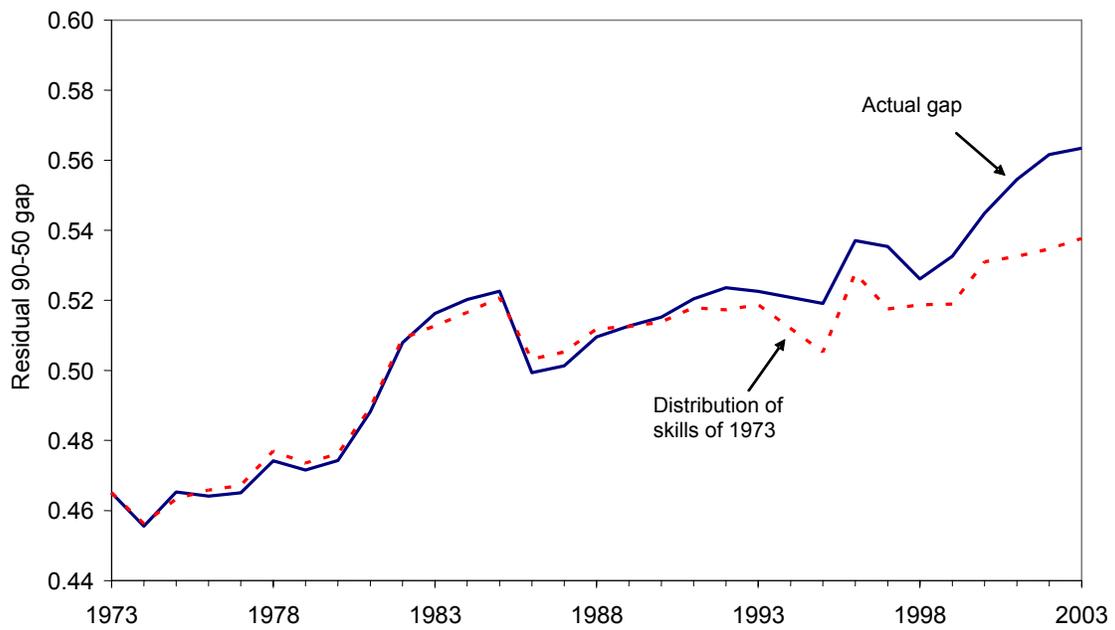
Appendix Figure 1c: Residual 50-10 wage gap for women, holding distribution of skills at their 1973 level
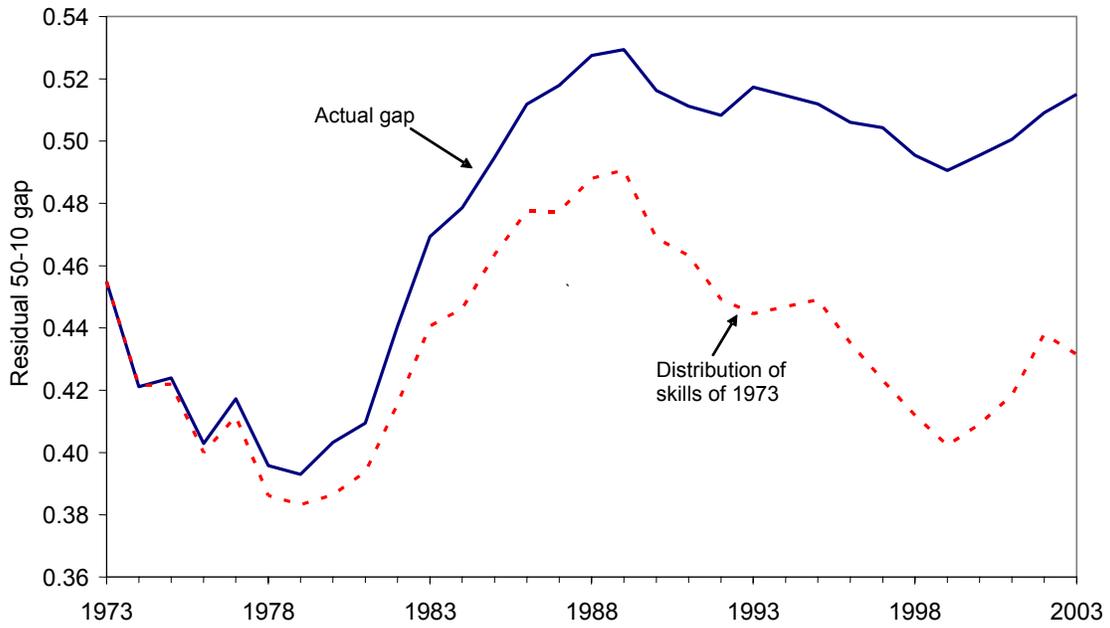

