8-2014

QUALITATIVE AND QUANTITATIVE ANALYSIS OF WAGE DIFFERENTIALS USING NONPARAMETRIC METHODS

Eric Makela
Clemson University, emakela@g.clemson.edu

QUALITATIVE AND QUANTITATIVE ANALYSIS OF WAGE DIFFERENTIALS USING NONPARAMETRIC METHODS

A Dissertation
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy
Economics

---

by
Eric Makela
August 2014

Accepted by:
Thomas Mroz, Committee Chair
Curtis Simon
Patrick Warren
Robert Tollison

ABSTRACT

The vast empirical literature on wage differentials reflects extreme interest in the subject from economists, social scientists, and policymakers alike. A review of the literature reveals two noticeable shortcomings. First, nearly all past studies have assumed a linear relationship between earnings and worker characteristics without consideration as to the potential specification error this imposes on their estimates. Second, researchers have generally focused on simple mean effects rather than estimating the effects of characteristics and wage rates on earnings distributions as a whole. Using data from the American Community Survey, this dissertation addresses both issues with a simple nonparametric approach, and applies the method to measure wage differentials between federal and private sector workers, as well as wage differences between males and females.

DEDICATION

My dissertation is dedicated to all of those who have donated their time, money, and effort for my benefit. I could list many names here, but will abstain in fairness to those I will never know.

ACKNOWLEDGMENTS

I need to acknowledge all those who made my dissertation possible, starting with Tom Mroz and Curtis Simon, as well as the other members of my committee. My thanks also goes out to all those who attended the labor seminar for their helpful comments and kind words of encouragement. A big acknowledgement also goes out to all my friends, and in particular Lyuda Sonchak, without whom I would not have made it through graduate school.

TABLE OF CONTENTS

LIST OF TABLES

List of Tables (Continued)

Table Page

CHAPTER 3

LIST OF FIGURES

Figure Page

CHAPTER 2

List of Figures (Continued)

Federal and Gender Wage Differentials: A Review

Eric Makela

July 2014

Abstract

The massive literature on wage differentials reflects great interest in the subject among social scientists. In fact, independent Google Scholar searches for "federal wage gap" and "gender wage gap" return over one million cumulative results. Due to technological and econometric advances, both differential literatures have developed in a predictably similar way. This paper summarizes the development of the federal and gender wage differential empirical literatures and provides a short discussion of some theories of the origin of these wage gaps.

1 Introduction

The empirical literature studying wage differentials is so vast one might mistake it for a book of federal rules and regulations. This dissertation studies two of the primary wage differential topics: gender wage gaps and public wage gaps. The purpose of this chapter is to familiarize readers with the evolution of empirical techniques used to measure wage differentials, as well as to introduce some of the theories regarding the origin of these wage gaps. Although this list of research is in no way complete, it should serve to demonstrate the advances of empirical work over the past 4 decades and help to identify the contribution of the following papers to this massive collection.

Section 2 provides a concise summary of the main wage differential theories. Section 3 outlines the empirical progress of researchers in the areas of gender and public wage differentials, while section 4 identifies my dissertation's niche in this literature and concludes.

2 Wage Differential Theory

A number of theories exist on why individuals with otherwise similar observable characteristics earn different wages for seemingly the same work. Compared to the size of the empirical literature, relatively few studies offer theoretical models to identify potential sources of wage differentials. Indeed, the primary focus of most research is on proper measurement of otherwise unexplained earnings gaps rather than providing a coherent explanation of the differential.

Probably the most notable theoretical explanation of wage differentials is that of pure discrimination, either positive (in the case of federal employees) or negative (in the case of women).[1] Compensating differentials are cited by Moore and Raisian

[1] For a complete discussion of discrimination as the driving force behind wage differentials, see Becker (1971).

(1991) as the potentially primary source of wage gaps. For example, suppose some particular characteristic of federal jobs, such as strict behavior policies or round-the-clock oversight, causes them to be less desirable. To entice workers to accept such job characteristics, federal agencies will have to pay higher wages to compensate their workers for the additional mental strain these policies impose. Studies using microdata rely on the assumption that, within reported occupation categories, federal and private jobs are equal. Should this not be the case, it may be that the measured wage differential is simply compensation for more undesirable job characteristics taken on by federal workers.[2]

The efficiency wage theory suggests that wage differentials arise not due to discrimination or systematic differences in job attributes between two groups, but rather because there are gains to productivity and efficiency from paying particular groups of workers more than others. If we consider differences in earnings by gender, firms could rationalize paying their male employees more than comparable females if, for example, males are more costly to recruit and train. This search interpretation is discussed by Bowlus (1997), who concludes that a portion of the raw gender wage gap is attributable to differences in labor market behavior between men and women. Similar theories regarding public employees are discussed by Gunderson (1979) and Bender (1998). In particular, if federal employees were given access to sensitive information, it may be optimal for agencies to have the lowest turnover rate possible, resulting in higher wages for most skill groups in the federal sector.

Numerous competing theories are cited in the literature as potential sources of observed wage gaps. Fogel and Lewin (1974) proposes that public wage differentials might be the result of difficulties applying the “prevailing wage” principle, which requires public entities to provide their employees “equal pay for equal work”. Not

[2]Or by men in the case of gender wage differentials.

only does this method of wage-setting fail because wage comparability surveys include only relatively large firms, but also because many jobs that exist in the public sector are nonexistent in the private labor market. Gunderson (1979) and Borjas (1980) offer some political explanations of the public-private wage gap. Borjas posits that public workers might earn more due to the dual mandate of their positions: they seek to provide public goods *as well as* produce political support for legislators. Gunderson cites the lack of a profit constraint in federal agencies as another source of differential pay practices. Short-run market disequilibria and union bargaining are also quoted as potential sources of both gender and public wage differentials, although there is little empirical evidence to support these arguments.

While some studies attempt to identify the source of wage differentials, others posit that the lack of optimal data renders these attempts futile. In particular, Rapaport (1995) tests for bias in estimates of the gender wage gap for teachers where wages are reportedly set in a nondiscriminatory fashion. Controlling for all personal characteristics used in the wage-setting process, the study still finds evidence of wage differentials. Rapaport concludes that, given her specification tests, her positive estimates could be the result of two potential data issues: measurement error or proxy error.[3]

The only consensus among labor economists regarding the source of wage differentials seems to be that the question will not soon be resolved. It should be noted that my studies cannot contribute to the theoretical literature, and my empirical results

[3] Gender wage differentials would be measured if there are systematic tendencies for one sex to misrepresent their human capital. For example, if females are more likely to overestimate their education or experience, regression analysis would measure a positive wage premium for men even if none existed (measurement error). Also, most micro datasets cannot perfectly measure human capital. Suppose a data source reports the number of years of education an individual has but cannot observe the quality of the education. If males tended to attend better quality schools than females, estimates would confirm a wage differential even if it could be accounted for by school quality (proxy error).

could be interpreted as one, or some combination, of the theories presented above.

3 Evolution of the Empirical Wage Differential Literature

This section will briefly describe many of the past studies in this literature. The empirical contribution of past research on wage differentials are classified into one of six categories based on their primary estimation method: linear Blinder-Oaxaca, dummy variable approach, sample selection models, indirect methods, quantile regression, and nonparametric models. Likely due to a multitude of reasons, the wage differential literature took flight after the developments of Blinder (1973) and Oaxaca (1973). Throughout the 1970's and 1980's academic journals were rife with papers measuring wage differentials using various forms of the Blinder-Oaxaca decomposition. Shortly after the development of the primary sample selection correction method by Heckman (1979), the differential literature evolved to include many variations of Heckman's correction. During the late 1990's and 2000's, some attention has shifted to estimating the *distributional* effects of gender and public employment as estimated by quantile regression methods. Most recently, researchers have used a quantile regression modification of the Blinder-Oaxaca decomposition. Although the wage differential literatures evolved along much the same timeline, I will present literature reviews of my two issues separately.

3.1 The Federal-Private Wage Gap

While conducting a review of this literature, a few sets of results seem particularly noteworthy. First, under the realm of public employment in the United States, federal employees are consistently estimated to earn significant wage premiums when compared to private sector employees. Second, state government workers typically earn either zero wage premium or a negative premium, depending on the sample used in the study. Last, local government workers are nearly always estimated to earn lower wages than private sector workers after controlling for individual characteristics.[4] To simplify comparison and to provide an accurate portrayal of where my paper fits in the literature, this section is restricted to studies of wage comparisons between the private sector and the central (federal) government. Some relatively recent reviews of the literature were performed by Bender (1998) and Gregory and Borland (1999). Readers are referred to these for complete reviews of the public wage differential literature prior to 2000.

In one of the earliest empirical studies of the earnings of public employees, Fogel and Lewin (1974) set the benchmark for this literature. Using primarily wage survey data, the authors compare mean earnings of public and private employees controlling only for differences in occupation. The study concludes that workers in some occupations receive federal pay premiums while those in other occupations do not. Despite the inconclusiveness of their results, the study pinpoints a primary difficulty in all public wage comparison studies: the necessary assumption that public and private jobs of the same occupational classification are, in fact, equal in terms of job attributes. The authors conclude that empirical studies based on microdata are unlikely to solve this issue soon.

The next empirical comparison of public and private wages came via Smith (1976) and Smith (1977).[5] Smith (1976) uses Census samples from 1960 and 1970 to estimate linear Blinder-Oaxaca models to compare the earnings of federal and private sector

[4]The magnitudes of differential estimates vary notably depending on the sample and analytic methods used, but these patterns tend to hold throughout the literature.

[5]Hammermesh (1973) studied the effects of government employment on union wages using a variety of different methods. His study is excluded from this list because the results are confounded by the union/nonunion dichotomy and are therefore not directly comparable to most other studies in this area.

workers. She found that on average, federal employees earn more. In 1960, approximately 60-65 percent of the absolute differential is not accounted for by differences in worker characteristics, and is labeled by Smith as an "Economic rent" accruing to those individuals in federal jobs. After refining her sample slightly, Smith (1977) separates males and females to estimate federal wage differentials for each sex. Not only did the females in her sample receive larger federal wage premiums than the males, but a greater portion of the premium is not explained by measured characteristics. Sharon Smith's primary contribution was to set a benchmark for Blinder-Oaxaca decompositions of public wage premiums, and her work prompted similar research in the United States and other countries.[6]

Subsequent Blinder-Oaxaca analyses generally have attempted to measure wage premiums more accurately by adding theoretically relevant explanatory variables to the wage equation.[7] For example, Quinn (1979) utilizes the 1969 wave of the Retirement History Study to measure wage premiums for workers in public administration occupations. Quinn's hypothesis was that public workers are more productive because they generally have more firm-specific human capital than workers in the private sector, who tend to change employers more often; the RHS survey allowed Quinn to control for both years of occupational experience *and* tenure at their current employer. His results suggest that federal public administrators earn an unexplained wage premium of about 20 percent. However, Quinn also notes that federal workers would earn more even if the two groups were paid equally because the federal workforce tends to have more education, more training, longer job tenure, better health, and also tends to be present in areas where wage rates are high.

Like Quinn (1979), many studies in this area have restricted their sample of public

[6] See Gunderson (1979) for an early analysis of public and private wages in Canada.

[7] For a complete listing of studies using linear Blinder-Oaxaca decompositions models, readers are referred to the "Double equation technique" section of Bender (1998).

workers to include only those in "public administration" occupations. Belman and Heywood (1988) propose that past differential estimates have been biased upward as a result of differences between wages in public administration and other government positions. Although their sample includes all levels of public employees in their analysis, the authors estimate a 1.1 percent wage premium when all public workers are included in the sample and a 6.6 percent wage *penalty* when public administrators, which constitute approximately one third of all non-private observations, are removed from the sample.

In an effort to further refine measurement of public wage differentials, Belman and Heywood (1990) suggest that when government agencies set their wages, they aim to set them comparable to those of workers in the largest private firms. Following the same B-O methodology, Belman and Heywood (1990) document how the inclusion of firm size in the wage equations drastically alters estimates of the public wage premium. In particular, the inclusion of firm size reduces federal wage differential estimates by 49 percent for females and by 25 percent for men. In general, Belman and Heywood demonstrate how sometimes arbitrary sample restrictions or omitted variable bias might cause past estimates of public wage differentials to overestimate the true effect of government on an individual's earnings. The trend in the empirical literature is toward B-O models with an increasing number of explanatory variables; in general, the more personal characteristics are accounted for in the model the smaller the estimate of the federal wage differential.

An important update of empirical methods came when Belman and Heywood (1989) noted that previous studies failed to account potential non-random selection into public jobs. Following methods developed by Heckman (1979), Belman and Heywood (1989) estimate a two-stage selection model correcting for non-random sorting into employment sectors based on observables. Comparing results with and without the selection term, it is clear that addition of the first-stage correction alters estimated marginal returns significantly. After controlling for selection into government work, the authors estimate white males in public jobs to receive a wage *penalty* between 4 and 19 percent of the average private wage, depending on the particular subsample of interest.[8] In another study, Gyourko and Tracy (1988) correct for the endogeneity of both employment sector and union status in their two-stage wage regression model. The authors measure a federal wage premium between 18.8 and 28.9 percent, depending on if selection is based on worker characteristics or differences in returns to those characteristics between sectors, and conclude that wage comparability legislation has done little to shrink the gap between federal and private compensation.

Around the same time period of the selection literature, some researchers were attempting to find indirect evidence that public employees receive a better overall compensation package. As Long (1982) points out, most studies of public wage differences are inconclusive, at least in terms of the magnitude of the estimates, because "the lack of individual data on fringe benefits has made it impossible to compare total compensation levels between similar workers in the public and private sectors." Long tests for overpayment by government agencies by comparing job separation behavior in the public and private sectors. His evidence clearly suggests that nearly every type of government employee, particularly federal workers, are significantly less likely to quit or be fired from their jobs. Along similar lines, Venti (1987) provides evidence that the skill groups which receive the largest federal wage premium are more likely to accept federal jobs, although these workers are also the least likely to be offered one. Venti uses 1982 CPS data to simulate the length of federal job queues; he estimates that, given the prevailing federal wage differentials, almost three times as many men and over six times as many women would be willing to work federal jobs as will be

[8]The sample includes public employees at all levels of government.

hired.

Recently, researchers have changed their focus from measuring the average effect of government employment to attempting to measure the *distributional* effects of government on individual earnings. Estimation of public wage differentials via quantile regression has grown in popularity over the past two decades, and many studies have been conducted applying this technique to foreign data.[9] Further, using data from the British SCELI survey, Bender (2003) shows it is possible to apply quantile regression methods to decompose raw wage differentials at deciles along the distribution of earnings. Critics of this approach might argue that although it is an upgrade over normal quantile regression and could be useful in identifying patterns of systematic wage differentials, its results are extremely difficult to interpret. Perhaps for this reason, to my knowledge no studies to this effect have been performed using data from the United States.

3.2 The Gender Wage Gap

While theories of the gender wage gap had been proposed early in the timeline of economic literature, it was at the heart of the women's rights movement that empirical estimates of the gap were first being published. Consensus among early empirical works, among them Sanborn (1964) and Fuchs (1971), was that women in the 1950's earned an unadjusted wage approximately 60 percent that of men, and that adjusting for productivity-related characteristics increases this estimate to 82 percent at most. Although not based on modern econometric methods, these studies established a few patterns that will persist throughout the male-female wage differential literature: 1) Wage gaps are largest for self-employed workers and smallest for government workers

[9] For some examples, see Mueller (1998), Melly (2005), and Bargain and Melly (2008). Poterba and Rueben (1994) study the distribution of wage premiums accruing to state and local government employees in the United States.

2) Wage gaps increase with age 3) A non-negligible portion of the wage gap can be explained by occupational choice, as females are much more likely to be employed in relatively low-salary occupations.

Much like the public-private wage differential literature, the mass of empirical work on the gender wage gap came following the seminal work of Oaxaca (1973) and Blinder (1973). Oaxaca's study used data from the 1967 Survey of Economic Opportunity to calculate the first two-equation estimate of the female wage differential.[10] He concludes that gender discrimination accounts for between 23 and 28 percentage points of the total differential for whites and between 20 and 25 percentage points of the total black differential.[11] Almost simultaneously, Blinder (1973) published his findings on the gender wage differential obtained from the Panel Study of Income Dynamics. His results are similar to those of Oaxaca, estimating that only 13.8 percent of the 45.6 percent raw wage differential could be attributed to the superior characteristics of the male workforce, while the remaining 86.2 percent was due to differences in estimated skill prices.

Utilizing datasets with a rich array of personal characteristics, researchers soon began testing hypotheses which generally stated that a portion of the gender wage gap was due to innate differences in personality and work preferences which led men and women to prepare for the labor market differently. Using data from the NLSY 1972 cohort, Daymont and Andrisani (1984) find that job preferences, such as the need to make money, be a leader, or helping others, leads the sexes to sort into different career paths. When included in the first-stage regression of a B-O decomposition,

[10] Readers are referred to Fortin et al. (2011) for a review of the current state of empirical methods based on the Blinder-Oaxaca decomposition.

[11] Subsequent studies point out that Oaxaca's discrimination coefficient is simply the portion of the raw differential that cannot be attributed to the productive characteristics in the wage equation. Barsky et al. (2002) discusses several issues *not* pertaining to the decomposition method itself, but rather potential misspecification of the wage equation.

these preferences account for 3.4 percentage points of the 12.9 percent total wage gap, a larger portion than is attributed to differences in estimated coefficients.[12]

One criticism of early research in the area is that it is impossible to perfectly control for education quality and occupation. Wood et al. (1993) attempts to correct this problem by investigating differences in male and female lawyer earnings using a special dataset of University of Michigan Law School graduates. Because the survey consists only of graduates, education quality is considered equal for all observations and, since every individual has a law degree, all are expected to be employed as lawyers. In the first year of work post-graduation, the majority of the 2.3 percent earnings gap could be explained by hours worked and job setting (size of firm, public/private, nonlegal, etc), while only an insignificant 0.1 percent was attributed to discrimination. After 15 years of work experience, the data show a raw yearly earnings gap of 48.5 percent; of this, 3.4 percentage points are due to increased family responsibilities on the part of females, 12 percentage points are due to males working longer hours, 15 percentage points are attributed to job setting, and 12.4 percentage points are "unexplained" or caused by potential discrimination on the part of employers. Throughout their specifications, females are estimated to receive higher wages due to superior human capital assets. Using a similarly-constructed dataset of University of Chicago MBA graduates, Bertrand et al. (2010) provides evidence that controlling for personal work preferences causes estimates of the gender earnings gap to become insignificantly different from zero. The work of Wood et al. (1993) and Bertrand et al. (2010) provide evidence for the family vs. career hypothesis of the gender earnings gap, which suggests that young workers tend to have relatively equal earnings, while the gap in wages grows as some females take time off and/or reduce their work hours to take on additional family responsibilities.

[12] Fortin (2008) finds similar results from more recent cohorts of the NLSY.

Other researchers have highlighted the importance of occupational choice and the general change in the nature of work as important considerations in measurement of the gender wage gap. In particular, Goldin and Polachek (1987) analyze changes in the male-female earnings ratio from 1890 to 1970, finding that occupations requiring schooling saw the most marked increase in the wage gap over that time period. Perhaps more importantly, the authors find a noticeable decline in the returns to male-dominated traits such as strength, and conclude that this phenomenon likely plays a role in the decline of measured male-female earnings ratios over time. Along similar lines, Gronau (1988) estimates that, because males and females tend to sort into different job categories, they acquire generally unequal levels of measurable human capital. From this perspective, wage differentials arise not entirely from discrimination and differing levels of human capital, but also from preferences for certain kinds of work, which Gronau (1988) shows to be crucial in explaining the sex differential in training and wages. This and all aforementioned studies have all been based on the traditional Blinder-Oaxaca approach of estimating monetary returns to worker traits.

However, as noted by Hellerstein et al. (1999), without some measure of worker productivity it is impossible to determine if measured wage differentials reflect discrimination, differences in worker productivity, or other confounding factors. Their study utilizes the new Worker Establishment Characteristics Database (WECD) which allows workers to be matched with employer characteristics such as firm size and output. After many robustness checks, the authors conclude that, while females are estimated to have slightly lower marginal productivity than males, the wage differential is greater than the productivity gaps, suggesting other factors must be at play in the wage-setting process. Deriving estimates from the WECD merged with the Longitudinal Research Datafile, Hellerstein et al. (2002) test whether market competition for profits serves to reduce or eliminate the gender wage gap, finding evidence of increased profitability in firms that hire a greater percentage of females. Meng (2004) uses a unique employee-employer matched dataset to estimate negative "firm effects" on the gender wage gap, to the magnitude of 2 to 5 percent, suggesting that within-firm pay comparability practices narrow the measured gap between male and female earnings.

