

Selection, Heterogeneity and the Gender Wage Gap*

Cecilia Machado[†]

October 15, 2009

Abstract

Observed measures of the gender wage gap may be biased by selection because of differences in the characteristics of the workforce relative to the overall population. This problem is particularly severe for women, who have non-participation and missing information on wages at still very high levels. Existing approaches generally either impose a unique parametric correction procedure or assume the sign of average selection is known. In the case of female selection, however, both positive and negative rules likely co-exist, rendering the sign of average selection unknown. This paper proposes an alternative estimator which recovers a local measure of the gap in models with unobserved heterogeneity in selection. The local measure applies to the subpopulation of “always employed,” who are the women that remain employed in the presence of a young child. This is a relevant subpopulation for measuring the gender wage gap as the “always employed” women are similar to men in labor force attachment. In CPS data from 1976 to 2005, I show that the gender wage gap has narrowed substantially for this group, a more than twofold reduction from a $-.573$ to $-.267$ gap points (female minus male log wages).

*I thank Douglas Almond, Janet Currie, Lena Edlund and Edward Vytlačil for advice, and Tiago Berriel, Bo Honoré, Yinghua He, Mariesa Herrmann, Dennis Kristensen, David Lee, Bentley MacLeod, Marcelo Moreira, Serena Ng, Alexei Onatski, Cristian Pop-Eleches, Ricardo Reis, Yona Rubinstein, Johannes Schmieder, Till von Wachter, Reed Walker, and participants at the Columbia Applied Micro Colloquium, Applied Micro Workshop and Econometrics Lunch for useful discussions and comments. All errors are my own.

[†]ccm2116@columbia.edu

1 Introduction

The narrowing of the gender gap in recent decades has been one of the most striking changes in the US labor market. Whereas in 1979 the US gender log wage differential was $-.459$, it shrank to $-.227$ in 1998 (Blau and Kahn, 2006). Those estimates, however, may be biased by selection, since observed wages are sensitive to the characteristics of the individuals who opt into employment. In the US, the missing information problem is quite severe, as female full-time participation, albeit increasing, still averages just 50% in recent years.

Correction methods for the selection problem date back to Gronau (1974) and Heckman (1974). Recently, non-parametric bounds estimators of wage distribution parameters have been proposed under an alternative set of theoretical restrictions (Manski, 1990). Surprisingly, the two approaches have arrived to starkly different conclusions about the evolution of the gender wage gap.

Mulligan and Rubinstein (2008), using a parametric selection model (among other methods), found no reduction in the gender gap in the US: their selection-corrected measure has remained stable at around $-.338$ since the 1970's. They find that selection into employment has switched from negative to positive, with observed gains masking important changes in the composition of the workforce. In contrast, Blundell et al. (2007) use bounding procedures and find an improvement in the pay gap in the UK from 1978 to 1998: a substantial change that ranges from $.23$ to $.28$ log points. A key assumption in their bounding procedure is the imposition of a positive selection rule, as the estimated change in the gap becomes uninformative once this assumption is relaxed.

The stance taken on the self-selection process is critical because female non-participation remains high. In the case of parametric selection models, a monolithic rule, which does not vary with other unmeasured characteristics of women, is implicitly assumed. In the case of non-parametric bounds, the sign of average selection is generally imposed.

When it comes to female wages and employment decisions, however, it is questionable whether a unique or “homogeneous” selection mechanism can be assumed, or if the sign of the average selection effect is *a priori* known. While positive selection is the gold standard in traditional labor supply models, negative selection is also plausible for women. And if both positive and negative selection rules co-exist in different parts of the labor market, it is unclear which selection mechanism dominates.

Neal (2004) documents heterogeneity in female selection by race: while black women were positively selected into work, white women were negatively selected. He finds that the employment

decision of black women likely reflected low market wages, as non-working black women were generally under government assistance programs. White women’s decision, on the contrary, likely reflected high reservation wages, as non-working white women were more likely to be married and have high household income. More generally, heterogeneity in selection can also depend on unobserved characteristics: if couples match based on skills and out-of-work income rises with skills, one might conjecture that the employment decision of high skilled women reflects high reservation wages, while for low skilled women low potential earnings may be more important (Blundell et al., 2007).

This paper departs from the existing approaches when differing selection rules co-exist, and the average sign of the selection effect is unknown. Specifically, I consider the case in which the heterogeneous rules depend on the same unobservables that govern self-selection.

I first generate an example with simulated data to show that under unobserved heterogeneity in the selection rule both the parametric selection model and the non-parametric bounds that sign the average selection by assumption might fail to recover meaningful information on wages. I then show that a local measure of wages can still be recovered under a class of models that satisfy a monotonicity condition on the participation response. The local measure of wages can be recovered for the “always employed,” and draw the analogy to the “always takers” in treatment effect models (Angrist et al., 1996). This subpopulation corresponds to individuals who participate in the labor market regardless of the value of the instrument. In the case of the gender wage gap, the “always employed” women may be a particularly relevant subpopulation to compare to men, as male labor force attachment is high and stable.

Using CPS data, and the presence of a child younger than six as the instrument for female employment, I estimate an improvement in the US gender gap for this group from 1976 to 2005 of .30 log wage points, a more than twofold improvement, from -.573 to -.267 points. In the same data, I also replicate the finding in Mulligan and Rubinstein (2008) and in Blundell et al. (2007), showing that the methods, rather than differences in the sample or in the choice of the instrument, are the driving force for the disparate results. As unobserved heterogeneity in selection is a relevant feature of female employment decisions, the estimator proposed in this paper constitutes an alternative parameter to be incorporated in studies of the gender wage gap.

The results are maintained under a series of robustness exercises, such as using alternative covariate controls, accounting for male selection and redefining the instrument by different ages of the young child present. Finally, following (Angrist et al., 1996), I assess the bias of the estimator

when the identifying assumptions fail, and show that relaxing some of them yield qualitatively similar results.

This paper is organized as follows. The next section outlines the selection problem with missing outcomes, provides an example of unobserved heterogeneity in the selection rule and identifies a class of models under which a local measure of potential wages can be recovered. Section 3 presents the baseline estimates and section 4 contains the robustness exercises. A discussion on the validity of the assumptions necessary to identify the local gender gap is found in section 5. Section 6 concludes.

2 Heterogeneity and Selection with Missing Outcomes

Let \mathcal{Y} denote the potential wages received by individuals in the market place and Y the observed wages. The selection problem arises because $Y = \mathcal{Y}$ only for those found to be employed. Since individuals who choose employment are plausibly different from the ones who do not, the observed distribution of wages does not generalize to the entire population. Thus, a simple OLS regression of observed wages on a gender dummy will mask selection effects in both the male and female population.

How big is the selection problem? Denote by $E \in \{0, 1\}$ the employment status and $G \in \{0, 1\}$ the gender indicator, with $G = 1$ for women. Conditional on covariates X , the unobserved mean of potential wages can be decomposed as:

$$E(\mathcal{Y}|X, G) = E(Y|E = 1, X, G)\Pr(E = 1|X, G) + \underbrace{E(\mathcal{Y}|E = 0, X, G)}_{\text{selection effect}} \underbrace{\Pr(E = 0|X, G)}_{\text{non-participation effect}}. \quad (1)$$

In words, the parameter of interest $E(\mathcal{Y}|X, G)$ is a weighted average of wages among participants and non-participants, and the unobservability of $E(\mathcal{Y}|E = 0, X, G)$ poses the main challenge in any estimation strategy.

If $E(Y|E = 1, X, G) \geq E(\mathcal{Y}|E = 0, X, G)$, the selection effect is said to be positive, and individuals with higher wages are more likely to be working. Similarly, selection is negative when $E(Y|E = 1, X, G) \leq E(\mathcal{Y}|E = 0, X, G)$. The non-participation effect only magnifies the role of $E(\mathcal{Y}|E = 0, X, G)$: for higher non-participation rates, the more missing information there is on wages of individuals out of employment.

Either positive and negative rules can be justified by standard labor supply models. But without

prior information on the selection effect, $E(\mathcal{Y}|E = 0, X, G)$ could assume a wide range of values. A seemingly conservative approach is to put bounds on $E(\mathcal{Y}|X, G)$. With \mathcal{Y} bounded, the best and worse case scenarios can be constructed by taking $E(\mathcal{Y}|E = 0, X, G)$ to be either its extreme lowest or highest values $\underline{\mathcal{Y}}$ and $\overline{\mathcal{Y}}$:

$$\begin{aligned} E(Y|E = 1, X, G)\Pr(E = 1|X, G) + \underline{\mathcal{Y}}\Pr(E = 0|X, G) \\ \leq E(\mathcal{Y}|X, G) \leq \\ E(Y|E = 1, X, G)\Pr(E = 1|X, G) + \overline{\mathcal{Y}}\Pr(E = 0|X, G). \end{aligned} \quad (2)$$

Applying bounds to the evolution of the gender wage gap is further challenged by differences in selection and participation across genders and across time. This can be seen by letting X refer to time T , with $T \in \{t_1, t_2\}$:

$$\begin{aligned} \Delta(t_2) - \Delta(t_1) = \left\{ E(\mathcal{Y}|T = t_2, G = 1) - E(\mathcal{Y}|T = t_2, G = 0) \right\} \\ - \left\{ E(\mathcal{Y}|T = t_1, G = 1) - E(\mathcal{Y}|T = t_1, G = 0) \right\}, \end{aligned} \quad (3)$$

where each term in the expression has bounds given by (2). Even if wide (e.g., including zero), the resulting bounds from (3) highlight the magnitude of the missing data problem. The lack of information on $E(\mathcal{Y}|E = 0, X, G)$ becomes more important the higher the non-participation rate, and this is particularly troublesome in studies of the gender wage gap, as female participation rate is still relatively low, despite increasing over the past decades. Existing attempts to sign $\Delta(t_2) - \Delta(t_1)$ must impose more structure into a selection model.

2.1 Existing Approaches

There are three main approaches in the literature on selection with missing outcomes. The first approach uses information on the observed covariates X and restrictions motivated by economic models to impute values for the missing data (Neal, 1996; Johnson et al., 2000; Neal, 2004; Blau and Kahn, 2006; Olivetti and Petrongolo, 2007). For example, if non-participants earn less than median wages, imputing zero wages for those with missing information does not bias the estimate of median wages (Johnson et al., 2000). Alternatively, if selection into employment is purely random after we control for a very detailed set of observed covariates, one can match similar individuals

based on X , and impute the wages of non-participants by the mean wage of participants (Olivetti and Petrongolo, 2007).

A second approach acknowledges that selection can be based on unobservables and models the self-selection process. In general, the correction procedure amounts to including an extra term in the wage equation, the control function, which is either known, as in parametric models (Heckman, 1974; Mulligan and Rubinstein, 2008), or, when unknown, estimated by semi-parametric methods (see Vella (1998) for a detailed discussion on this literature). Identification under the control function approach requires an exclusion restriction¹, that is, an instrument Z that shifts employment but is unrelated to wages, and the estimated model generally falls into a variant of a single index partially linear model.

The third approach, while still accounting for unobserved selection, does not explicitly model self-selection. This bounding approach to the selection problem has equations similar to (2) as its starting point, and narrows bounds using restrictions motivated by econometric or economic theory. For example, Manski (1990) shows how the availability of an instrument can reduce the width of the bounds. Intuitively, an instrument shifts participation and reduces the weight placed on the unobserved wages of non-participants. In Blundell et al. (2007), the availability of an instrument is combined with a positive selection assumption and an additivity restriction in the wage equation for an empirical assessment of the gender wage gap in the UK.

A common denominator in all three branches of the literature is the imposition of some structure into the selection process. In imputation methods, selection on unobservables is assumed away as missing wage information is filled in based on observed covariates. Parametric, semi-parametric methods and the bounding approach all account for unobserved selection, but either assume a unique rule (or correction procedure), which is invariant to unobservable characteristics, or assume knowledge of the sign of average selection in order to derive informative answers from the bounds.