Appendix Figure 1d: Residual 90-50 wage gap for women, holding distribution of skills at their 1973 level
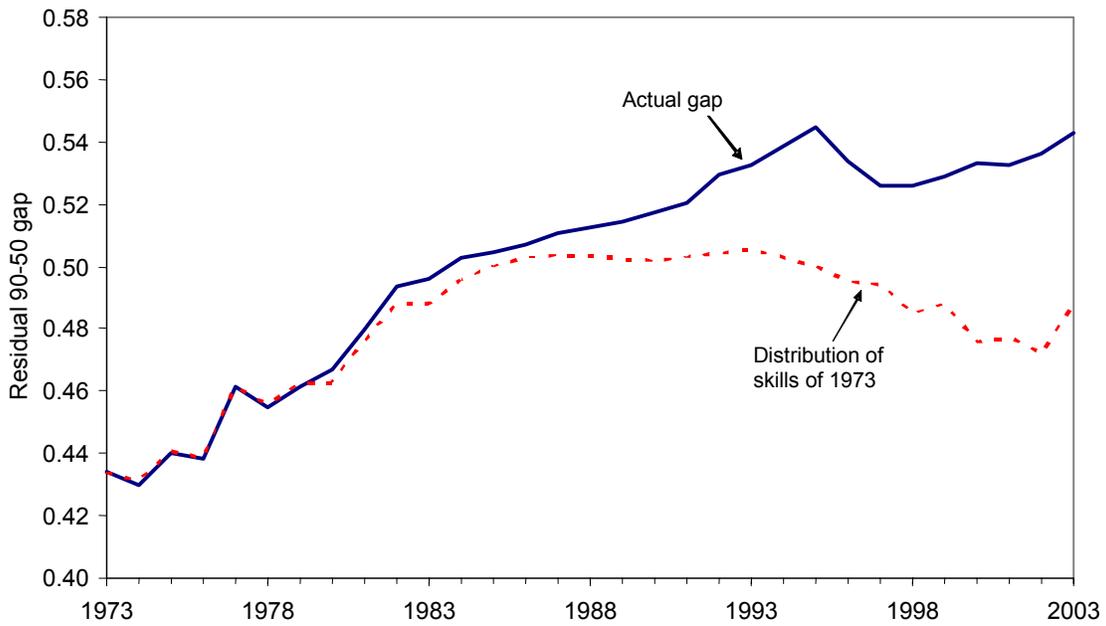
Appendix Figure 2a: Variance of log hourly wages for all workers and matched workers only, men


March, all

March, matched

May/ORG, all
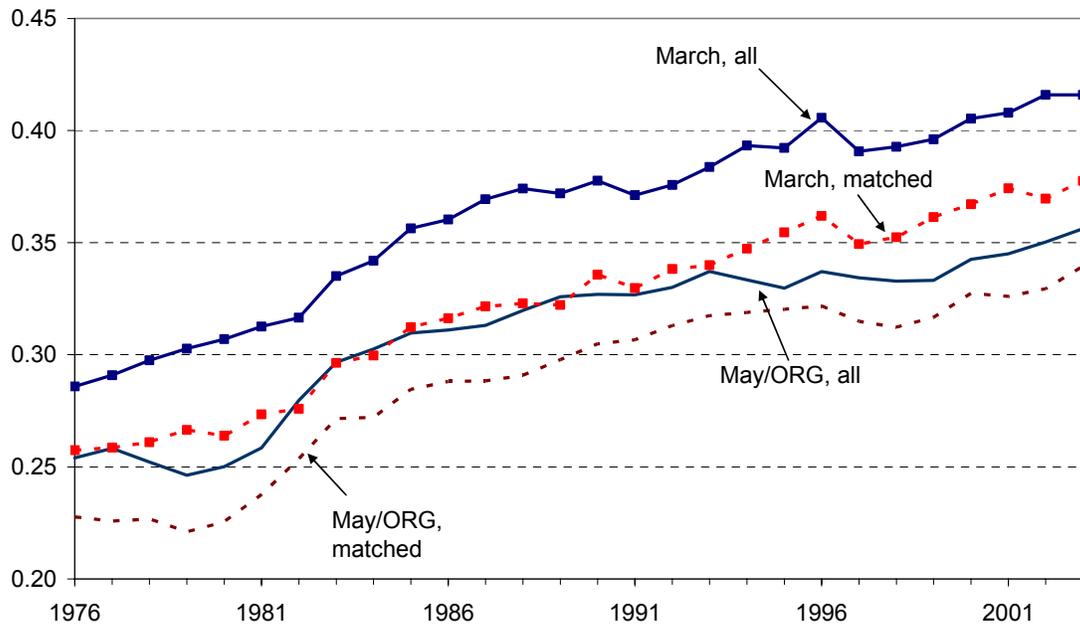
May/ORG, matched

Appendix Figure 2b: Variance of log hourly wages for all workers and matched workers only, women