As with other differential literatures, there has been a shift towards estimation of differences in the distribution of male and female earnings. The recent development of quantile regression has fostered a number of studies from research economists outside the United States.[13] Most recently, Boudarbat and Connolly (2013) offers insights on the wage gap distribution in Canada by performing decomposition analysis at various points in the total wage distribution. Their results indicate that women at the bottom of the wage distribution have more valuable characteristics than men in the same position, yet earn 2.6 percent less; at the 90th percentile, males are more likely to be employed in high-wage industries, which partially accounts for the 8.2 percent raw wage differential.[14]

In general, recent papers attempting to accurately assess the gender earnings gap fall into one of three categories: those attempting to assess the distribution of the wage gap, those utilizing linked employer-employee datasets to find indirect evidence of the wage gap based on firm performance, and those with specialized datasets intended to better control for differences in occupation and human capital. The following papers fall into the first category.

[13] See Mueller (1998) and Galego and Pereira (2010) for some examples. Arulampalam et al. (2007) compares the distribution of the male-female earnings gap across European countries.

[14] Barón and Cobb-Clark (2010) applies the same quantile decomposition method to Australian data.

4 Conclusion

Although this summary of the literature is in no way complete, I hope it will serve as a general guide to its development. As data and computing power have become increasingly available, so have the capabilities of researchers to produce accurate estimates of wage differentials.

The magnitudes of wage differential estimates vary greatly by sample and analytical methodology. Estimates of the federal wage differential range from around 2 to 40 percent.[15] Estimates of the gender wage gap have ranged from 0 to about 30 percent.[16] Both differentials have been shown to shrink over the period from 1970 to 2010 likely due to social changes, wage comparability legislation, and superior data, which allow researchers to control for more of the characteristics and work preferences that play a role in both job selection and wage offers.

One aspect of the differential literature which remains to be fully investigated is the effect of functional form assumptions on measured wage gaps.[17] The following papers in my dissertation provide three important updates of the current literature. First, I relax all functional form assumptions to obtain nonparametric estimates of the federal and gender wage differentials using recent data from the American Community Survey. Second, I compare these estimates with those obtained using a classical Blinder-Oaxaca decomposition model to test for potential specification error. Last, I am able to estimate counterfactual wage densities to assess the impact of wage rates on the distribution of earnings. My results suggest that past studies may have *overestimated* the true wage differential due to errors in functional form, although since no studies have been conducted using the ACS no direct comparison can be

[15] For studies where a representative sample of both males and females are included.

[16] Using data from the last 50 years.

[17] The only exception being Rapaport (1995), who discusses the effects of potential measurement error and proxy error in differential estimates.

made. In general empirical studies are attempting to find more accurate measures of true wage differentials, and the following papers offer a potentially important starting point.

References

Arulampalam, Wiji; Booth, Alison, and Bryan, Mark. Is there a glass ceiling over europe? exploring the gender pay gap across the wage distribution. 2007.

Bargain, Olivier and Melly, Blaise. Public sector pay gap in france: new evidence using panel data. 2008.

Barón, Juan and Cobb-Clark, Deborah. Occupational segregation and the gender wage gap in private- and public-sector employment: A distributional analysis. 2010.

Barsky, Robert; Bound, John; Charles, Kerwin, and Luptin, Joseph. Accounting for the black-white wealth gap: A nonparametric approach. *Journal of the American Statistical Association*, 2002.

Becker, Gary. *The Economics of Discrimination*. 1971.

Belman, Dale and Heywood, John. Public wage differentials and the public administration "industry". 1988.

Belman, Dale and Heywood, John. Government wage differentials: a sample selection approach. *Applied Economics*, 1989.

Belman, Dale and Heywood, John. The effect of establishment and firm size on public wage differentials. *Public Finance Quarterly*, 1990.

Bender, Keith. The central government-private sector wage differential. *Journal of Economic Surveys*, 1998.

Bender, Keith. Examining equality between public- and private-sector wage distributions. *Economic Inquiry*, 2003.

Bertrand, Marianne; Goldin, Claudia, and Katz, Lawrence. Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, 2010.

Blinder, Alan. Wage discrimination: reduced form and structural estimates. *Journal of Human Resources*, 1973.

Borjas, George. Wage determination in the federal government: The role of constituents and bureaucrats. 1980.

Boudarbat, Brahim and Connolly, Marie. The gender wage gap among recent post-secondary graduates in canada: a distributional approach. *Canadian Journal of Economics*, 2013.

Bowlus, Audra. A search interpretation of male-female wage differentials. *Journal of Labor Economics*, 1997.

Daymont, Thomas and Andrisani, Paul. Job preferences, college major, and the gender gap in earnings. 1984.

Fogel, Walter and Lewin, David. Wage determination in the public sector. *Industrial and Labor Relations Review*, 1974.

Fortin, Nicole. The gender wage gap among young adults in the united states: The importance of money versus people. *Journal of Human Resources*, 2008.

Fortin, Nicole; Lemieux, Thomas, and Firpo, Sergio. Decomposition methods in economics. *Handbook of Labor Economics*, 2011.

Fuchs, Victor. Differences in hourly earnings between men and women. 1971.

Galego, Aurora and Pereira, Joao. Evidence on gender wage discrimination in portugal: Parametric and semi-parametric approaches. *The Review of Income and Wealth*, 2010.

Goldin, Claudia and Polachek, Solomon. Residual differences by sex: Perspectives on the gender gap in earnings. 1987.

Gregory, Robert and Borland, Jeff. Recent developments in public sector labor markets. *Handbook of Labor Economics*, 1999.

Gronau, Reuben. Sex-related wage differentials and women's interrupted labor careers-the chicken or the egg. *Journal of Labor Economics*, 1988.

Gunderson, Morley. Earnings differentials between the public and private sectors. *The Canadian Journal of Economics*, May 1979.

Gyourko, Joseph and Tracy, Joseph. An analysis of public- and private-sector wages allowing for endogenous choices of both government and union status. *Journal of Labor Economics*, 1988.

Hammermesh, Daniel. *The effect of government ownership on union wages*. Princeton University, 1973.

Heckman, James. Sample selection bias as a specification error. 1979.

Hellerstein, Judith; Neumark, David, and Troske, Kenneth. Wages, productivity, and worker characteristics: Evidence from plant-level production functions and wage equations. 1999.

Hellerstein, Judith; Neumark, David, and Troske, Kenneth. Market forces and sex discrimination. *Journal of Human Resources*, 2002.

Long, James. Are government workers overpaid? alternative evidence. 1982.

Melly, Blaise. Public-private sector wage differentials in germany: Evidence from quantile regression. *Empirical Economics*, 2005.

Meng, Xin. Gender earnings gap: the role of firm specific effects. 2004.

Moore, William and Raisian, John. Government wage differentials revisited. *Journal of Labor Research*, 1991.

Mueller, Richard. Public-private sector wage differentials in canada: evidence from quantile regressions. *Economics Letters*, 1998.

Oaxaca, Ronald. Male-female wage differentials in urban labour markets. *International Economic Review*, 1973.

Poterba, James and Rueben, Kim. The distribution of public sector wage premia: New evidence using quantile regression methods. 1994.

Quinn, Joseph. Wage differentials among older workers in the public and private sectors. *Journal of Human Resources*, 1979.

Rapaport, Carol. Apparent wage discrimination when wages are determined by nondiscriminatory contracts. *American Economic Review*, 1995.

Sanborn, Henry. Pay differences between men and women. 1964.

Smith, Sharon. Pay differentials between federal government and private sector workers. *Industrial and Labor Relations Review*, 1976.

Smith, Sharon. Government wage differentials. *Journal of Urban Economics*, 1977.

Venti, Steven. Wages in the federal and private sectors. In Wise, David, editor, *Public Sector Payrolls*. 1987.

Wood, Robert; Corcoran, Mary, and Courant, Paul. Pay differences among the highly paid: The male-female earnings gap in lawyers' salaries. *Journal of Labor Economics*, 1993.

Measuring the Federal Wage Premium: A Nonparametric Approach

Eric Makela

July 2014

Abstract

Recently, the earnings of federal workers have come under renewed scrutiny. Past research has consistently concluded that positive wage premiums have accrued to federal workers over recent decades, but has all been based on highly restrictive statistical assumptions. This paper utilizes recently-developed econometric techniques which do not require functional form specification to measure the federal wage differential. I find the average wage premium accruing to federal employees is approximately 10.8% for men and 22.4% for women. I provide evidence that the compressed nature of the federal earnings distribution is due to the wage structure rather than the distribution of productive characteristics in the federal sector. Evidence also suggests the popular linear Oaxaca-Blinder decomposition is ill-equipped to measure federal wage differentials for specific groups of interest, and that the more pliable nonparametric approach should be utilized when possible.

1 Introduction

According to the Bureau of Labor Statistics, federal employees have comprised greater than 2 percent of all workers every year since records became available in 1939.[1] U.S. public employment totals climbed to their highest level in history in 2009 amidst growing fiscal consternation at federal, state, and local levels of government. Now more than ever, the systematic differential of wage incentives between the public and private sectors is an important topic of study. While empirical research on the public earnings premium in the United States has been scant in recent years,[2] numerous past studies have attempted to measure differential wages paid to federal, state, and local government workers. Results have varied greatly depending on the dataset, explanatory variables, and functional form specification.[3]

All past research on the topic of federal wage differentials has relied on linear functional form assumptions.[4] The purpose of this paper is to apply recently developed techniques to measure the yearly earnings differential between private sector and federal workers. The nonparametric matching method proposed here does not require specification of a functional form, and thus will not introduce potential bias resulting from misspecification. I estimate propensity score functions which allow us to weigh observations in order to control for differences in the joint distribution of explanatory variables between sectors. In this way, I study the federal earnings premium using a program evaluation approach requiring less restrictive assumptions than those made in previous research on the topic.

[1]Calendar year 2007 being the lone exception, when federal workers constituted 1.99 percent of all working people.

[2]See Borjas (2002) for a noteworthy exception.

[3]Estimated federal differentials vary widely among specification and group of focus. See Smith (1977), Belman and Heywood (1989), and Belman and Heywood (1990) for examples. Also, see Bender (1998) for a summary of the literature.

[4]Fogel and Lewin (1974) is the sole exception, although occupation is the only human capital variable the authors are able to control for.

The study finds that the average federal salary for males is \$6,718 greater than the average yearly salary in the private sector. If male federal employees were compensated according to the value of their skills in the private market, earnings for the average federal worker would fall by 10.8 percent; if men working in the private sector were paid federal wages, the average private sector worker is estimated to earn approximately 7.5 percent more. I estimate large federal premiums for unskilled workers and negative premiums for workers with at least a bachelor's degree. Although the linear and nonparametric models produce similar point estimates, I provide evidence that the assumptions underlying linear estimates are likely untrue, suggesting the nonparametric treatment effects approach should be used when data allow.

The paper is organized as follows. The next section summarizes the data and variables used in the paper. Section three discusses the public labor market literature and describes previous methods used to analyse the federal earnings premium. Section four introduces notation and estimates average federal treatment effects, comparing results with those derived from linear models. Section five presents earnings density estimates to supplement previous treatment effect calculations, and section six investigates earnings premium estimates for specific cases of interest in the federal workforce.

2 Data and summary statistics

Data used in the analysis are from the American Community Survey, a yearly survey of households conducted by the U.S. Census Bureau.[5] The dataset allows identification of 3-digit occupational code, as well as differentiating between federal, state, and local government employees. State and local government workers are excluded from the sample for consistency purposes, although a similar analysis of the wage structure for

[5] Alexander, Doeken, Genadek, Ruggles, Schroeder, and Sobek (2010) of the IPUMS project.

these workers would also be enlightening. The study is restricted to observations aged 25 to 65 who have completed at least nine years of education, were working at the time of the survey, and were not enrolled in school. The sample also excludes employees who have worked fewer than 40 weeks over the past year, who report working fewer than 30 hours in a typical workweek, or who work for nonprofit organizations.[6] Also excluded from the dataset are 141,048 observations who are not matched with either private or federal government counterparts.[7] Restricting the sample as such allows focus on earnings of workers whose time is fully invested in their occupation.

Summary statistics for the dataset can be found in appendix table A.1. Annual earnings are adjusted to 2011 dollars using the Consumer Price Index. The average male federal salary is $6,718 more than the average male salary in the private sector. Females in federal jobs do relatively better than men, earning an average income $16,229 higher than females employed in the private sector. High school graduates comprise 16.78 percent of male federal employees, compared to 31.05 percent of private males. Also, a federal employee chosen at random is more than twice as likely to have attained an advanced degree than a random worker in the private sector.[8] The federal government tends to employ older workers, as a significantly greater percentage of their workforce has at least 25 years of workforce experience.[9] Male and

[6] Also, the sample excludes military personnel and individuals living in institutions or other group quarters. Self-employed workers are also omitted from the sample. Observations reporting hourly wages below $5.75 are also excluded, as these are either mistakenly coded or extreme outliers in the light of all other exclusions.

[7] 46,703 males and 94,345 females were not matched with observations from the opposite sector, reflecting a fair amount of sector-specific specialization.

[8] Education categories are clearly defined in the ACS survey. Observations with some college are defined as those who report having between zero and four years of college, but less than a Bachelors degree. The advanced or professional degree category is comprised of individuals who report having attained a Master's, professional, or Doctoral degree.

[9] Experience is calculated separately for each education category. For high school dropouts, experience = age-18. For high school graduates, experience = age-20. For those with some college, experience = age-21. For workers with a Bachelor's degree, experience = age-22. For those with an advanced or professional degree, experience = age-25.

female scientists constitute a greater proportion of the federal workforce, while service occupations (except for protective services) are much more prevalent in the private sector. While the data suggest federal workers are likely paid higher wages due to their superior education and experience assets, it is unclear that federal jobs always require superior or specialized skills. For example, over 19 percent of both male and female federal workers are classified as mail and message distributors, thus postal workers represent a sizable portion of federal employment.

Table 1 documents mean earnings by education and sector. It is clear that a larger percentage of the federal workforce has attended at least some college, while private firms tend to hire more blue-collar workers. Those federal workers with less than a Bachelor's degree are paid more on average than similarly educated workers in the private sector, while those with at least a Bachelor's degree earn less on average. It is also clear from the table that the federal government generally hires more a more educated workforce.

Table 2 shows the distribution and mean earnings by 5-year experience groups. The table provides evidence that a randomly chosen federal employee is likely to have more workforce experience than a random employee at a private firm. It is also clear from the table that all experience groups earn sizeable average federal pay premiums, and that the federal earnings gap is larger for more experienced workers.

Appendix tables A.2 and A.3 document the mean earnings and mean federal differential by occupation. Positive mean federal wage differentials exist for male federal employees in 30 of the 40 defined occupational categories, and for females in 37 of these occupations. Police and protective service personnel receive the largest average federal premium, while lawyers in federal jobs earn significantly less than those in the private sector.

The intent of this paper is to provide a thorough analysis of the effects of differen-

Table 1: Distribution and mean earnings, by sex, education and sector						
	Federal			Private		
	Observations		Mean	Observations		Mean
	Total	Percent	earnings	Total	Percent	earnings
			Males			
HS Dropout	698	1.13	\$49,318	86,758	4.98	\$39,170
HS Graduate	10,329	16.78	\$55,066	540,793	31.05	\$47,594
Some college	16,414	26.67	\$63,831	507,741	29.15	\$59,707
Bachelor's degree	19,405	31.52	\$85,277	423,712	24.33	\$98,737
Advanced degree	14,710	23.90	\$111,823	182,576	10.48	\$148,309
Total	61,556	100.00	\$80,425	1,741,580	100.00	\$73,707
			Females			
HS Dropout	801	1.20	\$40,182	47,628	3.27	\$26,876
HS Graduate	15,676	23.49	\$47,950	451,248	30.98	\$33,000
Some college	23,605	35.37	\$54,586	540,210	37.08	\$42,044
Bachelor's degree	15,079	22.59	\$71,994	303,968	20.87	\$65,092
Advanced degree	11,580	17.35	\$95,610	113,685	7.80	\$94,836
Total	66,741	100.00	\$63,905	1,456,739	100.00	\$47,676

Table 1: The sample consists of 1,803,136 males and 1,523,480 females. All data are self-reported to the U.S. Census Bureau.

tial wage rates in the federal and private sectors on the federal wage differential while controlling for differences in basic human capital, and to compare the nonparametric treatment effects with those obtained from other estimation methods. While previous research has often assumed the effects of federal employment are homogeneous for all skill groups,[10] it has been well documented that this is likely a false assumption. This suggests a more pliable, less restrictive estimation method should be used.

3 Previous methodologies

The earnings of public employees were the focus of many studies by labor economists in the years following the work of Hammermesh (1973), Ehrenberg (1973), and Fo-

[10]Borjas (2002) is the most recent notable paper to make this assumption.

Table 2: Distribution and mean earnings, by sex, experience and sector						
	Federal			Private		
	Observations		Mean	Observations		Mean
Experience	Total	Percent	earnings	Total	Percent	earnings
			Males			
0-5 years	1,507	2.45	$59,363	38,202	2.19	$57,156
6-10 years	4.845	7.87	$61,926	204,663	11.75	$52,700
11-15 years	5,114	8.31	$74,816	188,743	10.84	$66,744
16-20 years	6,911	11.23	$79,129	235,546	13.52	$72,928
21-25 years	7,712	12.53	$82,414	236,228	13.56	$78,796
26-30 years	11,058	17.96	$86,336	272,249	15.63	$81,969
31-35 years	11,673	18.96	$82,883	251,164	14.42	$79,721
36 + years	12,736	20.69	$84,319	314,785	18.07	$78,367
Total	61,556	100.00	$80,425	1,741,580	100.00	$73,707
			Females			
0-5 years	1,555	2.33	$54,569	41,382	2.84	$47,791
6-10 years	4,117	6.17	$55,466	156,612	10.75	$42,265
11-15 years	4,075	6.11	$62,865	126,977	8.72	$49,115
16-20 years	5,892	8.83	$62,754	163,391	11.22	$49,426
21-25 years	7,773	11.65	$63,458	174,352	11.97	$50,732
26-30 years	11,272	16.89	$67,825	222,545	15.28	$50,504
31-35 years	13,253	19.86	$65,822	208,273	14.30	$48,613
36 + years	18,804	28.17	$63,596	363,207	24.93	$44,967
Total	66,741	100.00	$63,905	1,456,739	100.00	$47,676

Table 2: The sample consists of 1,803,136 males and 1,523,480 females. All data are self-reported to the U.S. Census Bureau.

gel and Lewin (1974). Studies of federal compensation in the United States, from Smith (1976) to Borjas (2002), have consistently found a positive and significant wage premium for federal employees. Readers are referred to Bender (1998) and Gregory and Borland (1999) for a complete survey of the literature pertaining to the earnings of workers in the central government prior to 2000. While approaches vary, the theoretical basis of the public labor market literature is derived from the notion that public production and thus compensation decisions are made differently by government administrators than by private business managers. Numerous explanations of government wage differentials have been offered in the literature, some of which include productivity differentials, lack of budget constraints in central governments, and difficulty in pay comparability due to job differentiation between government and private work. Political constraints are also cited as a potential reason for this phenomenon, suggesting that political pressures may induce a misallocation of resources between the government and private markets.[11] Various reasons for this are discussed in Bator (1958) and Demsetz (1964); since my empirical results cannot identify the source of measured wage differentials, the presence of any form of government failure will not be not be discussed in this paper. Regardless of cause, a plethora of studies have confirmed the existence of U.S. federal wage premiums through a variety of methods.