Where female employment and wages are concerned, such assumptions may be especially tenuous. While positive selection is the norm for male selection, negative selection is also plausible for women. Neal (2004) provides evidence of different female selection rules by race. More generally, selection rules can also differ by unobserved characteristics. For example, if couples match based on skills or some other unobserved productivity component, and out-of-work income rises with skills, one might conjecture that the employment decision of more skilled women reflect high reservation

¹The exclusion restriction is not necessary in parametric models, but, in practice, identification without an instrument is weak, as the correction term is often a linear function of the variables entering the outcome equation directly (Vella, 1998).

wages rather than low potential earnings. Where both positive and negative rules co-exist, the sign of average selection is also debatable.

The approach pursued in this paper departs from the previous literature in that it does not impose any structure on (or recover information about) the selection rule. It instead targets identification of mean wages under unobserved heterogeneity in the selection rule. By unobserved, I mean that the positive and negative selection rules can also depend on the unobservable determinants of wages and employment. The next section provides a stylized example illustrating those core features.

2.2 Unobserved Heterogeneity in Selection: An Example

Let female wages and employment be generated by the model:

$$\begin{aligned} Y_i &= \mathcal{Y}_i \text{ if } E_i = 1 & (4) \\ E_i &= \mathbf{1}[\alpha + \gamma Z_i \geq \epsilon_i] \\ Z_i &\perp\!\!\!\perp (\mathcal{Y}_i, \epsilon_i), \end{aligned}$$

with parametrization and data generating process (DGP) given by:

$$\begin{aligned} \epsilon_i &\sim N(0, \sigma_\epsilon^2), \quad \sigma_\epsilon = 4 & (5) \\ Z_i &\sim \text{Binomial}(0.5) \\ \alpha &= 2.5, \quad \gamma = -5 \\ \mathcal{Y}_i &= \begin{cases} (\rho_N * \sigma_N / \sigma_\epsilon) * \epsilon_i + \zeta_i & \text{if } \epsilon_i \geq 0 \text{ (type N)} \\ (\rho_P * \sigma_P / \sigma_\epsilon) * \epsilon_i + \zeta_i & \text{if } \epsilon_i < 0 \text{ (type P)} \end{cases} \\ \zeta_i &\sim N(0, 1) \\ (\rho_N, \sigma_N) &= (0.5, 4) \\ (\rho_P, \sigma_P) &= (-0.5, 4). \end{aligned}$$

The model in (4) is a standard labor supply model where \mathcal{Y}_i and ϵ_i are the unobservables that jointly determine employment and wages. In the terminology of those models, ϵ_i corresponds to the difference in the unobservables of the reservation and market wage equations². Under a joint

²This can be seen by letting the reservation wages Y_i^R be given by $Y_i^R = -\alpha - \gamma Z_i + \xi_i$. Thus, $E_i = 1$ if $\mathcal{Y}_i \geq Y_i^R$ or, alternatively, $\alpha + \gamma Z_i \geq \xi_i - \mathcal{Y}_i = \epsilon_i$.

normality assumption, the correlation between \mathcal{Y}_i and ϵ_i determines the sign of selection: if the correlation is positive, selection is said to be negative, and vice-versa³.

Heterogeneity in the selection rule is captured by the co-existence of two rules, N and P, corresponding to a negative and a positive selection mechanism, respectively. This heterogeneity is unobserved because it depends on ϵ_i , being of type N if $\epsilon_i \geq 0$ and of type P if $\epsilon_i < 0$. Since $\epsilon_i \sim N(0, \sigma_\epsilon^2)$, the two types are equally likely. If $\rho_N = \rho_P$, we return to the conventional setup.

Table (1) displays summary statistics for 1,000 datasets, with 10,000 observations each, generated by the above model and DGP. Panel A of the table reports some statistics on wages and employment that would be observed by researchers. Mean wages among the employed are 1.596, and are based on information about the 50% of the population that chooses to participate. The instrument Z decreases employment by 0.469 percentage points (0.734-0.265).

Panel B displays information on \mathcal{Y} by the underlying selection rule, which could only be recovered under knowledge of the true model. The average of potential wages is 1.595 and is very close to the observed mean wages among employed. Individuals with selection rule N, are less likely to be employed than type P, as higher values of ϵ_i meet the the employment threshold equation less frequently. As would be expected, type N individuals are negatively selected into employment, with $E(\mathcal{Y}|E = 0, N) \geq E(\mathcal{Y}|E = 1, N)$, as $\rho_N > 0$. Type P individuals are positively selected.

Because the two selection rules are symmetric, type N and P women have the same distribution of wages. This can be seen in figure (1), which plots the distribution of wages for one of the 1,000 datasets. The two types only differ with respect to ϵ_i , and, consequently, on the likelihood of employment. This model captures the idea that some women will fare better in terms of unobserved reservation wages, even though their market wages are on average the same. The women with $\epsilon_i \geq 0$ have their employment decision guided strongly by high reservation wages (relative to market wages), and among themselves, the ones that work have the lowest market wages (negative selection). The symmetric reasoning holds for women with $\epsilon_i < 0$. As an example, ϵ_i could be seen as an unobserved characteristic that make some women fare better in the marriage market and enjoy high reservation wages⁴.

Panel C presents estimates on \mathcal{Y} under two common approaches used to recover female wages (and the gender gap): a parametric selection model that assumes joint normality of unobservables

³When comparing \mathcal{Y}_i and $-\epsilon_i$ - the difference between market and reservation wages - the sign of the correlation and selection go on the same direction.

⁴If ϵ_i proxies some unobserved skill valued both in marriage and labor market, market wages could also differ by ϵ_i . This feature can be incorporated by letting $|\rho_N| \geq |\rho_P|$, but only strengthens the results presented in this section.

and a unique selection rule, as in Mulligan and Rubinstein (2008), and non-parametric bounds on mean and median wages, assuming the sign of the average selection is known, as in Blundell et al. (2007).

The parametric selection model includes the Mills ratio as a control in the wage equation, and estimates the mean wages to be 0.605, which deflates its true measure of 1.595 by more than half. Based on that estimation framework, one would conclude that selection is positive, as wages of the employed, 1.596, are higher than the estimated measure of $E(\mathcal{Y})$, even though positive selection is no more likely than negative selection.

Placing non-parametric bounds on mean and median wages combines the availability of the instrument Z with the assumption of either a positive or negative average selection mechanism. Information on the construction of the bounds is contained in appendix D.1. Since only 50% of sample chooses employment, best and worse case bounds on $E(\mathcal{Y})$ are very wide, and range from a low of 0.119 to a high of 3.229. Imposing average positive selection reduces the upper bound to 1.280, and the resulting bounds range from 0.119 to 1.280. Similarly, under negative selection the lower bound becomes 2.468, with the upper bound of 3.229 being maintained. Bounds under either positive or negative selection assumptions also miss $E(\mathcal{Y})$, as average selection in this example is zero. With respect to median median wages, bounds are relatively tighter, but still range from 0.501 to 1.825. As in the case of mean wages, the positive selection assumption reduces upper bound to 1.129, and misses median wages, which is 1.465 in this example. The negative selection assumption is rejected in this case, as upper and lower limits cross.

This example shows that under unobserved heterogeneity in the selection rule, both the parametric correction and the non-parametric bounds that sign the average selection may fail to recover $E(\mathcal{Y})$ (or $\text{Med}(\mathcal{Y})$). Nonetheless, a meaningful parameter can still be recovered for models of this type, a local measure of potential wages, even when selection rule is unobserved. The next section describes the necessary conditions for the identification of this parameter.

2.3 A Local Measure of Potential Wages

This section is built upon the potential outcome notation of causal models, as in Rubin (1974) and Heckman (1990). In fact, the selection problem with missing outcomes is a particular case of treatment effect models, with the outcome being observed only for the individuals that opt in. In contrast to those models, however, the goal is to recover information on a parameter for the entire population, rather than inferring the causal effect of a treatment.

I adhere to two conventions in the literature. First, I do not explicitly model observed covariates, and all is taken to be conditional on $X = x$. Second, I abstract from general equilibrium effects, even though they might be a relevant concern when extrapolating the results to universal participation.

For a binary Z , define E_1 and E_0 as the potential participation status when Z is externally set to 1 and 0 respectively. The model reads:

(AI) Existence of an instrument:

$$\text{Independence:} \quad Z \perp\!\!\!\perp (\mathcal{Y}, E_0, E_1).$$

$$\begin{aligned} \text{Nontrivial } Z: \quad & \Pr(E = 1|Z = z) - \Pr(E = 1|Z = z') \neq 0, \quad z \neq z', \\ & \text{with } \Pr(E = 1|Z = z) > 0, \quad \forall z \in \{0, 1\}. \end{aligned}$$

(AII) Exclusion Restriction:

$$\begin{aligned} Y &= \mathcal{Y} \text{ if } E = 1 \\ E &= ZE_1 + (1 - Z)E_0. \end{aligned}$$

(AIII) Monotonicity: Either $E_1 \leq E_0$ or $E_1 \geq E_0$ for all individuals.

Under (AI)-(AIII), a local measure of potential wages, $E(\mathcal{Y}|E_1 = 1, E_0 = 1)$, can be identified. For $E_1 \leq E_0$:

$$\begin{aligned} E(Y|E = 1, Z = 1) &= E(\mathcal{Y}|E_1 = 1) & (6) \\ &= E(\mathcal{Y}|E_0 = 1, E_1 = 1)\Pr(E_0 = 1|E_1 = 1) \\ &+ E(\mathcal{Y}|E_0 = 0, E_1 = 1)\Pr(E_0 = 0|E_1 = 1) \\ &= E(\mathcal{Y}|E_0 = 1, E_1 = 1)\Pr(E_0 = 1|E_1 = 1) \\ &= E(\mathcal{Y}|E_0 = 1, E_1 = 1). \end{aligned}$$

For $E_1 \geq E_0$, a similar reasoning shows that $E(\mathcal{Y}|E_1 = 1, E_0 = 1)$ is identified by $E(Y|E = 1, Z = 0)$. In the terminology of Angrist et al. (1996), $(E_0 = 1, E_1 = 1)$ are the “always takers,” i.e., the women who always work. In contrast to the treatment effect literature, which uses the IV to identify the treatment effect among “compliers”, the IV is used here to identify potential wages for the subsample of individuals who do not change their employment decision and remain working no matter the value of Z .

Assumption (AIII) rules out the existence of either $(E_0 = 0, E_1 = 1)$ or $(E_0 = 1, E_1 = 0)$, and the local estimator of $E(\mathcal{Y}|E_0 = 1, E_1 = 1)$ is sensitive to the excluded type: it is given by $E(Y|E = 1, Z = 1)$ when $E_1 \leq E_0$ and by $E(Y|E = 1, Z = 0)$ when $E_1 \geq E_0$. Technically, the two formulations are equivalent, as Z can be arbitrarily relabeled⁵, but in practice the monotonicity direction should be inferred, so its corresponding estimator can be applied.

A simple check comes by noting that the model represented by (AI)-(AIII) is alternatively represented by the model in (4), which was previously used in the simulation exercise⁶. Under (AIII), the direction of monotonicity can be recovered by verifying whether the instrument decreases or increases the employment probability, ie, whether $\Psi = \Pr(E = 1|Z = 1) - \Pr(E = 1|Z = 0)$ is negative or positive. A negative Ψ rules out the $(E_0 = 0, E_1 = 1)$ type individuals, and monotonicity holds in the decreasing direction, with $E_1 \leq E_0$. Similarly, a positive Ψ rules out the $(E_0 = 1, E_1 = 0)$ type individuals, and monotonicity holds in the increasing direction, with $E_1 \geq E_0$ ⁷.

Moving back to table (1), panels A and D show that $E(Y|E = 1, Z = 1)$ matches $E(\mathcal{Y}|E_1 = 1, E_0 = 1)$. This is the case because $\Psi < 0$, which implies monotonicity in a decreasing direction, with $E_1 \leq E_0$.