March, all

March, matched

May/ORG, all

May/ORG, matched
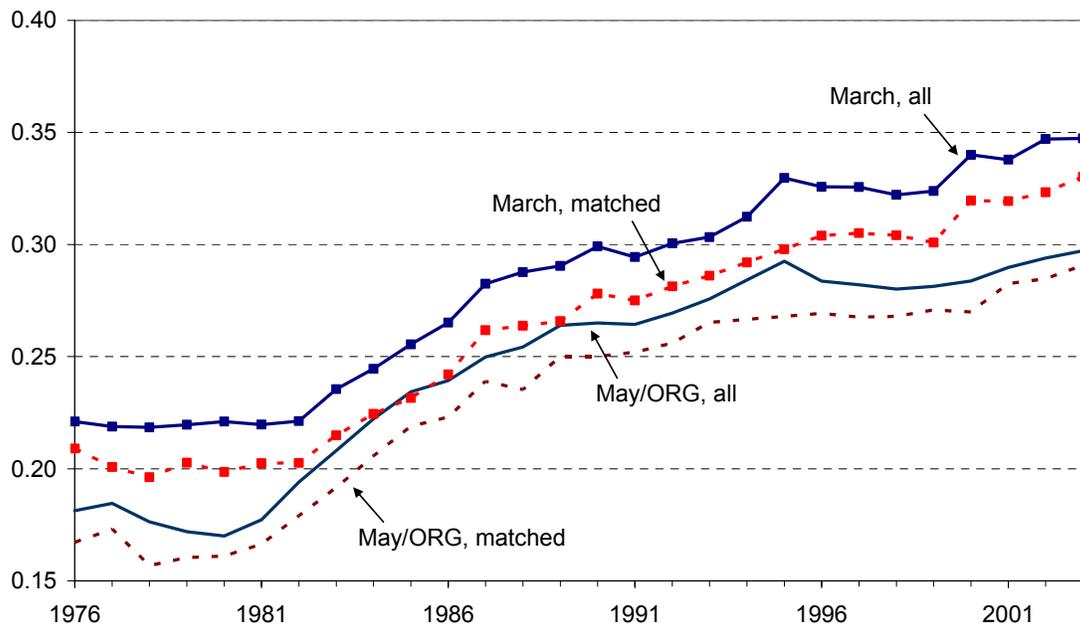
Appendix Figure 3a: Estimated measurement error variances in the March and May/ORG CPS, men

March, hourly

March, non-hourly

May/ORG, non-hourly

May/ORG, hourly
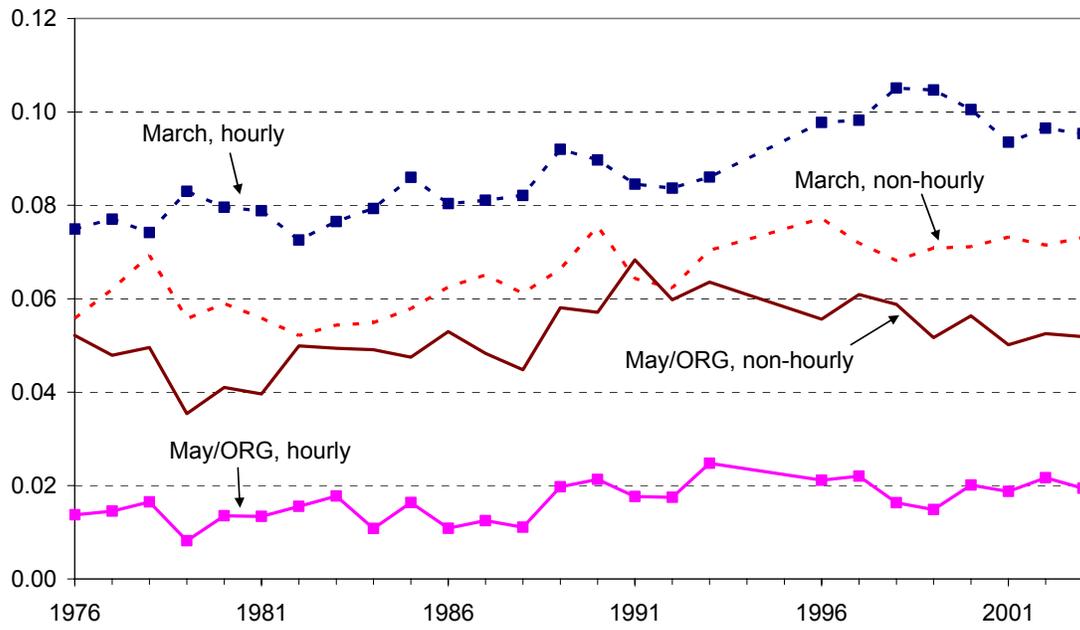
Appendix Figure 3b: Estimated measurement error variances in the March and May/ORG CPS, women