3.1 Dummy variable approach

The most simple empirical approach to estimating public wage differentials is through use of a dummy variable. Using this method, both private and public sector wages are estimated simultaneously according to the equation

$$ln(w_i) = X_i\beta + S_i\delta + \epsilon_i, \tag{1}$$

where $ln(w_i)$ is the log of yearly earnings for person i, X_i is a vector of the productive characteristics of person i, and ϵ_i a person-specific error term. S_i is an indicator variable equal to one if the individual is working in a public job, thus the coefficient estimate $\hat{\delta}$ is said to measure the effects of government enterprise on an individual's observed wage rate. For U.S. federal government workers, this approach consistently yields positive coefficient estimates for the public sector dummy,[12] suggesting rents

[11] Refer to the theoretical sections of Gunderson (1979) and Borjas (1980) for discussion.

[12] The size of wage differentials does change significantly by sample. For instance, large wage differentials to the magnitude of 37% are observed for women, much larger than that for men.

Table 3: Log wage premiums using the dummy variable approach, males only						
Explanatory variable		(1)	(2)	(3)	(4)	(5)
	Federal	0.02*	0.09*	0.09*	0.09*	0.08*
	HS Grad	0.22*	0.15*	0.15*	0.13*	0.15*
Education	Some college	0.43*	0.28*	0.30*	0.29*	0.27*
	Bachelor's degree	0.87*	0.55*	0.44*	0.64*	0.54*
	Advanced degree	1.23*	0.79*	0.61*	0.90*	0.78*
	6-10 yrs	0.20*	0.22*	0.21*	0.22*	0.21*
	11-15 yrs	0.34*	0.34*	0.23*	0.34*	0.43*
	16-20 yrs	0.47*	0.43*	0.25*	0.43*	0.57*
Experience	21-25 yrs	0.49*	0.40*	0.35*	0.49*	0.66*
	26-30 yrs	0.55*	0.53*	0.37*	0.53*	0.72*
	31-35 yrs	0.56*	0.53*	0.42*	0.54*	0.73*
	36+ yrs	0.49*	0.48*	0.42*	0.49*	0.70*
	Married	0.22*	0.18*	0.18*	0.18*	0.18*
	Occupation		x	x	x	x
	State of work		x	x	x	x
Specification	Educ x Exper			x		
	Educ x Occup				x	
	Exper x Occup					x

Table 3: The * symbol indicates that the coefficient is statistically significant at the 99 percent level. There are 61,556 male federal employees in the eleven year sample, along with 1,741,580 male private sector workers. The federal dummy variable equals one if the respondent was currently employed by the federal government at the time of survey. The married dummy variable equals one if a respondent is married with their spouse present and zero otherwise.

may be accruing to federal employees. The dummy variable approach is based on the econometric assumption that productive characteristics such as education and experience are rewarded equally in both sectors, and there exists no correlation between S and ϵ.

Five specifications of (1) are estimated and results recorded in tables 3 and 4, with explanatory variable *Federal* representing the estimate of $\hat{\delta}$ in equation (1). The

Public wage differences are also much greater on the low end of the wage distribution, and for racial minorities. Wage differential estimates are lower and often negative for state and local government employees.

Table 4: Log wage premiums using the dummy variable approach, females only						
Explanatory variable		(1)	(2)	(3)	(4)	(5)
	Federal	0.26*	0.21*	0.21*	0.22*	0.21*
	HS Grad	0.22*	0.14*	0.13*	0.18*	0.14*
Education	Some college	0.46*	0.26*	0.43*	0.39*	0.26*
	Bachelor's degree	0.88*	0.55*	0.63*	0.73*	0.54*
	Advanced degree	1.21*	0.77*	0.84*	0.99*	0.77*
	6-10 yrs	0.15*	0.16*	0.28*	0.17*	0.19*
	11-15 yrs	0.26*	0.26*	0.34*	0.26*	0.37*
	16-20 yrs	0.33*	0.31*	0.32*	0.32*	0.47*
Experience	21-25 yrs	0.37*	0.40*	0.35*	0.35*	0.54*
	26-30 yrs	0.39*	0.36*	0.41*	0.37*	0.58*
	31-35 yrs	0.39*	0.37*	0.45*	0.37*	0.57*
	36+ yrs	0.36*	0.34*	0.48*	0.35*	0.53*
	Married	0.02*	0.01*	0.01*	0.01*	0.01*
	Occupation		x	x	x	x
	State of work		x	x	x	x
Specification	Educ x Exper			x		
	Educ x Occup				x	
	Exper x Occup					x

Table 4: The * symbol indicates that the coefficient is statistically significant at the 99 percent level. There are 66,741 female federal employees in the eleven year sample, along with 1,456,739 female private sector workers. The federal dummy variable equals one if the respondent was currently employed by the federal government at the time of survey. The married dummy variable equals one if a respondent is married with their spouse present and zero otherwise.

dependent variable in each regression is the log of yearly earnings. In all specifications, male federal employees are estimated to earn between 2 and 9 percent more than their private sector counterparts, while results suggest females earn a federal premium between 21 and 26 percent. Due to the large sample size, all coefficients are highly significant. Moore and Raisian (1991) apply this technique using the 1979 and 1983 CPS samples, estimating federal hourly wage premiums between 7 and 13 percent. In his somewhat recent study of public wage premiums over time, Borjas (2002) finds the male federal differential fell throughout the 1980's, then fluctuated between 0 and

9 percent during the 1990's while the female differential remained steady between 20 and 29 percent throughout the same period.

3.2 Linear Oaxaca-Blinder approach

If one thinks productive characteristics might be valued differently in the private sector than in government work, the dummy variable approach will provide biased estimates of the true effect of government employment on an individual's earnings. Another common empirical method of deriving wage differentials was introduced by Blinder (1973) and Oaxaca (1973), and involves estimating wage equations separately for each group. Sharon Smith (1976) was the first to use the Oaxaca-Blinder decomposition method for study of federal employees in the United States, finding evidence of large federal wage differentials for hourly workers. Numerous subsequent studies[13] use this approach to find smaller unexplained wage differentials for federal employees in the U.S. and Canada. This method can be applied by first estimating

$$ln(w_{is}) = X_{is}\beta_s + \epsilon_{is}. \tag{2}$$

In the above equation, s denotes sector of employment; linear estimation of (2) is straightforward and allows productive characteristics to be valued differently in the private and government markets. Estimation gives the mean log wage in sector s as $ln(\bar{w}_s) = \bar{X}_s\hat{\beta}_s$, where $\bar{X}_s$ are the mean attributes of sector s employees and $\hat{\beta}_s$ are the estimated skill prices for those attributes in sector s. Slight manipulation of the estimates $\bar{X}_s$ and $\hat{\beta}_s$ gives the familiar linear version of the Oaxaca-Blinder decomposition of earnings premiums

$$\bar{w}_f - \bar{w}_p = \Sigma(\hat{\beta}_f - \hat{\beta}_p)\bar{X}_f + \Sigma(\bar{X}_f - \bar{X}_p)\hat{\beta}_p, \tag{3}$$

[13] See Smith (1977), Gunderson (1979), and Freeman (1987) for some examples.

where f denotes the federal sector and p denotes the private sector. The difference between the average federal and average private wage is now split into two portions. The second term on the right will amount to the nominal inter-sector difference in wages that can be attributed to higher average levels of education, experience, or other productive characteristics. This leaves the first term on the right side to equal the portion of the federal wage premium attributable to differential skill prices in the federal and private markets, which is the primary statistic of interest in the public labor market literature.

Note that the first term on the right side is a measure of the wage advantage enjoyed by federal employees that is attributed to differences in skill valuation between sectors. In other words, if federal agencies value a particular set of characteristics more than private firms, workers with those characteristics will receive a federal wage premium. In the program evaluation literature, such preferential treatment is called a treatment effect. If we consider federal employees the recipients of federal "treatment", the first term in decomposition (3) can be interpreted as a linear estimate of the average effect of treatment on the treated. Note also that equation (3) utilizes the private sector as its choice of weight, resulting in a measure of the average treatment effect on the treated, or ATT. Gregory and Borland (1999), among others, note that one may decompose wage differentials using a variety of weighting methods. An alternative would be to rearrange terms to estimate the average treatment effect on the untreated, or ATU. This is analogous to estimating the wage differential as $\bar{w}_f - \bar{w}_p = \Sigma(\hat{\beta}_f - \hat{\beta}_p)\bar{X}_p + \Sigma(\bar{X}_f - \bar{X}_p)\hat{\beta}_f$, thus the ATU is measured by the first term on the right hand side of the equation.[14]

[14]It is also possible to decompose the mean wage gap using any linear combination of the two decompositions described above. In particular, $ln(\bar{w}_f) - ln(\bar{w}_p) = \Sigma(\bar{X}_f - \bar{X}_p)\hat{\beta^*} + \bar{X}_f(\hat{\beta}_f - \hat{\beta^*}) + \bar{X}_p(\hat{\beta^*} - \hat{\beta}_p)$, where $\hat{\beta^*}$ is some linear combination of the federal and private skill price estimates. The issue with this approach is that there is no consensus on exactly what $\hat{\beta^*}$ should be. I ignore this problem by setting $\hat{\beta^*}$ equal to $\hat{\beta}_p$ or $\hat{\beta}_f$, resulting in the

I estimate (3) using the sample of individuals defined previously; results are reported in table 5. The raw male differential[15] is approximately 22.9 percent of mean private sector earnings. Estimated average treatment effects are somewhat greater for the treatment group than for the control. Estimates of average treatment effects are displayed in bold text to highlight their importance. Results indicate the average male in the federal workforce receives a skill price premium of about 10.9 percent, while the average male in the private market would receive a premium of 7.9 percent should he receive a federal job. While women earn less on average than men in both sectors, they earn significantly higher federal pay premiums, a result consistent with the findings of past research. The average female in the federal workforce receives a wage premium of 22.5 percent that is not accounted for by measurable productive characteristics. Sharon Smith (1976) argues that these unexplained premiums are equivalent to economic rents accruing to federal employees, although from an empirical viewpoint they amount to the portion of the federal wage gap that cannot be attributed to variables in the empirical model. Results still indicate there exists a significant unquantifiable difference in compensation practices between sectors.

Attempts at decomposing federal wage differentials have led researchers to vastly different estimates of the true wage premium. For example, Smith (1976) finds unexplained federal wage premiums of 30-32 percent, while Belman and Heywood (1990) find much smaller unexplained premiums. Chronologically, there has been a trend toward smaller unexplained residuals. However, no estimation procedure or set of explanatory variables will cause a misspecified model to produce the correct results.

This paper utilizes a different approach from that repeated in the literature. Rather than simply adding explanatory variables to a likely misspecfied model, I

---

twofold decompositions presented above.

[15] $ln(\bar{w}_f) - ln(\bar{w}_p)$

Table 5: Linear Oaxaca-Blinder decomposition				
	Males only		Females only	
$\bar{w}_f - \bar{w}_p$	0.229		0.381	
$(\bar{X}_f - \bar{X}_p)\hat{\beta}_p$	0.120	(0.04)	0.156	(0.03)
$(\hat{\beta}_f - \hat{\beta}_p)\bar{X}_f \rightarrow$ Linear ATT	**0.109**	**(0.02)**	**0.225**	**(0.02)**
$(\bar{X}_f - \bar{X}_p)\hat{\beta}_f$	0.150	(0.02)	0.186	(0.02)
$(\hat{\beta}_f - \hat{\beta}_p)\bar{X}_p \rightarrow$ Linear ATU	**0.079**	**(0.02)**	**0.195**	**(0.02)**
N	1,803,136		1,523,480	

Table 5: Regressions include five education dummy variables, eight experience dummy variables, and forty occupation dummy variables.

measure the federal wage premium using a nonparametric method which imposes no functional relationship between variables. Although less statistically efficient than a correctly-specified linear model, the nonparametric specification will provide superior measures of wage differentials if the relationship between skills and earnings is not truly linear.

4 Empirical Model

The primary empirical contribution of this paper is to obtain the first federal earnings differential measure which does not assume *a priori* knowledge of the wage equation's functional form. This study views the measurement of federal wage differentials through the lens of program evaluation, where federal workers are seen as the treatment group and private sector workers the control.[16] Using this methodology, I am

[16] I rationalize assignment of these groups as follows. In the private market, competition for profits should result in wage rates equal to the marginal productivity of the worker. Any worker paid more than his/her marginal product will be released from the firm or have their wages lowered; any worker paid below his/her marginal product would eventually be bid away by competing firms. The absence of a profit motive and hard budget constraint in the federal government is likely to lead to a misallocation of resources as well as rent-seeking behavior on the part of federal agencies and their employees (see Bator (1958) and Demsetz (1964) for discussion). Thus the private sector was chosen as the control group because observed wage rates are more likely to reflect true productivity

able to relax two restrictive assumptions made in previous research. First, because tables 1-4 suggest federal treatment does not affect the earnings of all workers equally, these effects are modeled as heterogeneous. I also eliminate the assumption of linear functional form. Because the true relationship between worker characteristics and salary is unknown, this removes any potential specification bias from the results.

We are able to relax homogeneity and functional form assumptions by using propensity score matching techniques. As in the program evaluation literature, we are then able to estimate the average effects of federal employment while removing possible selection bias. Note the impact of federal employment on an individual's earnings can be written

$$\tau_i = w_i(fed) - w_i(pri), \tag{4}$$

where $w_i(fed)$ is the worker's federal wage and $w_i(pri)$ is the worker's wage in the private sector. Clearly, one of these wage rates can be observed in cross-sectional data, while the other cannot, thus estimating the employer effect in the case of individuals is not possible in this setting.

It is possible, however, to estimate average treatment effects for specific groups of interest. While the federal workforce utilizes labor with a wide variety of skills, some skill groups are sparsely populated, suggesting federal jobs are not intended for workers of *all* kinds. Due to this, we focus the study on the average treatment effect on the treated and untreated rather than the federal effect on the labor force as a whole. The average effect of treatment on the treated, or ATT, is identified by Caliendo and Kopeinig (2008) as the most prominent statistic in the program evaluation literature. In the current setting, denote the ATT as

$$\tau_{ATT} = E(w(fed)|S = 1) - E(w(pri)|S = 1), \tag{5}$$

via measurable output.

where $S = 1$ implies the group having federal characteristics and $S = 0$ implies the group having private sector characteristics. We might also be interested in the effect treatment would have on the untreated population, in this case private sector workers. The average treatment effect on the untreated, or ATU, is

$$\tau_{ATU} = E(w(fed)|S = 0) - E(w(pri)|S = 0), \tag{6}$$

and measures how much the average private sector worker's earnings would change if he/she were to acquire a federal job. Similarly, (5) measures how the average federal employee's wages would change if he/she were to leave federal employment for private sector work.

Note that this approach is analogous to the linear Oaxaca-Blinder model, only with the added benefit of removing linearity assumptions, avoiding possible misspecification of the relationship between earnings and worker characteristics. I also relax the original model's homogeneity assumption by letting the wage effects of federal employment vary among workers. Relaxation of these assumptions leads to a more accurate and serviceable model with which to measure federal earnings premiums.

4.1 A hypothetical example

Allow us to continue by investigating a simple economy with q different types of workers who are paid an average wage of w_q. Let us also denote expected wages in the private and federal sectors respectively as $E_{pri}(w)$ and $E_{fed}(w)$. It follows that mean, or expected, wages in each sector are

$$E_{pri}(w) = \sum_{j=1}^{J} \pi_{j,pri} w_{j,pri} \tag{7}$$

$$E_{fed}(w) = \sum_{j=1}^{J} \pi_{j,fed} w_{j,fed}, \tag{8}$$

Table 6: Expected wage calculations in the federal and private sectors				
Type	Federal wage	Ratio of federal workforce	Private wage	Ratio of private workforce
1	\$10,000	$\frac{1}{3}$	\$10,000	$\frac{1}{3}$
2	\$20,000	$\frac{1}{3}$	\$20,000	$\frac{1}{2}$
3	\$30,000	$\frac{1}{3}$	\$30,000	$\frac{1}{6}$
	$E_{fed}(w) = \$20,000$		$E_{pri}(w) = \$18,333$	
	$\tau_{ATT} = \$0$		$\tau_{ATU} = \$0$	

Table 6: Federal workers earn more on average in this example, but not because of differential pay.

where the weights π represent the fraction of each worker type in that sector. As shown in table 1, mean federal earnings are significantly higher than mean private earnings. Without conditioning on worker characteristics, a simple comparison of these statistics is misleading. Specifically, proper comparison of expected earnings requires weights π to be equal across sectors. By conditioning on observable characteristics, I am able to evaluate the federal earnings premium, or the difference between expected earnings in the federal sector and expected earnings in the private sector.

To better frame the previous discussion, consider a hypothetical nation in which workers differ on only one dimension. For each type of worker, wages are equal in the federal and private sectors of the economy: type 1 workers earn \$10,000, type 2 workers earn \$20,000, and type 3 workers earn \$30,000.

The approach taken in this paper can be easily summarized by table 6. A naive comparison indicates that federal workers earn more on average than workers in the private sector. However, we know that conditional wages are the same regardless of employer. The initial analysis is flawed because mean earnings in each sector are calculated using different weights, meaning we are not measuring the effect of differential wages, but rather the effect of greater skill employed by the federal government.

However, it is straightforward to assess the federal premium in terms of treatment effects by imposing identical skill compositions in the two sectors. In the program evaluation literature, researchers attempt to eliminate bias by matching observations with similar observable characteristics. By matching the skill characteristics in the two sectors, the effect of treatment, which in this case is employment by a federal agency, is isolated. In particular, the estimate of τ_{ATT}, or the average treatment effect on the treated, is simply the weighted mean federal wage gap where weights are determined by the workforce composition of the federal sector. Likewise, τ_{ATU} assesses the mean federal wage gap using the private workforce composition as weights, and is said to measure the average treatment effect on the untreated. Note that this statistic is equal to zero in both cases, indicating the nonexistence of federal earnings premiums in our hypothetical world. Note that analysis by either dummy variable or linear Oaxaca-Blinder approaches in this case will also result in an estimate of zero government wage premium when accounting for worker type.

The previous example illustrates the spirit of the nonparametric approach proposed here. It is clear however, that individuals vary on many dimensions. This is the primary shortcoming of Belman et al. (1994),[17] who employ the methodology proposed here using the occupational composition of public and private workforces in the Wisconsin Wage Survey as a means of constructing weights, finding evidence of wage premiums for state and local government employees. A small sample size limits their ability to study additional distributional effects, hence this paper offers a vast improvement over the work of Belman et al. (1994), as I am able to control for other vital observable characteristics in the estimation process via superior data. While use of this method is non-existent in the federal labor market literature, similar

[17] This is also a prevalent issue in Fogel and Lewin (1974), who were the first to use the approach in studying public wage differentials.

approaches have proven fruitful to differential estimation in other arenas.[18] Because use of this reweighting technique for study of federal wage differentials is currently lacking, the purpose of this paper is to fill the gap of knowledge by demonstrating its use in this area.

4.2 Propensity score matching

In order to remove selection bias from the treatment effect calculation, we must account for differences in productive characteristics between the treated and untreated groups. To do this, we must first estimate

$$\theta(X) = Pr(S = 1|X). \tag{9}$$

$\theta(X)$ is a propensity score which measures the likelihood of an individual being a federal worker given his/her exact characteristics X, and its use in the program evaluation literature is well documented.[19]

Propensity score $\theta(X)$ can be estimated in one of two ways depending on the method of matching used. The sample can be split into cells based on categorical variables X then the proportion of observations in each cell calculated. This method allows the researcher to easily reweight observations such that the distribution of productive characteristics is equal in both sectors. To calculate average treatment effects, observations are then matched based on propensity score. The primary advantage of this characteristic matching method is that each observation in the federal sector is perfectly matched with private sector observations with the same observable productive assets. Characteristic matching therefore offers the best nonparametric approximation of average treatment effects because all information is used and only

[18] DiNardo, Fortin, and Lemieux (1996) study the effects of minimum wage and other labor market institutions on wage distributions, while Butcher and DiNardo (2002) use a similar approach in their comparison of immigrant and native wage distributions in the United States.

[19] See Caliendo and Kopeinig (2008).

comparable observations are matched with each other. Although there are costs associated with this approach,[20] the significant benefits outweigh the cost increased data requirements for my analysis.