In the context of selection with missing outcomes, I will refer to $(E_0 = 1, E_1 = 1)$ as the “always employed,” $(E_0 = 0, E_1 = 0)$ as the “never employed,” $(E_0 = 1, E_1 = 0)$ as the “switchers,” and $(E_0 = 0, E_1 = 1)$ as the “defiers”⁸. The defiers will be assumed away by imposing $E_1 \leq E_0$, as the instrument employed in the empirical analysis that follows decreases the probability of participation.

⁵That is, if $E_{Z=1} \geq E_{Z=0}$, one can redefine Z' to be 1 when $Z = 0$ and 0 when $Z = 1$ to have $E_{Z'=1} \leq E_{Z'=0}$

⁶The equivalence of both formulation has been established by Vytlačil (2002)

⁷This can be formally seen by examining the expressions for $\Pr(E_0 = 0, E_1 = 1)$ and $\Pr(E_0 = 1, E_1 = 0)$. If $\gamma < 0$, then:

$$\begin{aligned} \Pr(E_0 = 1, E_1 = 0) &= \Pr(\alpha < \epsilon_i, \alpha + \gamma \geq \epsilon_i) \\ &= \Pr(\alpha < \epsilon_i \leq \alpha + \gamma) \\ &= F_\epsilon(\alpha + \gamma) - F_\epsilon(\alpha) \\ &= \Pr(E = 1|Z = 1) - \Pr(E = 1|Z = 0) \\ \Pr(E_0 = 0, E_1 = 1) &= \Pr(\alpha \geq \epsilon_i, \alpha + \gamma < \epsilon_i) = 0. \end{aligned}$$

Similarly, if $\gamma > 0$:

$$\begin{aligned} \Pr(E_0 = 1, E_1 = 0) &= 0 \\ \Pr(E_0 = 0, E_1 = 1) &= \Pr(E = 1|Z = 0) - \Pr(E = 1|Z = 1). \end{aligned}$$

Thus, since $\Psi = \Pr(E = 1|Z = 1) - \Pr(E = 1|Z = 0)$ is either positive or negative, either $(E_0 = 1, E_1 = 0)$ or $(E_0 = 0, E_1 = 1)$ will be assumed away by monotonicity.

⁸Note that the “switchers” and “defiers” here correspond to the “defiers” and “compliers” of Angrist et al. (1996).

3 The Gender Wage Gap under Heterogeneous Selection Rules

3.1 Data

The data used in this paper comes from the Annual Demographic File (ADF) of the Current Population Survey (CPS) from 1976 to 2005 and follows the sample restriction typically employed in studies of the gender gap: I focus on white non-hispanic adults between ages of 25 and 44. The age restriction is tighter than in previous studies⁹ because the instrument employed in this paper, which is fertility related, affects women of childbearing age. I define participation by two employment variables: any work and full-time-full-year work (35+ hours per week and 50 weeks or more) during the year. The outcome variable is log hourly wages. More details on the construction of this sample is found in appendix A.

The instrument Z is a binary indicator for a presence of a child less than 6 years-old in the family. The bulk of the variation in this variable comes before age 44, as only 2% of women between 45-54 have a child younger than 6 years old. Moreover, although this variable is originally multivalued in the CPS survey, roughly 90% of my sample has either no children or only one child below the age of six, motivating the classification of the binary instrument. The choice of the instrument, although questionable, follows the previous literature. In Heckman (1974), one of the seminal works on female selection, number of children is used as an explanatory variable in the shadow price function. More recently, Mulligan and Rubinstein (2008) have used number of children younger than 6 interacted with marital status as variables determining employment, which are excluded from the market wage equation. A discussion about this instrument, and its relation to assumptions (AI)-(AIII), is found in section 5 of this paper.

Summary statistics for the data are displayed in table (2). Female participation increases from 65% to 80% over the period of analysis, a trend that is followed by the full-time full-year (FT) rate, at lower levels. Still by 2005, FT wages are only observed for 50% of women in the sample. The very high degree of missing wage information in the FT sample justifies having any employment as an alternative participation variable, bearing in mind that hourly wages for part-time workers could be smaller on average, and that the fraction of part-timers should be higher in the female population.

Relative to women, male employment rates are substantially more stable, though over this

⁹In Mulligan and Rubinstein (2008), the sample encompasses ages 25-54, and in Blau and Kahn (2006) it includes ages 18-65

period FT wages are not observed for more than 20% of men. Since the extent of missing information is greater for women, and a valid instrument for male labor supply is hard to find, I proceed by assuming that the observed wages of men proxy the latent distribution of potential wages for all men, but I will relax this assumption in section 4.

The race, ethnicity and age restrictions on the sample makes the universe of men and women very similar in observables aside from two other characteristics, which are marital status and education. Although the fraction married is similar for both male and female populations, the education distribution and its evolution between 1976 and 2005 does not display a similar pattern across gender. For instance, the fraction with a college degree or more increases 5 percentage points for men, a 16% change, whereas it increases by 16 points for women, an almost twofold change. The empirical analysis that follows takes X to be education and stratifies results by 4 groups: less than high school, high school graduates, some college and college graduate or more.

3.2 Estimation

The parameter of interest is the gender wage gap between the “always participating” women and men of similar characteristics. This local estimator of the gap depends on how the instrument changes women’s participation: whether in an increasing or decreasing direction. A first step inspects whether

$$\Psi_{xt} = \Pr(E = 1|T = t, X = x, G = 1, Z = 1) - \Pr(E = 1|T = t, X = x, G = 1, Z = 0) \quad (7)$$

is positive or negative. Under the monotonicity assumption (AIII), a negative Ψ_{xt} implies $E_1 \leq E_0$, rendering $E(Y|E = 1, Z = 1)$ as the local estimator of women’s wages. The opposite applies when Ψ_{xt} is positive.

Abstracting from selection effects in the male population, the second step is a simple OLS regression where the instrument Z enters interacted with gender:

$$Y_i = \beta_{0xt} + \beta_{1xt}G_i + \beta_{2xt}G_i * Z_i + u_i. \quad (8)$$

The local measure of the gap is then given by:

$$\Delta(x, t) = E(\mathcal{Y}|E_1 = 1, E_0 = 1, G = 1, X = x, T = t) - E(\mathcal{Y}|G = 0, X = x, T = t) = \begin{cases} \beta_{1xt} & \text{if } \Psi_{xt} > 0 \\ \beta_{1xt} + \beta_{2xt} & \text{if } \Psi_{xt} < 0 \end{cases}$$

3.3 Results

Table (4) displays the first stage results stratified by the four education groups for the first and last years of the sample. The presence of a child under six decreases female employment, both for any work or full time work, with the effect being stronger at higher levels of education (where participation levels are higher). Overall, results indicate that the presence of a child younger than six decreases participation, and the sensitivity is slightly smaller for years 2001-05 relative to 1976-80.

Since $\hat{\Psi}_{xt}$ is negative for all education groups and periods, implying monotonicity in a decreasing direction, the local measure of the gender wage gap is recovered by $\hat{\beta}_{1xt} + \hat{\beta}_{2xt}$. Results for equation (8) are displayed in table (5). Each panel of the table, one for each education group, has four regressions, which differ according to years, 1976-1980 vs 2001-2005, and to the participation classification, any employment or FT.

Education-wise comparison of the gender gaps indicate that they get smaller (in absolute value) as education increases. Whether they measure pure discrimination or other unobserved gender differences, women in the high end of the education distribution are found less subject to a penalty. Nonetheless, the local measure of the wage gap has decreased for all education groups between 1976 and 2005. The largest improvement in the gap, a twofold reduction, has occurred for the group with a college degree or more. For them, note that the local gap has closed by 20 percentage points, whereas the observed (or uncorrected) gap, displayed in the last line of the panel, indicates only a 10 percentage point reduction.

The above results can be summarized by weighting each education gap $\Delta(x, t)$ by corresponding education proportions. Since the education distribution varies over time and by gender (see table (2)), alternative weighting schemes can be employed. I consider four types of weights and display results in table (6). The first weight, the female variable weight, uses the female education proportions in each time period, $p_{xt}^{FV} = \Pr(X = x|G = 1, T = t)$, and computes the average gap

by:

$$\Delta(t) = \sum_{x=1}^4 \Delta(x, t) * p_{xt}^{FV}. \quad (9)$$

The observed evolution of the gap, without selection corrections, provides a modest proxy for the gap of “always employed women”: .306 versus .238 points for the ones with any employment and .258 versus .182 for the ones in full time full year work. As changes in this average gap reflect changes in each conditional gap as well as changes in the education composition of the female workforce, the next two weights in table (6) hold education fixed using either its 1976-80 or its 2001-2005 proportions. The female fixed 1976-1980 weight uses $p_{xt}^{7680} = \Pr(X = x|G = 1, T = 1976-1980)$ and the female fixed 2001-2005 weight uses $p_{xt}^{0105} = \Pr(X = x|G = 1, T = 2001-2005)$. These alternative weighting schemes show that although part of the improvement is due to changes in the educational composition of the female workforce, the bulk of the change is due to a uniform reduction in the gender gap for each education category. Taking the education proportions in the male population as weights, $p_{xt}^{MV} = \Pr(X = x|G = 1, T = t)$, the gains are slightly smaller and reflect the fact that the education proportions in the female population are skewed towards the groups with the highest gains.

3.4 Comparison to Previous Gender Gap Estimates

Putting the results of the previous section into perspective, figure (2) compares $\Delta(t)$ to other measures estimated in literature: the observed (or “uncorrected”) evolution of the gap and the gap from a parametric selection model. Appendix C outlines how those two measures were obtained in my sample.

The local gap portrayed in the figure has participation defined as full time work and weighs the education groups by p_{xt}^{FV} . The initial and final estimates of the local gap in the figure correspond to the numbers in the first line of table (6), panel B, columns (1)-(3): an improvement of 25 log wage points, from -0.521 to -0.263. The observed gap displays a similar trend, with the measured improvement being lower, at 18 points, from -0.447 to -0.265. In contrast, the measure from a parametric selection model shows no improvement of the pay gap, which remains around -.30 points from 1976 to 2005. Note that these numbers closely track the estimates in the literature¹⁰.

I also estimate non-parametric bounds on the median wage gap in my sample following the

¹⁰Blau and Kahn (2006) measure the observed differential as being -.459 in 1979 and -.227 in 1998 using PSID data. Mulligan and Rubinstein (2008) measure the observed differential to be -.414 in 1975-79 and -.254 in 1995-1999 using the CPS ADF data. Their parametric selection estimator of the gap is -.337 in 1975-79 and -.339 in 1995-1999.

procedure and assumptions in Blundell et al. (2007), who have used data from the UK. Details about the computation of the bounds are contained in appendix D. Two key features are worth noting. First, bounds pertain the median, rather than the mean, wage gap, as bounds on the mean wage gap require wages to have a bounded support, which is likely not the case. Second, for purpose of comparison to Blundell et al. (2007), I maintain the results stratified by education groups and assume that the changes in the education gap is the same for all ages. I replicate their findings for the US, and find that a positive selection assumption¹¹ is key to determining that the relative wages of women have increased. This result can be seen in the appendix figure (D.1), which plots bounds on the changes of the gender wage gap between 1976 and 2005. A positive number indicates that gap in 2005 is lower relative to 1976, and a negative number indicates it is higher.

What do all these estimates reflect? On the one hand, selection considerations are very important for measuring female wages, as their participation levels are low. This would, in principle, make the observed evolution of the gap a poor proxy for its “true” evolution. On the other hand, any attempt to correct for selection needs to impose some structure into a selection model. Not surprisingly, estimates using different strategies and assumptions yield conflicting answers: Mulligan and Rubinstein (2008) find that the gender wage gap has remained stable, whereas Blundell et al. (2007) find an improvement. Moreover, I am able to replicate both these findings in a single US data set illustrating that it is the method that drives the differing results, rather than any difference in sampling, time period, or choice of instrument.

