Wage Decompositions and Sample Selection: Evidence for the United States

Iván Fernández-Val[*] Franco Peracchi[†] Aico van Vuuren[‡]
Francis Vella[§]

February 1, 2018

Abstract

We analyze the patterns and determinants of female and male hourly real wage growth for 1976 to 2016 at various quantiles of the wage distribution. Employing non-parametric estimation methods, which account for the presence of selection bias we decompose wage growth at different quantiles into structural, composition and selection effects. Male real wages at the median and below have decreased despite an increasing skill premium and an increase in the educational attainment levels. This is primarily due to large decreases of the wages of the lowly educated. Wages at the upper quantiles of the distribution have increased drastically due to a large and increasing skill premium. For females the wage declines are less dramatic and occur at only the lower quantiles. The increases at the upper quantiles are also substantial and reflect increasing skill premia. These changes have resulted in a substantial increase in wage inequality. Finally, the increasing participating rate of women has enlarged female wage inequality.

[*]Boston University
[†]Georgetown University
[‡]Gothenburg University
[§]Georgetown University

1 Introduction

The dramatic increase in earnings inequality in the United States is an intensely studied phenomenon (see, for example, Katz and Murphy 1992, Murphy and Welch 1992, Juhn, Murphy and Pierce 1993, Katz and Autor 1999, Lee 1999, Lemieux 2006, Autor, Katz and Kearney 2008, Acemoglu and Autor 2011, and Murphy and Topel 2016). While contributions to this literature employ a variety of measures of earnings and inequality and utilize different data sets, there is a general agreement that earnings inequality has increased drastically since the early 1980s. This has alarmed economists and social commentators as substantial, and increasing inequality is associated with a variety of economic and social concerns. While the finding that earnings inequality has risen is robust to the measure of earnings employed we highlight that an individual's annual earnings reflect both the number of hours that a person works and the average hourly wage rate received for those hours. Thus understanding the determinants of each of these two factors is important for uncovering the sources of earnings inequality. The focus of this paper is to examine the evolution and determinants of hourly wage growth of males and females in the United States over the period 1976 to 2016.

Figure 1 presents the profiles of selected quantiles of real hourly wage rates for males and females using data from the U.S. Current Population Survey (CPS) for the survey years from 1976 to 2016. They are based on individuals aged 24 to 65 years at the survey date who worked a positive number of hours the previous year and are neither in the Armed Forces nor self-employed.[1] The profiles differ greatly by gender and by location in the wage distribution. While they each generally display some relationship with the business cycle the overall trends are different. For example, consider the median male wage rate. It increases slightly between 1987 and 1990 but then decreases steadily until 1994, reaching a minimum of 17.6 percentage points below its initial level. It rebounds between 1994 and 2002, but falls between 2007 and 2013. In 2016 it is 13.6 percentage points below its 1976 level. The female median wage, despite some small occasional dips, increases by nearly 25 percentage points over the sample period. While the decrease in the male median wage is concerning, an examination of the male wage profile at the 25th quantile is even more alarming.

[1] Section 2 provides a detailed discussion of the data employed in our analysis and our sample selection process.

This shows the same cyclical behavior as the median but a greater decline, especially between 1976 and 1994, and in 2016 it remains 18.2 percentage points below its 1976 level. In contrast, despite some cyclical fluctuations, the female wage at the 25th quantile increases by about 17 percentage points over our period. The profiles at the 10th quantile provide a similar story to those at the 25th and 50th quantiles.

Descriptions of wage inequality involve contrasts between the lower and upper parts of the wage distribution and thus we also examine the upper part of the wage distribution. At the 75th quantile there is a small increase in the male real wage over the whole period although as recently as the late 1990's it was below the 1976 level. However, the past 15 years have seen a notable increase. The profile at the 90th quantile for males resembles the 75th, although the periods of growth have produced larger jumps. For females there has been strong and steady growth at both the 75th and 90th quantiles since 1980 with an increasing gap appearing between each of these profiles and the median wage. Figure 1 suggests that the documented increases in inequality capture low or negative growth for wages below the median combined with extraordinary wage increases in the upper part of the distribution. The period of drastic growth differs by gender.

These respective wage profiles suggest that conventional measures of inequality, such as the Gini coefficient or 90th/10th quantile ratios, will indicate drastically increasing inequality for both males and females. This is confirmed by Figure 2 which reports the time series behavior of the 90th/10th quantile ratios of wages. For both males and females there is a widening gap with increases over our sample period of 54.6 and 37.9 percent respectively. However, rather than immediately focusing on these ratios we consider it insightful to first examine what is driving the wage profiles. Figure 1 suggests that they are either influenced differently by the same market factors and/or, alternatively, different factors. For example, the cyclicality of the profiles suggests they each respond to business cycle forces, although the strength of this relationship varies by location in the wage distribution. However, the behavior of the profiles at the upper and lower end of the wage distribution suggest that workers in these associated labor markets are undergoing contrasting experiences. Understanding what is generating these trends is the goal of this paper.

Our investigation is guided by previous empirical work which has identified the possible determinants of the growth in wage inequality. First, the labor economics literature (see for example, Juhn, Murphy and Pierce 1993, Katz and Autor 1999, Autor, Katz and Kearney 2008, Acemoglu and Autor 2011, and Murphy and Topel 2016) has highlighted the role of increasing skill premia and increasing returns to higher education. The effect by which the prices of individual's characteristics contribute to wages is known as the "structural effect". This may also capture factors such as the declining real minimum wage and the reduction in the bargaining power due to lower unionization rates of the workforce (see, for example, DiNardo, Fortin and Lemieux 1996 and Lee 1999) noting that these considerations are likely to influence the lower part of the wage distribution. Second, the nature of the workforce has changed drastically over the past 41 years suggesting that workers' characteristics have contributed to wage movements. This is primarily reflected by the increases in educational attainment. However, the large changes in the labor force participation rates of females may have also produced important changes in the composition of the labor force. The movements in hourly wages rates attributable to these changing characteristics are known as the "composition effects".

Recent work (see, for example, Angrist, Chernozhukov and Fernàndez-Val, 2006 and Chernozhukov, Fernández-Val and Melly, 2013) has estimated these structural and composition effects in general forms and under general conditions. However, they typically have not incorporated a role for selection bias resulting from the non randomness of the employment outcome (see Heckman, 1974, 1979). Thus, they have not allowed for wage movements capturing the changing nature, in terms of unobservables, of the workforce. This "selection effect", is potentially important as the drastic changes over this sample period in the participation rates and in the average number of annual hours worked of both males and females, shown in Figure 5, suggest the workforce, in terms of unobservables, has changed over time. Figure 4 shows the fraction of wage earners who work either full-time or full-year. For males, it fluctuates cyclically between 80 and 90 percent, and is about 85 percent on average. For women, despite the positive labor market trends during our period, it is still less than 75 percent in 2016. Moreover, just as the returns to observed characteristics may vary over time, and across different parts of the wage distribution the return to these unobservables may act similarly. Mulligan and Rubenstein (2008) find that selection into full-time employment played an important role in explaining variation in inequality both within and across gender. Arellano and Bonhomme (2017) provide decompositions of wage growth which incorporate a role for selection.

While accounting for selection effects in the estimation of the conditional mean, as is done in Mulligan and Rubenstein (2008), is generally straightforward, it is challenging when estimating the determinants of wages at different points of the wage distribution. One recent approach to modeling the income changes which incorporates selection is Arellano and Bonhomme (2017) who adopt a copula approach to modeling the wage distribution and a binary work decision. Fernández-Val, van Vuuren and Vella (2018) propose an estimation strategy for nonseparable models with a censored selection rule via the use of distribution regressions while accounting for sample selection via an appropriately constructed control function.

We employ the Fernández-Val, van Vuuren and Vella (2018) approach, hereafter FVV, to nonparametrically estimate the relationship between the individual's wages and their characteristics while accounting for selection. In contrast to Arellano and Bonhomme (2017) which studies selection with the more commonly encountered binary selection rule, FVV requires a censored selection rule. While this requires more information in the data in this instance the required information, namely hours of work, is provided. The following section discusses the data. Section 3 provides a discussion of our empirical model, the estimation procedure and the objects of interest and their implication for wage inequality. Section 4 provides the empirical results. Section 5 provides some additional discussion of the empirical results and concluding comments are offered in Section 6.

2 Data

We employ micro-level data from the Annual Social and Economic Supplement (ASEC) of the Current Population Survey (CPS), or March CPS for each of the 41 survey years from 1976 to 2016 which report annual earnings for the calendar years from 1975 to 2015.[2] We begin with the 1976 survey as information on weeks worked and usual hours of work per week last year are not available before that survey. To ensure that most individuals have completed their educational process we restrict attention to those aged 24–65 years in the survey year (aged 23–64 in the previous year). These criteria produce an overall sample of 1,794,466 males and 1,946,957 females, with an average annual sample size of 43,767 males and 47,487

[2]We downloaded the data from the IPUM-CPS website maintained by the Minnesota Population Center at the University of Minnesota (Flood et al. 2015), which provides standardized data from 1963 to 2016.

females. There are large year-to-year fluctuations in annual sample sizes, which range from a minimum of 30,767 males and 33,924 females in 1976 to a maximum of 55,039 males and 59,622 females in 2001.

We define annual hours worked last year as the product of weeks worked last year and usual hours of work per week. Most of those reporting zero hours worked last year report themselves as not in the labor force (i.e., they report themselves as doing housework, unable to work, at school, retired or other) in the week of the March survey. We define hourly wages as the ratio of reported annual earnings and annual hours worked in the year before the survey. Hourly wages are unavailable for those not in the labor force. However, this definition of hourly wages creates an issue for the Armed Forces, the self-employed, and unpaid family workers as their annual earnings and annual hours tend to be poorly measured. For example, annual weeks of work of those in the Armed Forces are recorded as zero till 1989, while about 60 percent of the self-employed report zero annual earnings. To bypass this problem, we exclude the Armed Forces, the self-employed, and unpaid family workers from our sample and focus attention on the sample consisting of civilian dependent employees with positive hourly wages and people out of the labor force last year. This sample contains 1,551,796 males and 1,831,220 females (respectively 86.5 percent and 94.1 percent of the original sample of people aged 24–65), with an average annual sample size of 37,849 males and 44,664 females. The subsample of civilian dependent employees with positive hourly wages contains 1,346,918 males and 1,276,125 females, with an average annual sample size of 32,852 males and 31,125 females.

The March CPS differs from the commonly employed Outgoing Rotation Groups of the CPS, or ORG CPS, which contains information on hourly wages in the survey week for those paid by the hour, and on weekly earnings from the primary job during the survey week for those not paid by the hour. Lemieux (2006) and Autor, Katz and Kearney (2008) argue that the ORG CPS data are preferable for two reasons. First, workers paid by the hour (more than half of the U.S. workforce) may recall their hourly wages better. Second, the March CPS data lack a point-in-time wage measure. No clear evidence is available about differences in reporting accuracy between hourly wages, weekly earnings and annual earnings. In addition, many workers paid by the hour also work overtime, so their effective hourly wage depends on the importance of overtime work and the wage differential between straight time and overtime. Furthermore, the failure of the March CPS to provide a point-in-time wage measure may actually be an advantage as it smooths out intra-annual variations in hourly wages.

There are concerns with using earnings data from either of the CPS files. First, defining hourly wages as the ratio of earnings (annual or weekly) to hours worked (annual or weekly) may induce a "division bias" (see, e.g., Borjas 1980). Second, CPS earnings data are subject to measurement issues, including: (i) top-coding of earnings (see Larrimore et al. 2008), (ii) mass points at zero and at values corresponding to the legislated minimum wage (DiNardo, Fortin and Lemieux 1996), (iii) item non-response, especially in the tails of the earnings distribution, which has been rising sharply over time (see, e.g., Meyer, Mok and Sullivan 2015), (iv) earnings response from proxies, typically another household member, and (v) earnings imputation procedures by the Census Bureau (see Lillard, Smith and Welch 1986, and Bollinger et al. 2017). Using distribution regression and quantile regression methods, as is done in this paper, somewhat mitigates the first two issues. Since there is no consensus on how to best address the remaining three issues we retain proxy responses and imputed earnings.

To explain the variation in hourly wage rate and annual hours of work we employ a number of conditioning variables including an individual's age and categorical variables for educational attainments (less than high-school, high-school graduate, some college, and college or more), race (white and non white), region of residence (North-East, South, Central, and West), and marital status (married with spouse present, and not married or spouse not present). We also employ household composition variables, including the number of members and the number of children in a household, and indicators for the presence of children under 5 years of age and the presence of other unrelated individuals.

3 Model and Objects of Interest

Our primary objective is to investigate how the observed changes over time in individual hourly wage rates can be attributed to changes in the composition of the labor market, as reflected by the individual characteristics of those in the labor market, and to movements in the prices associated with these characteristics. We also explore the role of "selection effects". We allow these relationships to vary by location in the wage distribution.

Two interesting features of Figure 1 should be captured by our modelling approach. First, the trends for males and females appear quite different. This is accommodated by estimating separate models by gender. Second, the trends of the hourly wage rates differ greatly depending on their location in the hourly wage distribution. Our modelling approach should incorporate this aspect of the data. In addition, Figure 3 suggests that the composition of the labor force has changed for both males and females over the sample period. This introduces the possibility that the observed wage movements partially reflect changes in the observable and unobservable characteristics of the work force. While the changes in the observable characteristics can be accounted for by including the appropriate conditioning variables the movements in the unobservables require that we account for "selection". We estimate a model which incorporates each of these features and investigate their respective contributions.

3.1 Model

We begin by estimating the following model based on the Heckman (1979) sample selection model, where the censoring rule generating the selection process in the data has been adapted to incorporate that the number of hours worked, and not only the binary work/not work rule, are observed. The model has the form:

$$W = g(X, \varepsilon), \text{ if } H > 0, \quad (1)$$

$$H = \max\{h(X, Z, \eta), 0\}, \quad (2)$$

where W denotes the hourly wage rate, H denotes annual hours worked, and X and Z are vectors of observable explanatory variables. The functions g and h are unknown and ε and η are respectively a vector and a scalar of potentially mutually dependent unobservables. We assume that ε and η are independent from X and Z in the total population and the function h is increasing in its last argument. The model is a nonparametric representation of the tobit type-3 model previously estimated under a variety of parametric and semi-parametric settings (see, for example, Amemiya 1984, Vella 1993, Honoré et al. 1997, Chen 1997, and Lee and Vella 2006). The most general treatment, and that most suited to our objectives and employed here, is provided by FVV.

Assuming that X and Z are independent of ε in the total population does not rule out dependence in the selected sample. However, FVV show that conditional on $V \equiv F_{H|X,Z}(H|X,Z)$ and $H > 0$, ε is independent from X and Z. That is $F_{H|X,Z}(H|X,Z)$, is the appropriate control function for our model.[3] The intuition behind this result is that it can be shown that $V = F_\eta(\eta)$ and thus V can be interpreted as the index for the error term of the hours equation. Moreover, this control function can be estimated via a distribution regression of H on X and Z.

Below we denote the support of the selected sample of X and the control function V by $\mathcal{XV}$. This set reflects the combination of the observed characteristics and the control function that are observed for those individuals that have a positive number of working hours. We also introduce the set $\mathcal{V}(x)$ which is the support of V conditional on a particular level of $X = x$ among the selected sample.

3.2 Local Effects

As previous studies partially attribute the increase in inequality to the skill premium we examine how the returns to education have evolved and how they can be estimated while accounting for the changing participation rates and hours of work. We estimate what FVV define as "local objects". These are objects of the outcome variable W conditional on the value of the control function V and for a given set of observed characteristics X. Local objects are defined for the total population and are not restricted to the selected sample. Two important local objects are based on the mean and the distribution of the outcome variable. These are defined as the local average structural function (LASF) and the local structural distribution function (LSDF).