Alternatively, $\theta(X)$ can be estimated using a standard discrete choice model, leading to a different class of matching estimators. Although observations are still matched based on their estimated propensity score, so-called nearest-neighbor matching matches observations from the federal sector with those observations in the private sector with the closest propensity score. Note that both methods will result in observations with the same characteristics having equal propensity scores. Nearest-neighbor matching, however, does not allow treated individuals to be matched with more than one nontreated person. When the treatment group is significantly smaller than the control, nearest-neighbor matching will eliminate all non-matched private observations from the sample, thus potentially important information is lost. In cells which have more federal than private sector workers, federal employees will be matched with private observations that may have very different characteristics. By imposing a functional form on the propensity score estimate, discrete choice models reduce the curse of dimensionality which plagues characteristic matching. However, the fact remains that nearest-neighbor matching can result in less than optimal matches. For comparability purposes, the effects of federal treatment on yearly earnings is estimated using both characteristic and nearest-neighbor matching.

In order to interpret the results, some statistical assumptions are needed. Rosenbaum and Rubin (1983) formalize the assumptions needed for strong ignorability of treatment assignment, which allow results to be interpreted as unbiased. The first of

[20]Notably, the characteristic matching method suffers the curse of dimensionality, meaning that additional explanatory variables potentially exacerbate the common support requirement. However, I am still able to control for personal characteristics which affect an individual's productivity and consequent earnings.

these conditions is conditional independence of outcomes. Formally, this assumption can be summarized as

$$(w(fed), w(pri)) \perp S|X,$$

where $\perp$ denotes independence. The conditional independence assumption requires that outcomes (earnings in this case) are independent of treatment after controlling for all observable characteristics X, and are thus there is no selection into employment sector based on unobserved characteristics. While Pfeifer (2008) provides evidence that individuals who are more risk-averse will tend to sort themselves into the public sector, it is not plausible to collect and correctly interpret such data. Since measures of innate skill, risk aversion, or attitude toward government are unavailable, there is little choice but to accept this point as a limitation of the model.

The second condition needed for strong ignorability is that of common support. In order for proper inference to be made regarding average treatment effects on a population, there must be sufficient characteristic overlap between the treatment and control groups. Thus it is necessary that

$$0 < Pr(S = 1|X) < 1,$$

meaning that each federal worker must have at least one comparable private sector observation. Should these conditions hold, treatment assignment can be considered strongly ignorable, which allows the researcher to eliminate selection bias based on observables by accounting for characteristic differences between the treatment and control groups.

4.3 Empirical estimation

Estimation of treatment effects requires matching based on characteristics or propensity scores. Define expected earnings in sector s conditional on propensity score as

$E_s(w|\theta(X))$. The conditional gap between federal and private wages is then

$$Q(X) = E_{fed}(w|\theta(X)) - E_{pri}(w|\theta(X)), \tag{10}$$

thus $Q(X)$ measures the gap between mean federal and mean private wages of matched groups. Note that while observations with equal characteristics will receive the same propensity score, treatment effect estimates will vary depending on the type of matching used.

Treatment effects are calculated by measuring the impact on wages while keeping workforce characteristics constant. Allow $g(X|S)$ to be the distribution of characteristics in sector S. The average treatment effect on the treated is then estimated as

$$\begin{aligned} \tau_{ATT} &= E(w(fed)|S=1) - E(w(pri)|S=1) \\ &= \int Q(X)g(X|S=1)dX. \end{aligned} \tag{11}$$

The federal composition of workforce skill is thus used to estimate a weighted mean treatment effect. The average effect of treatment on the untreated is

$$\begin{aligned} \tau_{ATU} &= E(w(fed)|S=0) - E(w(pri)|S=0) \\ &= \int Q(X)g(X|S=0)dX, \end{aligned} \tag{12}$$

and measures the increase in wages that would be realized by an average member of the control group (the private sector).

Estimation of τ_{ATT} and τ_{ATU} results in a point estimate of the wage effect of federal employment for the average worker in the treated and untreated groups, respectively. To calculate standard errors in the characteristic matching model, a simple bootstrap was used. Both statistics were calculated in 1000 repetitions of sample size N drawn with replacement.[21] Bootstrapped standard errors are clustered by the state in which the job is performed.

Although the nonparametric approach requires significantly more data, this complication eliminates a potentially worrisome property of linear O-B estimation methods. First, it eliminates the need to specify a functional form in the wage equation. Also, as noted by Barsky et al. (2002), linear regression minimizes mean square error for each sector's wage estimation. However, MSE is not necessarily minimized in estimation of counterfactuals, thus over- or underestimating both expected counterfactual wages and the contribution of explanatory variables. By estimating treatment effects using nonparametric matching, I eliminate any specification error in both wage regressions and counterfactuals.

4.4 Federal wage differentials

This section examines average federal treatment effects for men in the federal and private sectors. Raw wage gaps are seen in the first column of results in tables 7 and 8. Men in federal jobs earn approximately 22.9 percent more on average than men employed by private firms. Table 8 shows that by comparison, federal wages for average females are just over 38 percent higher than private sector wages. As noted previously, these estimates should be viewed as naive approximations of true federal wage premiums. If the federal government hires a more educated and experienced workforce, or if their workers perform jobs which are on average more valuable than those in the private sector, a positive differential should be expected. Estimates based on (11) and (12), corresponding to the second and third column of results respectively, and give the mean weighted federal earnings gap.

[21] Nearest-neighbor matching estimates were derived through 10 replications using the "psmatch2" Stata command.

When characteristic matching is used, I estimate federal employment increases earnings of the average male federal worker by 10.8 percent while the wage effect is only 7.5 percent for the average worker in the private sector. Another interpretation of these average treatment effects is possible. Should a random male federal employee lose his job and accept a new job in the private sector, he is expected to earn 10.8 percent less than previously. Likewise, if a random male working in the private sector were to obtain a federal job, it is estimated he will earn 7.5 percent more than he made in his private sector job.

Nearest-neighbor matching estimates slightly different effects. Notably, the federal treatment wage effect on the average federal employee is estimated to be 9.2 percent, while the effect for the average private worker is 9.8 percent. Because there exist some skill groups which are more numerous in federal work, nearest-neighbor matching with replacement produces noticeably larger standard errors. As noted by Lechner (2002), unnecessary inflation of the estimated variance will result in this case because a potentially small number of private sector observations will be repeatedly used in matching. In other cases where the number of private workers outnumber comparable federal employees, unmatched private observations are simply removed from the estimation sample, causing further inflation of estimated standard errors.

Table 7: Federal treatment effects, males only

Matching method	Raw differential	τ_{ATT}	τ_{ATU}
Characteristic matching	0.229	0.108	0.075
		(0.011)	(0.019)
Nearest-neighbor matching	0.229	0.092	0.098
		(0.037)	(0.019)

Table 7: Bootstrapped standard errors are displayed in parentheses. Nearest-neighbor matching utilizes a probit model to estimate propensity scores. Earnings are measured in 2011 dollars as the log of yearly labor income. Average log earnings are 10.947 in the private sector and 11.176 in the federal sector.

Table 8: Federal treatment effects, females only			
Matching method	Raw differential	τ_{ATT}	τ_{ATU}
Characteristic matching	0.381	0.224	0.189
		(0.015)	(0.019)
Nearest-neighbor matching	0.381	0.204	0.187
		(0.032)	(0.017)

Table 8: Bootstrapped standard errors are displayed in parentheses. Nearest-neighbor matching utilizes a probit model to estimate propensity scores. Earnings are measured in 2011 dollars as the log of yearly labor income. Average log earnings are 10.559 in the private sector and 10.940 in the federal sector.

Table 8 documents a similar series of estimates obtained for women. Matching based on observable characteristics, the calculated ATT suggests females in federal work receive an average premium of approximately 22.4 percent; nearest-neighbor matching predicts a slightly smaller effect. An average woman working in the private sector is estimated to earn between 18.7 and 18.9 percent more should she receive a federal job. Results suggest that, in general, characteristic matching offers more precise estimates of average treatment effects in this setting, and should be used when the data allow.

Although the evidence presented here is compelling, it would be remiss to not state the possibility that the explanatory variables do not adequately capture all wage-determining characteristics, and it is these omitted variables through which the earnings differential arises. If this is the case, we can infer only that my estimates measure the effects of all factors *other than* education, experience, and occupation. For example, if federal agencies tend to employ people who tend to be more hardworking than those in the private sector, average treatment effects would be overestimated. If this is not the case, the estimated federal treatment effect can be attributed solely to the difference in conditional wages between sectors, meaning that a randomly selected individual from the workforce will likely earn a significant wage premium in federal work.

4.5 Comparison with linear estimates

Having estimated the federal earnings premium for men and women using nonparametric models, it is appropriate to compare the findings to those obtained by more standard estimation procedures. Results of both models for males are shown in table 9, and hint that while the methods vary in their apportionment of measured wage premiums, point estimates of average treatment effects are similar in all three models.

In comparison, nonparametric characteristic and nearest-neighbor matching methods imply that 40-48 percent of the raw earnings gap can be attributed to differential treatment of federal employees. The linear OB model predicts that 10.9 percent of the mean differential is attributable to preferential treatment in federal work, thus little specification error can be observed. Nonparametric models predict the gains to be had for a private sector worker receiving a federal job are slightly less than estimated by the linear model.

Table 9: Performance of linear and nonparametric models, males only

Model	τ_{ATT}	$\frac{100 \times \tau_{ATT}}{0.229}$	τ_{ATU}	$\frac{100 \times \tau_{ATU}}{0.229}$
Raw differential = 0.229				
Linear	0.109	47.6	0.079	34.5
	(0.024)	-	(0.023)	-
Nonparametric (characteristic)	0.108	47.2	0.075	32.8
	(0.011)	-	(0.019)	-
Nonparametric (nearest-neighbor)	0.092	40.2	0.098	42.8
	(0.037)	-	(0.019)	-

Table 9: Percent columns refer to the portion of the raw earnings differential attributed to treatment effects. The linear model includes 5 education, 8 experience, and 40 occupation dummies. The nonparametric models includes full interaction terms in addition to these variables.

Table 10: Performance of linear and nonparametric models, females only

Model	τ_{ATT}	$\frac{100 \times \tau_{ATT}}{0.381}$	τ_{ATU}	$\frac{100 \times \tau_{ATU}}{0.381}$
Raw differential = 0.381				
Linear	0.225	59.1	0.195	51.2
	(0.029)	-	(0.026)	-
Nonparametric (characteristic)	0.224	58.8	0.189	49.6
	(0.015)	-	(0.019)	-
Nonparametric (nearest-neighbor)	0.204	53.5	0.187	49.1
	(0.032)	-	(0.017)	-

Table 10: Percent columns refer to the portion of the raw earnings differential attributed to treatment effects. The linear model includes 5 education, 8 experience, and 40 occupation dummies. The nonparametric models includes full interaction terms in addition to these variables.

Result comparison for females are displayed in table 10. Again, linear and characteristic matching models produce very similar point estimates. A generally higher portion of the raw differential is attributed to differences in skill prices for women than for men. Also noticeable is that estimates of τ_{ATT} are, with the exception of the male nearest-neighbor estimate, greater than τ_{ATU}, suggesting potential positive selection into federal work for members of both sexes.

The linear specification estimated here utilizes only explanatory variables used in the nonparametric models, although this need not be so. It is certainly the case that linear models may include a far greater number of explanatory variables without the exponential increase in data requirements which limits the effectiveness of nonparametric approaches. In this way, linear models may reduce the incidence of omitted variable bias. It is unclear as to the validity of the functional form assumption, thus in order to determine if this is the case, a more in-depth investigation is needed.

5 Counterfactual wage densities

Although previous sections suggest measurement of federal differentials through differencing of weighted means, it is possible to generalize the decomposition techniques to estimate entire counterfactual wage densities. The density of wages in each sector can be written as the integration of the conditional wage density over the distribution of productive assets in that sector. Allow $f_s(w)$ to represent the wage density for workers in sector s. Formally, earnings densities for the private and federal workforces respectively, are

$$f_{pri}(w) = \int h_{pri}(w|X)g(X|pri)dX \tag{12}$$

$$f_{fed}(w) = \int h_{fed}(w|X)g(X|fed)dX. \tag{13}$$

Above, $h_s(w|X)$ is the wage density conditional on productive characteristics X and $g(X|s)$ is the density of productive characteristics in sector s.

Under the assumption that alterations to the distribution of productive characteristics will not induce changes to the conditional wage density in either sector, we can imagine what the density of earnings in each sector would be if they were paid the skill prices given in the opposite sector. Optimally, the researcher would like to estimate

$$f_{pri}^{fed}(w) = \int h_{fed}(w|X)g(X|pri)dX \tag{14}$$

$$f_{fed}^{pri}(w) = \int h_{pri}(w|X)g(X|fed)dX. \tag{15}$$

Interpretation of these statistics is relatively straightforward. (14) is an estimate of the private sector wage density if workers were instead paid federal wages, while (15) estimates the federal wage density if its employees were paid their private market values.

However, researchers cannot observe these counterfactuals directly since an individual worker's earnings in the opposite sector are unknown. But these densities can be estimated empirically using the characteristic reweighting method.[22] In particular, I can estimate

$$f_{pri}^{fed}(w) = \int \theta_{pri}(X) h_{fed}(w|X) g(X|fed) dX \tag{16}$$
$$\theta_{pri}(X) = \frac{g(X|pri)}{g(X|fed)}.$$

Equation (16) is identical to (14) and can be interpreted as such. Similarly, it is possible to estimate (15) as

$$f_{fed}^{pri}(w) = \int \theta_{fed}(X) h_{pri}(w|X) g(X|pri) dX \tag{16}$$
$$\theta_{fed}(X) = \frac{g(X|fed)}{g(X|pri)}.$$

This visual exercise allows evaluation of the effect of the wage structure on the entire workforce. In a sense, I am able to assess whether the shape of earnings densities are due to the skills in that workforce or the skill prices offered in that sector. All that is needed is a reliable method of density estimation.

5.1 Kernel density estimation

Allow $\hat{g}(w_{0s})$ to represent the empirical density estimate in sector s at yearly salary w_0. The earnings density is estimated by the kernel function

$$\hat{g}(w_{0s}) = \frac{\theta}{h} \sum_{i=1}^{n} K\left(\frac{W_{is} - w_{0s}}{h}\right), \tag{15}$$

where h is the bandwidth and g is the univariate density of wages. The θ above are observation-specific weights, which depend on w_{0s} and the type of kernel function used in estimation.

[22] Also note that such counterfactual densities could be estimated using the linear O-B approach. In this case, we might calculate $w_{pri}^{fed} = X_{fed}\hat{\beta}_{pri}$ and $w_{fed}^{pri} = X_{pri}\hat{\beta}_{fed}$, then estimate the density of these counterfactual wages. It is assumed here that there is no variability of wages when conditioning on explanatory variables X.

An important consideration in kernel density estimation is the choice of bandwidth. Although the dataset is not lacking in observations, a proper bandwidth is needed to maintain a balance between bias and variance in the density estimates. Large bandwidths result in kernel estimates with low variance and high bias; selection of a small bandwidth will reduce the bias at the expense of higher variance. To select the bandwidth for each density estimate, this paper utilizes the plug-in technique proposed by Sheather and Jones (1991). Also relevant is the choice of kernel function. The Epanechnikov kernel is used here, although tests of alternative kernel functions suggest this issue is of little empirical importance in this instance.

Kernel estimates of federal and private sector male earnings densities are shown in figure 1a. Federal workers do not appear to be paid very low wages, and the variance of yearly earnings appears to be much lower in the federal sector. The wage gap is also apparent in the figure, as federal male workers earn an average of $80,425 while average earnings are only $73,707 in the private sector, suggesting a naive federal earnings advantage of $6,718 for male workers. Figure 1b documents the female wage densities for the two sectors, which appear visually similar to their male counterparts aside from the noticeably larger mean wage gap.[23]

5.2 Wage density estimates

As is clear from figure 1, both mean earnings and earnings densities in the federal and private sectors differ significantly. Many potential reasons exist for this disparity.[24] It has been documented that federal employees are on average more educated and experienced than their private sector counterparts, accounting for a portion of this differential.

[23] Average female earnings in the federal sector are $63,905 and average earnings are $47,676 in the private sector. The mean federal wage differential for females is then $16,230.

[24] See Bender (1998) for a summary of the theoretical literature.

Figure 1: Federal and private earnings densities

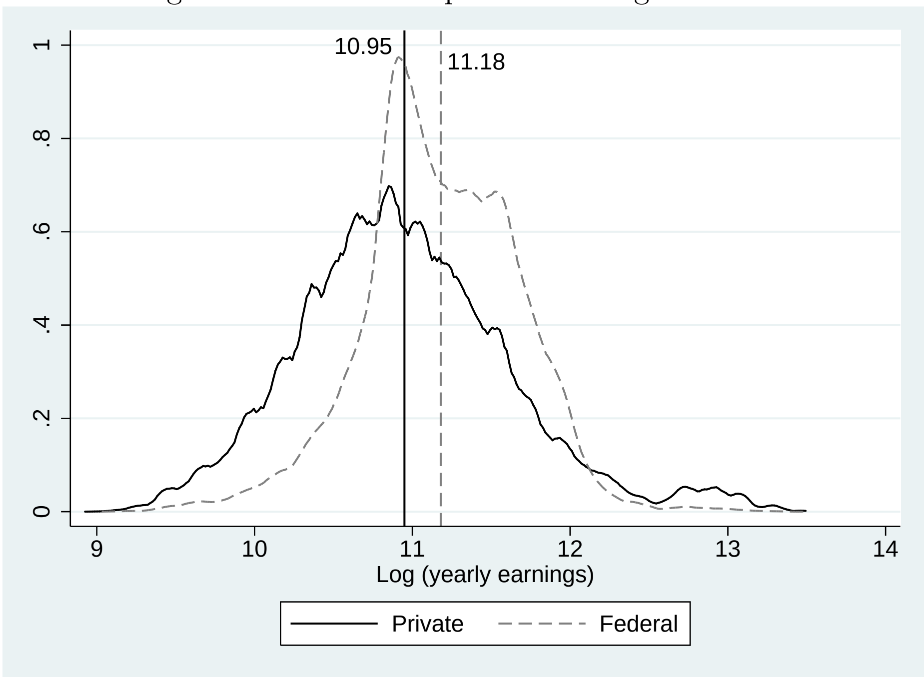

(a) Males

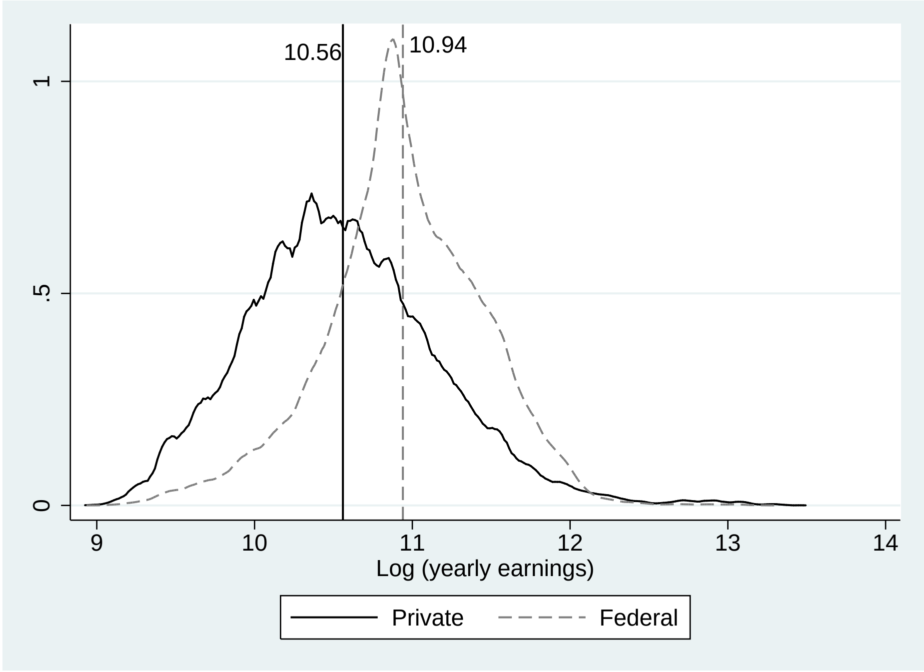

(b) Females

We wish to employ a nonparametric model to demonstrate exactly how these characteristics affect the prevailing wage densities. In order to perform this visual experiment, observation weights $\theta(X)$ must first be calculated. Observations are divided and sorted into cells according to their personal characteristics. The analysis focuses on 5 education groups, 8 experience groups and 40 occupations, resulting in a total of 5 x 8 x 40 = 1600 cells, or bins in each sector. The weight for each cell can be simplified conceptually to

$$\theta_{pri}(X) = \frac{\alpha_{pri}(n)}{\beta_{fed}(n)}, n \in 1, 2, ..., N \tag{16}$$

$$\theta_{fed}(X) = \frac{\beta_{fed}(n)}{\alpha_{pri}(n)}, n \in 1, 2, ..., N, \tag{17}$$

where n represents the bin number and N is the number of bins. $\alpha_{pri}(n)$ is the fraction of private sector workers in education, experience and occupation bin n, while $\beta_{fed}(n)$ is the fraction of federal employees in the same bin. In order to obtain a consistent, fully nonparametric estimate of propensity score function $\theta(X)$, it is necessary that $\frac{n}{N} \to 0$ as $n \to \infty$ and $N \to \infty$.