The local measure of the gender gap also needs to impose some structure into the selection problem, with assumptions given by (AI)-(AIII). Taken at face value, these assumptions are no different than the ones already assumed in the parametric selection model, for instance. The proposed method departs from the literature in that it does not impose either a unique selection mechanism, or knowledge of the sign of average selection. These considerations are particularly important in the presence of unobserved heterogeneity, as should be the case for women and the labor market.

The trade-off, obviously, is that the method developed in this paper estimates the gap for a local sample: women who do not change participation (and remain employed) in the presence of a young child. Nonetheless, this is plausibly the group most comparable to men, as the latter have higher attachment to the labor force and seldom leave their jobs when they have children. In this

¹¹Positive selection is imposed through a stochastic dominance assumption, rather than the median restriction, as defined in Blundell et al. (2007).

“apples to apples” comparison, substantial progress in closing the gender wage gap is indeed found.

4 Robustness

This section examines the sensitivity of the estimates of the local measure of the gap to: 1) covariates other than education; 2) male selection; and 3) an alternative, but related, instrument exploring the age of the youngest child present. For the purpose of comparison to the estimates in figure (2), I consider participation to be full time employment, and use the female variable weights p_{xt}^{FV} when averaging is necessary.

4.1 Including Covariates

The analysis in previous sections has used education as the single covariate in X . However, selection effects may vary along other characteristics, such as age and marital status. In this section, I follow Card (1996) and Lee (2009) and incorporate all available covariates in a “skill” index. The index is used to sort workers into groups of similar characteristics and the local measure of the gender wage gap is computed within the groups.

The procedure is as following. For each period and gender, I estimate a wage equation¹² and use the model to predict wages for the entire sample, whether working or not. I then compute the four quartiles of the predicted wage distribution and sort observations into each quartile. Finally, for each period, I compute the local gender gap between men and women that have the same rank on its own predicted wage distribution. Although the predicted wage distribution varies by gender, this exercise aims to recover gaps under the assumption that the skills are approximately the same, but are rewarded differently in the market place.

Table (7) presents the summary statistics of the the sample by the four quartiles of predicted wages, stratified by period and gender. It confirms that education is an important variable in the classification of skill types. For example, in the 1976-1980 period, the female composition in the lower quartile has education attainment below high school, whereas the upper quartile has women with at least some college education. Nonetheless, it also shows that the other covariates also have explanatory power in the skill classification, and that education is not its unique determinant. A

¹²The explanatory variables are: five education dummies (some high school, high school graduates, some college, college graduates and more than college, relative to less than high school education), three age dummies (ages 30 to 34, 35 to 39 and 40 to 44, relative to ages 25 to 29) and four marital status dummies (widowed, divorced, separated and never married, relative to married).

child younger than 6 years old decreases female full time employment for all skill groups and time periods, in line with the last two lines of panel A¹³.

Estimates of the local gap are displayed in figure (3). The local gender wage gap is lower, in level, for the lowest skill group, and possibly reflects minimum wage policy limiting disparities in the low end of the wage distribution. The trend towards wage equality between 1976 and 2005 is verified for all skill groups. But, most strikingly, the figure also shows continuous improvement of the gender wage gap throughout the 1990s for the two upper quartiles of the skill distribution. For purpose of comparison, figure (4) displays the uncorrected gap by the skill types, and confirms the finding of slowing convergence of the gender gap after 1990, as documented by Blau and Kahn (2006). Taken together, the two figures show that selection effects have masked substantial improvements in the gap in 1990s.

4.2 Male Selection

The approach taken so far has assumed that the observed wage distribution for men proxies the distribution of potential wages. In the US, however, male selection into work may challenge this assumption as one quarter of wage information is missing for the full-time employed sample, as seen in tables (2) and (3). If average selection is positive for men, and the ones participating have the highest wages, the results in previous sections constitute an exaggerated estimate of the wage gap, as the average potential wages of men should be lower.

In principle, the presence of a young child can also be used as an instrument for male participation. In fact, the summary statistics in table (2) show that the full time employment of men increases by 9% when a young child is present. Relative to women, the sensitivity of male participation to the presence of a young child is smaller and goes in the opposite direction. Probits for male participation confirm this evidence and are available upon request.

Since θ_{xt} has a negative sign for women and a positive sign for men, the following variant of equation (8) is estimated:

$$Y_i = \beta_{0xt} + \beta_{1xt}G_i + \beta_{2xt}G_i * Z_i + \beta_{3xt}(1 - G_i) * Z_i + u_i, \quad (10)$$

¹³Results are available upon request.

where the local measure of the gap between “always participant” men and women is given by:

$$\begin{aligned} \Omega(x, t) &= E(\mathcal{Y} | E_1 = 1, E_0 = 1, G = 1, X = x, T = t) \\ &\quad - E(\mathcal{Y} | E_1 = 1, E_0 = 1, G = 0, X = x, T = t) = \beta_{1xt} + \beta_{2xt}. \end{aligned}$$

Results for (10) are summarized in figure (5). Relative to the measure of $\Delta(t)$, which only accounts for female selection, $\Omega(t)$ pictures a reduced wage gap. Male selection considerations becomes more important towards the end of the sample period, as the wedge between $\Delta(t)$ and $\Omega(t)$ gets wider. For the 2001-2005 period, the point estimates of $\Delta(t)$ and $\Omega(t)$ are -0.263 and -0.229, and depart by .034 log wage points.

4.3 Child Age and Female Participation

The presence of a child younger than six has a substantial impact on women’s participation decisions. As seen in table (2), full time employment of women with a young child is lower by 25 percentage points, almost half of the participation of women with no young child present. Since younger children require more maternal input, participation effects could differ by the age of the child. Most importantly, because the local gap recovers the pay penalty among “always employed” women, the subpopulation that does not change participation when a child younger than, say, 1 is present should be even more like men in terms of unobservables, such as job commitment. As the age of the youngest child decreases, and having Z as an indicator for his/her presence, the local measure of the gap should be smaller.

I investigate this hypothesis by utilizing an alternative dataset, the June CPS survey. The June CPS has a fertility supplement, available every other year, with information on the birth month and birth year of the last child. I make restrictions similar to the ones in the ADF sample, which are detailed in appendix B. The new instruments considered, Z_j , are binary indicators of the presence of a child less than j , with $j \in [1, 6]$, and are considered one at a time.

One shortcome of this alternative dataset is its small sample size, as wage information in the June sample is only recorded for individuals in the outgoing rotation groups (ORG). Therefore, in this section, I will aggregate the June data into more sparse time groups, allocating approximately one quarter of total observation into four periods: 1979 to 1982, 1983 to 1987, 1988 to 1992, and 1994-2002. Summary statistics are presented in table (3). The age of the youngest child alters participation in a similar manner, and the difference in participation between women with and

without a child present is relatively constant (around 25%) for all values of j . In fact, full time employment can remain insensitive to the age of the child if women with young children return to part-time jobs.

Figure (6) summarizes the local gender wage gap using the alternative Z_j 's as instruments¹⁴. The figure suggests that as the age of the young child decreases, the local gender wage gap gets smaller. This result is in line with the conjecture that “always employed” women with a very young child share similar unobserved job characteristics with men, and the wage difference among them is tighter. The standard errors of the estimates, however, do not allow the inference that the local gap using Z_1 is statistically different from the one using Z_6 for the later periods in the sample.

5 Discussion

5.1 The Local Gap and External Validity

The instrumental variable approach proposed in this paper recovers a local measure of the gender wage gap, the gap in pay between women who would choose to participate whether or not a young child is present and similar men. One criticism of using instrumental variables to recover a local parameter regards the particular and unobserved group of individuals to which the estimates refer to. In the case of the gender gap, however, the subpopulation of always participating women should be the most relevant comparison to men, as, in general, women’s attachment to work in the presence of a child is lower than men’s. This should make the “always participants” very close to men in important unobserved wage determinants, such as job commitment and motivation.

Nonetheless, in order to inspect the external validity of the estimate to the total population of women, a useful exercise backs out the proportion of “switchers,” “always employed” and “never employed.” Noting that:

$$\begin{aligned}\Pr(E = 1|Z = 1) &= \Pr(E_1 = 1) = \Pr(E_1 = 1, E_0 = 1) + \Pr(E_1 = 1, E_0 = 0) \\ \Pr(E = 1|Z = 0) &= \Pr(E_0 = 1) = \Pr(E_1 = 1, E_0 = 1) + \Pr(E_1 = 0, E_0 = 1),\end{aligned}$$

¹⁴The estimation results for equations (7) and (8) are available upon request.

and using the monotonicity condition $E_1 \leq E_0$, the proportion in each group is given by:

$$\begin{aligned}\Pr(E_1 = 0, E_0 = 1) &= \Pr(E = 1|Z = 0) - \Pr(E = 1|Z = 1) \\ \Pr(E_1 = 1, E_0 = 1) &= \Pr(E = 1|Z = 1) \\ \Pr(E_1 = 0, E_0 = 0) &= 1 - \Pr(E = 1|Z = 0).\end{aligned}$$

From table (2) the proportion of “always employed” averages 27.1% for the entire sample, and becomes quite sizable by 2001-05. For that last period, it reaches 34.7% of the sample. Insofar as the “always employed” have higher wages than the “never employed” and the “switchers,” the unconditional average wage gap could be higher than the local measure provided by this paper.

The rise in the proportion of “always employed” in the female population raises concern about the comparability of this group in time. In principle, it is possible that improvements in the wage gap reflect better unobserved characteristics of the marginal women who become “always employed.” It is unclear, however that all women joining the “always employed” pool have the highest wages. If, for example, “always employed” have the highest wages to start with, as more women join the pool from lower parts of the wage distribution, average wage for the group should fall, rather than increase. This paper aims to correct for employment selection in the presence unobserved heterogeneous rules, but does not address the reason for the falling gender wage gap, whether if it is due to lower discrimination or better unobserved characteristics of the “always employed”.

5.2 The Identifying Assumptions

As in treatment effect models, solving the problem of selection with missing outcomes generally require an instrument that shifts participation and is excluded from the outcome equation. The literature, however, has treated the IV assumption “quite casually”, viewing identification as an intrinsic econometric problem, with little emphasis on its economic content (Moffit, 1999).

The discussion is often neglected because it is very hard to defend all the necessary assumptions for a valid instrument, particularly in a non-experimental setting. The use of the presence of a young child as an IV in the female selection problem is no different. Although very little can be done when the instrument fails the independence condition (AI), which is the case if fertility decisions are not exogenous to employment and wages (Angrist and Evans, 1998; Browning, 1992; Willis, 1998), the failure of assumptions (AII) and (AIII) are not necessarily catastrophic. In this section, I follow

Angrist et al. (1996) and discuss the sensitivity of the local measure of wages to the exclusion and monotonicity restrictions.

5.2.1 Monotonicity

The monotonicity condition rules out the existence of type $(E_1 = 1, E_0 = 0)$ women, and would be violated if there are women who do not work, but start working once a child is present, to complement household income, for instance. By manipulation of the terms in equation (6), the bias of the local wage estimator is given by:

$$\begin{aligned} E(Y|E = 1, Z = 1) - E(\mathcal{Y}|E_1 = 1, E_0 = 1) = \\ \left\{ E(\mathcal{Y}|E_1 = 1, E_0 = 1) - E(\mathcal{Y}|E_1 = 1, E_0 = 0) \right\} \Pr(E_0 = 0|E_1 = 1). \end{aligned}$$

The sign of the bias depends on whether $E(\mathcal{Y}|E_1 = 1, E_0 = 1)$ is smaller or greater than $E(\mathcal{Y}|E_1 = 1, E_0 = 0)$, and in principle could go either way, but should be a decreasing function of the proportion of “defiers” among “defiers” and “always employed”, $\Pr(E_0 = 0|E_1 = 1)$.