The LASF is equal to:

$$\mu(x, v) = \mathbb{E}[g(x, \varepsilon) \mid V = v],$$

i.e. the mean of the outcome variable in the case that everybody with a control function equal to v would have had characteristics x. The LDSF is equal to:

$$G(y, x, v) = F_{g(x,\varepsilon)|V}(y|v) = \mathbb{E}[1\{g(x, \varepsilon) \leq y\} \mid V = v]$$

[3]This result is closely related to that of Imbens and Newey (2009).

i.e. the distribution of the outcome variable in the case that everybody with a control function equal to v would have had characteristics x. The LASF and LSDF are identified provided the combination of x and v considered belongs to the support of the selected sample, $(x, v) \in \mathcal{XV}$. This also implies that for a given level of x, these objects are identified on the observed set $\mathcal{V}(x)$.[4] The LASF can be estimated using a series regression of the outcome variable on X and V and the LDSF can be estimated using series distribution regression methods. This is explained in greater detail in the appendix.

From the LSDF we derive the local quantile structural function (LQSF):

$$q(\tau, x, v) := \inf\{y \in \mathbb{R} : G(y, x, v) \geq \tau\}.$$

The LQSF is the left-inverse function of $w \mapsto G(w, x, v)$ and corresponds to the quantiles of the outcome variable for a given set of observed characteristics X conditional on a particular value of V.

3.3 Global Effects and Wage Decompositions

FVV derive global counterparts of the local effects by integrating over the control variable. An example of a global effect at $x \in \mathcal{X}$ is:

$$\theta(x|x_0) = \int_{\mathcal{V}(x_0)} \theta(x, v) dF_{V|X,H>0}(v|x_0), \tag{3}$$

where $\theta(x, v)$ can be any of the local objects defined above, $F_{V|X,H>0}(v|x_0)$ is the distribution of V conditional on the outcome $X = x_0$, and $\mathcal{V}(x_0)$ is the corresponding support. Since the distribution of $F_{V|X,H>0}(v|x_0)$ is identified, identification of $\theta(x|x_0)$ requires the identification of $\theta(x, v)$ over $\mathcal{V}(x_0)$. From above we know that local objects for a given level of x are identified on the set of $\mathcal{V}(x)$. Hence, the identification restriction of this global object equals $\mathcal{V}(x_0) \subset \mathcal{V}(x)$. That is, individuals with a positive probability of being in the selected sample with observed characteristics x_0 also have a positive probability of being selected if their observed characteristics are equal to x. The appropriateness of this restriction depends on the situation. For example, it seems reasonable to assume that all low educated

[4]FVV provide the proofs of these and related identification results. The proofs are constructive in that they show that these objects are identical to the objects of the outcome variable conditional on $X = x$, $V = v$ and conditional on the selected sample.

workers with a positive probability of working will also have a positive probability of working when assigned a higher level of education.

An interesting global object is the average structural function (ASF) conditional on $X = x_0$ in the selected population:

$$\mu^s_{x_0}(x) := \mathbb{E}[g(x, \varepsilon) \mid H > 0, X = x_0].$$

This represents the average of the wage in the selected population of individuals with X equal to x_0 when their observed characteristics are set equal to x. The average treatment effect of changing X from x_0 to x_1 in the selected population is:

$$\mu^s_{x_0}(x_1) - \mu^s_{x_0}(x_0). \tag{4}$$

Similarly, one can consider the distribution structural function (DSF) in the selected population as in Newey (2007). That is:

$$G^s_{x_0}(y, x) := \mathbb{E}[1\{g(x, \varepsilon) \leq y\} \mid X = x_0, H > 0], \tag{5}$$

which gives the distribution of the potential outcome $g(x, \varepsilon)$ at y in the selected population among individuals with characteristics x_0 when their characteristics equal x. This is a special case of the global effect (3) with $\theta(x, v) = G(y, x, v)$. We construct the quantile structural function (QSF) in the selected population as the left-inverse of the mapping $y \mapsto G^s(y, x)$. That is:

$$q^s_{x_0}(\tau, x) := \inf\{y \in \mathbb{R} : G^s_{x_0}(y, x) \geq \tau\}.$$

The QSF gives the quantiles of $g(x, \varepsilon)$. Unlike $G^s_{x_0}(y, x)$, $q^{s_{x_0}}(\tau, x)$ cannot be obtained by integration of the corresponding local effect, $q(\tau, x, v)$, because we cannot interchange quantiles and expectations. The τ-quantile treatment effect of changing X from x_0 to x_1 in the selected population is:

$$q^s_{x_0}(\tau, x_1) - q^s_{x_0}(\tau, x_0)$$

which measures how the τ-th quantile of the outcome variable changes among individuals in the selected sample with observed characteristics x_0 when their observed characteristics are set equal to x_1. We employ these quantile treatment effects to evaluate the impact of education.

We also employ these global effects to generate counterfactual distributions constructed by integration of the DSF with respect to different distributions of the explanatory and control variables. We then perform wage decompositions (e.g., DiNardo, Fortin and Lemieux, 1996, Chernozhukov, Fernández-Val and Melly, 2013, Firpo, Fortin and Lemieux, 2011, and Arellano and Bonhomme, 2017) to examine the underlying forces of changes in the hourly wage distribution. We focus on the global effects for the selected population. To simplify notation, we use a superscript s to denote these functionals, instead of explicitly conditioning on $H > 0$. The decompositions are based on the following expression for the observed distribution of W:

$$G_Y^s(w) := \int_{\mathcal{XZV}} F_{W|X,V}^s(y|x,v) dF_{X,Z,V}^s(x,z,v),$$

where we use that W is independent of Z conditional on X, V and $H > 0$. Using Bayes' rule, we rewrite this equation as:

$$G_Y^s(w) = \frac{\int_{\mathcal{XZV}} G(y,x,v)\mathbf{1}(h(x,z,v) > 0) dF_{X,Z,V}(x,z,v)}{\int_{\mathcal{XZV}} \mathbf{1}(h(x,z,v) > 0) dF_{X,Z,V}(x,z,v)}.$$

We construct counterfactual distributions by combining the component distributions G and $F_{X,Z,V}$ as well as the selection rule h from different populations that can correspond to different time periods or demographic groups. Thus, let G^t, and F_{X_k,Z_k,V_k} denote the distributions in groups t, r, and k, and let $\mathcal{XZV}_k$ denote the support $\mathcal{XZV}$ in groups k. Then, the counterfactual distribution of W where the wage structure G is as in group t, the joint distribution of the characteristics and the control function is as in group k and the selection rule is as in group r is defined as:

$$\begin{aligned} G_{Y_{\langle t|k,r\rangle}}^s(w) &:= \frac{\int_{\mathcal{XZV}_r} G^t(y,x,v)\mathbf{1}(h^r(x,z,v) > 0) dF_{X_k,Z_k,V_k}(x,z,v)}{\int_{\mathcal{XZV}_r} \mathbf{1}(h^r(x,z,v) > 0) dF_{X_k,Z_k,V_k}(x,z,v)} \\ &= \frac{\int_{\mathcal{Z}_r} \int_{\mathcal{XV}_r(z)} G^t(y,x,v)\mathbf{1}(h^r(x,z,v) > 0) dF_{X_k,V_k|Z_k}(x,v) dF_{Z_k}(z)}{\int_{\mathcal{Z}_r} \int_{\mathcal{XV}_r(z)} \mathbf{1}(h^r(x,z,v) > 0) dF_{X_k,Z_k,V_k}(x,z,v)}, \end{aligned}$$

where the second line is conditional probability. We know that $G^t(y,x,v)$ is identified for all combinations of x and v for which $(x,v) \in \mathcal{XV}_t$. Note that the integration above is over the observed combinations of x and v of the population r, conditional on the outcome for Z equal to z. Hence, a necessary condition for the identification of $G^s_{Y_{\langle t|k,r\rangle}}(w)$ is that $\mathcal{XV}_r(z) \subseteq \mathcal{XV}_t$. Note that $\mathcal{XV}_r(z)$ is a subset of the combination of x and v in the total population of r, *i.e.* the identification set of the population r, or $\mathcal{XV}_r(z) \subseteq \mathcal{XV}_r$. Hence, a sufficient condition of the above integral to be identified is that $\mathcal{XV}_r \subseteq \mathcal{XV}_t$. That implies that the observed combinations of x and v in population r must be a subset of the population t. That is, individuals with a certain combination of x and v in the selected sample for population r should also be in the selected sample for population t. The opposite is not necessary. That is, an individual with characteristics x and v who is in the selected sample in population t does not necessarily have to be in the selected sample of population r would he or she have belonged to that population instead. Under this definition, the observed distribution in group t is $G^s_{W_{\langle t|t,t\rangle}}$. We can write:

$$G^s_{Y_{\langle t|k,r\rangle}}(w) := \frac{\int_{\mathcal{Z}_k}\int_{\mathcal{XV}_k(z)} G^t(w,x,v)\mathbf{1}\{(x,v)\in\mathcal{XV}_r(z)\}dF^s_{X_k,Z_k}(x,z)}{\int_{\mathcal{Z}_k}\int_{\mathcal{XV}_k(z)} \mathbf{1}\{(x,v)\in\mathcal{XV}_r(z)\}dF^s_{X_k,Z_k}(x,z)}.$$

and since $(x,v)\in\mathcal{XV}_r(z)$ if and only if $v > \mathbb{P}(H_r \leq 0|X_r = x, V_r = v)$, we can write this as:

$$G^s_{Y_{\langle t|k,r\rangle}}(w) := \frac{\int_{\mathcal{Z}_k}\int_{\mathcal{XV}_k(z)} G^t(w,x,v)\mathbf{1}\{v > \mathbb{P}(H_r \leq 0|X_r = x, V_r = v)\}dF^s_{X_k,Z_k}(x,z)}{\int_{\mathcal{Z}_k}\int_{\mathcal{XV}_k(z)} \mathbf{1}\{v > \mathbb{P}(H_r \leq 0|X_r = x, V_r = v)\}dF^s_{X_k,Z_k}(x,z)}.$$

This equation has its sample analog as:

$$\widehat{G}^s_{Y_{\langle t|k,r\rangle}}(w) := \frac{\sum_i \widehat{G}^t(w,X_{ik},V_{ik})\mathbf{1}\{V_{ik} > \widehat{\mathbb{P}}(H_r \leq 0|X_r = X_{ik}, V_r = V_{ik})\}}{\sum_i \mathbf{1}\{V_{ik} > \widehat{\mathbb{P}}(H^r \leq 0|X^r = X_{ik}, V^r = V_{ik})\}}.$$

We can decompose differences in the observed distribution between group 1 and 0 using counterfactual distributions:

$$G^s_{W_{\langle 1|1,1\rangle}} - G^s_{W_{\langle 0|0,0\rangle}} = \underbrace{[G^s_{W_{\langle 1|1,1\rangle}} - G^s_{W_{\langle 1|1,0\rangle}}]}_{(1)} + \underbrace{[G^s_{W_{\langle 1|1,0\rangle}} - G^s_{W_{\langle 1|0,0\rangle}}]}_{(2)} + \underbrace{[G^s_{W_{\langle 1|0,0\rangle}} - G^s_{W_{\langle 0|0,0\rangle}}]}_{(3)},$$

where (1) is a selection effect due to the change in the distribution of the control variable given the explanatory variables, (2) is a composition effect due to the change in the distribution of the explanatory variables, and (3) is a structural effect due to the change in the conditional distribution of wages given the explanatory variables and control variable.

4 Empirical Results

4.1 Estimation of Hours Equation

The left panel of Figure 5 reports the movements in the employment rates of males and females, defined as the percentage of people in our sample who report positive annual hours of work and positive annual earnings in the year before the survey. For males, the employment rate fluctuates cyclically around a downward trend, falling during recessions (especially the Great Recession of 2008–2011) and increasing during expansions, peaking at 90.0 percent in 1976, reaching a minimum of 82.1 percent in 2012, and rebounding to 83.1 percent in 2016. For females, it increases almost steadily from 1976 to 2000, from a low of 56.5 percent in 1976 to a high of 75.3 percent in 2001. It decreases somewhat ending at 70.0 percent in 2016.

The right panel of Figure 5 plots movements in average annual hours worked of wage earners separately by gender. For males, average annual hours fluctuate cyclically around a slightly upward trend, falling during recessions (especially the Great Recession) and increasing during expansions, from 2032 hours in 1976 to a minimum of 1966 hours in 1983 to a maximum of 2160 hours in 2001, ending at 2103 hours in 2016. For females, average hours increases steadily from 1515 hours in 1976 to 1792 hours in 2000. The increase in average annual hours continues after 2000, though at a slower rate and with some fluctuations, reaching 1838 hours in 2016. The number of average hours for females is highest at the end of the sample. These patterns suggest movements along both the intensive and the extensive margins of labor supply. The large variation in the average annual hours worked, combined with the drastic movements in the participation rates, suggest that the large changes in the hourly wage distribution reported in Figure 1 may partially reflect the changes in the composition of the work force. Moreover, these compositional changes may capture both observable and unobservable factors.

To implement the FVV approach, we first obtain the control function, given as the conditional CDF of hours, by estimating the annual hours of work outcome via distribution regression separately for males and females for each year separately. We do this via logistic regression as outlined in the appendix where the conditioning vector includes age and aged squared, education dummy variables for highest educational attainment reported as less than high school, high school, some college, or college or more, dummy for marital status (married or not), an indicator variable for race reported as non white, a dummy denoting the region of residence, a dummy variable denoting the presence of other individuals, the number of children in the household, the number of household members, a dummy variable denoting a child aged under 5 years present in household. We also include the interaction of the educational indicator functions with age and age squared. Each of these variables appear in the hourly wage models with the exception of those capturing family composition. These variables, along with the variation in hours, are the basis of identification.[5] As our focus is on the estimation of the wage equations we do not discuss the estimation results for the hours equations.

4.2 Wage Equations and the Impact of Education

We now estimate the model explaining the variation in individuals hourly wage rates. We estimate the model separately by gender for each of the cross sections by distribution regression over the sub sample of workers reporting a positive wage. The conditioning variables are those included in the hours equation with the exception of the household composition variables. We also include the control function and its square and the remaining conditioning variables are all interacted with the control function.

Substantial empirical evidence suggests that the widening wage gaps partially reflect changes in the impact of education on wages (see for example, Autor, Katz and Kearney, 2008). To explore this possibility we estimate the impact of education on earnings, while accounting for selection of the hours decision but treating the individual's education level as exogenous. The source of the "endogeneity/selection" to which we refer below is that related to the hours of work decision.

Education is captured by three dummy variables denoting that the individual has acquired, as their highest level of schooling, "high school", "some college" or "college or more". The excluded group is "less than high school". Starting in 1992, the CPS

[5]For a detailed discussion of issues related to identification in this model the reader is referred to FVV.

measures educational attainment by the highest year of school or degree completed instead of the previously employed “highest year of school attendance”. Although the educational recode by the IPUM-CPS aims at maximizing comparability over time, there is a discontinuity between 1991 and 1992 in how those with a high school degree and some college are classified. Thus, in Figure 3, we focus on the fraction of wage earners with a high school degree or less and with a college degree or more. The figure illustrates the dramatic changes in the educational attainment of the labor force. The fraction of male workers with at most a high-school degree fell steadily from 64.1 percent in 1976 to 38.4 percent in 2016, while the fraction with at least a college degree rose from 20.4 percent in 1976 to 35.2 percent in 2016. The trends for females are even more striking with the fraction of workers with at most a high-school degree falling from 69.1 percent in 1976 to 29.7 percent in 2016. At the same time the fraction with at least a college degree rose from 16.1 percent in 1976 to 40.1 percent in 2016. While male workers were more educated than female workers at the beginning of the sample period, this is reversed at the end of our sample period.