Kernel density estimates using this procedure are displayed in the following figure. In addition to the federal and private sector densities (the solid and dashed lines, respectively), two alternative wage densities are estimated, denoted by the dotted line. Figure 2a estimates the male private earnings density under the counterfactual assumption that private sector workers are valued as they are by the federal government. Figure 2b predicts the density of earnings that would prevail for male federal sector workers if the productive characteristics of its workers were valued as they are by the private market.

Observed mean log earnings for male federal workers in the sample are 11.18, while mean log earnings for private sector employees are 10.95. The vertical line represents the mean value for each alternative earnings density. Figure 2a suggests that given

Figure 2: Federal, private, and counterfactual kernel density estimates, males

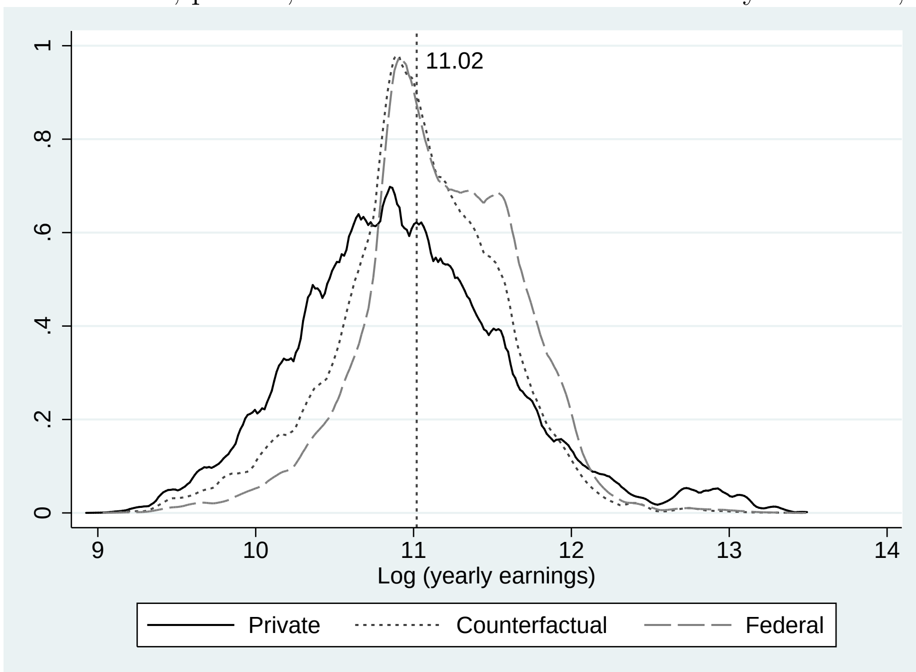

(a) Private characteristics, federal compensation

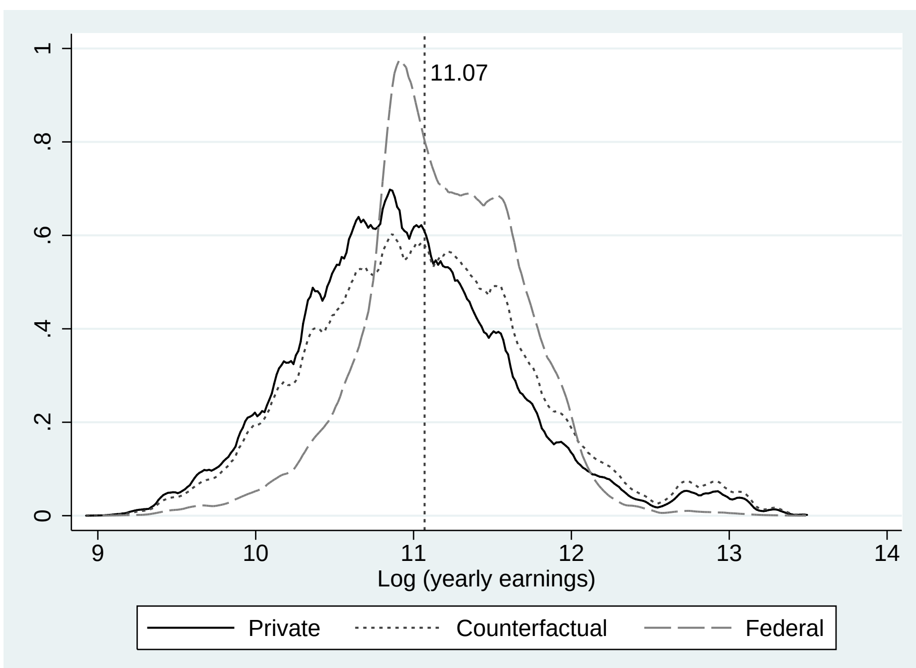

(b) Federal characteristics, private compensation federal skill prices, the density of earnings in the private sector shifts to the right; mean log earnings rise from 10.95 to 11.02.[25] Figure 2b implies that if federal workers were paid as their private sector counterparts, the wage density would more closely match that of the true private wage density. Mean log earnings for federal employees would fall from 11.18 to 11.07, but are more widely dispersed.[26] This implies that the low variance of wages in the federal sector is not due to its workers having similar characteristics, but rather is due to the federal compensation scale. These estimates correspond to average male treatment effects from table 7.

Note that the gap between the federal and counterfactual lines in figure 2a represents treatment effects for men in the private workforce. Given federal wages, there appears to be greater mass near the mean of the private wage density. Given private wages, the wage density for federal sector workers is more dispersed, displaying higher variance with more mass at the tails. We call the gap between the private and counterfactual wage densities in 2a to be the effect of treatment on the untreated. observed gap to solely to differences in conditional wage functions between sectors. Analogously, the gap between the countrefactual and federal lines in figure 2b signifies the effects of federal treatment on the treated population. If the conditioning variables adequately capture all wage-determining characteristics, these gaps can be attributed solely to differences in conditional wage functions between sectors. If we assume that unobserved characteristics also play a role in the wage-setting process, the aforementioned gaps signal the effects of all variables other than education, experience, and occupation.

Estimation results for female workers are displayed in figure A.1. When the private

[25] The standard deviation of log earnings in the private sector falls from 0.68 to 0.52 under the federal pay scale.

[26] If federal workers were compensated with private sector wages, the standard deviation of log earnings in that sector would rise from 0.49 to 0.71, and would thus display more variance than wages in the private sector.

workforce is given federal compensation, mean log earnings rise from 10.56 to 10.75. When federal sector female workers are paid private skill prices, mean log earnings fall from 10.94 to 10.72. Similar to the results for men, the distribution of wages displays much lower variance under the federal wage structure. Again, those individuals near the top of the earnings distribution tend to do relatively worse when given federal earnings. Not only do the figures offer visual confirmation of federal wage premiums, but provide further evidence that the effects of federal employment are far from homogeneous.

6 Further investigation

All results thus far have controlled for differences in the observable characteristics education, experience, and occupation, although making different assumptions regarding the relationship between these variables. Of course, there may be reasons to focus on more narrowly defined groups. One might wish to investigate federal wage competitiveness for a very particular type of worker; alternatively, a private firm may need to determine the wage it must offer in order to compete with the federal government for the best employees. This section explores relative wage profile estimates obtained from the three empirical models discussed in the paper.

Possibly the most noteworthy occupational group of federal employees are mail carriers. Comprising more than 20 percent of the male federal workforce, individuals in this occupation earn significant average wage premiums as shown in tables A.2 and A.3. With over 500,000 full time employees and over $45 billion in salaries and benefits paid in 2012,[27] the U.S. Postal Service is the largest federal entity currently participating in a consumer market.

Figure 3a displays estimated federal log earnings differentials for male postal work-

[27]Data from the United States Postal Service.

Figure 3: Estimated federal wage premiums by age, males only

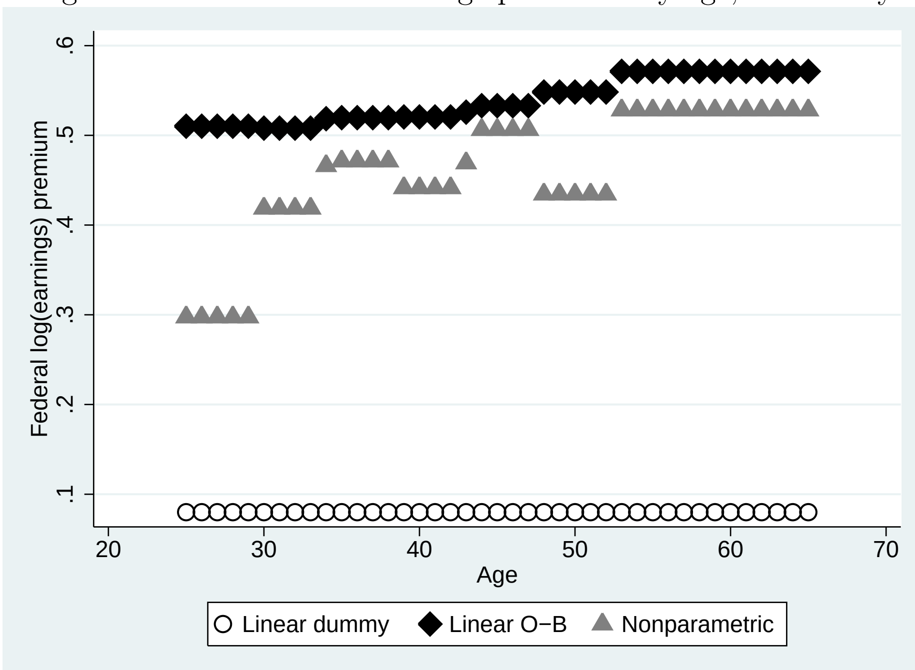

(a) Communication, mail, and message distributors, high school degree

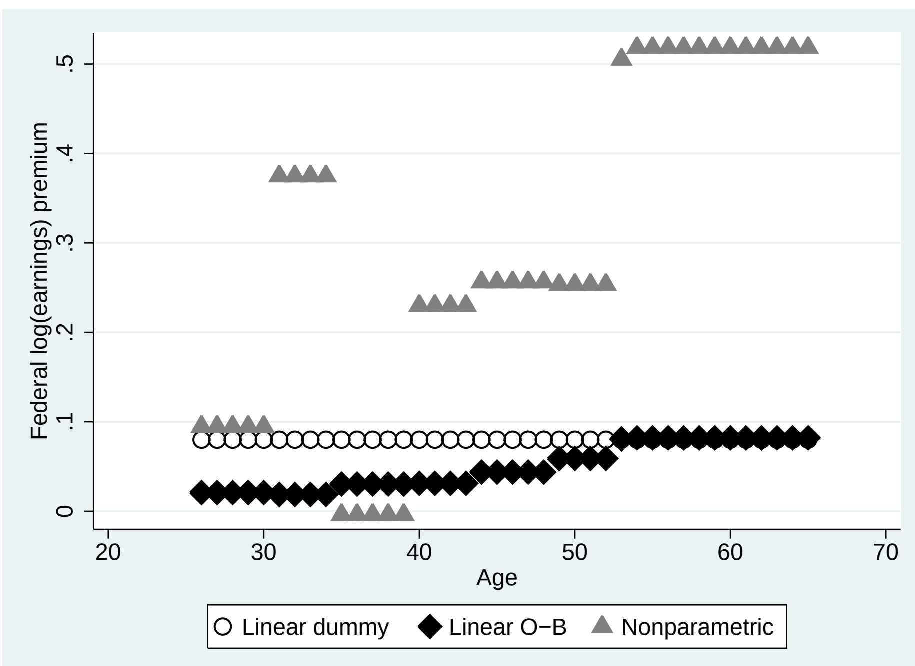

(b) Health services, some college

ers with a high school degree.[28] The dummy variable model, adopted from specification (5) in table 3, estimates that all male federal employees receive a log wage premium of 0.08 regardless of education, age, or occupation. Using the linear O-B approach, I find federal earnings premiums between 50 and 60 percent for mail carriers. Nonparametric models suggest linear O-B will overestimate the federal wage premium for this skill group, while the dummy variable approach will vastly underestimate the premium. While the magnitudes of estimates vary significantly by cohort using the nonparametric approach, it is clear that the two parametric models are misspecified in this instance, and that workers in this skill group receive greater monetary compensation from federal agencies than the private market.

Figure 3b demonstrates that differential estimates for some skill groups vary even more significantly by age and analytical method. While the linear models perform comparably for health service workers, the nonparametric model suggests the gap between federal and private sector earnings grows substantially by age, an observation which may lead to testable hypotheses regarding occupation and sector-specific human capital accumulation or labor market attachment.

Results for women are found in appendix table A.1, and document more variability in federal differential estimates than those for men. In particular, nonparametric models estimate some female mail carriers earn a wage premium surpassing 50 percent; lawyers are estimated to earn significantly less in federal jobs.

Visual results in this section further the earlier assertion that parametrically-specified models might measure differentials incorrectly, either under or overestimating the true differential for some skill groups. Researchers for their part must exercise caution when performing quantitative analysis of treatment effects, particularly if there are numerous groups of interest for which treatment effects might differ.

[28]These workers are classified as "Communications, mail, and message distribution."

More implications regarding the variation in differential earnings across ages for both skill groups are possible. It may be that mail carriers and health service workers in the federal and private markets reach their earnings peak at different ages. Another possibility is that unobserved cohort-specific skills cause the relative productivity of these groups to differ between the federal and private sector. Irrespective of cause, linear decomposition methods predict the relative earnings trajectory for all federal workers as compared to their private counterparts to be the same, simply adjusted based on educational attainment. In order to test theoretical hypotheses, it is vital to have the most accurate measures of federal wage differentials that are possible. The nonparametric model proposed in this paper can be viewed as a potentially important upgrade of current techniques.

7 Conclusion

The preceding analysis has measured the federal earnings premium using a variety of methods. I first examine the dummy variable approach, which concludes that male federal employees earn a wage premium between 2 and 9 percent while females earn a federal premium between 21 and 26 percent. However, summary statistics suggest that the effects of federal employment are likely heterogeneous across skill groups, a fact which invalidates inference using this approach.

Second, I estimate federal earnings gaps using the linear O-B model. These estimates suggest that the average federal male employee earns 10.9 percent more than he would in the private sector, while the average private sector male worker earns 7.9 percent less than he would were he given a federal job. These point estimates are similar to those obtained from the nonparametric matching model. However, subsequent analysis provides evidence that the two models do not perform quite so comparably in all situations. Even within narrowly-defined skill groups, differential estimates vary widely between methods, suggesting that the functional form imposed by parametric models may result in specification error.

To remove the specification error, I estimate two nonparametric models. When all data is utilized via characteristic matching, males are estimated to earn a federal premium of 10.8 percent while federally employed females receive an average wage advantage of 22.4 percent. When nearest-neighbor matching models are used, estimates of the federal earnings premium decrease to 9.2 and 20.4 percent for males and females, respectively. The nonparametric models demonstrated here makes no assumption regarding the functional form and will thus never introduce specification error into the results. Matching methods also allow for heterogeneity of treatment effects, with treatment in this case being employment by a federal agency. The benefit of relaxing these assumptions is that the approach provides more reliable and accurate estimates than parametric models when the functional form is unknown, while the cost is a simple increase in data requirements. Of course, researchers generally will not be able to compare multiple methods, nor have they a predetermined notion of whether the estimated relationship is truly linear. With the vast data archives available to statisticians and social scientists today, there would appear to be little need to make such restrictive assumptions when more advanced statistical methods are available.

It should also be noted that wage income measures only a portion of the likely federal earnings differential. The U.S. federal government spent $69.4 billion on federal employee pension contributions in the 2010 fiscal year, and there is speculation that federal pensions represent a disproportionately high percentage of total compensation as compared to an average private sector pension. Federal workers also receive health coverage in higher proportions than their private counterparts. It is therefore likely that this paper has underestimated the true federal earnings premium.

A common theme among all results is a positive and significant wage premium accruing to federal workers in the sample. Some skill groups, such as postal workers and police, earn wage premiums much higher than the average federal worker while other groups, such as lawyers and managers, tend to earn less in federal work than they otherwise would in the private market. At most, we may attribute the remaining unexplained wage gap to differential wage determination practices in the federal and private sectors. At a minimum, we can attribute this gap to all wage-determining factors other than education, experience, and occupation. The nonparametric matching model's generality leaves much in this topic to be investigated, and its lack of restrictive assumptions makes it an attractive alternative to parametric specifications when the functional form is unknown.

References

Alexander, J. Trent; Doeken, Ronald; Genadek, Katie; Ruggles, Steven; Schroeder, Matthew B., and Sobek, Matthew. Integrated public use microdata series: Version 5.0. *Minneapolis: University of Minnesota*, 2010.

Barsky, Robert; Bound, John; Charles, Kerwin, and Luptin, Joseph. Accounting for the black-white wealth gap: A nonparametric approach. *Journal of the American Statistical Association*, 2002.

Bator, Francis. The anatomy of market failure. *Quarterly Journal of Economics*, 1958.

Belman, Dale and Heywood, John. Government wage differentials: a sample selection approach. *Applied Economics*, 1989.

Belman, Dale and Heywood, John. The effect of establishment and firm size on public wage differentials. *Public Finance Quarterly*, 1990.

Belman, Dale; Franklin, Thomas, and Heywood, John. Comparing public and private earnings using state wage surveys. *Journal of Economic and Social Measurement*, 1994.

Bender, Keith. The central government-private sector wage differential. *Journal of Economic Surveys*, 1998.

Blinder, Alan. Wage discrimination: reduced form and structural estimates. *Journal of Human Resources*, 1973.

Borjas, George. Wage determination in the federal government: The role of constituents and bureaucrats. 1980.

Borjas, George. The wage structure and the sorting of workers into the public sector. *NBER Working Papers*, 2002.

Butcher, Kristen and DiNardo, John. The immigrant and native-born wage distributions: Evidence from united states censuses. *Industrial and Labor Relations Review*, 2002.

Caliendo, Marco and Kopeinig, Sabine. Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, 2008.

Demsetz, Harold. The exchange and enforcement of property rights. *Journal of Law and Economics*, 1964.

DiNardo, John; Fortin, Nicole, and Lemieux, Thomas. Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. *Econometrica*, 1996.

Ehrenberg, Ronald. The demand for state and local government employees. *The American Economic Review*, 1973.

Fogel, Walter and Lewin, David. Wage determination in the public sector. *Industrial and Labor Relations Review*, 1974.

Freeman, Richard. Longitudinal analyses of the effects of trade unions. *Journal of Labor Economics*, 1987.

Gregory, Robert and Borland, Jeff. Recent developments in public sector labor markets. *Handbook of Labor Economics*, 1999.

Gunderson, Morley. Earnings differentials between the public and private sectors. *The Canadian Journal of Economics*, May 1979.

Hammermesh, Daniel. *The effect of government ownership on union wages*. Princeton University, 1973.

Lechner, Michael. Some practical issues in the evaluation of heterogeneous labour market programmes by matching methods. *Journal of the Royal Statistical Society A*, 2002.

Moore, William and Raisian, John. Government wage differentials revisited. *Journal of Labor Research*, 1991.

Oaxaca, Ronald. Male-female wage differentials in urban labour markets. *International Economic Review*, 1973.

Pfeifer, Christian. Risk aversion and sorting into public sector employment. *IZA Discussion Papers*, 2008.

Rosenbaum, Paul and Rubin, Donald. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 1983.