5.2.2 Exclusion Restriction

The exclusion restriction would fail if Z , aside from shifting employment deteriorates women’s earnings prospect. The literature on the family gap (Waldfogel, 1998) has found that women with children earn lower wages, which could reflect compositional changes among working moms combined with changes in earnings.

In this case, the model in section 2 could be modified to:

$$\begin{aligned} Y &= Z\mathcal{Y}_1 + (1 - Z)\mathcal{Y}_0 \text{ if } E = 1 \\ E &= ZE_1 + (1 - Z)E_0, \end{aligned}$$

where \mathcal{Y}_1 and \mathcal{Y}_0 are now the potential wages when Z is externally set to 1 and 0 respectively. Maintaining monotonicity and modifying (AIII) to $Z \perp\!\!\!\perp (\mathcal{Y}_0, \mathcal{Y}_1, E_0, E_1)$, it can be shown that $E(Y|E = 1, Z = 1) = E(\mathcal{Y}_1|E_1 = 1, E_0 = 1)$. That is, the instrument still recovers a local measure of wages for “always employed” when Z is set to 1. If earnings decrease in the presence of a child, $E(Y|E = 1, Z = 1)$ provides a lower bound for \mathcal{Y} among all “always employed.”

6 Conclusion

In the early 1970s women started to enter the labor market in terms similar to men: they migrated to “high-powered” professions (Goldin and Katz, 2002) and increased their employment attachment (Goldin, 2006). Acceptance towards the new role of women in society gained force with female rights movements, and has been corroborated by anti-discrimination laws for equal pay. Taken altogether, women have gradually converged to be as men in the workplace, looking for “careers” rather than “jobs,” this being part of the larger phenomenon which Claudia Goldin termed as the “quiet revolution” that transformed women’s lives.

While improvements in female workforce characteristics have been partially responsible for the closing gender wage gap, selection effects, made evident by the increase in female employment rates, could cast doubts in the relative improvement of women’s wages. This paper explores the idea that likely both positive *and* negative selection rules have guided women’s employment decisions. Under unobserved heterogeneity in the selection rule, I show that the usual selection correction procedures might fail to recover the parameters they target, but a local measure of the gender wage gap can still be recovered. The local gap pertains the subpopulation of “always employed” women, and is a relevant comparison group to men, as they exhibit high labor force attachment.

Whether capturing “discrimination” or some other unobserved gender difference in productivity, I conclude that the gender wage gap has indeed decreased in the US for this comparable group of men and women. Empirically, I find that the observed (uncorrected) gap provides a good proxy for the local gender wage gap. I do not conclude, however, that selection effects and corrections are unimportant. On the contrary, they are so extreme in the female employment case, which is made evident by the different gap measures estimated in the literature, that a more conservative approach, as proposed in this paper, is called for. In the presence of heterogeneity in selection, focusing on the local gender wage gap is altogether less distorting than usual selection corrections.

References

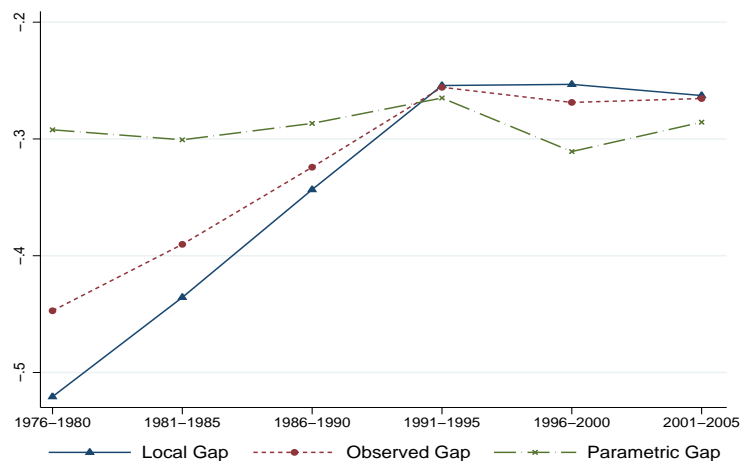
- J. D. Angrist and W. Evans. Children and their parents labor supply: Evidence from exogenous variation in family size. *American Economic Review*, June 1998.
- J. D. Angrist, G. D. Imbens, and D. B. Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–454, June 1996.
- F. Blau and L. Kahn. The us gender pay gap in the 1990s: Slowing convergence. *Industrial and Labor Relations Review*, 60(1):45–66, 2006.
- R. Blundell, A. Gosling, H. Ichimura, and C. Meghir. Changes in the distribution of female and male wages accounting for employment composition using bounds. *Econometrica*, 75(2):323–363, March 2007.
- M. Browning. Children and household economic behavior. *Journal of Economic Literature*, 30(3): 1434–1475, September 1992.
- D. Card. The effect of unions on the wage structure of wages: A longitudinal analysis. *Econometrica*, 64(4):957–979, July 1996.
- C. Goldin. The quiet revolution that transformed women’s employment, education, and family. *AEA Papers and Proceedings, Richard T. Ely Lecture*, 92(2):1–21, May 2006.
- C. Goldin and L. F. Katz. The power of the pill: Oral contraceptives and women’s career and marriage decisions. *Journal of Political Economy*, 110(4):730–770, August 2002.
- R. Gronau. Wage comparison - a selectivity bias. *Journal of Political Economy*, 82:1119–1143, 1974.
- J. Heckman. Shadow price, market wages, and labor supply. *Econometrica*, 42(4):679–694, July 1974.
- J. Heckman. Varieties of selection bias. *American Economic Review Papers and Proceedings*, 80 (2):313–318, May 1990.
- W. Johnson, Y. Kitamura, and D. Neal. Evaluating a simple method for estimating black-white gaps in median wages. *American Economic Review Papers and Proceedings*, 90:339–343, May 2000.

- D. Lee. Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Review of Economic Studies*, 76(3):1071–1102, June 2009.
- C. Manski. Nonparametric bounds on treatment effects. *The American Economic Review Papers and Proceedings*, 80(2):319–323, 1990.
- R. Moffit. *Handbook of Labor Economics*, chapter New Developments in Econometric Methods for Labor Market Analysis, pages 1367–1397. Elsevier, 1999.
- C. Mulligan and Y. Rubinstein. Selection, investment and women’s relative wage over time. *Quarterly Journal of Economics*, 2008.
- D. Neal. The role of pre-market factors in black-white wage differences. *Journal of Political Economy*, 104:869–895, October 1996.
- D. Neal. The measured black-white wage gap among women is too small. *Journal of Political Economy*, 112(1):S1–S28, 2004.
- C. Olivetti and B. Petrongolo. Unequal pay or unequal employment? a cross-country analysis of gender gaps. 2007.
- D. Rubin. Estimating causal effects of treatments in randomized and non-randomized studies. *Journal of Educational Psychology*, 66:688–701, 1974.
- F. Vella. Models with sample selection bias: A survey. *Journal of Human Resources*, 33(1):127–169, Winter 1998.
- E. Vytlacil. Independence, monotonicity and latent index models: An equivalence result. *Econometrica*, 70(1):331–341, January 2002.
- J. Waldfogel. Understanding the “family gap” in pay for women with children. *Journal of Economic Perspectives*, 12(1):137–156, Winter 1998.
- R. Willis. Children and their parents’ labor supply: Evidence from exogenous variation in family size. *American Economic Review*, 88(3):450–477, June 1998.

Figure 1: Simulated Data, Distribution of Wages




Figure 2: Gender Wage Gap, Local vs Observed (Uncorrected) vs Parametric




Data: March CPS, ADF, 1976-2005. Sample delimitation is as in notes to table (2), with further details in appendix A. The figure plots the gender wage gap (female minus male wages) in log wage points. The local gap averages estimates of (8), $\beta_{1xt} + \beta_{2xt}$, using female education proportions in each time period. The observed (uncorrected) and parametric gaps average estimates of (C.1) and (C.3), κ_{1xt} and θ_{1xt} , using the same weights. In all estimates, employment is defined by full time full year work. The instrument used to estimate the local and parametric gaps is an indicator for the presence of a child younger than six years old.

Figure 3: Local Gender Wage Gap, by Percentiles of Predicted Wage



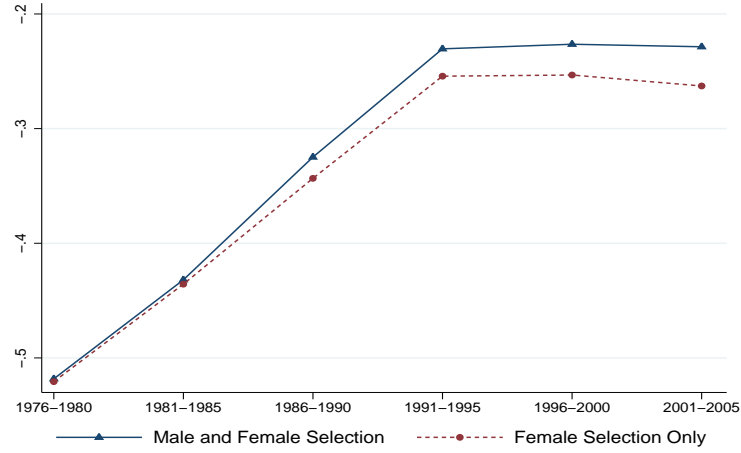
Data: March CPS, ADF, 1976-2005. Sample delimitation is as in notes to table (2), with further details in appendix A. The figure plots the local gender wage gap (female minus male wages) in log wage points, by quartiles of predicted wages. The quartiles are generated as following: for each period and gender, I estimate a wage equation - the dependent variables are: five education dummies (some high school, high school graduates, some college, college graduates and more than college, relative to less than high school education), three age dummies (ages 30 to 34, 35 to 39 and 40 to 44, relative to ages 25 to 29) and four marital status dummies (widowed, divorced, separated and never married, relative to married) - and use the model to predict wages for the entire sample, whether working or not; I then compute the four quartiles of the predicted wage distribution and sort observations into each quartile. The local gender gap is estimated between men and women that have the same rank in its own predicted wage distribution. I take X to be the four quartiles and estimate (8). In all estimates, employment is defined by full time full year work. The instrument used to estimate the local gaps is an indicator for the presence of a child younger than six years old. The lines stratify results by the four quartiles of the predicted wage distribution.

Figure 4: Observed (Uncorrected) Gender Wage Gap, by Percentiles of Predicted Wage




Data: March CPS, ADF, 1976-2005. Sample delimitation is as in notes to table (1), with further details in appendix A. The figure plots the observed (uncorrected) gender wage gap in log wage points, by quartiles of predicted wages. The quartiles are generated as in notes to figure (3). The observed (uncorrected) gap is estimated by the coefficient on the female dummy in a log wage regression, as outlined in appendix C. In all estimates, employment is defined by full time full year work. The lines stratify results by the four quartiles of the predicted wage distribution.

Figure 5: Local Gender Wage Gap, Female and Male Selection



Data: March CPS, ADF, 1976-2005. Sample delimitation is as in notes to table (2), with further details in appendix A. The figure plots the local gender wage gap (female minus male wages) in log wage points. The local gap averages estimates of (8), $\beta_{1xt} + \beta_{2xt}$, using female education proportions in each time period. In both estimates, employment is defined by full time full year work. The instrument used to estimate the local gaps is an indicator for the presence of a child younger than six years old. Lines differ on whether the instrument is used to corrected for male and female selection into employment or to correct for female selection only.

Figure 6: Local Gender Wage Gap, by Age of Young Child



Data: June CPS, 1979-2002. Sample delimitation is as in notes to table (3), with further details in appendix B. The figure plots the local gender wage gap (female minus male wages) in log wage points. The local gap averages estimates of (8), $\beta_{1xt} + \beta_{2xt}$, using female education proportions in each time period. In all estimates, employment is defined by full time full year work. Lines differ with respect to instrument used in estimation: child younger than j , with $j \in \{1, \dots, 6\}$.