We examine the impact of education by estimating the various treatment effects introduced in Subsection 3.3. We use the sample analog of (4) for the estimation of the treatment effect:

$$\frac{1}{n_e}\sum_{i=1,E_i=e}\widehat{\mu}(X_i,e',\widehat{V}_i)-\frac{1}{n_e}\sum_{i=1,E_i=e}\widehat{\mu}(X_i,e_i,\widehat{V}_i),$$

where E_i is education level of individual i, e and e' are the actual and counterfactual education levels respectively and X_i and $\widehat{V}_i$ denote the other observed characteristics and the estimated control function for individual i. As discussed in the appendix, the function $\widehat{\mu}$ is the estimated LASF based on series regression of log wages on (X_i, E_i, V_i).[6] The quantile treatment effects can be estimated likewise although we first estimate the LSDF rather than the LASF to obtain global estimates of the structural distribution function, *i.e.* equation (5). We then invert that distribution to obtain the quantile treatment effects. As we use log wages the final outcomes can be interpreted as percentage changes. To satisfy the identification restriction we calculate for individuals with low levels of education, *i.e.* education level e, what they would have received extra in terms of wages if they would have had a high

[6] FFV provide conditions under which the confidence intervals can be estimated via bootstraps.

level of education, *i.e.* e'. This is the "opposite" of a treatment effect although we employ that terminology here for simplicity.

Figures 6 to 8 present the average treatment effects for males and females for various contrasts in education levels. The average treatment effects for the "high school" to "less than high school" comparison, shown in Figure 6, reveal that for our sample period the effect is increasing but reasonably flat. The figure gives the appearance of substantial volatility, particularly for females for the late 1990's and onwards. However, the apparent volatility occurs within a small range. As the recipients of this premium are likely to be located in the lower part of the wage distribution for both genders it appears that the high school premium is not increasing quickly. Moreover, as the wages of the non high school group have probably decreased, it is likely that the premium is not sufficient to ensure real wages do not fall. This appears to be the case for males as suggested by the trends discussed above.

Figure 7 provides the corresponding treatment effect estimates for the "some college" to "less than high school" comparison. There is a substantial change in the return to this qualification. For males the effect is 33.1 percent in 1976, reaches a peak of 50.0 percent in 2008 and is 44.9 percent in 2016. Although there are episodic declines, there is a reasonable increase over the sample period. For females there is a notable increase during the 1980's which corresponds to a period of wage increases for all of those above the lowest quantiles of the wage distribution. In general, the pattern for females is similar to that of males and the changes are of a similar magnitude. The estimates suggest that education is contributing to the rise in wages for those located higher in the overall wage distribution.

While Figures 6 and 7 are somewhat suggestive of a role for education in exacerbating inequality the evidence in Figure 8 which provides the treatment effect for the difference between "college" and "less than high school" is very dramatic. Over the sample period this premium increases from 49.9 to 89.1 percent for males and from 68.9 to 87.8 percent for females. For males there is a steady increase over the 40 years of our sample although there are several periods of drastic increases in this premium. Note, however, that there are the occasional episodes of decreases which either reflect sampling issues or cyclical influences. The story for females is somewhat similar although there is a large decline in the late 1970's before a drastic increase in both the 1980's and 1990's. For both males and females the drastic growth for this premium in the late 1900's is not observed for the 2000's although a slight increase is observed.

While these various average treatment effects are informative about general trends they fail to reflect the degree of heterogeneity in the role of education on wages. Accordingly we provide corresponding quantile treatment effects in Figures 9 to 11. We present the behavior of these quantile treatment effects at the 25th, 50th and 75th quantiles as these appear to be reasonably indicative of what is occurring in the lower, middle and upper parts of the wage distribution.

Figure 9 presents the "high school"/"less than high school" comparison. A number of features of these plots are notable. First, despite the evidence that there is growth in this premium for males and females from the period 1976 to 2000, the profiles are reasonably flat over large parts of the sample. Second, for males there are some difference across quantiles with those at the 75th quantile showing the strongest evidence of growth. Females, however, have relatively similar profiles at the three quantiles we examine. For both males and females there are relatively small differences across quantiles.

Figure 10 contrasts "some college" to "less than high school" and provides a clearer indication that this premium has increased over the sample with particularly strong evidence of large increases in the 1980's and part of the 1990's for females but once again there is not a great deal of variability across the quantiles. This is not suggestive of a great deal of heterogeneity in this premium even if it does suggest that education is contributing to the increase in inequality.

As with the average treatment effects the evidence from these quantile treatment effects for the college premium provides a clearer description of the impact of education over this period. These are provided in Figure 11. We focus on males first. At the first quartile there is a notable increase in the college effect until around 2000 before it flattens and displays both minor increases and decreases. At the median there is a more prolonged increase in the premium with continued growth. This growth reflects a more heterogeneous effect with the quantile effect at the first quartile starting at around 52.3 percent before ending at 81.1 percent while the median increases to as high as 93.4 in 2013 before ending at 89.6. The steady increase seen at the median is even clearer at the third quartile. Here there is a steady growth over the sample with periods of dramatic growth.

For females similar patterns appear although some additional features are interesting. At the median and below there is dramatic growth in the 1980's and parts of the 1990's and some levelling out over the 2000's. At the third quartile, however, there is a steady increase over the 80's and 90's and periods of equally fast growth in some post 2000 periods. For females, there is significant heterogeneity in the effects across quantiles with that at the third quartile clearly higher than those below.

The empirical results regarding these education treatment effects provide insight into the possible contributors to inequality. First, the returns to education have increased over our sample period and the premium for the higher education levels have increased dramatically. Second, the treatment effects, particularly for the college premium, show a greater degree of variability with the quantile treatment effects showing that the third quartile treatment effect for college education is very high for both males and females. Finally, the general trends and level of the average and quantile treatment effects are similar across gender although the trends in wages discussed above are very different by gender. This clearly suggests that the education premia is not the only contributing factor to these trends. While we do not focus on the other conditioning variables they are included in the decomposition exercise.

Before we proceed to the decomposition results, we note that these results are for the selectivity adjusted estimates. Although we do not report them here, the unadjusted estimates are similar even though they appear to display less variability. This may reflect the additional sampling induced by the estimation and introduction of the control function. Also, we estimated the different local effects corresponding to different values of the control function. These also display greater variability, particularly at the lower values of the control function. The results are available from the authors.

4.3 Decompositions

In conducting the decompositions we set the base years to 1976 for women and 2010 for men. Given the increasing participation rate of women the choice of 1976 seems the best choice for the base year. It assumes that individuals with a certain combination of x, v in the year 1976 should also have a positive probability to be working in any other year. For men, the best choice is at the depth of the financial crisis since that year had the lowest level of participation over the 40 years period. Nevertheless, the choice of the base year for men is somewhat harder to defend.

As policy discussions primarily typically focus on the median wage we commence with the decompositions at this wage level. This is presented in Figure 15. Recall that the median wage for men decreases by 13.4 percent over the period examined, while the female hourly wage increases by 20.1 percent. This is illustrated by the total changes reflected in these figures. For males there is a striking similarity between the total effect and the structural effect with the difference completely due to the composition effect. That is, there is no evidence of any selection effect over the sample period on the median wage. The composition effect is increasing over the sample period and this primarily reflects the educational composition of the workforce. The composition effect appears to increase the median wage by around 1.8 percent. In contrast the structural component appears to be strongly procylical and producing large negative effects on the wage. While there appear to be some upturns, coinciding with periods of improving economic conditions, the structural effect is negative over the entire period. This is a striking result and consistent with what is understood about the median wage over the past twenty years. Perhaps what is most striking is the size of the wage decreases that would have occurred in the absence of the large increase in educational attainment as it is this increase in education that is most likely captured in the composition effects.

While the male median wage decreases over our sample period the increase in the female wage is around 2 percent. Figure 15 reveals that this is almost entirely driven by the composition effect. Similar to the males, the structural effect is negative for the whole period, but its influence is far less substantial than that for males. However, the turning points appear to coincide across the genders. As with males, there is little evidence of selection bias at the median female wage. The contribution of the selection component is negative but appears to be negligible. This is in contrast to Mulligan and Rubenstein (2008) who find a large selection effect at the mean.

The decompositions for the median suggests that the well established decline in the median male wage is due to the prices associated with the male skill characteristics. What is most surprising is that the previous section established that the skill premiums, as reflected by the returns to schooling, have generally increased over the sample period and the composition effects reflect that the male labor force has become increasingly more skilled. This suggests that the returns to the lowly skilled have decreased drastically. A similar pattern is observed for the female median wage although the negative structural impact is more than offset by the composition effects.

To further explore this conjecture that the decrease in the male median wage is the result of the returns to education for the more lowly paid we explore the decompositions of wage growth at the 10th and 25th quantiles. We also provide the corresponding decompositions for females. These are reported in Figures 13 and 14. For males the decompositions at the 10th and 25th quantiles are remarkably similar. However, more importantly, they suggest the reductions in the male wage rate are the result of prices associated with the human capital of individuals located at these lower quantiles. That is, the penalty associated with a lack of education is resulting in large wage decreases. At each of these quantiles there is evidence of a negative structural component of around 25 to 30 percent in both the late 1990's and the late 2010's. While these effects are somewhat offset by the composition effects the overall effect on wages is negative and non trivial. They reveal a very dramatic reduction in the wage almost entirely due to how worker's characteristics in this part of the male wage distribution are valued. There is no sign of selection effects for males and we discuss this below.

The evidence at the 10th and 25th quantiles for females is drastically different than that for males. First, there are greater differences between the 10th and 25th quantiles. The negative structural effects for females are more evident at the 10th quantile. At the 25th quantile the negative structural effects are small and offset by the composition effects. For both the 10th and 25th quantiles the overall wage changes for females become positive in the late 1990's and generally become more positive, although with some dips, over the remainder of the sample noting the change in the female median wage changes also turned positive in the later 1980's. Recall that the growth of the male wage at the 10th, 25th and 50th quantiles is negative over the whole sample period.

Finally, we examine the decomposition of the wage growth at the 75th and 90th quantiles in Figures 16 and 17. The male wage at the third quartile shows a small increase over this period exacerbating the level of inequality based on comparisons involving differences between the lower and higher quantiles. For males, the structural component displays a pattern strikingly similar to that for females at the earlier quantiles discussed. That is, there is a large decrease initially before rebounding and then staying relatively flat from around the early 2000's onwards. Unlike the lower quantiles the negative changes resulting from the structural component are not sufficiently large to dominate the positive composition effects so from the early 2000's onwards the overall wage growth at the 75th quantile is positive.

The figure for the female hourly wage change at the 75th quantile highlights our introductory discussion that it is at the higher quantiles of the wage distribution that the more dramatic changes occur. The changes in the structural component are initially negative before turning positive at around the mid of the 1980's. Most interestingly, already from the beginning of the 1980's the structural component shows a positive trend and combined with the composition effect to produce a steadily increasing wage. However, while the decomposition at the 75th quantile is suggestive of the primary sources of inequality are the structural components at higher quantiles the figures for the 90th quantile are even more supportive of this perspective. For males the structural component is less negative than that of lower quantiles and this combined with the consistently positive composition effect is producing a wage gain over the whole period. For females the structural component is positive from the middle of the 1980's and, in contrast to all the other decompositions, has a larger positive effect than the composition effect. The two effects combine to produce a remarkable 40.7 growth in the wage.

Another intriguing feature of the decompositions are the selection effects for each gender. Recall that the selection effects captures the impact of the unobservable factors on the work/non-work decision. Moreover, since we look at changes in the wage distribution, our selection effects should be interpreted as a change of this impact. Given this interpretation, consider the role of selection for males. Figures 15 to 17 suggest that selection effects cannot explain the observed changes in the males' wage distribution. At the 50th, 75th and the 90th quantiles the selection effect is essentially zero suggesting that at these locations in the wage distribution there appear to be no effects on the wage operating through unobservables determining the hours decision. This result is not surprising as males in this area of the wage distribution have a strong commitment to the labor force. Moreover, not only is there likely to be relatively little movement on the extensive margin, there is also likely little movement on the intensive margin given males' level of commitment to full time employment. One might anticipate it is at the lower quantiles that one might uncover some type of selection effects as these individuals are likely to have a weaker commitment to full time employment and thus the movements at the extensive and intensive margins due to unobservables may be more important. However, the evidence in the figures does not support this.

Now focus on the selection effects for females. At the higher quantiles there is very minor support for the evidence that selection effects had any impact on the changes in the wage distribution. However, just as with the male sample, it is likely that the females in this area of the wage distribution have relatively strong commitment levels to employment. Thus even with the large increases in participation rates and average annual hours worked the role of unobservables does not appear to be important. However, further down the wage distribution the selection effects not only begin to appear but also appear to be large. This is particularly true at the 10th and 25th quantiles. For example, at the 10th quantile the selection effect contribution is 2.2 percent, while the total wage change is 8.5 percent. Note that this negative selection effect implies that the female wage is lowered by 2.2 percent due to the higher participation rates of women in the later years of our sample period. Hence, this story is in line with what in a standard sample selection model is called a positive sample selection effect. Such a positive sample selection effect is also found by Mulligan and Rubenstein (2008). The large selection effect provides strong evidence that these factors are contributing to inequality comparisons and that they are particularly important for comparisons which involve the lower parts of the wage distribution. Moreover, the changes in the selection effects are increasing inequality as they decrease wages. It should be noted that the trend uncovered as we move down the wage distribution suggests that these selection effects will become even more important as we move further down the female wage distribution. This reflects that it is these individuals, those displaying a weaker commitment to the labor force, that are associated with stronger selection effects. This result regarding selection is similar to the findings of Arellano and Bonhomme (2017) and the empirical evidence in FVV which both study the evolution of female wages in the British labor market for the period 1978 to 2000.

5 Discussion

A number of interesting findings arise from our empirical results. First, the trends of the skill premia associated with the returns to schooling drastically differ across the educational treatments. While completing high school and attaining some college increase wages, relative to not completing high school, there is relatively little evidence that the change in these premia have contributed to wage inequality. However, for both males and females there is compelling evidence that the college premium has increased dramatically over the sample period and increased wage inequality.

The decompositions reveal a number of important findings. First, the mechanisms operating at different points of the wage distribution and across gender appear to vary. For males, the dramatic fall in wages at the median and below appears primarily due to the low wages associated with lower education levels. This is reflected by the substantive and negative structural effects for this area of the wage distribution. As the education premium is high and more workers are receiving higher education the composition effects are positive and somewhat offset the negative structural effects. However, the large, and increasing educational premium, reflect that the "penalty" to not being educated has increased over time. The negative structural effects for males are not restricted to the lower part of the wage distribution. Even at higher parts of wage distribution the structural effects are negative suggesting that the "value" of labor has decreased. At higher parts of the wage distribution these effects are offset by the larger number of individuals with higher levels of education. Combining the evidence of the negative structural effects for males with the increasing education premia suggests that the high college premium reflects how poorly valued are the skills of the uneducated. For females, we see wage increases at each quantile we examined and these reflect, in part, positive and increasing composition effects. Moreover, while the structural effects are generally negative for the whole period at the 10th and 25th quantiles, they are typically positive at the quantiles we examined above the median. Most notably, the structural effects for females are not dampening wage growth to the same extent as for males and at some quantiles they are greatly increasing wages.

The evidence also suggests that at the lower parts of the female wage distribution the impact of selection is negative and can be substantial. These become more important as we move further down the wage distribution. Given the evidence at the 10th and 25th quantiles for females it is likely that evaluating selection effects at even lower quantiles would reveal even greater negative selection effects. Conversely, the large increases in wage growth we find at the 90th quantile for both males and females are likely to be greater at higher quantiles.

The micro nature of our investigation is unable to provide direct insight into the macro factors generating these wages profiles. Nor does it provide evidence on the role of institutional factors which disproportionately influence certain sectors of the work force. However it does seem that the mechanisms affecting the wages at the bottom are very different that those influencing those at the top. At the top of the distribution the evidence is supportive of an increasing skill premium. At the bottom it appears that such considerations associated with the protection of lower wage workers may be operative. These are most likely the decreases in the real value of the minimum wage and the reduction in unionization.