March, hourly

March, non-hourly

May/ORG, non-hourly

May/ORG, hourly
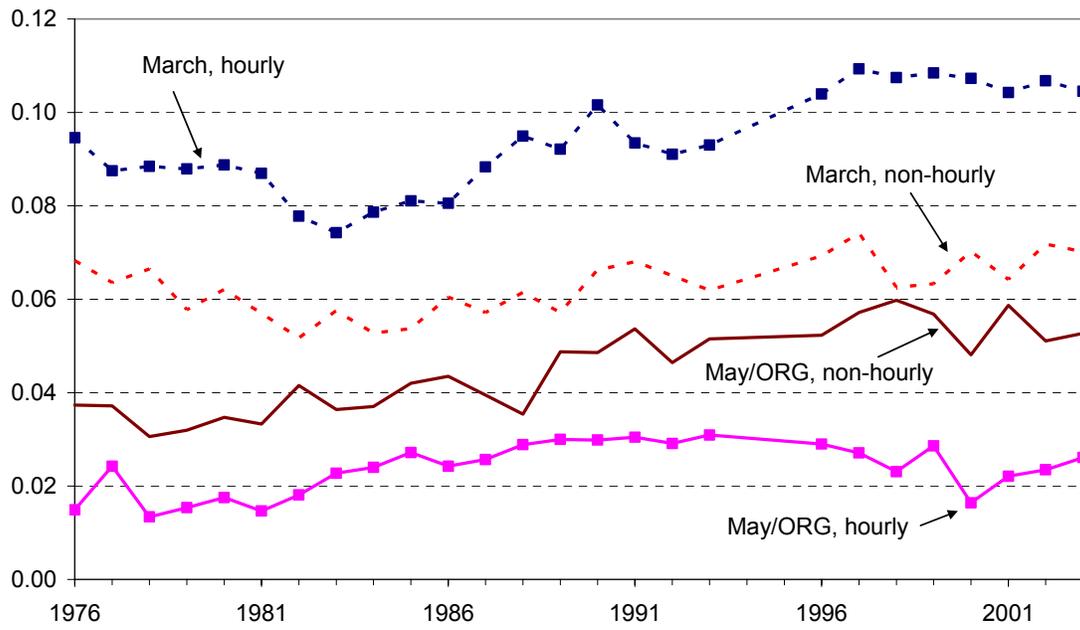
Appendix Figure 4a: Residual variance for men in the March CPS, holding distribution of skills at their 1975 level

Residual variance

Actual variance

Distribution of skills of 1975
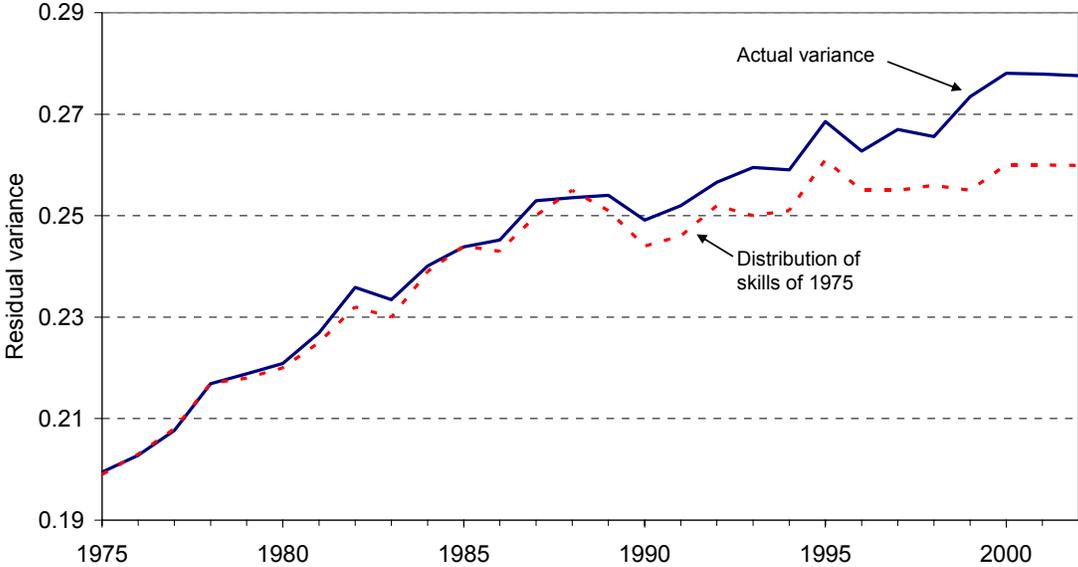
Appendix Figure 4b: Residual variance for women in the March CPS, holding distribution of skills at their 1975 level

Residual variance

Actual variance

Distribution of skills of 1975