Table 1: Simulated Data

	Mean	90% CI
Panel A: Observed Data		
$E(Y E = 1)$	1.596	(1.560,1.630)
$E(Y E = 1, Z = 1)$	2.468	(2.404,2.528)
$E(Y E = 1, Z = 0)$	1.280	(1.241,1.320)
$\Pr(E = 1)$	0.500	(0.492,0.507)
$\Pr(E = 1 Z = 1)$	0.265	(0.256,0.275)
$\Pr(E = 1 Z = 0)$	0.734	(0.724,0.745)
Panel B: Unobserved Data - by Selection Rule		
$E(\mathcal{Y})$	1.595	(1.569,1.623)
$\text{Med}(\mathcal{Y})$	1.465	(1.432,1.496)
$E(\mathcal{Y} \mathbb{N})$	1.595	(1.559,1.633)
$E(\mathcal{Y} \mathbb{P})$	1.595	(1.560,1.632)
$\Pr(E = 1 \mathbb{N})$	0.234	(0.224,0.244)
$\Pr(E = 1 \mathbb{P})$	0.765	(0.756,0.775)
$E(\mathcal{Y} E = 1, \mathbb{N})$	0.605	(0.552,0.657)
$E(\mathcal{Y} E = 0, \mathbb{N})$	1.898	(1.855,1.940)
$E(\mathcal{Y} E = 1, \mathbb{P})$	1.899	(1.858,1.941)
$E(\mathcal{Y} E = 0, \mathbb{P})$	0.604	(0.551,0.657)
Panel C: Estimation		
Parametric Estimator of $E(\mathcal{Y})$	0.605	(0.518,0.688)
LB on $E(\mathcal{Y})$	0.119	(-0.060,0.261)
UB on $E(\mathcal{Y})$	3.229	(2.966,3.571)
UB on $E(\mathcal{Y})$, + Selection	1.280	(1.241,1.320)
LB on $E(\mathcal{Y})$, - Selection	2.468	(2.404,2.528)
LB on $\text{Med}(\mathcal{Y})$	0.501	(0.441,0.562)
UB on $\text{Med}(\mathcal{Y})$	1.825	(1.756,1.897)
UB on $\text{Med}(\mathcal{Y})$, + Selection	1.129	(1.086,1.179)
LB on $\text{Med}(\mathcal{Y})$, - Selection	2.367	(2.289,2.444)

continued on next page

continued from previous page (Table 1)

	Mean	90% CI
Panel D: By Latent Response		
$E(\mathcal{Y} E_1 = 0, E_0 = 0)$	2.467	(2.422,2.510)
$E(\mathcal{Y} E_1 = 0, E_0 = 1)$	0.605	(0.579,0.629)
$E(\mathcal{Y} E_1 = 1, E_0 = 1)$	2.467	(2.424,2.510)
$\Pr(E_1 = 0, E_0 = 0)$	0.266	(0.258,0.274)
$\Pr(E_1 = 1, E_0 = 0)$	0.468	(0.460,0.477)
$\Pr(E_1 = 1, E_0 = 1)$	0.266	(0.259,0.273)
Panel E: Mean Squared Error of Estimators		
Parametric Estimator of $E(\mathcal{Y})^\dagger$	0.983	(0.813,1.156)
UB on $E(\mathcal{Y})$, + Selection [†]	0.100	(0.078,0.120)
LB on $E(\mathcal{Y})$, - Selection [†]	0.764	(0.666,0.866)
$E(Y E = 1, Z = 1)^\ddagger$	0.001	(0.000,0.003)

Data generated from 1,000 replications the model in (4) and (5) with 10,000 observations each. The parametric selection estimator of $E(\mathcal{Y})$ corresponds to the constant of a regression that includes the mills ratio, generated from a first stage probit of E on Z , as a control. UB and LB abbreviate upper and lower bound. The construction of upper and lower bounds in Panel C is detailed in appendix D.1. †: mean squared error with respect to $E(\mathcal{Y})$. ‡: mean squared error with respect to $E(\mathcal{Y}|E_1 = 1, E_0 = 1)$.

Table 2: Summary Statistics March CPS (means)

	Female			Male		
	1976-2005	1976-80	2001-05	1976-2005	1976-80	2001-05
% married	0.709	0.791	0.671	0.666	0.776	0.616
% widowed	0.00905	0.0102	0.00873	0.00221	0.00171	0.00286
% divorced	0.114	0.0910	0.118	0.0868	0.0593	0.0916
% separated	0.0278	0.0260	0.0243	0.0192	0.0178	0.0171
% never married	0.140	0.0819	0.178	0.226	0.145	0.273
% less than hs	0.0901	0.162	0.0549	0.0960	0.147	0.0704
% hs grad	0.381	0.466	0.284	0.350	0.361	0.317
% some college	0.257	0.181	0.307	0.245	0.204	0.276
% college or more	0.273	0.191	0.354	0.309	0.288	0.336
number children ≤ 6	0.421	0.457	0.408	0.409	0.477	0.383
% no child ≤ 6	0.689	0.659	0.702	0.704	0.650	0.725
% 1 child ≤ 6	0.214	0.238	0.202	0.197	0.238	0.180
% 2 children ≤ 6	0.0853	0.0900	0.0836	0.0867	0.0987	0.0825
% 3+ children ≤ 6	0.0119	0.0125	0.0123	0.0124	0.0135	0.0122
% employed	0.763	0.648	0.795	0.946	0.965	0.931
ln(hourly wage)	1.861	1.726	2.022	2.206	2.245	2.278
% employed w/o child ≤ 6	0.822	0.737	0.841	0.934	0.955	0.915
% employed w/ child ≤ 6	0.632	0.477	0.686	0.975	0.984	0.972
% full time	0.444	0.304	0.511	0.764	0.758	0.782
ln(hourly wage)	1.970	1.851	2.081	2.267	2.296	2.323
% ft w/o child ≤ 6	0.522	0.391	0.581	0.738	0.732	0.753
% ft w/ child ≤ 6	0.271	0.137	0.347	0.825	0.806	0.858
observations	469050	67813	92418	419208	59052	81639

Data: March CPS, ADF, 1976-2005. Sample restricted to white non-hispanic individuals between ages 25 and 44. More details on sample delimitation are contained in appendix A. % indicates the mean of a binary variable. All means are computed using the March Supplement weight. An individual is categorized as employed if it reports any wage in previous year. Full time workers are the employed individuals who have worked 50 or more weeks, and 35 or more hours per week, in previous year. Wages are log hourly wages deflated by the CPI-U, in the 1982-1984 base.

Table 3: Summary Statistics June CPS (means)

	Female			Male		
	1979-2002	1979-1982	1994-2002	1979-2002	1979-1982	1994-2002
% married	0.710	0.758	0.671	0.675	0.730	0.627
% absent or separated	0.0360	0.0345	0.0377	0.0249	0.0248	0.0287
% widowed or divorced	0.124	0.115	0.128	0.0859	0.0751	0.0942
% never married	0.130	0.0933	0.163	0.214	0.170	0.250
% less than hs	0.105	0.143	0.0763	0.108	0.140	0.0842
% hs grad	0.413	0.477	0.323	0.365	0.373	0.337
% some college	0.244	0.190	0.310	0.229	0.196	0.273
% college or more	0.238	0.190	0.291	0.298	0.290	0.306
% any child ≤ 6	0.325	0.339	0.324	.	.	.
% any child ≤ 5	0.279	0.291	0.274	.	.	.
% any child ≤ 4	0.233	0.244	0.228	.	.	.
% any child ≤ 3	0.180	0.190	0.177	.	.	.
% any child ≤ 2	0.124	0.130	0.122	.	.	.
% any child ≤ 1	0.0629	0.0660	0.0623	.	.	.
% full time	0.442	0.374	0.474	0.796	0.799	0.780
ln(hourly wage)	1.928	1.871	1.980	2.195	2.224	2.184
% ft w/o child ≤ 6	0.522	0.458	0.551	.	.	.
% ft w/ child ≤ 6	0.275	0.208	0.313	.	.	.
% ft w/o child ≤ 5	0.509	0.445	0.537	.	.	.
% ft w/ child ≤ 5	0.268	0.199	0.307	.	.	.
% ft w/o child ≤ 4	0.497	0.432	0.526	.	.	.
% ft w/ child ≤ 4	0.259	0.191	0.297	.	.	.
% ft w/o child ≤ 3	0.484	0.419	0.513	.	.	.
% ft w/ child ≤ 3	0.250	0.179	0.291	.	.	.
% ft w/o child ≤ 2	0.471	0.405	0.502	.	.	.
% ft w/ child ≤ 2	0.234	0.161	0.272	.	.	.
% ft w/o child ≤ 1	0.456	0.389	0.488	.	.	.
% ft w/ child ≤ 1	0.226	0.159	0.263	.	.	.
observations	59270	14689	14697	50826	12238	13006

Data: June CPS, 1979-2002. Sample restricted to white non-hispanic individuals between ages 25 and 44. More details on sample delimitation are contained in appendix B. % indicates the mean of a binary variable. All means are computed using weights from the earner study. An individual is categorized as employed if holding a job by the time of the survey. Full time workers are the employed individuals who work more than 35 hours per week. Wages are log hourly wages deflated by the CPI-U, in the 1982-1984 base. Age of youngest child is computed using the birth month and birth year of the child available in the fertility supplement.

Table 4: Changes in Female Employment Induced by the Presence of Any Child ≤ 6

Entries in each triplet are: ($\Psi_{xt}, N_{xt}, \%E_{xt}$)				
$x \backslash t$	E : Any Employment		E : Full Time, Full Year	
	(1)	(2)	(3)	(4)
	1976-80	2001-05	1976-80	2001-05
less than high school	-0.197*** 10759 0.517	-0.133*** 4831 0.583	-0.137*** 10759 0.193	-0.162*** 4831 0.315
high school graduate	-0.258*** 31448 0.630	-0.150*** 27025 0.761	-0.236*** 31448 0.292	-0.215*** 27025 0.482
some college	-0.285*** 12551 0.680	-0.161*** 29395 0.812	-0.294*** 12551 0.322	-0.243*** 29395 0.514
college or more	-0.316*** 13055 0.774	-0.186*** 31167 0.841	-0.372*** 13055 0.411	-0.275*** 31167 0.563
pooled (all education groups) [†]	-0.273*** 67813 0.648	-0.166*** 92418 0.795	-0.258*** 67813 0.304	-0.244*** 92418 0.511

Data: March CPS, ADF, 1976-2005. Sample delimitation is as in notes to table (1), with further details in appendix A. Each triplet refers to an education (x) and time (t) specific sample. The entries are, respectively, the change in probability $\Psi_{xt} = \Pr_{xt}(E = 1|Z = 1) - \Pr_{xt}(E = 1|Z = 0)$, where Z indicates the presence of child younger than 6, the number of observations, and the average employment rate. [†] controls for education specific employment rates. *** denotes statistical significance at the 1% level.