Our focus thus far has been on wage rates. We now directly examine the issue of inequality to examine how an established measure of inequality has changed and how these changes can be decomposed. The 90/10 ratios for hourly wages for males and females have increased by 54.6 and 37.9 percent respectively. Our evidence illustrates that hourly wage growth at different locations in the wage distribution is affected differently by the various factors but we are unable to directly infer from that evidence the respective components of the changes in inequality. For males recall that wages at the 10th quantile were due to the large negative structural effects. There was no evidence of selection effects and a small positive composition effect offset the decreases from around the early 1990's. In contrast the large gains at the 90th quantile reflected a steadily increasing composition effect and a structural effect which contributed both negatively and positively over the sample period. There are no signs of changes in selection effects. Figure 19 provides a decomposition of the changes in the 90/10 ratio for males. Given the evidence above, it is not surprising that the large increase in the 90/10 ratio for males over the period is almost entirely due to structural effects, although it is important to note that this reflects negative structural effects at the 1st decile and not positive effects at the ninth decile. The composition effects increase the ratio and this reflects the positive composition effects for males at the ninth decile.

For females the evidence is more complicated. There is an increasingly positive composition effect at the first decile of wages but a negative structural effect produces a decline in wages. At the 9th decile there is a steadily increasingly composition effect and a structural effect which are generally increasing the change in wages. The large increase in the total effect for the majority of the sample period reflects the sum of these two positive effects. Selection at the first decile indicates that the change in the selection effects on wages are negative while there is no sign of selection at the 9th decile. This result seems sensible given that the increase in the labor force participation rate and the number of hours worked of women who entered the workplace suggests that these are less productive. Figure 19 presents the change in the 90/10 ratio for females. Given the various issues related to wage growth, it is not surprising that the change in the 90/10 ratio appears to reflect almost entirely a structural effect. The composition effect has slightly decreased this measure of inequality due to the larger composition effects at the lower part of the wage distribution.

The change in the selection effects on the 90/10 suggest that they increase inequality and the contribution of this effect, in some periods, is relatively large as a fraction of the total change. Our conclusion regarding the impact of selection on inequality is consistent with our previous evidence regarding the role of selection on the wages of those located in various locations of the wage distribution.

While our focus is not upon across gender inequality we employ our earlier results to examine the trends in the male/female hourly wage ratio at different points of the wage distribution. These are reported in Figure 20. First, at all locations of the wage distribution females appear to be catching up to males. Moreover, the greatest gains appear to be at the median and below. This result should be treated with caution as it appears to be largely due to the reduction in the male wage and not large increases in the female wage. This is confirmed by a re-examination of Figure 1. Second, at all the quantiles reported in Figure 20 the improvement in the relative performance is almost entirely due to structural effects. As the earlier evidence suggested that the sign and the size of the structural effects on the individual gender specific wages varied by sample period and location in the wage distribution it is surprising that the impact on the gender wage ratios is so clear. The evidence suggests a relative improvement in the value of female labor at all points of the wage distribution. Third, the composition effects also have steadily increased the relative performance of females at all quantiles noting that the size of the effect diminishes as we move up the wage distribution until at the 9th decile the effect is small. Finally, the change in the selection effects is increasing gender inequality although the effects only appear at the lower parts of the wage distribution.

6 Conclusions

This paper documents the changes in female and male wages over the period 1976 to 2016 and decomposes these changes into structural, composition and selection components by implementing an estimation procedure for nonseparable models with selection. We find that male real wages at the median and below have decreased over our sample period despite an increasing skill premium and an increase in the educational attainment levels of the sample. The reduction is primarily due to large decreases of the wages of the lowly educated. Wages at the upper quantiles of the distribution have increased drastically due to a large and increasing skill premia and this has combined with the decreases at the lower quantiles to substantially increase wage inequality. For females the wage declines are less dramatic and occur at only the lower quantiles. The increases at the upper quantiles for females are also substantial and reflect increasing skill premia and these changes have resulted in a substantial increase in female wage inequality. As our sample period is associated with a large changes in the participation rates and hours of work of females we explore the role of ”selection” in wage movements. We find that the impact of selection is to decrease the wage growth of those at the lower quantiles with very little evidence of selection effects at other locations in the female wage distribution. These selection effects thus appear to increase wage inequality.

References

[1] Acemoglu D., and Autor D. (2011), “Skills, tasks and technology: Implications for employment and earnings”, *Handbook of Labor Economics*, Vol. 4B, pp. 1043–1171.

[2] Amemiya T. (1984), “Tobit models: A survey.” *Journal of Econometrics*, 24: 3–61.

[3] Angrist J., Chernozhukov V., and Fernàndez-Val I. (2006), “Quantile regression under misspecification, with an application to the U.S. wage structure.” *Econometrica*, 74: 539–563.

[4] Arellano M., and Bonhomme S. (2017), “Quantile selection models with an application to understanding changes in wage inequality.” *Econometrica*, 85: 1–28.

[5] Autor D. H., Katz L. F., and Kearney M. S. (2008), “Trends in U.S. wage inequality: Revising the revisionists.” *Review of Economics and Statistics*, 90: 300–323.

[6] Borjas G. (1980). “The relationship between wages and weekly hours of work: The role of division bias.” *Journal of Human Resources*, 15: 409–423.

[7] Chernozhukov V., Fernández-Val I., and Melly B. (2013), “Inference on counterfactual distributions.” *Econometrica*, 81: 2205–2268.

[8] DiNardo J., Fortin N. M., and Lemieux T. (1996), “Labor market institutions and the distribution of wages , 1973–1992: A semiparametric approach.” *Econometrica*, 64: 1001–1044.

[9] Fernández-Val I., van Vuuren A., and F. Vella (2018), “Nonseparable sample selection models with censored selection rules”, working paper, Boston University.

[10] Fortin N. M., Lemieux T., and Firpo S. (2011), “Decomposition methods in economics.” *Handbook of Labor Economics*, Vol. 4A, pp. 1–102.

[11] Foresi S., and Peracchi F. (1995), "The conditional distribution of excess returns: An empirical analysis." *Journal of the American Statistical Association*, 90: 451–466.

[12] Flood S., King M., Ruggles S., and Warren J. R. (2015). "Integrated Public Use Microdata Series, Current Population Survey: Version 4.0. [Machine-readable database]." University of Minnesota, mimeo.

[13] Heckman J. J. (1974), "Shadow prices, market wages and labor supply." *Econometrica*, 42: 679–694.

[14] Heckman J. J. (1979), "Sample selection bias as a specification error." *Econometrica*, 47: 153–161.

[15] Imbens G. W., and Newey W. K. (2009), "Identification and estimation of triangular simultaneous equations models without additivity." *Econometrica*, 77: 1481–1512.

[16] Juhn C., Murphy K. M., and Pierce B. (1993), "Wage inequality and the rise in returns to skill." *Journal of Political Economy*, 101: 410–442.

[17] Katz L. F., and Autor D. H. (1999), "Changes in the wage structure and earnings inequality." In O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, Vol. 3A, pp. 1463–1555.

[18] Katz L. F., and Murphy K. M. (1992), "Changes in relative wages, 1963–87: Supply and demand factors." *Quarterly Journal of Economics*, 107: 35–78.

[19] Koenker R., and Bassett G. (1978), "Regression quantiles." *Econometrica*, 46: 33–50.

[20] Larrimore J., Burkhauser R. V., Feng S., and Zayatz L. (2008). "Consistent cell means for topcoded incomes in the Public Use March CPS (1976–2007)." *Journal of Economic and Social Measurement*, 33: 89–128.

[21] Lee D. S. (1999), "Wage inequality in the United States during the 1980s: Rising dispersion or falling minimum wage?" *Quarterly Journal of Economics*, 114: 977–1023.

[22] Lemieux T. (2006), "Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill?" *American Economic Review*, 96: 461–498.

[23] Lillard L., Smith J. P., and Welch F. (1986). "What do we really know about wages? The Importance of nonreporting and Census imputation." *Journal of Political Economy*, 94: 489–506.

[24] Meyer B. D., Mok W. K. C., and Sullivan J. X. (2015), "Household surveys in crisis." *Journal of Economic Perspectives*, 29: 199–226.

[25] Mulligan C., and Y. Rubenstein (2008), "Selection, investment, and women's relative wages over time ", *Quarterly Journal of Economics*, 123: 10611110.

[26] Murphy K. M., and Topel R. H. (2016), "Human capital investment, inequality, and economic growth." *Journal of Labor Economics*, 34: S99–S127.

[27] Murphy K. M., and Welch F. (1992), "The structure of wages." *Quarterly Journal of Economics*, 57: 285–326.

[28] Newey, W.K. (2007), "Nonparametric continuous/discrete choice models", *International Economic Review*, 48: 1429–39.

[29] Vella, F. (1993), "A simple estimator for models with censored endogenous regressors." *International Economic Review*, 34: 441–457.

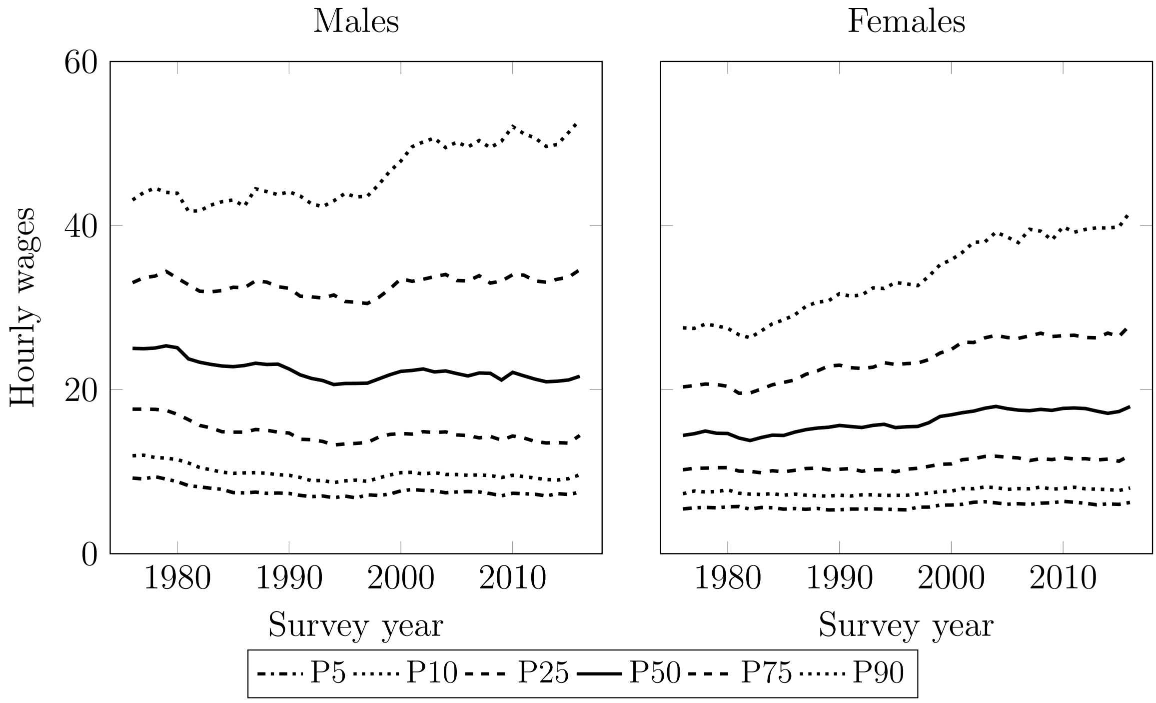

Figure 1: Percentiles of real hourly wages

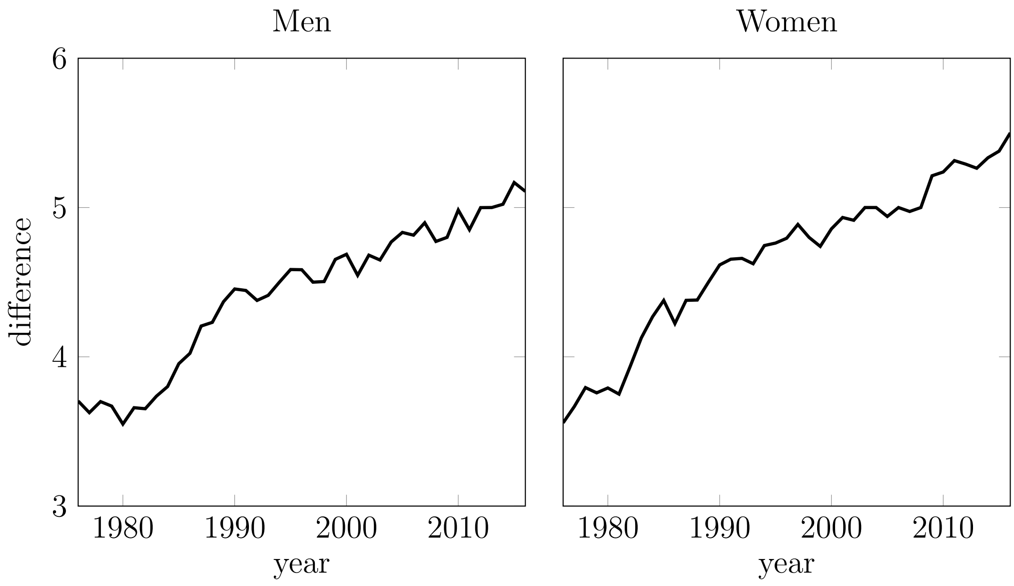

Figure 2: Development of the D9-D1 ratio.