Sheather, S.J. and Jones, M.C. A reliable data-based bandwidth selection method for kernel density estimation. *Journal of the Royal Statictical Association B*, 1991.

Smith, Sharon. Pay differentials between federal government and private sector workers. *Industrial and Labor Relations Review*, 1976.

Smith, Sharon. Government wage differentials. *Journal of Urban Economics*, 1977.

Appendix A

Table A.1: Descriptive statistics of the sample by sex and sector

	Male		Female	
	Federal	Private	Federal	Private
Mean yearly earnings (2011 $)	80,424.83	73,706.69	63,905.36	47,675.78
Std. dev. of yearly earnings (2011 $)	42,889.48	70,538.52	34,486.12	42,384.02
Education				
HS Dropout	1.13	4.98	1.20	3.27
HS Graduate	16.78	31.05	23.49	30.98
Some college	26.67	29.15	35.37	37.08
Bachelor's degree	31.52	24.33	22.59	20.87
Advanced degree	23.90	10.48	17.35	7.80
Experience				
0-5 years	2.45	2.19	2.33	2.84
6-10 years	7.87	11.75	6.17	10.75
11-15 years	8.31	10.84	6.11	8.72
16-20 years	11.23	13.52	8.83	11.22
21-25 years	12.53	13.56	11.65	11.97
26-30 years	17.96	15.63	16.89	15.28
31-35 years	18.96	14.42	19.86	14.30
36 + years	20.69	18.07	28.17	24.93
Occupation				
Managers	11.97	15.82	13.05	11.83
Accountants and auditors	2.59	1.76	3.63	3.21
Analysts, HR, and labor relations specialists	1.67	2.19	4.44	3.05
Inspectors and management support	2.85	1.15	3.26	1.50
Engineers	6.89	4.11	1.15	0.67
Mathematical and computer scientists	4.28	2.59	3.91	1.48
Natural scientists	5.21	0.58	2.48	0.37
Health diagnosis and treatment	1.99	1.54	5.09	6.21
Teachers (primary and secondary)	0.94	0.45	2.34	1.87
College and university instructors	0.24	0.19	0.17	0.27
Librarians, archivist and curators	0.24	0.02	0.43	0.05
Social scientists and planners	1.20	0.23	0.80	0.30
Lawyers	2.85	1.08	2.50	0.66
Artists and entertainers	0.77	1.69	0.90	1.68
Health technicians and specialists	0.64	0.43	1.86	2.93
Science technicians	2.30	1.49	1.07	0.57
Software developers	1.52	2.22	1.41	0.74
Other operators and technicians	1.85	0.63	1.63	1.33
Service sales representatives	0.11	1.43	0.13	1.50

Commodities sales representatives	0.31	7.63	0.85	6.84
Office supervisors	3.57	1.14	3.63	3.27
Receptionists and typists	0.31	0.33	1.60	3.08
Secretaries	0.21	0.17	5.28	7.69
Records and financial clerks	0.63	0.51	2.98	5.76
Communication, mail, and message distribution	19.62	3.27	19.03	3.02
Adjusters and investigators	1.12	1.53	3.21	5.19
Miscellaneous administrative support	0.99	0.54	4.13	3.91
Police, fire, and protective services	9.38	0.66	3.29	0.31
Food preparation and services	0.16	1.78	0.46	3.31
Health service occupations	0.19	0.24	0.90	4.04
Building and personal services	1.90	2.19	1.58	2.60
Farming, forestry, and fishing occupations	1.10	1.68	0.50	0.39
Automobile and electric mechanics and repairers	3.34	8.14	0.42	0.43
Construction workers	1.68	6.59	0.13	0.20
Extractive occupations	0.01	0.07	0.01	0.00
Precision production workers and supervisors	1.12	4.60	0.45	2.19
Machine operators	0.50	4.29	0.21	2.57
Welders and hand-assembly occupations	0.48	3.23	0.33	2.11
Transportation and material moving occupations	2.00	7.53	0.36	0.75
Helpers and other laborers n.e.c.	0.92	4.29	0.39	2.11
Observations	61,556	1,741,580	66,741	1,456,739

Table A.2: Male mean earnings, by occupation and sector

Occupation	Federal	Private	Differential
Managers	105,653	118,472	-12,819
Accountants and auditors	86,351	105,341	-18,990
Analysts, HR, and labor relations specialists	88,706	121,969	-33,263
Inspectors and management support	85,741	79,254	6,487
Engineers	97,447	94,247	3,200
Mathematical and computer scientists	91,519	83,376	8,143
Natural scientists	89,622	98,736	-9,114
Health diagnosis and treatment	141,131	153,634	-12,503
Teachers (primary and secondary)	67,736	57,529	10,207
College and university instructors	92,726	72,543	20,183
Librarians, archivist and curators	80,681	61,900	18,781
Social scientists and planners	113,192	118,559	-5,367
Lawyers	128,277	199,022	-70,745
Artists and entertainers	85,669	72,718	12,951
Health technicians and specialists	70,357	60,951	9,406
Science technicians	68,158	58,803	9,355
Software developers	93,570	98,158	-4,588
Other operators and technicians	112,307	74,709	37,598
Service sales representatives	94,154	132,486	-38,332
Commodities sales representatives	60,428	75,042	-14,614
Office supervisors	89,277	65,725	23,552
Receptionists and typists	51,051	49,439	1,612
Secretaries	65,913	50,136	15,777
Records and financial clerks	61,479	49,973	11,506
Communication, mail, and message distribution	54,892	44,669	10,223
Adjusters and investigators	64,775	55,064	9,711
Miscellaneous administrative support	59,523	55,973	3,550
Police, fire, and protective services	75,330	41,917	33,413
Food preparation and services	42,076	31,732	10,344
Health service occupations	54,919	40,227	14,692
Building and personal services	55,409	40,132	15,277
Farming, forestry, and fishing occupations	55,584	34,754	20,830
Automobile and electric mechanics and repairers	57,708	51,371	6,337
Construction workers	60,783	54,692	6,091
Extractive occupations	71,740	62,684	9,056
Precision production workers and supervisors	67,163	54,856	12,307
Machine operators	57,104	43,730	13,374
Welders and hand-assembly occupations	63,050	45,348	17,702
Transportation and material moving occupations	58,858	46,689	12,169
Helpers and other laborers n.e.c.	46,274	37,937	8,337

Table A.3: Female mean earnings, by occupation and sector

Occupation	Federal	Private	Differential
Managers	81,505	75,271	6,234
Accountants and auditors	69,233	62,053	7,180
Analysts, HR, and labor relations specialists	72,892	69,666	3,226
Inspectors and management support	72,938	57,132	15,806
Engineers	88,026	76,199	11,827
Mathematical and computer scientists	81,715	69,485	12,230
Natural scientists	77,284	76,509	775
Health diagnosis and treatment	82,690	67,774	14,916
Teachers (primary and secondary)	47,214	37,615	9,599
College and university instructors	67,932	55,461	12,471
Librarians, archivist and curators	74,292	53,084	21,208
Social scientists and planners	91,901	80,684	11,217
Lawyers	113,262	148,152	-34,890
Artists and entertainers	70,011	55,681	14,330
Health technicians and specialists	52,808	44,004	8,804
Science technicians	56,548	47,686	8,862
Software developers	83,921	80,877	3,044
Other operators and technicians	71,789	50,738	21,051
Service sales representatives	79,339	69,189	10,150
Commodities sales representatives	42,382	42,394	-12
Office supervisors	71,944	45,515	26,429
Receptionists and typists	40,384	31,424	8,960
Secretaries	46,092	37,554	8,538
Records and financial clerks	47,873	36,194	11,679
Communication, mail, and message distribution	48,331	33,566	14,765
Adjusters and investigators	54,309	38,392	15,917
Miscellaneous administrative support	46,232	33,764	12,468
Police, fire, and protective services	66,348	36,467	29,881
Food preparation and services	24,621	22,722	1,899
Health service occupations	36,209	28,089	8,120
Building and personal services	44,334	29,142	15,192
Farming, forestry, and fishing occupations	54,804	27,012	25,792
Automobile and electric mechanics and repairers	56,375	47,740	8,635
Construction workers	61,989	44,275	17,714
Extractive occupations	41,871	53,291	-11,420
Precision production workers and supervisors	48,349	35,269	13,080
Machine operators	41,612	30,076	11,536
Welders and hand-assembly occupations	50,089	32,800	17,289
Transportation and material moving occupations	45,628	34,407	11,221
Helpers and other laborers n.e.c.	35,891	25,597	10,294

Figure A.1: Federal, private, and counterfactual kernel density estimates, females

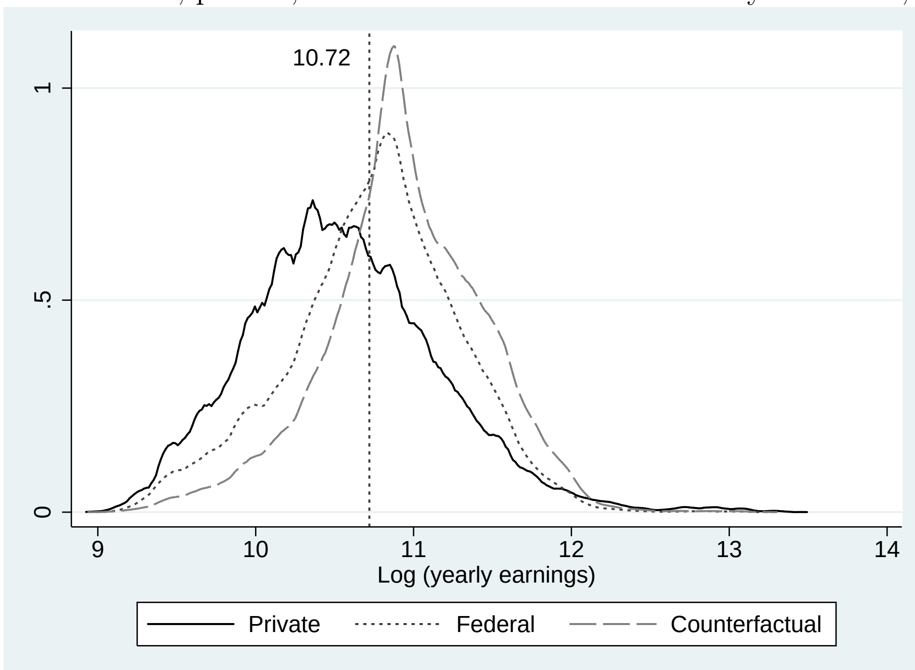

(a) Private characteristics, federal compensation

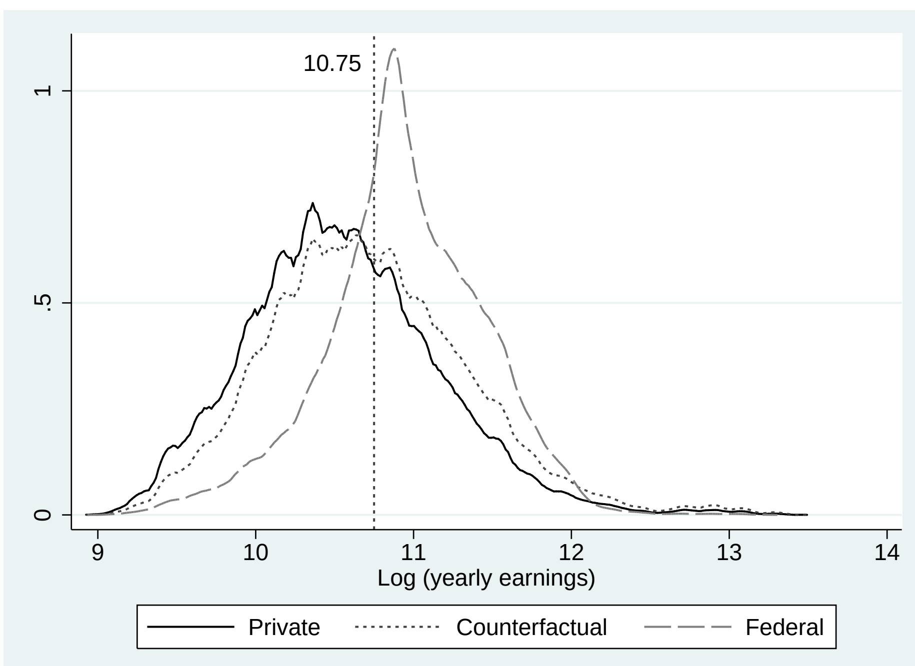

(b) Federal characteristics, private compensation

Figure A.2: Estimated federal wage premiums by age, females only

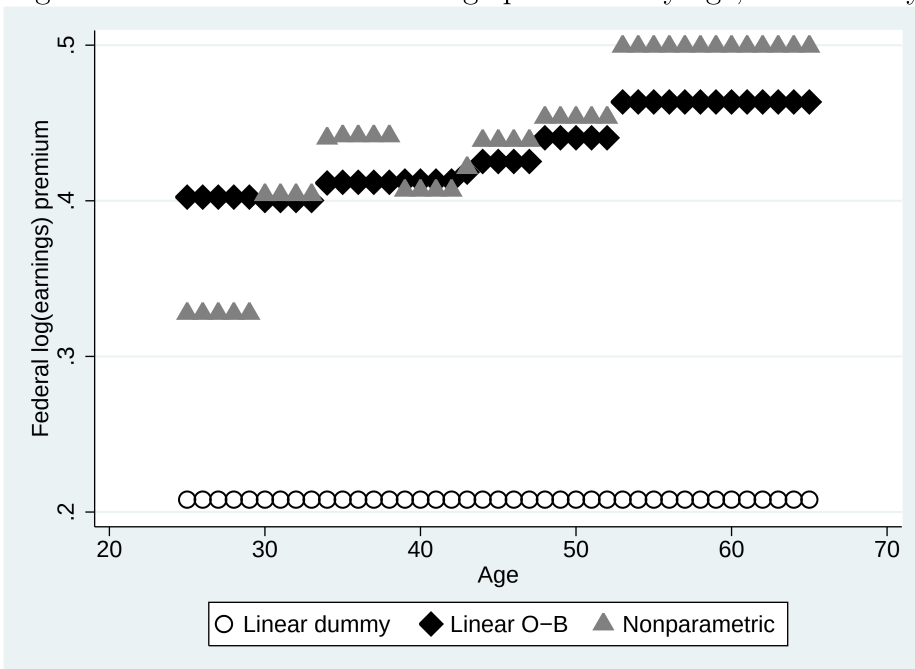

(c) Communication, mail, and message distributors, high school degree

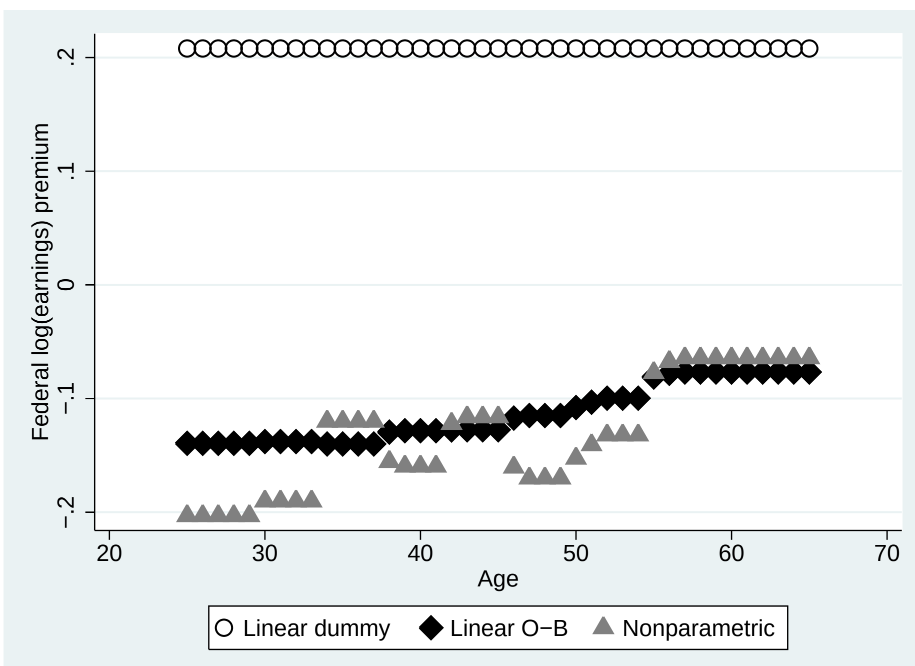

(d) Lawyers. professional degree

Appendix B

The purpose of this appendix is to provide a clear explanation of the empirical process, from sample creation to differential estimation. The following pages give precise instructions to obtaining my estimates, including variable categorization, sample exclusions, and coding (if deemed necessary). Anyone attempting to recreate my study may contact me at emakela@g.clemson.edu if they run into insurmountable issues.

Creating the sample

I begin with the full sample from the American Community Survey (henceforth ACS) from the years 2000-2011. Data was downloaded from IPUMS USA thanks to Alexander et al. (2010). Variable codes referenced in this appendix are those from IPUMS, and in some instances may differ from the coding structure in the raw ACS data files.

After importing the sample, I first remove all observations reported as living in group "institutions" (332558 observations deleted) or "other group quarters" (236691 observations deleted). Then, individuals over the age of 65 (3612641 observations deleted) and under the age of 25 (8058616 observations deleted) are removed. To remove any possible simultaneity in the effects of work and schooling on work productivity, I remove all observations currently attending school (762624 observations deleted). Those who report as having less than 9 years of education are also dropped from the sample (549852 observations deleted).

Since the study focuses on the work income, those observations who were not employed and currently working are removed from the sample, along with those employed by the armed forces (3563403 observations deleted). I then remove workers who report as self-employed (1062678 observations deleted), unpaid family workers (17634 observations deleted), and those who work for non-profit organizations

(724216 observations deleted). To conform with previous studies, I focus solely on full-time, full-year (FTFY) workers, thus I then remove individuals who work fewer than 40 weeks per year (570795 observations deleted) or fewer than 30 hours in a typical workweek (392305 observations deleted). Since veterans typically differ from non-veterans in terms of unmeasured human capital and would thus likely violate the model's statistical assumptions, I remove all veterans from the sample (716000 observations deleted). Also excluded are 3603 remaining observations who report as being employed outside the United States.

I then wish to remove those individuals who report wage income below the legal threshold, as these observations are either miscoded or misreported in light of all other restrictions.[29] The following code is used to convert earnings to 2011 dollars, then calculate hourly wages for each observation:

- replace incwage=incwage*cpi11
- gen ww=43.5*(wkswork2==4)+48.5*(wkswork2==5)+51*(wkswork2==6)
- gen totalhours=uhrswork*ww
- gen wage=incwage/totalhours

All observations reporting hourly wage income less than $5.75 (the lowest real minimum wage in the sample period) are then removed from the sample (118583 observations deleted).

In the spirit of abiding by the unconfoundedness assumption required to interpret my results, the next sample restriction seemed prudent. It is my opinion that there are likely to be systematic differences in unmeasured human capital stocks among different racial groups. To eliminate the possibility of this affecting my results, I construct my sample using only white observations, dropping all individuals who

[29] In real terms, the lowest minimum wage in the sample time frame was in 2006 when the real minimum wage (in 2011 dollars) was $5.75.

report their race as non-white (956787 observations deleted). For similar reasons, I also make a few sample restrictions based on occupation. These omissions are listed here:

- Social, recreation, and religious workers: drop if occ1990>173 & occ1990<177 (27378 observations deleted)
- Supervisors and proprietors of sales jobs: drop if occ1990==243 (170155 observations deleted)
- Farm operators and managers: drop if occ1990>472 & occ1990<477 (4875 observations deleted)

The final sample restrictions are made to adhere to the model's common support requirement. Because the wage effects of federal employment are clearly not equal for males and females, I conduct the analysis separately for each, meaning that I require common support for both sexes rather than the sample as a whole. For each sex, I calculate the ratio of workers in each education/experience/occupation group.[30] If there exists a federal or private skill group of either sex with no observations in the comparable skill group in the opposite sector, all individuals in this skill group are dropped from the sample. The (unsimplified) code for this restriction is as follows:

- bys sex sector: egen count_ss=count(logwage)
- bys sex sector ed exp occ: egen count_sseeo=count(logwage)
- gen w_eeo=count_sseeo/count_ss
- gen weeo_mg=w_eeo if male==1 & sector==1
- gen weeo_fg=w_eeo if male==0 & sector==1
- gen weeo_mp=w_eeo if male==1 & sector==0
- gen weeo_fp=w_eeo if male==0 & sector==0
- bys ed exp occ: egen mg_eeo=mean(weeo_mg)
- bys ed exp occ: egen fg_eeo=mean(weeo_fg)
- bys ed exp occ: egen mp_eeo=mean(weeo_mp)

[30] I will henceforth refer to each education/experience/occupation grouping a "skill group".