Table 5: Local Gap by Education

	<i>E</i> : Any Employment		<i>E</i> : Full Time, Full Year	
	(1)	(2)	(3)	(4)
	1976-80	2001-05	1976-80	2001-05
Panel A: Less than High School				
female, (a)	-0.572*** (0.013)	-0.326*** (0.021)	-0.502*** (0.016)	-0.306*** (0.021)
female*Z, (b)	-0.048* (0.026)	-0.058 (0.039)	-0.036 (0.038)	-0.043 (0.048)
constant	2.015*** (0.008)	1.819*** (0.012)	2.063*** (0.008)	1.858*** (0.012)
R ²	0.178	0.058	0.157	0.065
N	13298	7249	7231	4665
Local Gap, (a)+(b)	-0.620 (0.025)	-0.384 (0.037)	-0.538 (0.037)	-0.349 (0.046)
Observed (Uncorrected) Gap	-0.582	-0.337	-0.507	-0.312
Panel B: High School Graduate				
female, (a)	-0.562*** (0.007)	-0.300*** (0.008)	-0.459*** (0.007)	-0.262*** (0.008)
female*Z, (b)	-0.077*** (0.011)	-0.037*** (0.013)	-0.099*** (0.014)	-0.039** (0.016)
constant	2.225*** (0.004)	2.069*** (0.005)	2.250*** (0.004)	2.106*** (0.004)
R ²	0.210	0.062	0.199	0.058
N	40811	45127	25420	33161
Local Gap, (a)+(b)	-0.639 (0.011)	-0.337 (0.012)	-0.557 (0.014)	-0.301 (0.015)
Observed (Uncorrected) Gap	-0.581	-0.308	-0.475	-0.269

continued on next page

continued from previous page (Table 5)

	<i>E</i> : Any Employment		<i>E</i> : Full Time, Full Year	
	(1)	(2)	(3)	(4)
	1976-80	2001-05	1976-80	2001-05
Panel C: Some College				
female, (a)	-0.424*** (0.010)	-0.278*** (0.007)	-0.384*** (0.010)	-0.255*** (0.008)
female*Z, (b)	-0.105*** (0.020)	-0.013 (0.011)	-0.144*** (0.028)	-0.027** (0.012)
constant	2.239*** (0.005)	2.224*** (0.005)	2.295*** (0.005)	2.262*** (0.005)
R ²	0.129	0.053	0.132	0.058
Obs	20405	45847	13356	33278
Local Gap, (a)+(b)	-0.529 (0.019)	-0.290 (0.011)	-0.528 (0.027)	-0.282 (0.012)
Observed (Uncorrected) Gap	-0.453	-0.281	-0.407	-0.261
Panel D: College or More				
female, (a)	-0.365*** (0.008)	-0.281*** (0.008)	-0.361*** (0.009)	-0.276*** (0.008)
female*Z, (b)	-0.051*** (0.016)	0.109*** (0.010)	-0.049** (0.021)	0.074*** (0.012)
constant	2.382*** (0.005)	2.586*** (0.005)	2.437*** (0.005)	2.618*** (0.005)
R ²	0.095	0.036	0.102	0.042
Obs	26844	52170	19287	39848
Local Gap, (a)+(b)	-0.416 (0.015)	-0.173 (0.010)	-0.409 (0.020)	-0.202 (0.011)
Observed (Uncorrected) Gap	-0.378	-0.250	-0.367	-0.259

Data: March CPS, ADF, 1976-2005. Sample delimitation is as in notes to table (1), with further details in appendix A. Regression results from equation (8), stratified by education, time and participation (any employment and full time, full year). Standard errors in parenthesis. The measured gap is obtained by a wage regression that includes a gender dummy as explanatory variable. *** denotes statistical significance at the 1% level.

Table 6: Average Gap, Different Weights

	Local Gap		Observed Gap			
	(1)	(2)	(3)	(4)	(5)	(6)
	1976-80	2001-05	Difference (2)-(1)	1976-80	2001-05	Difference (5)-(4)
Panel A: Any Employment						
Female Variable Weights	-0.573 (0.00787)	-0.267 (0.00645)	0.306 (0.0102)	-0.519 (0.00412)	-0.281 (0.00398)	0.238 (0.00573)
Female Fixed Weights 76-80	-0.573 (0.00787)	-0.305 (0.00879)	0.268 (0.0118)	-0.519 (0.00412)	-0.297 (0.00490)	0.223 (0.00640)
Female Fixed Weights 01-05	-0.525 (0.00865)	-0.267 (0.00645)	0.258 (0.0108)	-0.470 (0.00442)	-0.281 (0.00398)	0.189 (0.00594)
Male Variable Weights	-0.549 (0.00793)	-0.272 (0.00662)	0.277 (0.0103)	-0.497 (0.00410)	-0.283 (0.00401)	0.213 (0.00573)
Panel B: Full Time, Full Year						
Female Variable Weights	-0.521 (0.0108)	-0.263 (0.00744)	0.258 (0.0131)	-0.447 (0.00460)	-0.265 (0.00413)	0.182 (0.00618)
Female Fixed Weights 76-80	-0.521 (0.0108)	-0.286 (0.0108)	0.234 (0.0152)	-0.447 (0.00460)	-0.273 (0.00510)	0.175 (0.00687)
Female Fixed Weights 01-05	-0.495 (0.0119)	-0.263 (0.00744)	0.232 (0.0140)	-0.418 (0.00472)	-0.265 (0.00413)	0.152 (0.00627)
Male Variable Weights	-0.506 (0.0109)	-0.266 (0.00775)	0.240 (0.0134)	-0.435 (0.00452)	-0.266 (0.00416)	0.168 (0.00614)

Data: March CPS, ADF, 1976-2005. Sample delimitation is as in notes to table (1), with further details in appendix A. The table summarizes the measures of the gender wage gap from table (5) using alternative education proportions as weights, as specified by equation (9). The female variables weights uses female education proportions in each time period. The female fixed weights holds fixed the female education proportion at either its 1976-80 or 2001-2005 level. The male variable weights uses male education proportions in each time period. Standard errors in parenthesis.

Table 7: Summary Statistics March CPS (means), by Predicted Wage Quartile

	1976-1980				2001-2005			
	pct ≤ 25	pct 25-50	pct 50-75	pct ≥ 75	pct ≤ 25	pct 25-50	pct 50-75	pct ≥ 75
Panel A: Female								
age	31.63	36.04	31.89	33.11	31.90	37.07	35.12	37.74
% married	0.854	0.927	0.641	0.668	0.478	0.778	0.729	0.740
% never married	0.0294	0	0.131	0.204	0.252	0.0861	0.185	0.164
% less than hs	0.586	0	0	0	0.194	0	0	0
% hs grad	0.414	1	0.303	0	0.573	0.558	0	0
% some college	0	0	0.697	0.213	0.233	0.442	0.526	0
% college or more	0	0	0	0.787	0	0	0.474	1
% no child ≤ 6	0.549	0.732	0.630	0.719	0.666	0.746	0.722	0.677
% full time	0.211	0.253	0.364	0.426	0.455	0.501	0.551	0.543
ln(hourly wage)	1.644	1.752	1.853	2.036	1.788	1.922	2.153	2.443
% ft w/o child ≤ 6	0.289	0.297	0.491	0.527	0.520	0.558	0.629	0.618
% ft w/ child ≤ 6	0.116	0.133	0.149	0.169	0.325	0.332	0.349	0.385
Panel B: Male								
age	29.08	32.29	34.96	36.80	32.21	35.63	36.81	37.33
% married	0.544	0.747	0.878	0.957	0.330	0.676	0.663	0.892
% never married	0.355	0.153	0.0496	0.00753	0.512	0.203	0.231	0.0655
% less than hs	0.332	0.242	0.000948	0	0.226	0	0	0
% hs grad	0.481	0.395	0.510	0.00109	0.601	0.476	0	0
% some college	0.109	0.272	0.195	0.248	0.173	0.408	0.597	0
% college or more	0.0776	0.0915	0.294	0.751	0	0.116	0.403	1
% no child leq 6	0.668	0.636	0.685	0.602	0.803	0.747	0.773	0.556
% full time	0.617	0.729	0.815	0.887	0.665	0.795	0.832	0.883
ln(hourly wage)	2.056	2.215	2.327	2.531	1.984	2.217	2.398	2.727
% ft w/o child ≤ 6	0.578	0.699	0.800	0.879	0.641	0.774	0.820	0.865
% ft w/ child ≤ 6	0.695	0.782	0.849	0.899	0.767	0.856	0.875	0.906

Data: March CPS, ADF, 1976-2005. Sample delimitation is as in notes to table (1), with further details in appendix A. For each period and gender, I estimate a wage equation - the dependent variables are: five education dummies (some high school, high school graduates, some college, college graduates and more than college, relative to less than high school education), three age dummies (ages 30 to 34, 35 to 39 and 40 to 44, relative to ages 25 to 29) and four marital status dummies (widowed, divorced, separated and never married, relative to married) - and use the model to predict wages for the entire sample, whether working or not. I then compute the four quartiles of the predicted wage distribution and sort observations into each quartile. Summary statistics are presented within each quartile. % indicates the mean of a binary variable. All means are computed using the March Supplement weight.

Appendix

A March CPS Data

The Annual Demographic File (ADF) is an annual supplement conducted every March in the CPS survey. It surveys all individuals in the monthly sample, and has detailed demographic information, as well as income and labor force participation in the previous calendar year.

In this paper, I use data from 1976 to 2005. The analysis starts in 1976 because hours worked per week is a categorical variable prior to 1976. I restrict the sample to non-hispanic whites between ages 25-44, and exclude individuals living in group quarters and the self-employed. I further eliminate observations with inconsistent reports of weeks worked per year, hours per week, and annual wages, as well as any allocation for these variables. Hourly wages are generated by dividing annual wages by annual hours (week per year times hours per week), and are deflated by the CPI-U (Consumer Price Index, All Urban Consumers) in the 1982-1984 base. The wage measure used in the paper is the log of hourly wages.

An individual is classified as employed if hourly wage information is available. The full time full year employment sample is further restricted to individuals working more than 35 hours per week and more than 50 weeks during the year.

Presence of a child younger than 6 is generated from “number of children in family under 6” and is available for both female and male individuals.

The other demographic information in the sample are marital status and education. I classify individuals into five marital status categories - married, widowed, divorced, separated and never married - and four education groups - less than high school (includes some high school), high school graduate, some college, and college graduate or more.

All estimates using the March CPS data employ the March Supplement weight.

B June CPS Data

The Marriage and Fertility Supplement is conducted on the month of June of the CPS survey in selected years. The questions asked in this supplement vary by survey, but the main fertility variable used in this paper, the year and month of birth of the youngest child, is asked in every fertility supplement, and is consistently available for all females of childbearing age (18-44). I use year and month of birth of youngest child to generate indicator variables of presence of young child less than x , with $x \in \{1, 6\}$.

I follow the same sample delimitation as in the March sample, keeping the non-hispanic whites between ages 25-44 and excluding individuals living in group quarters and the self-employed. Wages and employment come from the earner study, which is conducted every month in the CPS sample for the individuals in the outgoing rotation groups (ORG) 4 and 8. Therefore, relative to the ADF file, the sample in the ORG correspond to one fourth of monthly sample. The earner study has information on weekly wages and hours worked on the jobs individuals hold by the time of the survey. I eliminate individuals with inconsistent reports of weekly earnings and hours per week, as well as any allocation for these two variables. I generate hourly wages by the ratio of weekly earnings and hours per week, and deflate it by the CPI-U. The wage measure is log of hourly wages.

Employment is defined by individuals with hourly wage informations. The full time sample corresponds to those working more than 35 hours per week. Other demographics in the sample are marital status and education. The marital status classification is slightly different from the one used in the ADF sample, due to questionnaire differences, and has four categories: married, spouse absent or separated, widowed or divorced and never married. The education classification is the same as in the ADF sample: less than high school (includes some high school), high school graduate, some college, and college graduate or more.

Since the earner study starts in 1979, this is the first year in my June CPS sample. Survey years included in the sample follow the availability of the Marriage and Fertility Supplement and are 1979-88, 1990, 1992, 1994, 1995, 1998, 2000 and 2002.

All estimates using the June CPS data employ weights from the earner study.

C Observed and Parametric Selection Gender Gap

This section describes the other estimates of the gender gap pictured in figure (2). The observed gap was obtained from the following regression:

$$Y_i = \kappa_{0xt} + \kappa_{1xt} * G_i + u_i, \tag{C.1}$$

where x again subscripts four education categories: less than high school (includes some high school), high school graduate, some college, and college graduate or more. κ_{1xt} correspond to education specific gaps, and the observed gap in period t is recovered by weighting κ_{1xt} by female education proportions in each time period.