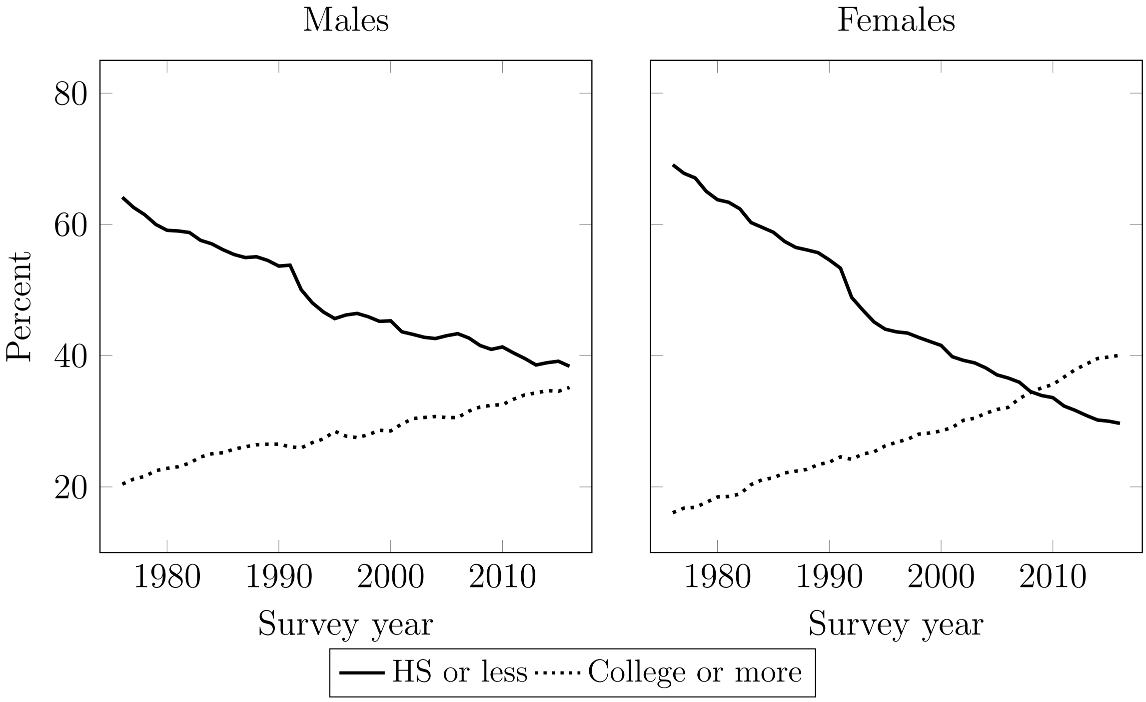

Figure 3: Distribution of wage earners by education level

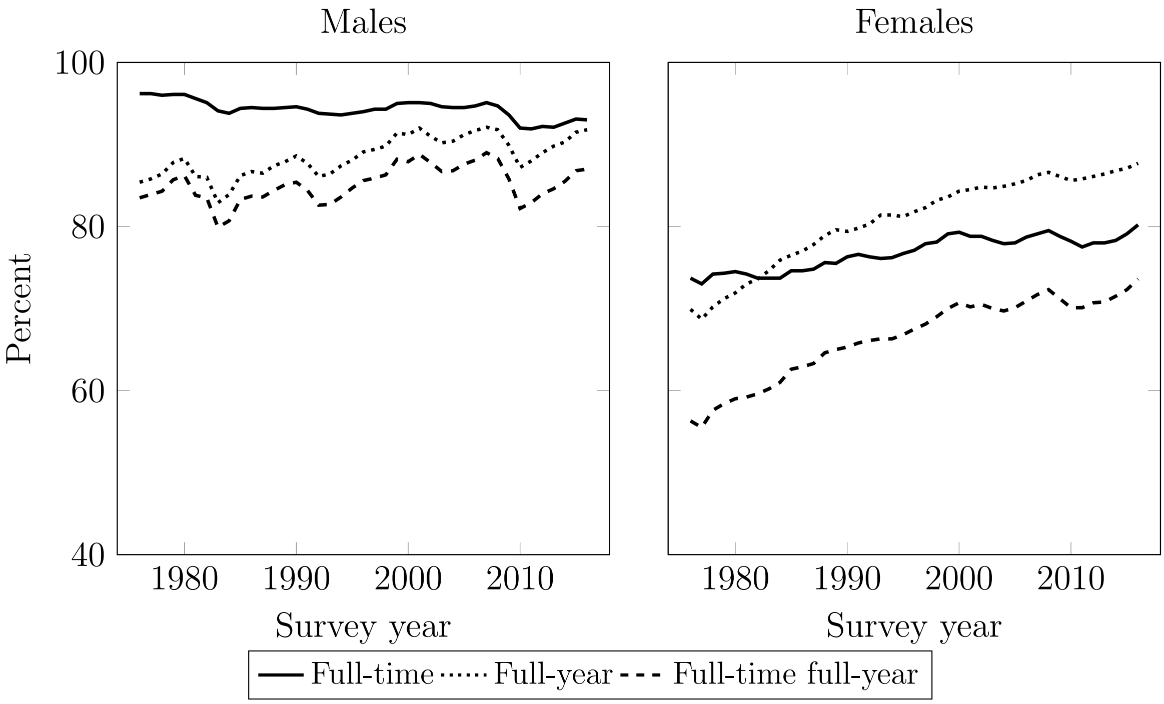

Figure 4: Fraction of wage earners working either full-time or full-year

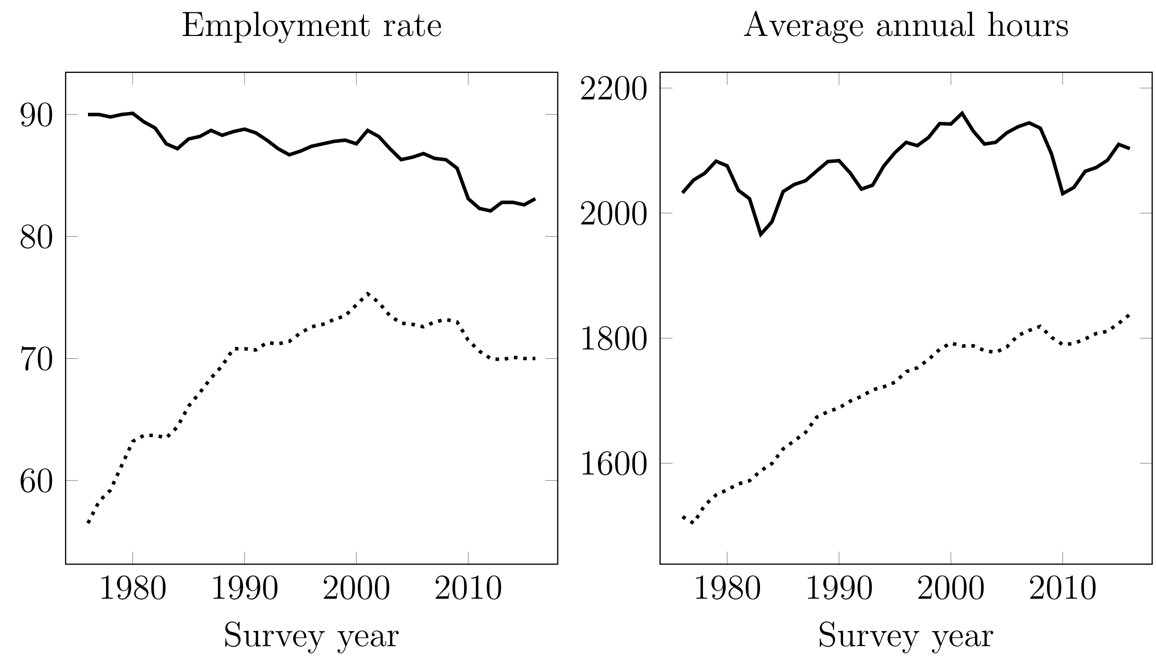

Figure 5: Employment rate and average annual hours worked of employees

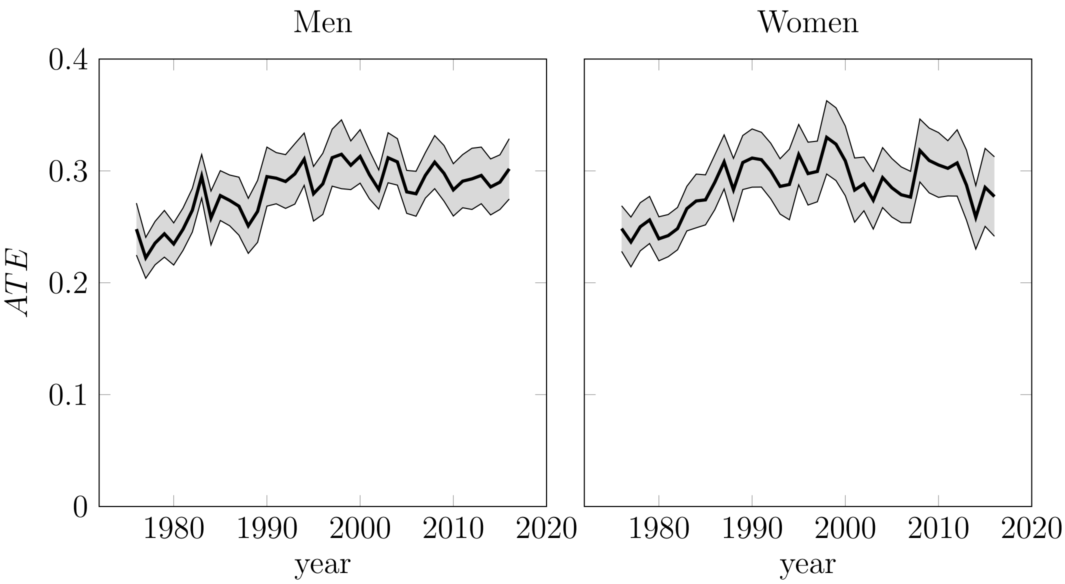

Figure 6: Average treatment effect of education, high school versus less than high school with correction.

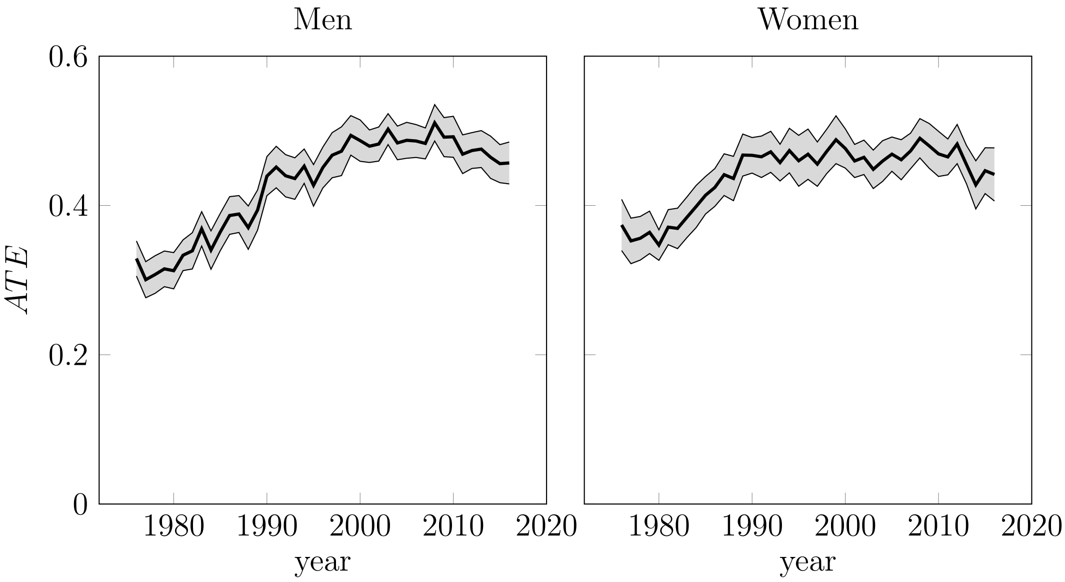

Figure 7: Average treatment effect of education, some college versus less than high school with correction.

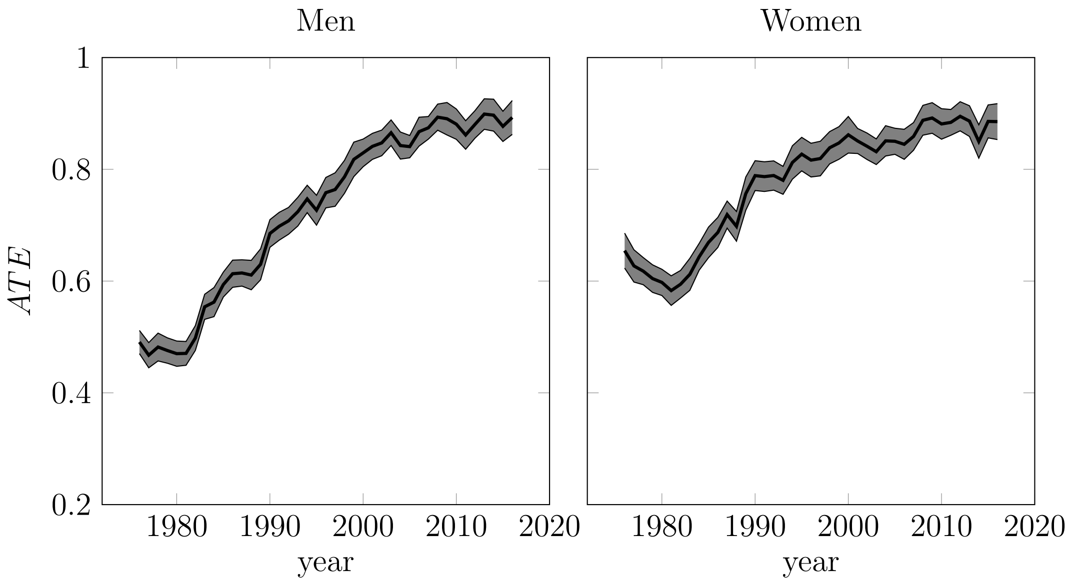

Figure 8: Average treatment effect of education, college versus less than high school with correction.

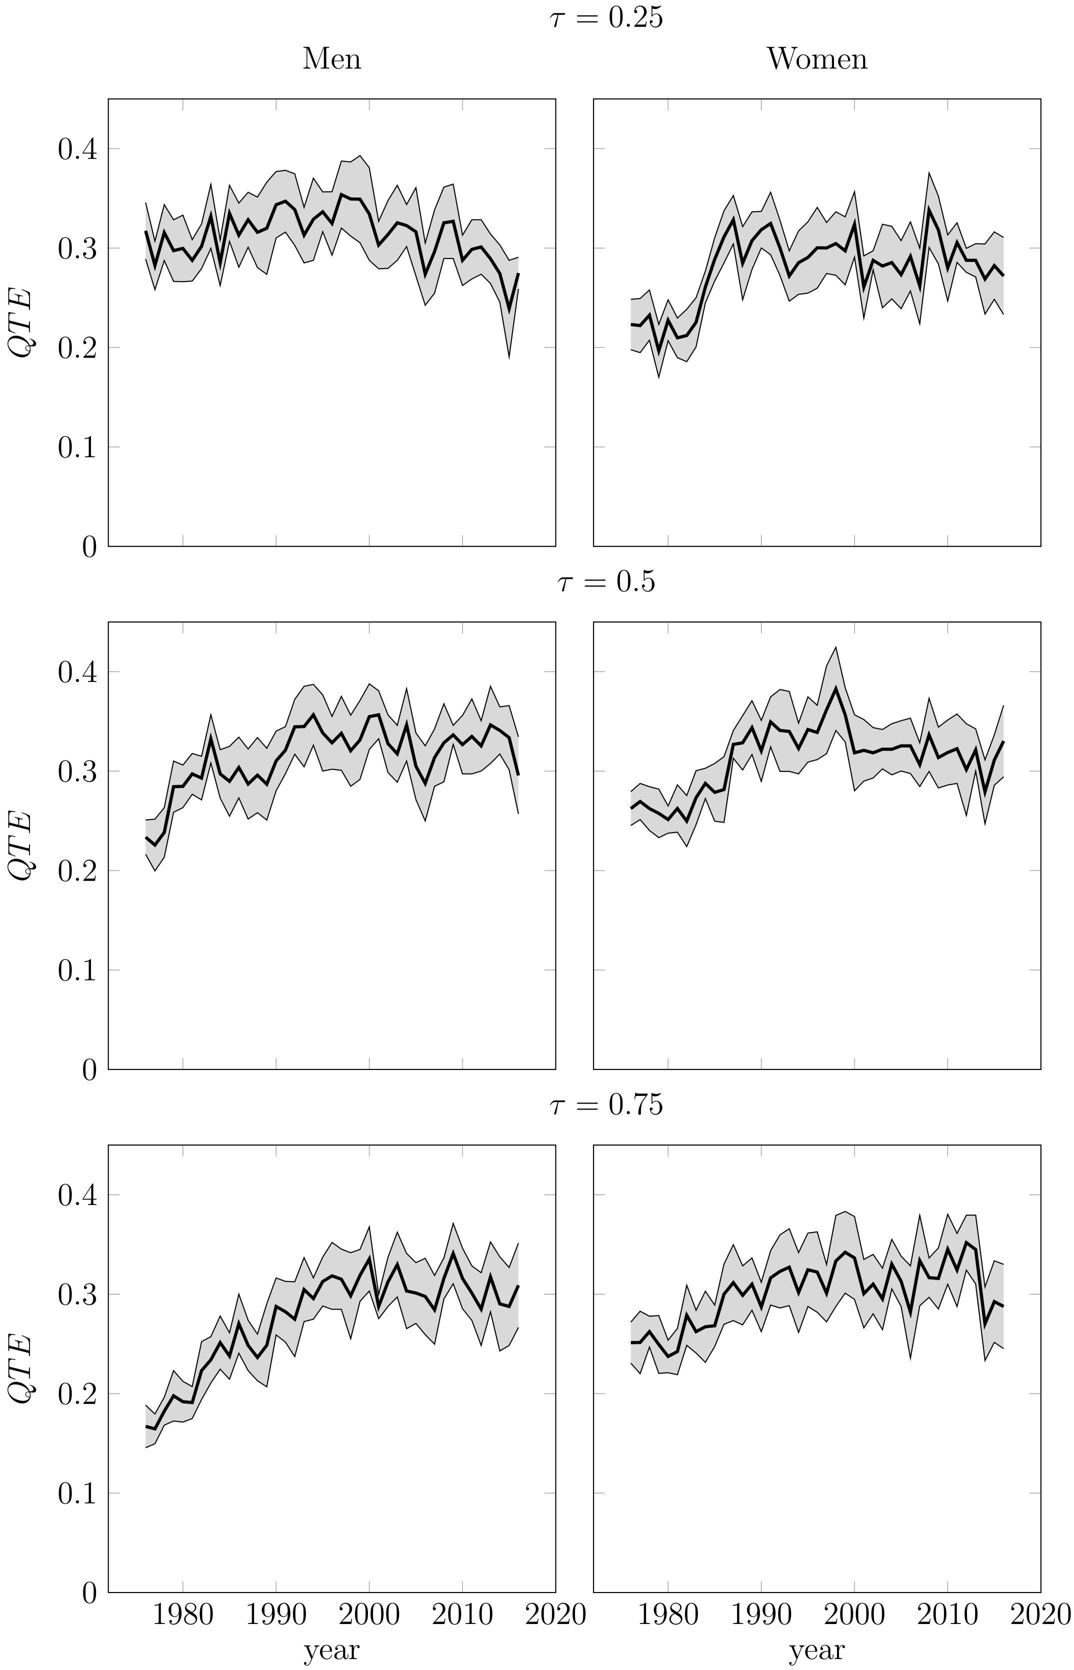

Figure 9: Quantile treatment effect of education, high school versus less than high school with correction.