- bys ed exp occ: egen fp_eeo=mean(weeo_fp)
- gen fedw_eeo_m=mg_eeo/mp_eeo
- gen fedw_eeo_f=fg_eeo/fp_eeo
- gen priw_eeo_m=mp_eeo/mg_eeo
- gen priw_eeo_f=fp_eeo/fg_eeo
- drop if fedw_eeo_m==. (46703 observations deleted)
- drop if fedw_eeo_f==. (94345 observations deleted)
- drop if priw_eeo_m==.
- drop if priw_eeo_f==.

As is shown above, 46,703 males are dropped from the sample for common support purposes, along with 94,345 females.

In sum, I begin with the entire ACS sample from the years 2000-2011, which includes a total of 26,075,611 observations. After all excluded observations are removed, the final sample consists of 3,198,319 private sector observations (1,741,580 male and 1,456,739 female) and 128,297 federal sector observations (61,556 male and 66,741 female).

Variable creation

The paper focuses on three primary explanatory variables. Education is clearly defined in the ACS. School dropouts are defined as all observations reporting fewer than 12 years of education, or as having attended grade 12 but not having attained a diploma. High school graduates include those individuals with a GED or alternative credentials. The "some college" category includes all workers who report attending at some college, but have not attained a bachelor's degree.[31] The advanced degree educational category includes all observations who report as having attended college for 5 years or more, or have been granted a master's, doctoral, or professional degree.

[31] This category includes workers who have been granted a 2 year associates degree of any kind.

Experience categories account not only for potential workforce experience, but also for estimated work intensity. Potential experience is calculated as:

- gen potexperience=age-18
- replace potexperience=age-20 if ed==2
- replace potexperience=age-21 if ed==3
- replace potexperience=age-22 if ed==4
- replace potexperience=age-25 if ed==5

where ed=1,2,3,4,5 corresponds to education categories HS dropout, HS graduate, some college, bachelor's degree, and advanced degree, respectively. Because workers of different skill groups might accrue human capital at varying rates, I use March CPS samples (years 2000-2013) to calculate the average total hours worked for workers based on their birth cohort, education, and sex:

- gen birthyr=year-age
- drop if birthyr>1986
- gen cohort=1*(birthyr<1940)+2*(birthyr>1939 & birthyr<1945)
  +3*(birthyr>1944 & birthyr<1950)+4*(birthyr>1949 & birthyr<1955)
  +5*(birthyr>1954 & birthyr<1960)+6*(birthyr>1959 & birthyr<1965)
  +7*(birthyr>1964 & birthyr<1970)+8*(birthyr>1969 & birthyr<1975)
  +9*(birthyr>1974 & birthyr<1980)+10*(birthyr>1979)
- bys cohort ed fed: egen wksw=mean(wkswork1)
- bys cohort ed fed: egen hrsw=mean(uhrswork)
- gen totalhrs=wksw*hrsw
- replace totalhrs=totalhrs/2000

To construct experience categories, I merge the compressed CPS dataset with the ACS microdata by cohort, education, and federal employment status. Experience is then calculated as the product of potential experience and the work intensity variable totalhrs. This variable is then converted into a categorical variable based partially on the even distribution of observations in each category:

- gen experience=potexperience*totalhrs
- gen exp=1*(experience<5)+2*(experience>5 & experience<10)+3*(experience>10 & experience<15)+4*(experience>15 & experience<20)+5*(experience>20 & experience<25)+6*(experience>25 & experience<30)+7*(experience>30 & experience<35)+8*(experience>35)

Worker occupation was classified as closely as possible following the 1990 census occupation classification system. Using occupational categories which are too wide and inclusive would risk invalidating differential estimates because this would increase the risk of systematic differences in job characteristics between federal and private sector workers. However, using occupation categories which are too narrowly defined would result in greater difficulty achieving common support and would likely mean the elimination of many more observations from the sample. The following table contains the precise occupational codes used to create the occupation variable in the paper. Excluded occupational categories are listed above.

Estimation

Once the sample is created, estimation is relatively straightforward. The code sample referenced above can be used to calculate the reweighting variables θ in the paper's empirical model.[32] More precisely, to estimate the counterfactual wage density in figure 2a, one needs only to estimate the wage density for males in the federal sector with observations reweighted by the function priw_eeo_m (seen just above).

The process is almost the same when calculating nonparametric average treatment effects. When evaluating τ_{ATT}, I am simply recording the difference between the average private wage and the average federal wage, where the private sector workforce is reweighted such that it has the characteristics of the federal workforce, with which

[32] When all explanatory variables are accounted for, fedw_eeo_f is equivalent to $\theta_fed(X)$ in (16) if we wish to measure the counterfactual wage density for females in federal work. Likewise $\theta_pri(X)$ is equal to priw_eeo_f if we wish to analyze the federal wage gap for females using the private sector as a base.

Table A2.1: Categorization of occupational codes

Occupation	Code
Managers	occ1990<23
Accountants and auditors	occ1990¿22 & occ1990<25
Analysts, HR, and labor relations	occ1990>24 & occ1990<28
Inspectors and management support	occ1990>27 & occ1990<40
Engineers and architects	occ1990>42 & occ1990<60
Mathematical and computer scientists	occ1990>60 & occ1990<69
Natural scientists	occ1990>68 & occ1990<84
Health diagnosis and treatment	occ1990>83 & occ1990<110
Teachers (primary and secondary)	(occ1990>153 & occ1990<164) *(ind1990!=850)
College and university instructors	(occ1990>153 & occ1990<164) *(ind1990==850)
Librarians, archivist and curators	occ1990>163 & occ1990<166
Social scientists and planners	occ1990>165 & occ1990<174
Lawyers and judges	occ1990>177 & occ1990<180
Artists and entertainers	occ1990>180 & occ1990<201
Health technicians and specialists	occ1990>201 & occ1990<210
Science technicians	occ1990>210 & occ1990<226
Software developers	occ1990==229
Other operators and technicians	occ1990>225 & occ1990<240 & occ1990!=229
Service sales representatives	occ1990>243 & occ1990<257
Commodities sales representatives	occ1990>257 & occ1990<300
Office supervisors and equip. operators	occ1990>302 & occ1990<309
Receptionists and typists	occ1990>313 & occ1990<325
Secretaries	occ1990==313
Records and financial clerks	occ1990>325 & occ1990<345
Communication, mail, and message dist.	occ1990>344 & occ1990<374
Adjusters and investigators	occ1990>374 & occ1990<379
Miscellaneous administrative support	occ1990>378 & occ1990<390
Police, fire, and protective services	occ1990>410 & occ1990<428
Food preparation and services	occ1990>428 & occ1990<445
Health service occupations	occ1990>444 & occ1990<448
Building and personal services	occ1990>447 & occ1990<470
Farming, forestry, and fishing occupations	occ1990>478 & occ1990<500
Automobile and electric mechanics	occ1990>500 & occ1990<550
Construction workers	occ1990>550 & occ1990<600
Extractive occupations	occ1990>600 & occ1990<620
Precision production workers	occ1990>620 & occ1990<700

Machine operators	occ1990>700 & occ1990<780
Welders and hand-assembly occupations	occ1990>780 & occ1990<800
Transportation and mat. moving occs.	occ1990>800 & occ1990<860
Helpers and other laborers n.e.c.	occ1990>860 & occ1990<900 + occ1990==405

it is being compared. For example, to calculate τ_{ATT} for female private sector workers, one would use the following code:

- sum logwage if male==0 & sector==1
- global i=r(mean)
- sum logwage if male==0 & sector==0 [aw=fedw_eeo_f]
- global j=r(mean)
- global ATT=$i-$j
- di $ATT

Linear treatment effect estimates are obtained using the *oaxaca* package, where τ_{ATT} is equivalent to the "unexplained" portion of the measured mean wage gap. Linear wage equations are estimated using dummy variables for education, experience, and occupation, and do not include interaction terms.

The information provided in this appendix should be adequate to properly reconstruct the dataset and produce the empirical results in the paper. If anyone attempting this task is having difficulty, please feel free to contact me.

A New Perspective on the Male-Female Wage Differential

Eric Makela

July 2014

Abstract

The vast empirical literature on male/female wage differentials reflects extreme interest in the subject. However, the majority of empirical studies have used Blinder-Oaxaca decomposition techniques which assume *a priori* knowledge of the wage equation's functional form. This paper utilizes a generalized form of the popular Blinder-Oaxaca model to investigate how differences in compensation affect the overall distribution of female earnings. Data shows that much of the difference between male and female wage densities would be eliminated in the absence of differential wage rates. The proposed model estimates that the average female in the private sector is paid a wage approximately 21.4 percent lower than the average male salary after controlling for observable differences in human capital. Differences in differential estimates between the generalized and linear B-O models are also compared, and suggest that errors of functional form may be present in linear specifications.

1 Introduction

Differences in earnings between men and women are a subject of interest to economists and policymakers alike. Although empirical research has documented the relative gains achieved by women in the labor market over recent decades, unexplained wage differences persist.

Review of the literature reveals two primary shortcomings. First, an overwhelming majority of empirical studies focus on simple mean effects, failing to account for the distributional effects of human capital and gender. Second, among studies attempting to identify wage gaps along the distribution of wages, strong functional form assumptions have been used. The primary focus of this paper is to provide a clear analysis of the effects of both human capital and skill price differences on the distribution of male and female earnings and, subsequently, the male-female wage gap. The model proposed here is comparable to those used in previous studies in that both techniques identify the effects of gender on earnings for the average worker in a population. However, the proposed empirical method differs substantially from previous methods in that it allows estimation of the effects of conditional wages and human capital on the distribution of earnings for an entire population, all without *a priori* specification of a functional form of the wage equation.

Results confirm the existence of differential compensation practices between men and women. In the private sector, the average female receives a wage penalty of approximately 21.4 percent. In a labor market in which all workers were paid as if they were males, the female earnings density is estimated to shift rightward, eliminating a majority of the raw wage gap and increasing the mass of women in the upper tail of the earnings density. The proposed nonparametric model estimates smaller gender wage differentials than a linear Blinder-Oaxaca decomposition utilizing the same explanatory variables, suggesting possible functional misspecification of the wage equation.[1] Results from three differential estimation methods are then compared, and hint that nonparametric models should be utilized when data allow.

The paper is organized as follows. The next two sections outline relevant empirical studies and provide a thorough description of the data. Section four highlights the empirical model used in estimation, and presents the quantitative and qualitative results. Section five concludes.

2 Background

This paper builds on the litany of empirical work on gender wage differentials. Very generally, the function

$$w_i = f(X_i)$$

is used to represent the relationship between an individual's wages w_i and the characteristics X_i which affect how productive that individual is. Under the human capital theory, any estimated gap in earnings between males and females after controlling for productive assets is generally interpreted as the effect of gender, or of discrimination on the basis of gender. An interpretation of gender wage gaps under the efficiency wage theory might suggest that, for any reason, firm-specific human capital is more valuable to men than for women, and that higher wage rates for males are paid to encourage greater attachment to the firm. Measurement of wage gaps has also been attributed to various errors in the empirical process. Rapaport (1995) states that if systematic measurement error exists between the two sexes, wage differentials might appear in the data when in reality none exists. She also notes that proxy error may

[1] If this misspecification is systematic over all past studies using the linear Blinder-Oaxaca technique, it is possible that *all* the previous literature on the subject of male/female wage differentials has overestimated the true unexplained wage gap.

cause incorrect measures of wage gaps. For example, if male lawyers generally attend more prestigious law schools than female lawyers and quality of the school is not measured in microdata, a wage gap between men and women will be measured even though it might be accounted for by superior data. For generality's sake, I approach the problem of gender wage differentials agnostic to the specific mechanism through which the wage gaps arise; I will later discuss various interpretations of my results.

Empirical work on gender wage gaps[2] builds primarily upon the basic decomposition techniques introduced by Blinder (1973) and Oaxaca (1973). Baker et al. (1995) make an important observation of early empirical results, noting that the choice of sample, analytical methodology and explanatory variables appear to be important determinants of the estimated magnitude of the gender wage gap. Inconsistencies in sample and methodology render many results incomparable, another facet of the current literature which has influenced this paper.

While the majority of empirical studies focus on mean effects or gender wage gaps in narrowly defined skill groups, a select few researchers have focused explicitly on the distribution of the male-female wage differential. In their empirical analysis, Baker et al. (1995) employ quantile regression to estimate the gender earnings gap at various points in the overall distribution of earnings. Although they do not control for other observable characteristics in the analysis, they are unable to document significant differences in the gender wage gap between the 10th and 90th percentiles, but based on decompositions within skill groups conclude that the gender wage gap is larger for highly-skilled women in Canada.

In a more recent study, Boudarbat and Connolly (2013) use quantile decomposition techniques developed by Firpo et al. (2009) to decompose gender wage gaps along

[2]See Filer (1985), Gronau (1988), Groshen (1991), Lazear and Rosen (1990), Wood et al. (1993), Blau and Kahn (1997), Hellerstein et al. (2002), Fortin (2008), Galego and Pereira (2010), and Bertrand, Goldin, and Katz (2010) for a small sample.

the overall distribution of wages. The study finds similar results to those of Baker et al. (1995), suggesting that highly educated females in the Canadian workforce face the largest gender wage gap while those at the bottom of the earnings distribution suffer least from discriminatory pay practices. Although the study is restricted to only those individuals who are college graduates, the authors conclude that university graduates are the driving force behind the widening male-female earnings gap.

The few studies focused on the distribution of the gender earnings gap are able to make inference only when relatively restrictive statistical assumptions are used.[3] More specifically, past results can only be interpreted as valid when the wage equation is correctly specified. Herein lies a primary fault with the existing literature: all past studies, to my knowledge,[4] have assumed a linear functional form of the wage equation. If the relationship between earnings and personal characteristics is not linear, past results have may be biased due to specification error, although the direction of this bias is not entirely clear.

This study is similar in spirit to Baker et al. (1995) and Boudarbat and Connolly (2013) in that it attempts to identify the effects of gender on the entire wage distribution. Rather than assuming *a priori* knowledge of the wage equation's functional form however, I adopt nonparametric techniques developed by DiNardo et al. (1996) to assess the gender wage gap both quantitatively and qualitatively. The empirical results presented here are the first nonparametric estimates of the gender wage gap, and thus represent a marked improvement over previous studies in terms of model pliability and required statistical assumptions.

[3]Also, to my knowledge, no studies have performed such an analysis using data from the United States.

[4]Perhaps with the exception of Rapaport (1995), who estimates one fully nonparametric specification. She finds that race and gender are still estimated to have a significant effect on teachers' earnings, results which closely mirror those obtained by her parametric specifications.

3 Data

Data used in the analysis are from the American Community Survey, years 2000-2011. The primary advantage of this data source is the large sample size, a data characteristic required in order to relax the functional form assumptions of past studies. The study is restricted to individuals aged 25 to 65 who have completed at least nine years of education, were working at the time of the survey, and were not enrolled in school. The sample also excludes employees who have worked fewer than 40 weeks over the past year, who report working fewer than 30 hours in a typical workweek, who work for nonprofit organizations, or who are identified as self-employed.[5] Also excluded from the dataset are 10,333 observations who are not matched with either male or female counterparts. Importantly, the study is limited to "white" individuals. By focusing on a single racial demographic, I eliminate the possibility that the results are biased in either direction by unobservable characteristics which are heterogeneous across racial backgrounds.[6]

The study focuses on three human capital variables: education, experience, and occupation. All variables used in the paper are defined explicitly in the *variable creation* section of Appendix B, where the precise coding for the education and occupation variables can be viewed. Each individual's experience is calculated as the product of potential workforce experience and a work intensity variable. Potential workforce experience is equal to an observation's age minus 18 minus a constant which varies depending on educational attainment.[7] Work intensity is calculated from a CPS

[5] Also, the sample excludes military personnel and individuals living in institutions or other group quarters. Self-employed workers are also omitted from the sample. Observations reporting wage income less than $5.75 per hour worked are excluded, as these must be mistakenly coded in light of other restrictions.

[6] See Appendix B for a step-by-step list of casewise exclusions.

[7] Potential experience is calculated separately for each observation. For high school dropouts, potential experience = age-18. For high school graduates, potential experience = age-20. For those with some college, potential experience = age-21. For workers with a Bachelor's degree, potential

sample separately by education and birth cohort as the average number of total hours worked per year divided by 2000.[8]

The intent of creating experience variables in such a way is to account for differences in experience accumulation due to differential labor supply decisions between men and women and among birth cohorts. Males typically work more hours in the paid workforce per year than females and would thus gain experience at a faster rate; the work intensity variable is intended to account for these difference. The result is that although women in the sample are on average older than men (44.7 years versus 42.7), calculated average workforce experience is very similar (24.7 years for women versus 24.5 years for men). Although I am unable to perfectly control for workforce experience using a large, cross-sectional dataset, I believe experience, as calculated here, effectively controls for differences in human capital stock arising from differential labor supply patterns.[9]

Earned income is the primary variable of interest in the paper. Table 1 shows that average yearly earnings for males are over $23,000 greater than that of females, and are also more dispersed than female earnings. On average, however, females in the sample work fewer weeks per year and fewer hours per workweek, thus comparison of yearly earnings is likely to be confounded by differences in the labor supply decision between men and women. To counter this effect, I use hourly wages as my measure of earned income. Hourly wages are computed using work information reported to the census bureau.[10] Even after accounting for differences in usual time worked, a large

experience = age-22. For those with an advanced or professional degree, potential experience = age-25.

[8]See Appendix B for exact details and justification of created variables.

[9]Treatment effects were also calculated using the "potential experience" variable as the individual's measure of workforce experience. When work intensity is included in the experience variable creation, the estimated wage effects of gender were slightly closer to those measured in previous literature, although the change was not large in significance.

[10]More precisely, the ACS provides researchers with total earned income, a continuous variable of usual hours worked in a week, and a categorical variable indicating the number of weeks worked

Table 1: Descriptive statistics of the workforce by sex, 2000-2011

Variable	*Male*	*Female*
Mean yearly earnings (2011 $)	71,289	47,936
Mean hourly wage (2011 $)	30.89	22.69
Std. dev. of hourly wage (2011 $)	26.34	16.72
Public	16.05	25.20
Education		
HS Dropout	6.23	3.43
HS Graduate	28.85	27.62
Some college	28.04	33.95
Bachelor's degree	24.17	21.90
Advanced degree	12.70	13.10
Experience		
0-10 years	13.12	13.64
11-15 years	11.59	10.19
15-19 years	12.26	11.37
20-24 years	13.21	12.97
25-29 years	14.39	14.61
30-34 years	14.62	16.24
35-39 years	11.82	12.57
40+ years	9.00	8.42
Observations	2,179,783	2,004,101

Table 1: Dollar figures are converted to 2011 dollars using the CPI. Public employees include local, state, and federal workers.

earnings gap exists between the two groups; males earn an estimated hourly wage of $30.89 while females earn an average of only $22.69 per hour of work, an average wage premium equal to 26.5 percent of the average male wage.

The demographic composition of the sample, as shown in table 1, does not easily lend itself to a story of persistent negative female earnings gaps, thus a simple examination of mean earnings does little to identify the source of differential wages between men and women. In particular, it has been shown in the literature that females in

in the past year. Number of weeks worked is imputed using the midpoint of each category as such: $Wks.Worked = 43.5 * (40 \leq weeks \leq 47) + 48.5 * (48 \leq weeks \leq 49) + 51 * (50 \leq weeks \leq 52)$. Hourly wages are then calculated by dividing total wage income by the imputed total number of hours worked in the last year. See Appendix B for detailed summaries of variable construction.

different skill groups will not face the same gender wage gap, suggesting heterogeneity in the effects of gender.[11] This heterogeneity is documented in table 2, which displays the mean female wage gap for each education/experience group in the sample. The data support the findings of past research measuring smaller gender wage gaps for public employees than private sector workers.[12] Among the youngest cohort of government workers with greater than a bachelor's degree, the mean hourly wage gap is approximately 2.4 percent. The most highly educated and highly experienced female skill group faces an estimated mean wage penalty of nearly 50 percent in the private sector.