The gap obtained from the parametric selection model is estimated in two stages. The first

stage estimates a participation probit for the female population:

$$\Pr(E = 1|T = t, X, G = 1) = \Phi(\alpha'_t X_i + \gamma_t Z_i), \quad (\text{C.2})$$

where X_i is a vector with education dummies and Z_i indicates the presence of a child less than 6. The above results is used to construct the mills ratio, $\lambda_i = \lambda(\hat{\alpha}_t X_i + \hat{\gamma}_t Z_i)$, and the second stage regression includes $\lambda_i * G_i$ as a control for the female selection bias:

$$Y_i = \theta_{0xt} + \theta_{1xt} * G_i + \theta_{2xt} \lambda_i * G_i + u_i. \quad (\text{C.3})$$

θ_{1xt} is weighted by female education proportions in each time period.

D Non-parametric Bounds on the Gender Gap

D.1 Bounds on Mean and Median Wages

When an instrument is available (and maintaining the assumption \mathcal{Y} has a bounded support), we can use condition $E(\mathcal{Y}) = E(\mathcal{Y}|Z)$ to restate (2) as:

$$\begin{aligned} \max_Z \left\{ E(Y|E = 1, Z) \Pr(E = 1|Z) + \underline{\mathcal{Y}} \Pr(E = 0|Z) \right\} \\ \leq E(\mathcal{Y}) \leq \min_Z \left\{ E(Y|E = 1, Z) \Pr(E = 1|Z) + \bar{\mathcal{Y}} \Pr(E = 0|Z) \right\}, \end{aligned} \quad (\text{D.1})$$

where I have abstracted from X and G . That is, Z shifts participation in a manner that is unrelated to the measure of wages, making it valid to consider the best bounds under each value of Z .

Positive selection - $E(Y|E = 1) \geq E(\mathcal{Y}|E = 0)$ - implies $E(\mathcal{Y}) \geq E(Y|E = 1)$, which can be seen by simple manipulation of (1). Substituting this tighter upper bound to (D.1) delivers:

$$\begin{aligned} \max_Z \left\{ E(Y|E = 1, Z) \Pr(E = 1|Z) + \underline{\mathcal{Y}} \Pr(E = 0|Z) \right\} \\ \leq E(\mathcal{Y}) \leq \min_Z \left\{ E(Y|E = 1, Z) \right\}. \end{aligned} \quad (\text{D.2})$$

Similarly, under negative selection:

$$\begin{aligned} \max_Z \left\{ E(Y|E = 1, Z) \right\} \leq E(\mathcal{Y}) \leq \\ \min_Z \left\{ E(Y|E = 1, Z) \Pr(E = 1|Z) + \bar{\mathcal{Y}} \Pr(E = 0|Z) \right\}. \end{aligned} \quad (\text{D.3})$$

I estimate each quantity in (D.1), (D.2) and (D.3) by its sample counterpart, taking $\underline{\mathcal{Y}}$ and $\bar{\mathcal{Y}}$ to the minimum and maximum of observed wages.

Moving away from the assumption of bounded support, I estimate bounds on median wage following a procedure similar to the one outlined in Blundell et al. (2007). The starting point is the CDF of wages $F(\mathcal{Y})$ which is itself bounded by 0 and 1:

$$\max_Z \left\{ F(Y|E = 1, Z) \Pr(E = 1|Z) \right\} \leq F(\mathcal{Y}) \leq \min_Z \left\{ F(Y|E = 1, Z) \Pr(E = 1|Z) + \Pr(E = 0|Z) \right\}. \quad (\text{D.4})$$

Let Z^l (Z^u) denote the argmax (argmin) of above lower (upper) bound. Bounds on median wages are given by:

$$Y^l \leq \text{Med}(\mathcal{Y}) \leq Y^u, \quad (\text{D.5})$$

where Y^l and Y^u solve the q^l and q^u quantiles of $F(\cdot|E = 1, Z^l)$ and $F(\cdot|E = 1, Z^u)$ with:

$$q^l = \frac{0.5 - \Pr(E = 0|Z^l)}{\Pr(E = 1|Z^l)} \quad \text{and} \quad q^u = \frac{0.5}{\Pr(E = 1|Z^u)}. \quad (\text{D.6})$$

Positive and negative selection assumptions change upper and lower bounds of (D.4) in the same spirit as they change bounds for $E(\mathcal{Y})$. Under positive selection, wages for those observed working stochastically dominate those on non workers and $F(Y|E = 1) \leq F(\mathcal{Y}|E = 0)$. Thus, lower bound in (D.4) is tightened to $\max_Z F(Y|E = 1, Z)$. Similarly, under negative selection upper bound in (D.4) becomes $\min_Z F(Y|E = 1, Z)$. I estimate each quantity in (D.4) and (D.6) by its sample counterpart and use the observed distribution of $Y|E = 1, Z$ when searching for quantiles.

Two features of the bounding procedure on median wages are worth noting. First, when the exclusion restriction is invalid, upper and lower limits in (D.6) may cross. In the case the instrument is valid, but sign of selection is incorrectly imposed, bounds may also cross, as seen in the example with simulated data, where the exclusion restriction is combined with the negative selection assumption. Second, the quantiles q^l and q^u can only be computed when participation is above 50%. For example, estimation of (D.6) requires that $\Pr(E = 0|Z^l) \leq 0.5$ and that $\Pr(E = 1|Z^u) \geq 0.5$.

D.2 Blundell et al. (2007) Bounds using US Data

I follow the approach in Blundell et al. (2007) and construct bounds on the median gender wage gap using US data, with same sample delimitation as in section 3. Participation is defined by any employment¹⁵, and conditioning covariates are time T , education X and age A . I let T denote the start and end of my sample, $t_1=1976-1980$ and $t_2=2001-2005$, X , the four education categories and

¹⁵Full time work rate is below 50% for women in start of sample period, and bounds on $\text{Med}(\mathcal{Y})$ cannot be recovered.

A , five year age groups (25-29, 30-34, 35-39, 40-44). The parameter of interest is the improvement of the gender gap between 1976-1980 and 2001-2005, and is given by:

$$\Delta(t_2, X, A) - \Delta(t_1, X, A) = \left\{ \text{Med}(\mathcal{Y}|t_2, X, A, G = 1) - \text{Med}(\mathcal{Y}|t_2, X, A, G = 0) \right\} - \left\{ \text{Med}(\mathcal{Y}|t_1, X, A, G = 1) - \text{Med}(\mathcal{Y}|t_1, X, A, G = 0) \right\}. \quad (\text{D.7})$$

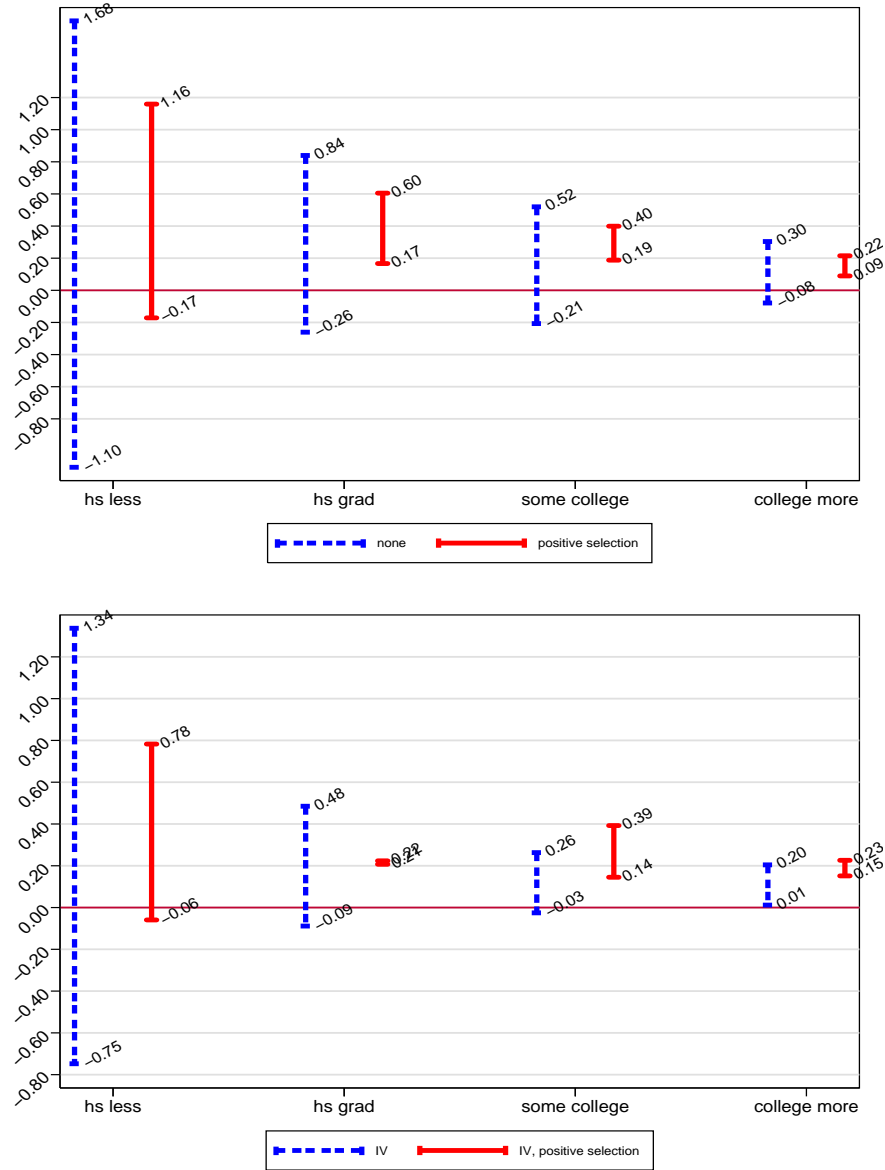
Bounds on each quantity of (D.7) are estimated following the procedure outlined in section D.1 of this appendix, where I consider four scenarios. The baseline result are bounds on $\Delta(t_2, X, A) - \Delta(t_1, X, A)$ that impose neither the exclusion restriction nor the sign of average selection. The second scenario imposes positive selection for both men and women, and for both time periods. The third uses the exclusion restriction, the presence of a child less than six years old, in the computation of female median wages, with no restrictions on the sign of average selection for both men and women. Finally, the fourth case combines the exclusion restriction for women with positive selection for both genders.

Results on $\Delta(t_2, X, A) - \Delta(t_1, X, A)$ are summarized by education categories using the “additivity” assumption of Blundell et al. (2007), in which changes in education differentials across time are assumed to be the same for all age groups, making it valid to search across A for the best values of $\Delta(t_2, X, A) - \Delta(t_1, X, A)$. That is:

$$\max_A LB \left\{ \Delta(t_2, X, A) - \Delta(t_1, X, A) \right\} \leq \Delta(t_2, X) - \Delta(t_1, X) \leq \min_A UB \left\{ \Delta(t_2, X, A) - \Delta(t_1, X, A) \right\} \quad (\text{D.8})$$

where LB and UB stand for lower and upper bounds. Results are displayed in figure (D.1). The baseline case is uninformative about the evolution of the gap for all education categories. For those with some college, for example, the gender gap could have improved as much as .52 log points or deteriorated by 0.21 log points. The positive selection assumption is shown to be quite crucial to sign the evolution of the gap: all individuals with high school degrees or more experience a substantial improvement in gap under this assumption. The exclusion restriction alone is barely able to sign changes in gap (only for individuals with a college degree or more), even when abstracting from sample variability. Finally, the combination of the exclusion restriction and positive selection is not valid, as bounds for those with some college and more than college cross.

Figure D.1: Bounds on Median Wage Gap, US



Data: March CPS, ADF, 1976-2005. Sample delimitation is as in notes to table (1), with further details in appendix A. The figure plots changes in the gender median wage differential from 1976 to 2005, as specified in (D.7) and (D.8), stratified by education groups. “none” denotes neither the use of an exclusion restriction nor imposing the sign of average selection when estimating the bounds. “positive selection” denotes use of a positive selection assumption for all periods and all genders. “IV” denotes the use of an exclusion restriction, a binary indicator for the presence of a child younger than six, in the bounds for the median wages of women. “IV, positive selection” combines the exclusion restriction for women with a positive selection assumption for both men and women. The bounds under “IV, positive selection” for individuals with some college or more cross.