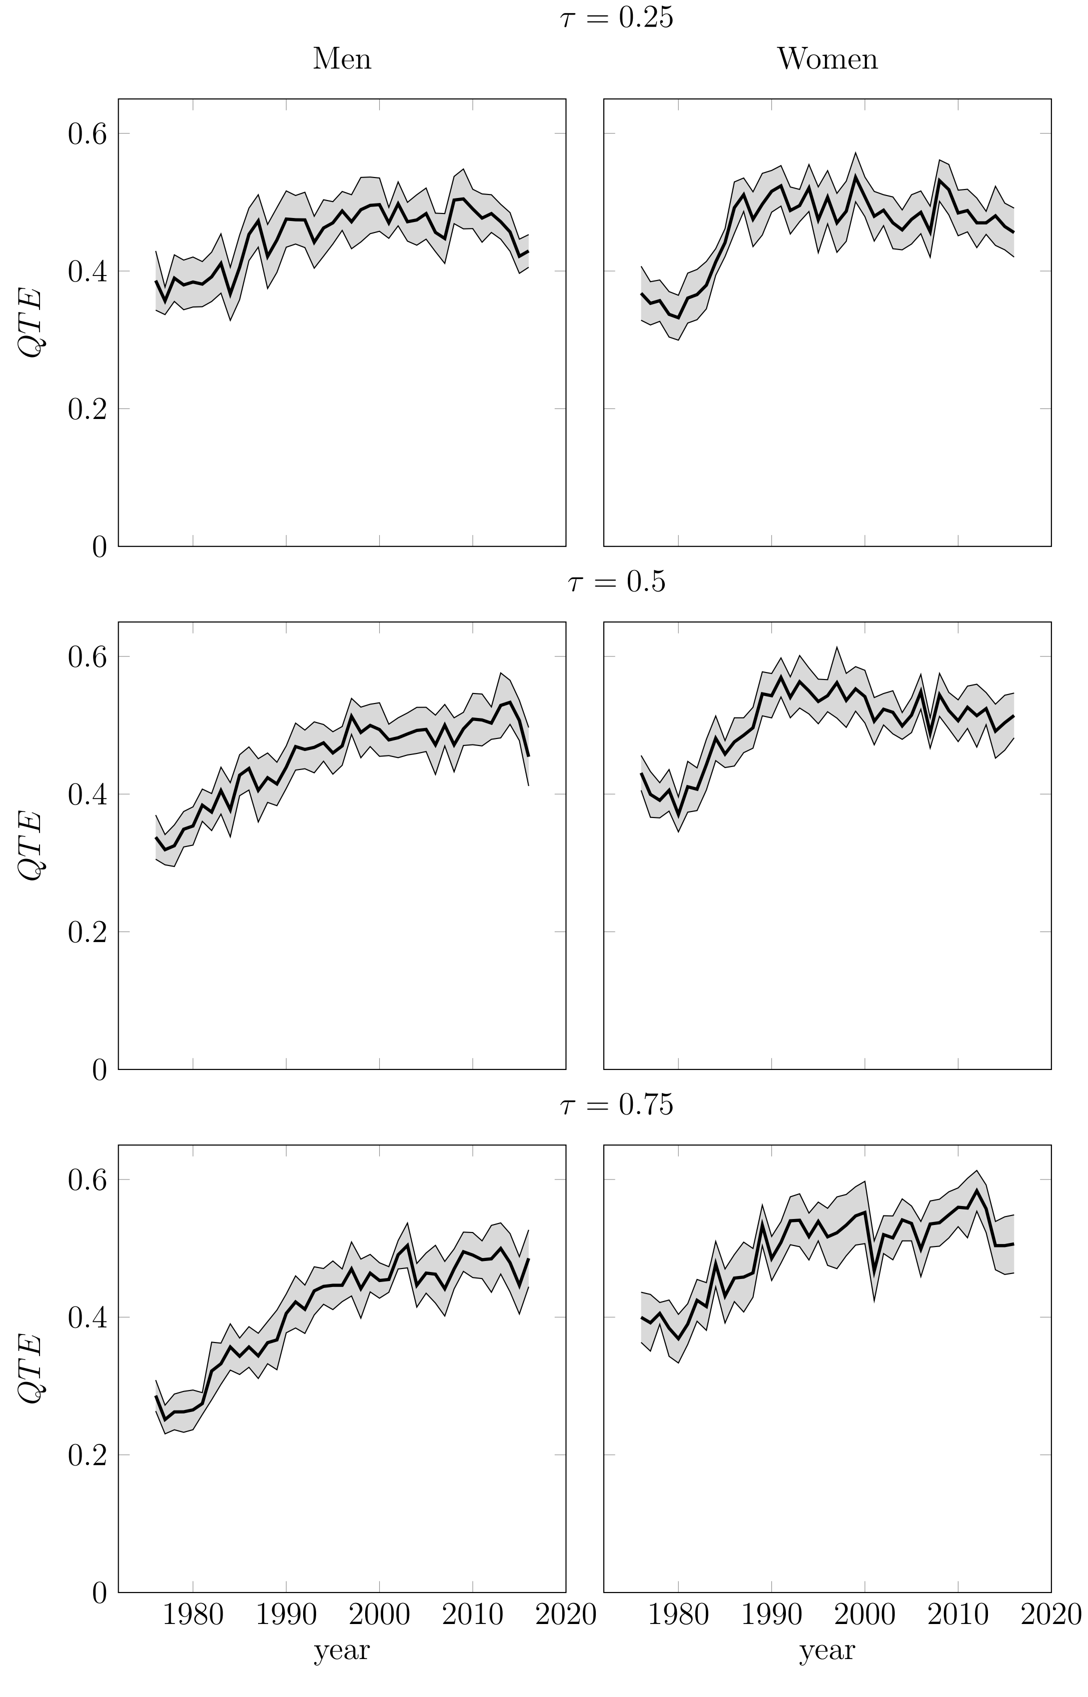

Figure 10: Quantile treatment effect of education, some college versus less than high school with correction.

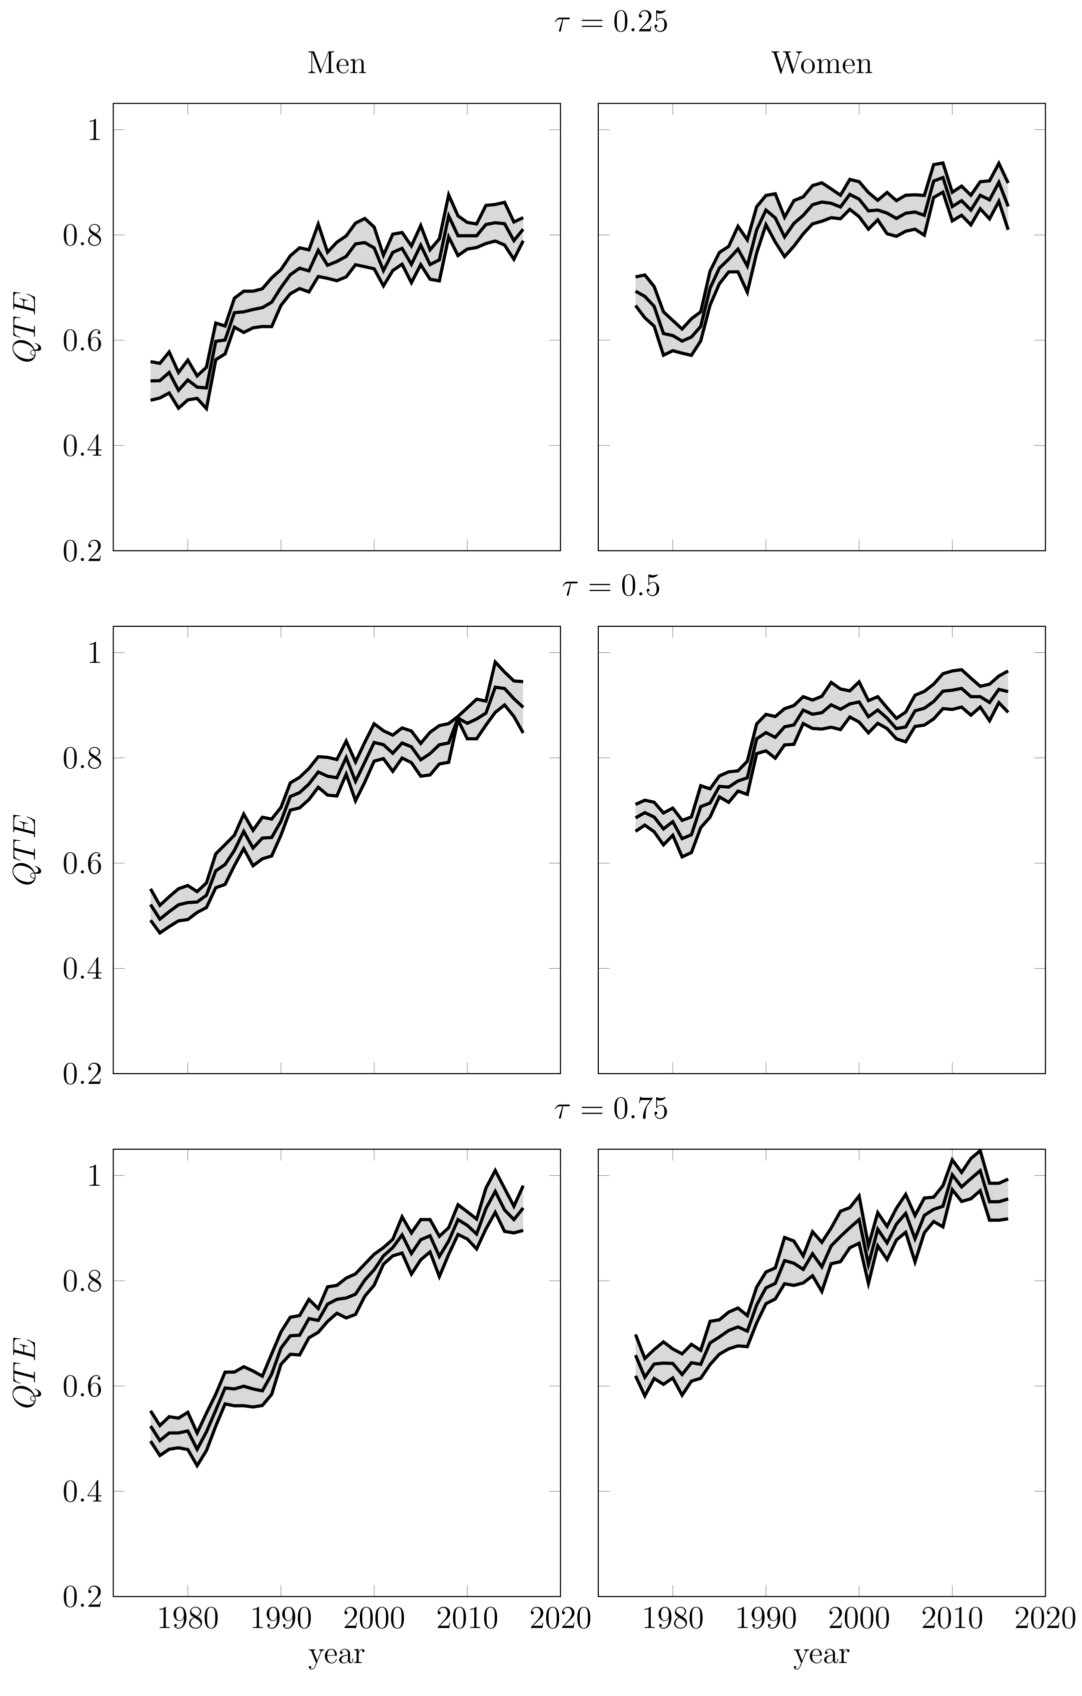

Figure 11: Quantile treatment effect of education, college versus less than high school with correction.

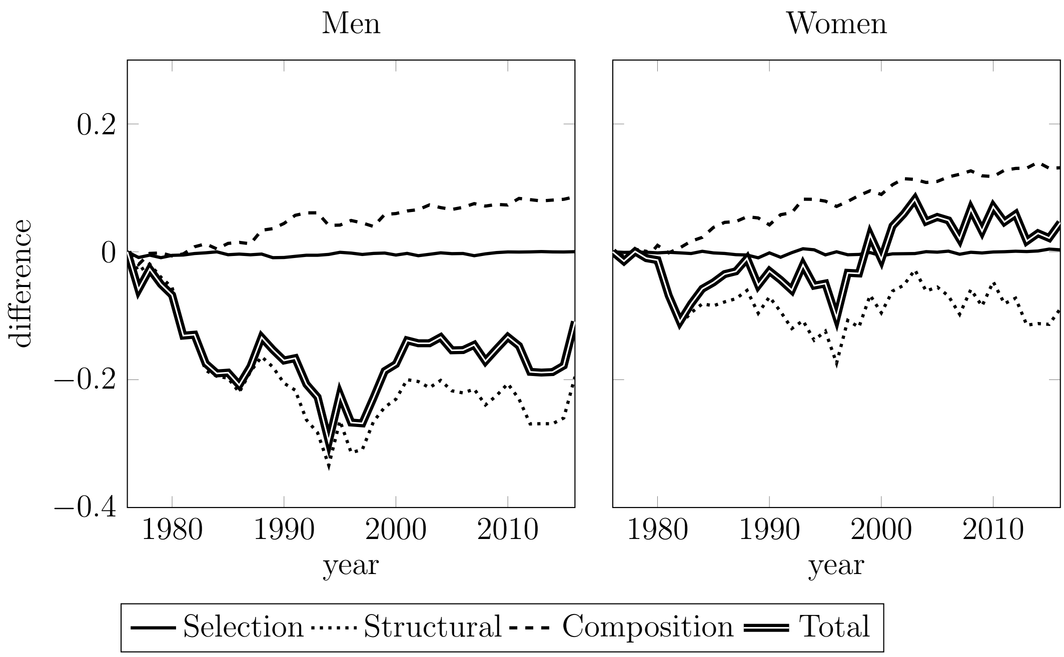

Figure 12: Decompositions at the mean.

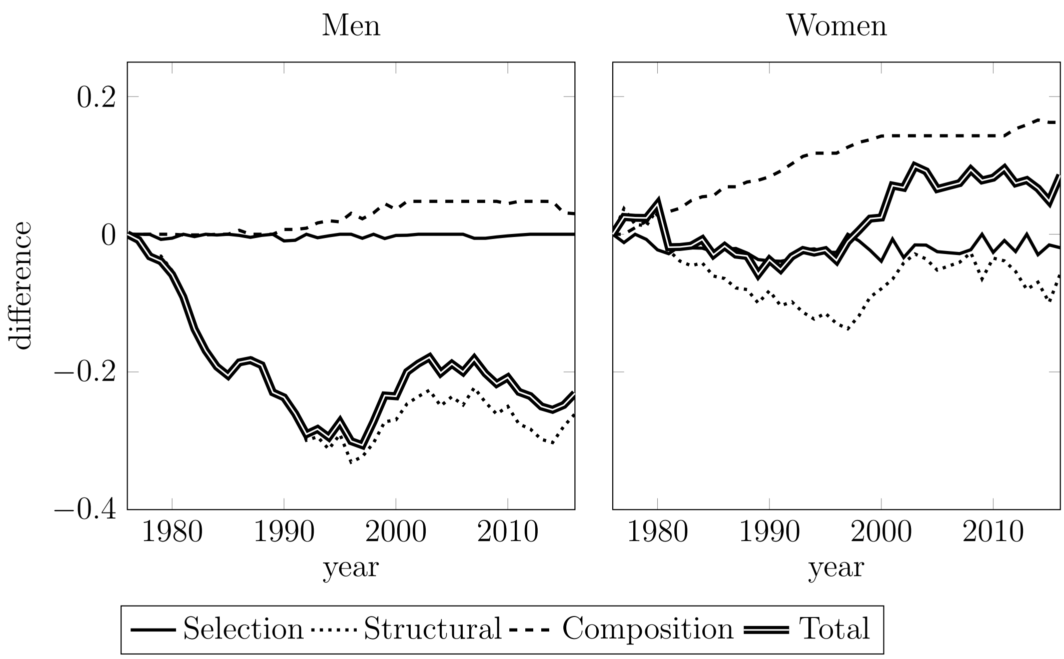

Figure 13: Decompositions at D1.

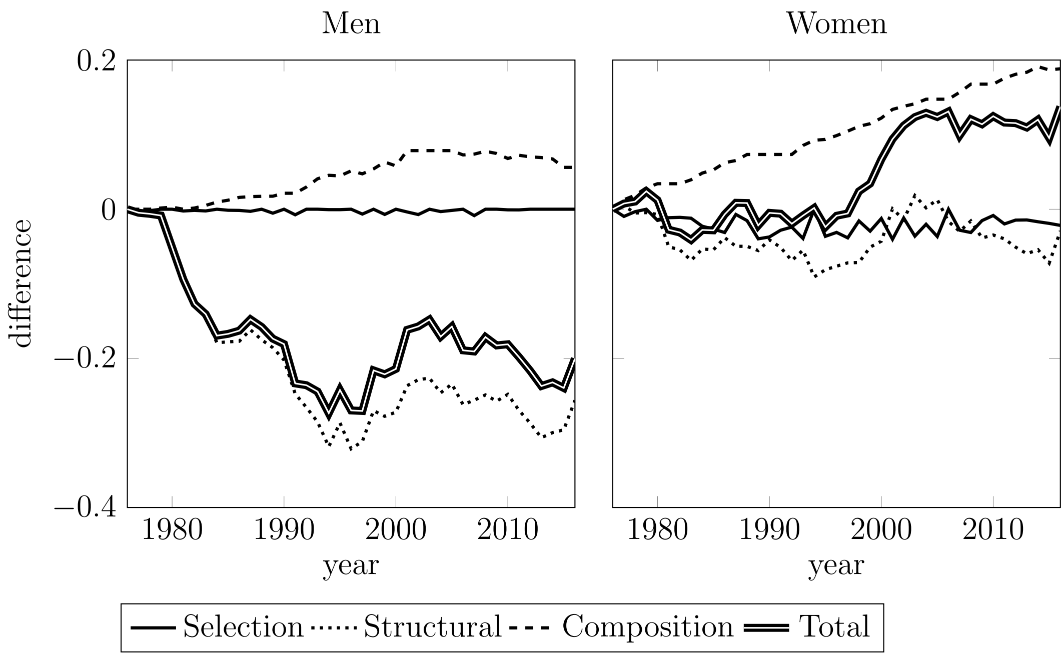

Figure 14: Decompositions at Q1.

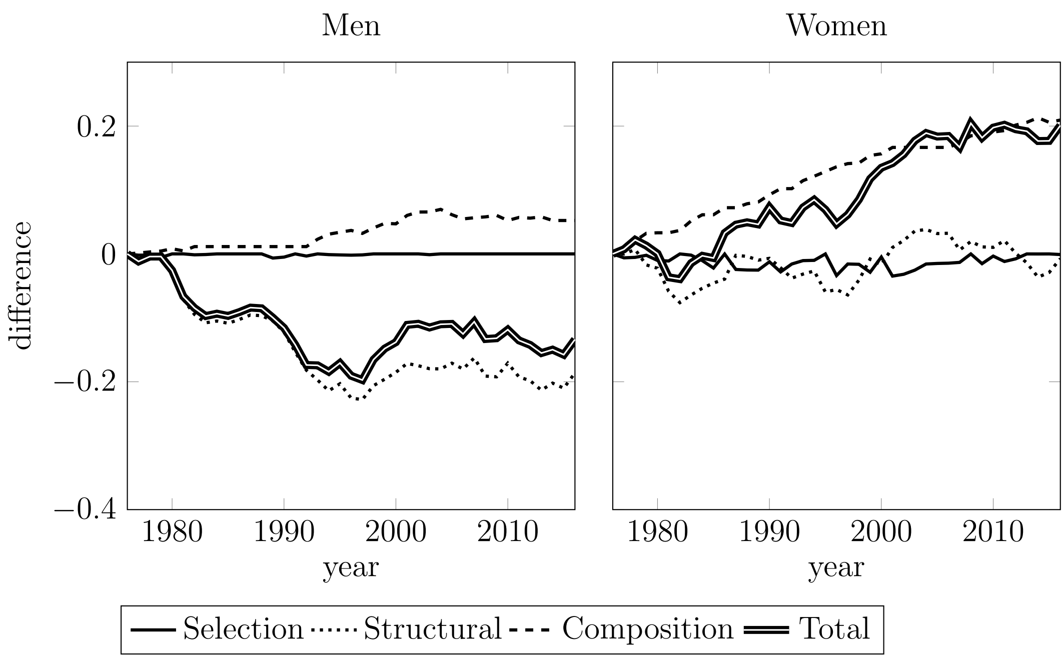

Figure 15: Decompositions at Q2.

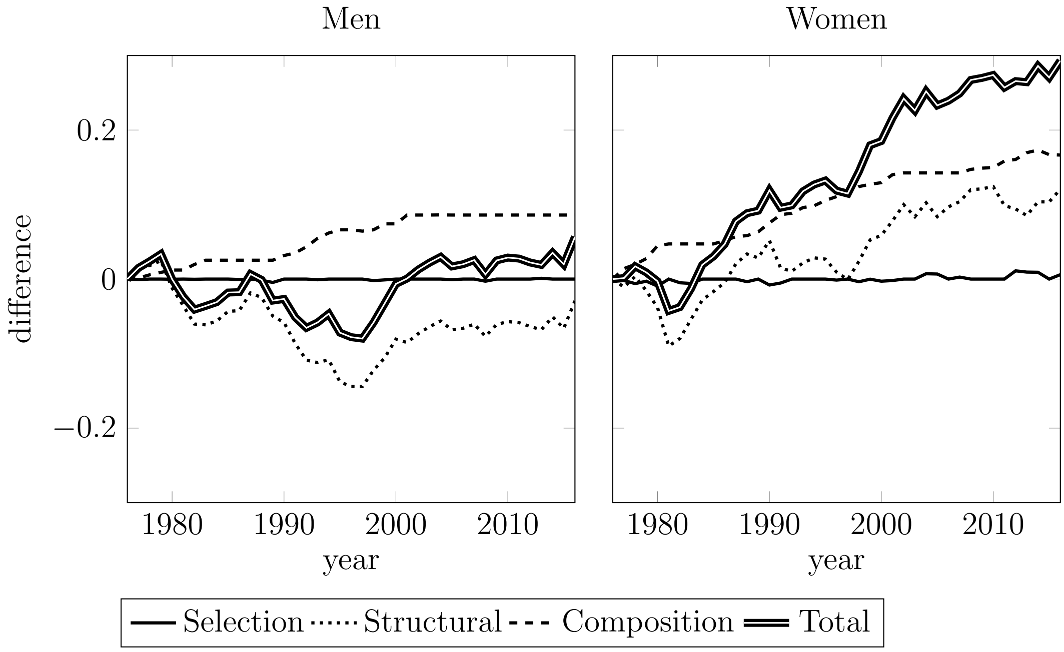

Figure 16: Decompositions at Q3.

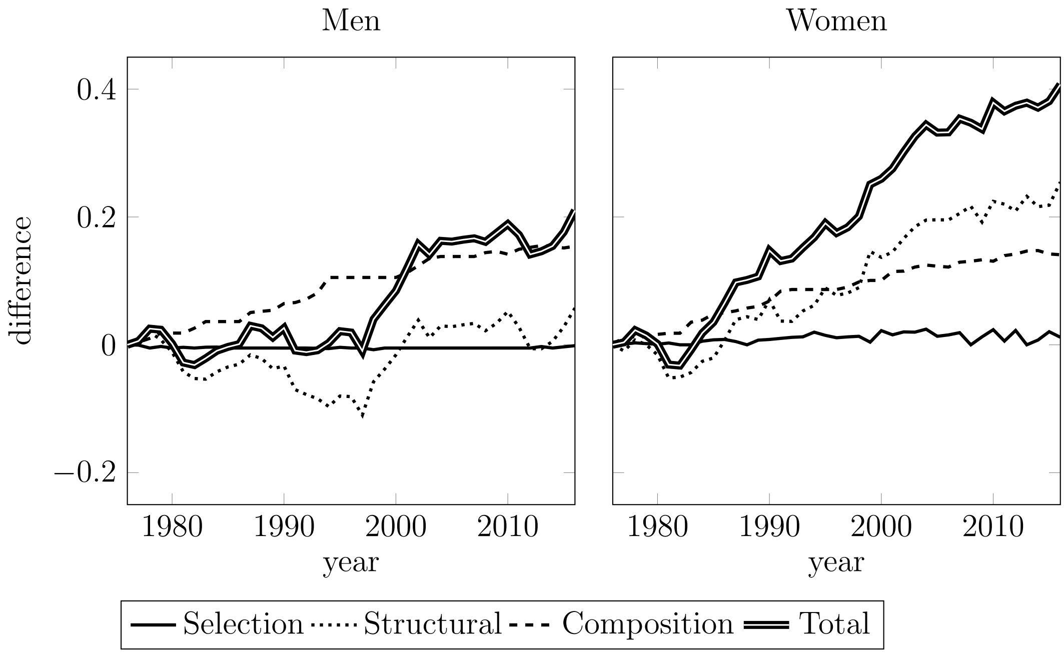

Figure 17: Decompositions at D9.

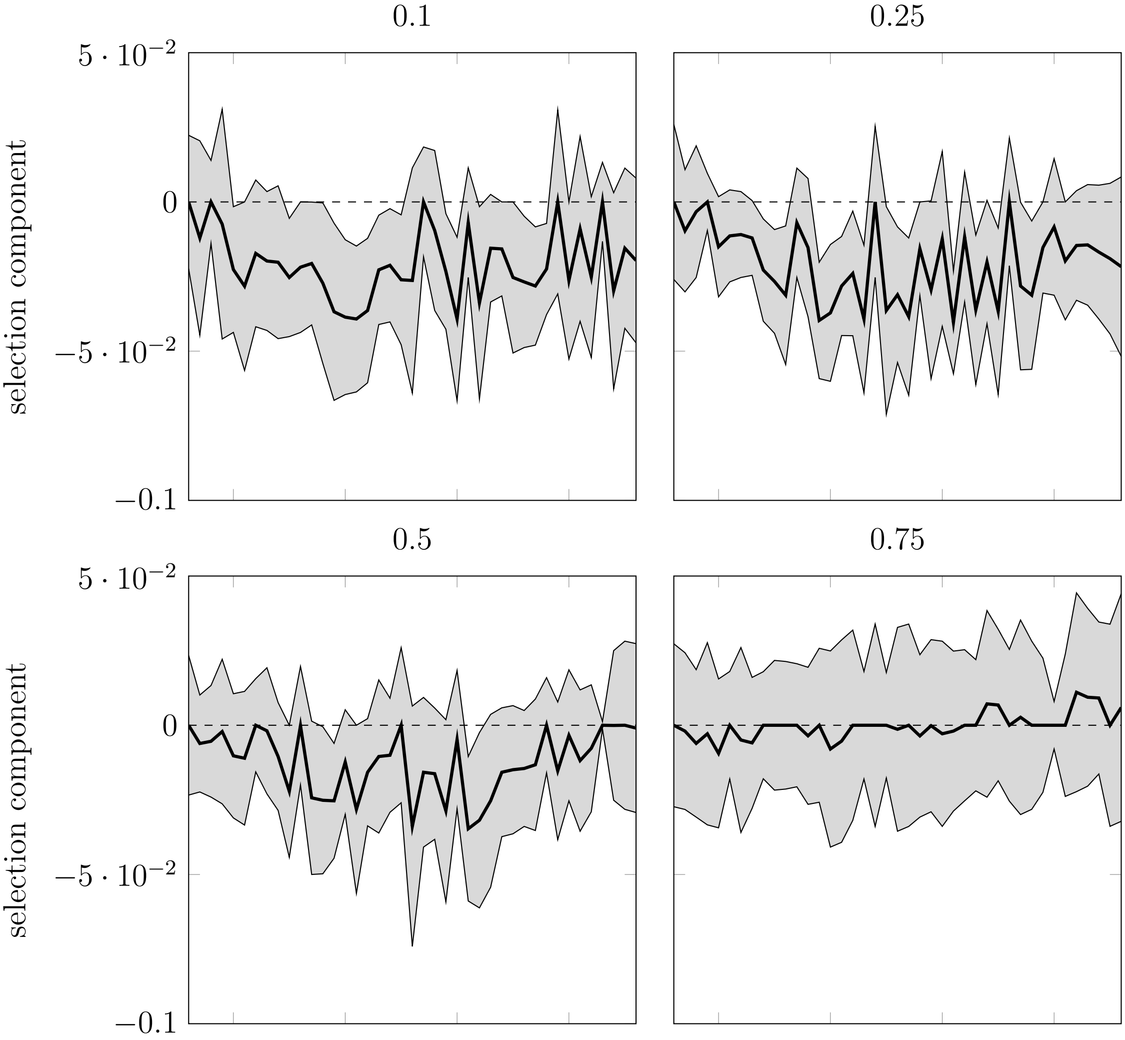

Figure 18: Selection component and 95-% confidence intervals for women.

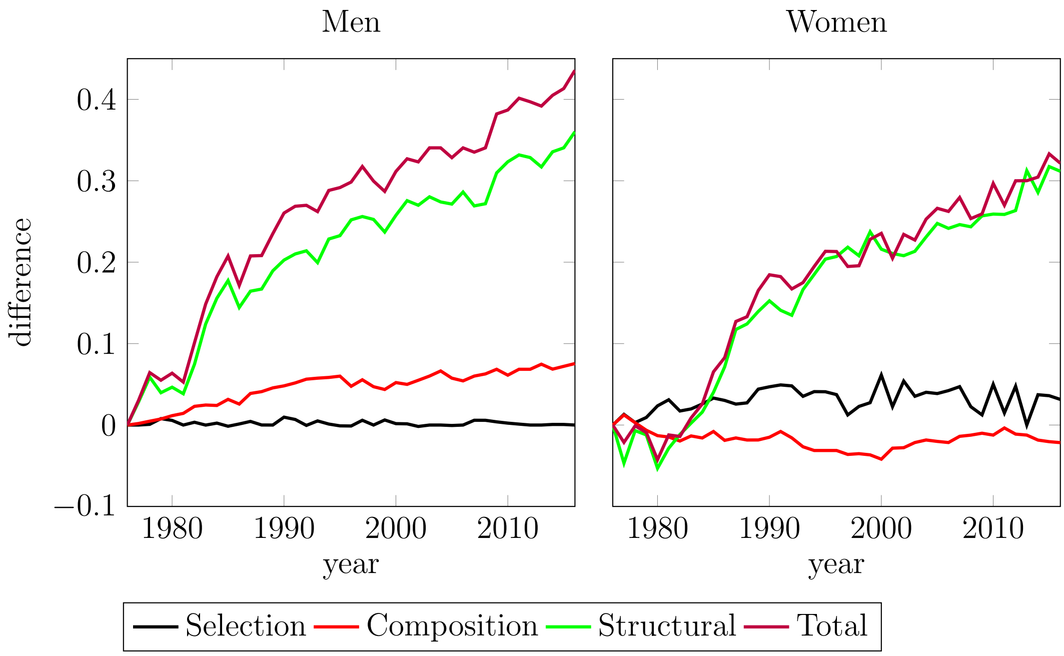

Figure 19: Decompositions at D9-D1 ratio.

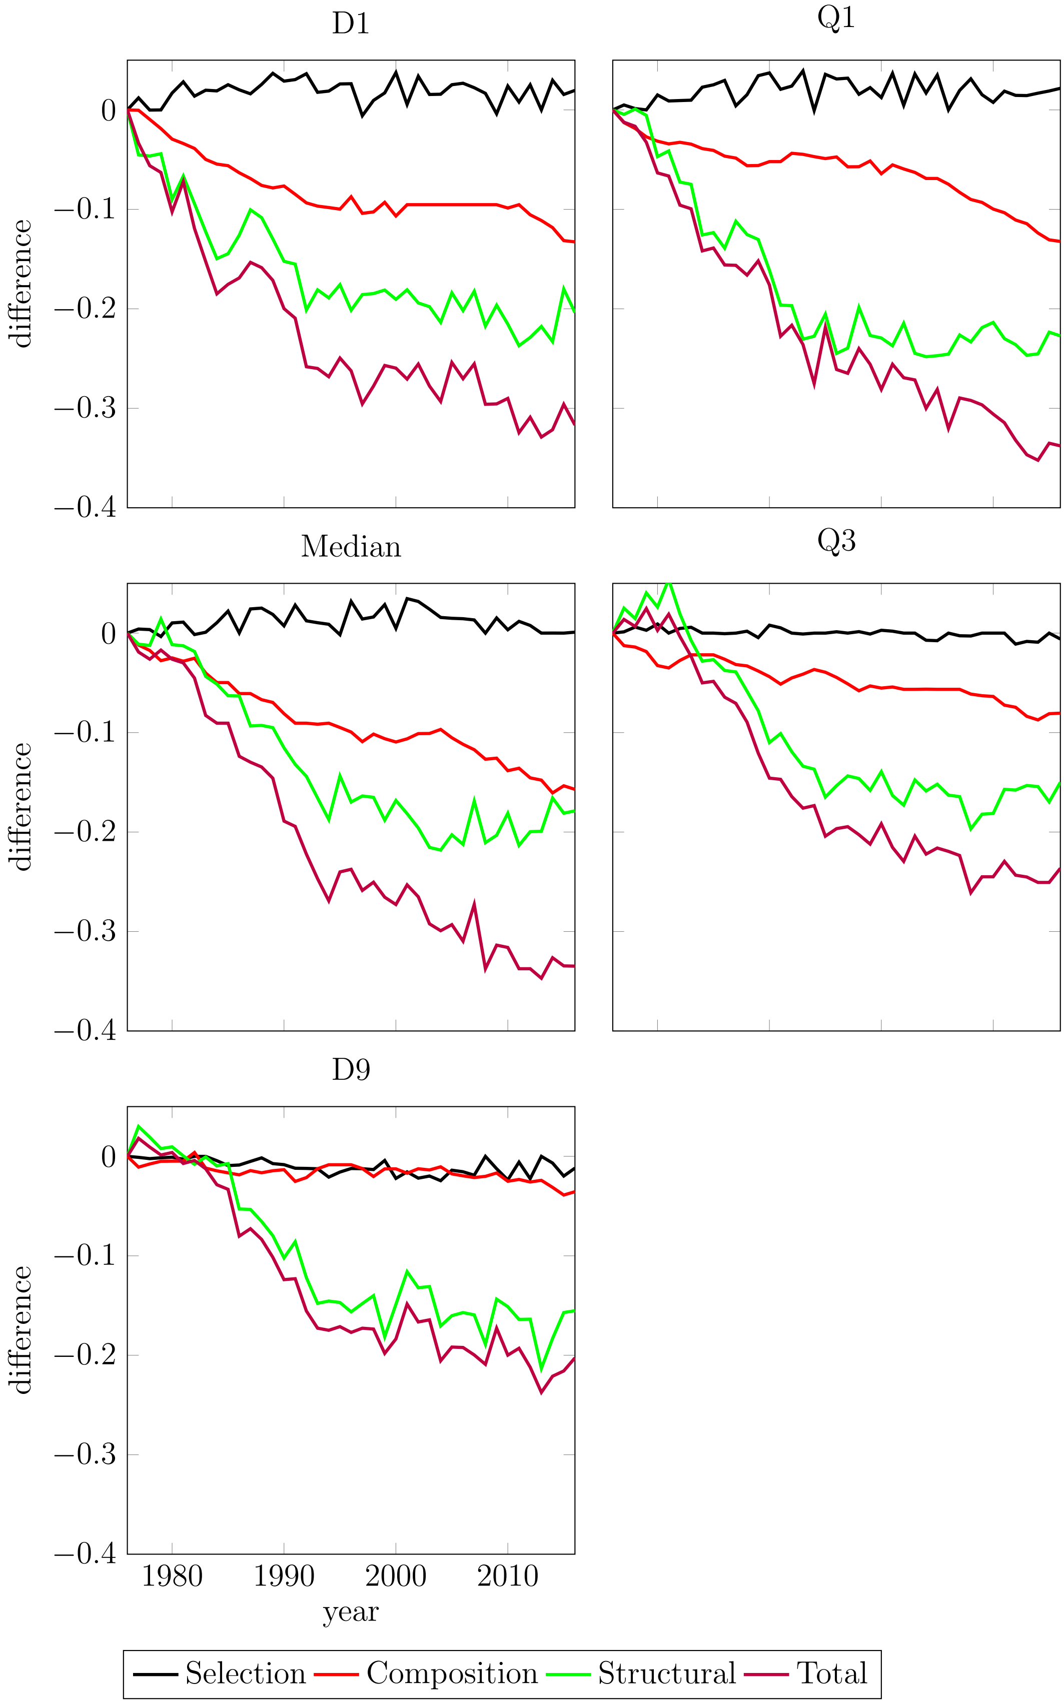

Figure 20: Decompositions of the wages of men divided by the wages of women.

Appendix

We briefly describe the FVV procedure and the estimation of the objects of interest. For further details see FVV. We estimate the control variable as the conditional distribution function of H given Z , via logistic distribution regression. That is:

$$\widehat{V}_i = \Lambda(P(Z_i)^\top \widehat{\pi}(H_i)), \quad i = 1, \dots, n,$$

where $\Lambda(u) = [1 + \exp(-u)]^{-1}$ is the logistic distribution function, $P(z)$ is a p-dimensional vector of transformations of z involving polynomials, and

$$\widehat{\pi}(h) = \arg\max_{\pi \in \mathbb{R}^p} \sum_{i=1}^{n} \big[1\{H_i \leq h\} \log \Lambda(P(Z_i)^\top \pi)) + \\ +1\{H_i > h\} \log\{1 - \Lambda(P(Z_i)^\top \pi)\} \big],$$

for $h \in \mathcal{H}_n$, the empirical support of H.

The estimator of the LASF is $\widehat{\mu}(x, v) = P(x, v)^\top \widehat{\beta}$, where $\widehat{\beta}$ is the least squares estimator and $P(x, v)$ is a p-dimensional vector of transformations of (x, v) involving polynomials and interactions:

$$\widehat{\beta} = \left[\sum_{i=1}^{n} P(X_i, \widehat{V}_i) P(X_i, \widehat{V}_i)^\top \right]^{-1} \sum_{i=1}^{n} P(X_i, \widehat{V}_i)^\top W_i.$$

The estimator of the LDSF is $\widehat{G}(y, x, v) = \Lambda(P(x, v)^\top \widehat{\beta}(w))$, where $\widehat{\beta}(w)$ is the logistic distribution regression estimator:

$$\widehat{\beta}(w) = \arg\max_{\beta \in \mathbb{R}^p} \sum_{i=1}^{n} \Big[1\{W_i \leq w\} \log \Lambda(P(X_i, \widehat{V}_i)^\top \beta)) + \\ +1\{W_i > w\} \log \Lambda(P(X_i, \widehat{V}_i)^\top \beta)) \Big].$$

Similarly, the estimator of the LQSF is $\widehat{q}(\tau, x, v) = P(x, v)^\top \widehat{\beta}(\tau)$, where $\widehat{\beta}(\tau)$ is the Koenker and Bassett (1978) quantile regression estimator:

$$\widehat{\beta}(\tau) = \arg\max_{\beta \in \mathbb{R}^p} \sum_{i=1}^{n} \rho_\tau(W_i - P(X_i, \widehat{V}_i)^\top \beta),$$

and $\rho_\tau(u) = [\tau - 1\{u < 0\}]u$ is the “check function”.

Estimators of the local derivatives are obtained by taking derivatives of the estimators of the local structural functions. The estimator of the LADF is:

$$\widehat{\delta}(x, v) = \partial_x P(x, v)^\top \widehat{\beta},$$

and the estimator of the LQDF is:

$$\widehat{\delta}_\tau(x, v) = \partial_x P(x, v)^\top \widehat{\beta}(\tau).$$

We obtain estimators of the generic global effects by approximating the integrals over the control variable by averages of the estimated local effects evaluated at the estimated control variable. The estimator of the effect (3) is

$$\widehat{\theta}_{x_0}(x) = \sum_{i=1}^{n} S_i K_i(x_0) \widehat{\theta}(x, \widehat{V}_i) / \sum_{i=1}^{n} S_i K_i(x_0),$$

for $K_i(x_0) = 1\{X_i = x_0\}$ when X is discrete, and $K_i(x_0) = k_h(X_i - x_0)$ when X is continuous, where $k_h(u) = k(u/h)/h$, k is a kernel and h is a bandwidth such as $h \to 0$ as $n \to 0$.

The estimator of the counterfactual distribution is:

$$\widehat{G}^s_{Y_{\langle t|k,r \rangle}}(y) = \sum_{i=1}^{n} \Lambda(P(X_i, \widehat{V}_i)^\top \widehat{\beta}_t(y)) 1\{\widehat{V}_i > \Lambda(P(Z_i)^\top \widehat{\beta}_r(0))\} / n^s_{kr},$$

where the average is taken over the sample values of $\widehat{V}_i$ and Z_i in group k, $n^s_{kr} = \sum_{i=1}^n 1\{\widehat{V}_i > \Lambda(P(Z_i)^\top \widehat{\beta}_r(0))\}$, $\widehat{\beta}_t(y)$ is the distribution regression estimator of step 2 in group t, and $\widehat{\beta}_r$ is the distribution regression estimator of step 1 in group r. Here we are estimating the components $F^s_{Y_t}$ by logistic distribution regression in group t and the component $F^s_{Z_k}$ by the empirical distribution in group k.

We do not have direct empirical analogs for the conditional distribution of the control variable $F^s_{V_r}$ and the support $\mathcal{V}_r(z)$ over all $z \in \mathcal{Z}_k$. We estimate these components using the empirical distribution of $\widehat{V}_i$ in group k, conditional on $\widehat{V}_i > \Lambda(P(Z_i)^\top \widehat{\beta}_r(0))$. This condition selects observations with the values of the explanatory and control variables from group k, which would have been selected with the parameters of the selection equation of group r. This estimation relies on:

$$F_V^s(v \mid Z = z) = 1\{v > F_C(0 \mid Z = z)\}\ v,$$

and the assumption that the ranking of the observations in the conditional distribution of the selection variable is invariant across groups.

Several explanations have been proposed for the rising earnings inequality in the U.S. and other developed economies: demand-side explanations focusing on skill-biased technological changes (see, e.g., Acemoglu 2002, Acemoglu and Autor 2011, Murphy and Topel 2016) or changes in the patterns of international trade, supply-side explanations focusing on the limited expansion of college graduates (Bound and Johnson 1992) or immigration of less-skilled workers, and institutional factors such as the de-unionization or the decline in the real value of minimum wages (Di Nardo, Fortin and Lemieux 1996 and Lee 1999).

The recent paper by Murphy and Topel (2016) provides a useful general equilibrium framework for interpreting the observed empirical evidence.

They argue that rising earnings inequality is "an equilibrium outcome in which endogenous human capital fails to keep pace with steadily rising demand for skills, driven by skill-biased technical change (SBTC) or other shifts in economic fundamentals, such as a decline in the price of capital, that favor highly skilled labor."

MT focus mainly on the supply side and distinguish three different margins on which the relative supply of skilled labor responds to a rise in its relative price.

The first margin is the choice of which type of human capital to invest in, proxied by the decision whether to attend or complete college.

The second is the decision of how much human capital of the chosen type to acquire.

The third is the decision of how intensively to apply the chosen amount of human capital to the market sector through effort, labor supply or occupational choice.

These different margins have different effects on earnings inequality.

Investments on the first margin, the "extensive margin", mitigate the impact of rising demand on the skill price, and thereby mitigate the resulting rise in inequality.

On the other hand, the choices on the other two "intensive margins", magnify the growth of inequality because they increase the quantity of human capital each worker employs: "rising returns to skills increase the incentives of able individuals to invest in human capital and, once it is produced, to use human capital more intensively."

MT argue that these forces are important in the light of the slowdown in educational attainments in the U.S.

They show that the supply of college graduates has flattened out after the 1960s among males, though not among females.

This shortfall of investment on the extensive margin not only contributes to inequality directly by driving up the price of skills, but also sets in motion supply responses on the intensive margin that cause further growths of inequality.

They conclude that remedies to the inequality problem lie on the supply side and consist of policies that encourage or enable the acquisition of skills or encourage the immigration of highly skilled individuals.

Other remarks:

- Because of substitutability, changes in the distribution of male and female wages are not independent.

- Lee (1999) argues that the decline in the real value of minimum wages accounts for most of the increased dispersion in the lower tail of the wage distribution observed during the 1980s, especially for women.

- Lemieux (2006) stresses the role of composition effects (increasing dispersion in the distribution of unobservable skills, especially during the 1980s and 1990s) and measurement errors (especially for wages paid by the hour), although measurement error does not appear to affect the trends over time. He also argues that "the March CPS does not provide a very accurate measure of wages for the majority of workers who are paid by the hour" and that "the ORG supplement of the CPS provides more accurate measures of hourly wages for much larger samples of workers than the March CPS." He points at two problems that he considers as open: (i) "why the overall growth in wage inequality is so concentrated in the 1980s", and (ii) "why wage inequality has mostly expanded in the upper end of the wage distribution".