Similar heterogeneity is observed in appendix tables A.1 and A.2 where I examine mean gender wage gaps by occupational category. On average, women receive a wage penalty of approximately 28.7 percent in the private sector and 20.7 percent in the public sector. Wage gaps in the private labor market range from 1.8 percent for workers in resource extraction occupations to 51.1 percent in health treatment categories. In every private sector occupation category, mean wages for women are below average male wages. In the public sector, females in technical operation jobs receive a wage penalty of around 44.7 percent, while those in farming and similar occupations actually receive a wage *premium* of about 1 percent.

Results in tables 2, A.1 and A.2 document a noticeable lack of homogeneity in the effects of gender on earnings when controlling for various measures of human capital. Among education groupings, the earnings gap is lowest among young workers, and generally rises with potential workforce experience. Within experience cohorts, gender wage gaps seem to rise with educational attainment for private sector workers and fall with educational attainment in the public sector. Both patterns suggest

[11]See Lazear and Rosen (1990) and Baker et al. (1995).

[12]See Wood et al. (1993) and Barón and Cobb-Clark (2010) for examples.

Table 2: Female log wage premiums, by education and experience

Experience	*HS Dropout*	*HS Grad.*	*Some College*	*Bach. Deg.*	*Adv. Deg.*
		A) Private sector only			
0-10 yrs	-0.201	-0.177	-0.140	-0.113	-0.133
11-15 yrs	-0.225	-0.214	-0.191	-0.151	-0.210
15-19 yrs	-0.250	-0.244	-0.231	-0.228	-0.271
20-24 yrs	-0.298	-0.277	-0.254	-0.288	-0.331
25-29 yrs	-0.331	-0.301	-0.282	-0.334	-0.382
30-34 yrs	-0.339	-0.317	-0.294	-0.386	-0.431
35-39 yrs	-0.332	-0.323	-0.287	-0.392	-0.483
40+ yrs	-0.327	-0.296	-0.276	-0.387	-0.496
		B) Public sector only			
0-10 yrs	-0.119	-0.178	-0.229	-0.100	-0.024
11-15 yrs	-0.225	-0.216	-0.267	-0.141	-0.092
15-19 yrs	-0.248	-0.274	-0.312	-0.200	-0.138
20-24 yrs	-0.309	-0.282	-0.344	-0.256	-0.165
25-29 yrs	-0.314	-0.286	-0.332	-0.268	-0.188
30-34 yrs	-0.256	-0.270	-0.303	-0.258	-0.191
35-39 yrs	-0.283	-0.257	-0.282	-0.248	-0.197
40+ yrs	-0.269	-0.218	-0.228	-0.223	-0.244

Table 2: Coefficient estimates are the result of regressions of log wages on a female dummy variable for each education/experience/sector category. All coefficients are significant at the 1% level.

likely heterogeneity in the effects of gender on individual wages. Further, the variety of coefficient magnitudes signal that the wage effects of being female are more complicated than the simple linear shift in the earnings profile suggested by previous research. These facts suggest that differential estimation using a simple dummy variable for gender will lead to biased results and improper inference. The method proposed in this paper allows for heterogeneity in gender's wage effects and will thus not introduce bias into wage gap estimates.

4 Methodology and results

While summary statistics have provided some evidence for the existence of gender earnings gaps, a simple comparison of mean earnings leaves many interesting questions unanswered. Adopting the methodology proposed by DiNardo et al. (1996) and detailed by Fortin et al. (2011), the remainder of this paper is dedicated to investigating the effects of human capital and differential pay practices on the distribution of female earnings. The proposed model requires only minimal statistical assumptions, offering a specification error-free alternative to past empirical models.

The attractive properties of this approach are twofold. First, by eliminating parametric assumptions, I eliminate any potential bias due to specification error in the wage equation. Although past studies generally include robustness checks, most economists will likely agree that the true relationship between wages and productive characteristics is unknown. Although the rationale for adopting the linear functional form is not given in most studies, I expect it is likely used as a "best guess" approximation due to its desirable small-sample properties and the ability to compare results with those of previous studies.

Secondly, the approach allows study of the distributional effects of human capital and conditional wages, rather than the simple mean effects obtained by most previous research.[13] In fact, the proposed model can be adapted to obtain measures of the male-female wage gap at any point along the distribution of earnings, allowing past results to be compared with estimates of the mean effect of the nonparametric model. The model's pliability allows study of the effects of gender on earnings along the entire wage distribution, and offers an important methodological update of the existing literature.

[13] See Baker et al. (1995), Barón and Cobb-Clark (2010), and Boudarbat and Connolly (2013) for noteworthy exceptions.

4.1 Kernel density estimation

In order to qualitatively assess the distributional effects of gender-specific wage rates on the distribution of earnings, I must first establish a reliable method to estimate wage densities. Allow $\hat{g}(w_{0s})$ to represent the empirical density estimate for sex s at hourly wage w_0. The earnings density is estimated by the kernel function

$$\hat{g}(w_{0s}) = \frac{\theta_i}{h} \sum_{i=1}^{n} K\left(\frac{W_{is} - w_{0s}}{h}\right), \tag{1}$$

where h is the bandwidth and g is the univariate density of wages. The θ_i above are observation-specific weights, which depend on w_{0s} and the type of kernel function used in estimation.

An important consideration in kernel density estimation is the choice of bandwidth. Although the dataset is not lacking in observations, a proper bandwidth is needed to maintain balance between bias and variance in the density estimates. Large bandwidths result in kernel estimates with low variance and high bias; selection of a small bandwidth will reduce the bias at the expense of higher variance. To select the bandwidth for each density estimate, this paper utilizes the plug-in technique proposed by Sheather and Jones (1991). Also relevant is the choice of kernel function. The Epanechnikov kernel is used here, although tests of alternative kernel functions suggest this issue is of little empirical importance in this instance.

Kernel density estimates of male and female earnings in both the public and private sectors are displayed in figure 1. Vertical lines represent mean hourly log earnings for each sex. Male wage densities lie distinctly above the female densities in both sectors. There appears to be greater mass on the lower end of the female wage density for private workers, resulting in the noticeably large mean gap in earnings in this sector. A two-sample Kolmogorov-Smirnov test for equality of distribution gives an estimated p-value of 0.000 in both sectors, suggesting that male and female hourly

Figure 1: Male and Female Earnings Densities, by Sector

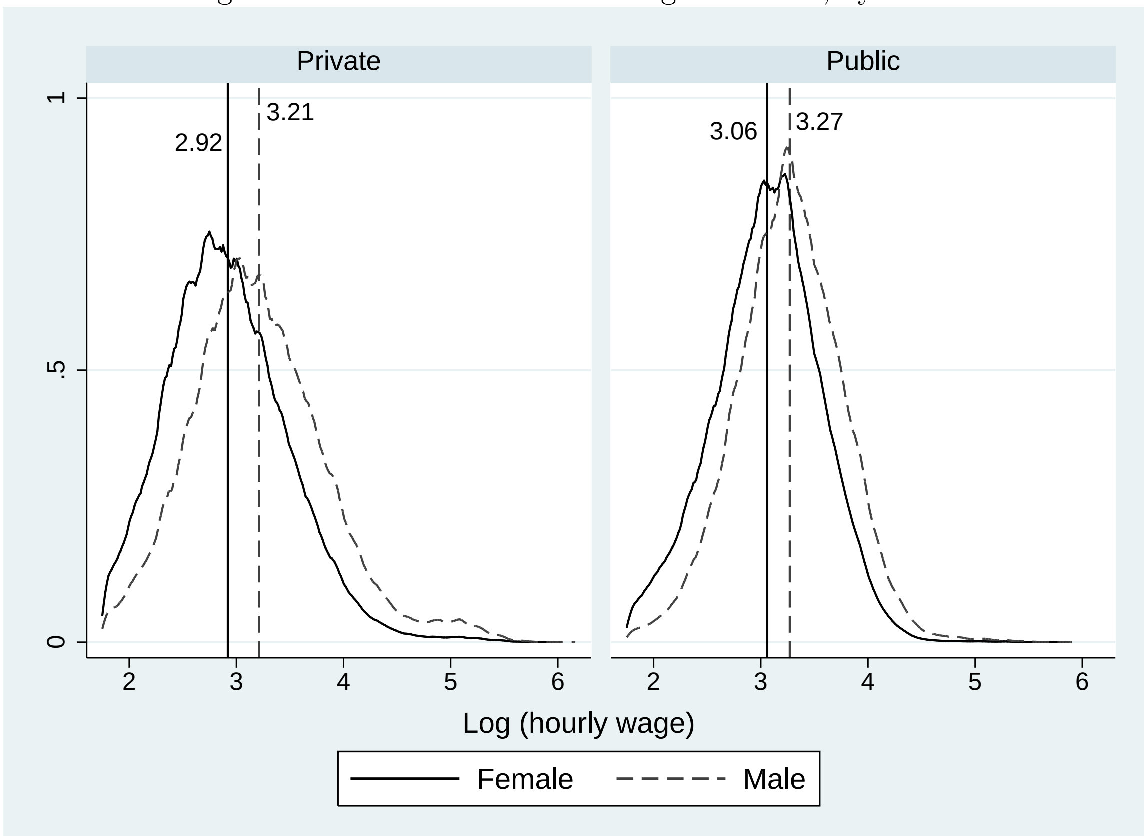

Figure 1: Total hours worked are calculated as the imputed value of weeks worked last year times usual number of hours worked per week. Hourly wages are then equal to total wage income divided by total hours worked.

wages would be unlikely to be drawn from the same distribution if wage rates were indeed random variables.[14] Similarly, a two-sample t-test rejects the null hypothesis that mean wages for men and women are equal.[15]

[14]Figure A.1 shows the male and female wage densities when public and private sector workers are combined. The qualitative results are similar, and a Kolmogorow-Smirnov test on these estimates again produces a p-value of 0.000.

[15]Tests were conducted for public employees, private employees, and a pooled sample of both groups. In each case, H_0 is rejected at the 1% significance level.

4.2 Empirical approach

"How would the earnings of males and females compare if both sexes were truly compensated equally?" This general question is the heart of most studies of the male-female wage differential, and is central in the current paper. In a general sense, I wish to compare the labor market outcomes of men and women. Specifically, I will estimate the density of earnings that *would* result if females were paid as if they were males, as well as the *average* effect being female on one's earnings after controlling for productive characteristics.

Denote the expected earnings of females and males respectively as $E_f(w)$ and $E_m(w)$. In a simple context, mean, or expected earnings of females, is

$$E_f(w) = \sum_{j=1}^{J} \pi_{j,female} w_{j,female}, \tag{2}$$

where $w_{j,female}$ is the mean earnings of women in skill group j and $\pi_{j,female}$ is the fraction of women in skill group j. Clearly, expected earnings for women depend not only on conditional earnings w but also on the number of women who have chosen to be a part of skill group j. It is therefore possible that the raw earnings gaps presented in the previous section are a function of π rather than differences in conditional earnings between men and women.

To investigate this possibility, a set of counterfactual earnings for women must be estimated. Obtaining this counterfactual is rather straightforward. Denote this estimate as

$$E_f^m(w) = \sum_{j=1}^{J} \pi_{j,female} w_{j,male}, \tag{3}$$

where subscripts denote the population from which characteristic weights π are taken and superscripts denote the population from which wages are taken. Thus (3) is an estimate of the mean earnings of females if they were paid male wages.[16]

While the reweighting method described above might be used to analyze differences in earnings between the *average* female and the *average* male, it is possible to generalize it to study the effects of gender on wages in the population as a whole. Define the distribution of earnings for females to be

$$f_f(w) = \int h_f(w|X)g(X|f)dX, \tag{4}$$

where $h_f(w|X)$ is the female conditional wage function and $g(X|f)$ denotes the distribution of productive characteristics X for females. To assess the impact of wage rates on the observed distribution of earnings, we can imagine replacing female wages with male wages in the above equation, thus yielding

$$f_f^m(w) = \int h_m(w|X)g(X|f)dX, \tag{5}$$

which is simply (4) with the conditional wage function for males substituted in for that of females.

Equation (5) may be estimated indirectly by reweighting (4) such that the population of interest is given male characteristics *and* female wages. (5) can thus be rewritten as

$$f_f^m(w) = \int h_m(w|X)g(X|m)\theta_f(X)dX, \tag{6}$$

which is the male wage distribution reweighted such that the sample has the characteristics of females. In the above equation, weight function

$$\theta_f(X) = \frac{g(X|f)}{g(X|m)}, \tag{7}$$

thus (6) simplifies to (5).

[16]Alternatively, one might estimate $E_m^f(w) = \sum_{j=1}^{J} \pi_{j,male} w_{j,female}$, which is simply average male earnings if men were paid according to the female wage scale.

In this paper, reweighting function $\theta_f(X)$ is estimated by taking the proportion of women with exact characteristics X and dividing by the proportion of men with the same characteristics.[17] For example, table A.1 shows that in the private sector, women are less likely than men to be employed as managers. Because the fraction of men who are managers is greater than the fraction of female managers, $\theta_f(X) < 1$, thus each male manager in the estimate of (6) will receive less weight than in the actual male wage distribution if occupation was the only characteristic variable in the vector X. Likewise, females are more likely than males to be employed in a health diagnosis occupation, thus each male will receive a weight greater than 1 when estimating counterfactual (6). Throughout the remainder of the paper, counterfactual estimates are obtained using the nonparametric method outlined above. Counterfactual wage densities obtained by (6) are interpreted as the female wage density that would be observed if females were paid male wages, assuming the female characteristic distribution is unaffected by changes in the wage structure.

Note that this reweighting process is similar in spirit to that of estimating treatment effects in the program evaluation literature. By comparing the mean wage of the counterfactual wage density described above with the mean wage of the observed female wage density, I am able to obtain an estimate of the average treatment effect on the treated. Further, by comparing the nonparametric ATT estimates with linear Blinder-Oaxaca estimates of the ATT, I am able to test for potential bias caused by specification error of the linear wage equations.

[17] θ may also be estimated using any standard probability model. Rosenbaum (1987) notes that estimating propensity scores by parametric specification results in superior estimates in some cases, as these model-based estimates "compensate to some degree for the difference between the population and actual sample proportions... thereby correcting for both systematic and chance imbalances." My paper abstracts from this issue to focus solely on the benefits full nonparametric specification as opposed to linear shifts in wage profiles in estimation of gender wage differentials.

4.3 Counterfactual wage densities

I begin the empirical portion of the paper by offering a more qualitative view of the effects of differential wage rates between sexes. Generally speaking, the nonparametric reweighting method outlined above is ideal for differential analysis because it not only allows researchers to identify mean effects, but also reflects changes would occur in the earnings distribution as a whole if male and female conditional wage functions were the same. It should also be noted that although this study focuses on average treatment effects as the quantitative measure of the real wage gap, it is possible to modify this pliable methodology slightly to estimate real wage gaps at any set of points along the male and female wage distributions.[18]

The question at the heart of the discussion of gender wage gaps is the extent to which differential compensation practices are the driving force behind observed wage differentials. Thus I am interested in the portion of the male-female wage gap which is *not* due to differences in observable productive characteristics, or more precisely, the effect of differential conditional wage functions. Note that to properly assess this statistic the two groups must have equal sets of characteristics. Qualitatively, the effect of differential CWFs can be seen by comparing the wage density estimated by (4) with the wage density estimated by (5).

These empirical estimates are relatively easy to obtain. The empirical version of (4) is simply the unweighted kernel density estimate (1). Counterfactual density (5) is also obtained via the kernel density estimator. The empirical version of (5) is estimated as the kernel density estimate of *male* wages, with observations each receiving analytical weight (7). The result is that counterfactual (5) simulates what the *female* wage density would be if they each received the average wages of men in

[18]For a short discussion, see Butcher and DiNardo (2002); see Gregory and Borland (1999) for general information on the reweighting method.

Figure 2: Estimates of female wage densities given male wages, private sector

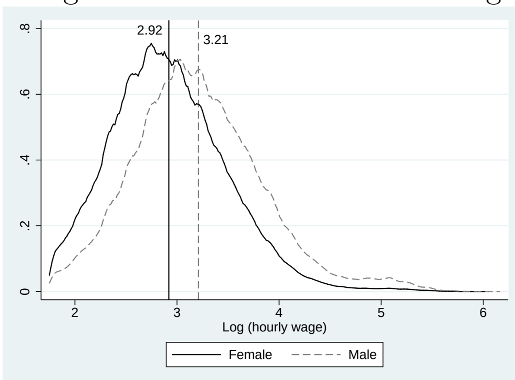

(a) Observed wage densities

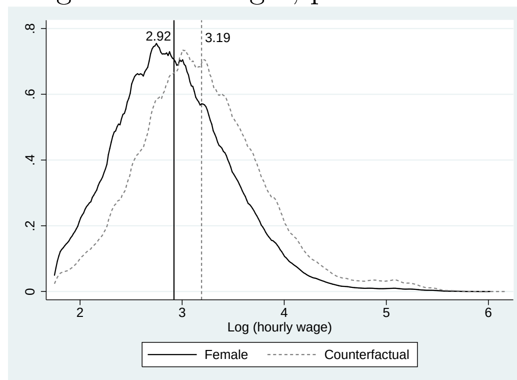

(b) $\theta = f(education)$

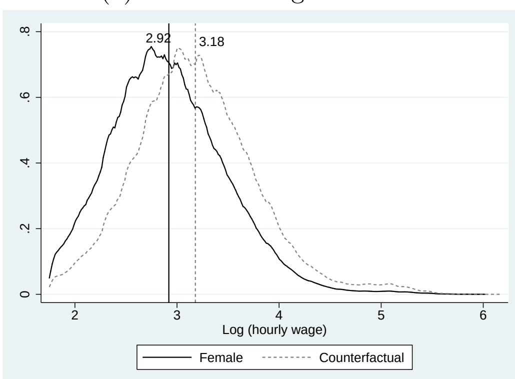

(c) $\theta = f(education, experience)$

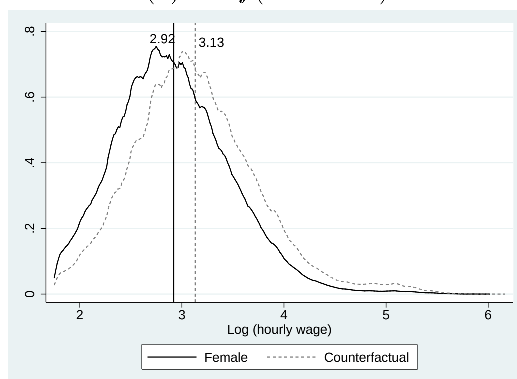

(d) $\theta = f(education, experience, occupation)$

Figure 2: Differences between the female and counterfactual wage densities are interpreted as the effect of differential compensation practices between men and women. Gaps between the male and counterfactual wage densities can be interpreted as a measure of the raw wage gap that can be explained by differences in observable human capital attributes.

their respective skill groups.

Figure 2 demonstrates the effects of differential wage rates on the observed female earnings distribution. Figure 2a shows the observed male and female wage densities for private sector workers; again, vertical lines mark the mean earnings for each population. Other subfigures present counterfactual estimates of what the female earnings density *would* be if women were paid wages equal to that of comparable males in the sample, controlling for different sets of human capital variables. Accounting only for differences in education, the estimated counterfactual appears very similar to that of the observed male wage density. Interpretation of figure 2b is relatively simple. If males and females in the private sector had equal distributions of educational attainment and were paid equal wages, their earnings densities would be almost identical and mean log hourly wages for females would be approximately 2 percent lower than that of males (3.21 vs 3.19 in figures 2a and 2b). The addition of experience to explanatory vector X appears to have a small effect on the estimated counterfactual, suggesting that differential experience has little influence on gender wage gaps.

Of the variables of interest, occupation is possibly the characteristic which is most responsible for the defining differences between male and female wage densities. The addition of occupation to the calculation of $\theta(X)$ shifts greater mass into the lower portion of the counterfactual wage density and lowers estimated mean log hourly earnings to 3.13. This leftward shift of the counterfactual density means that controlling for occupation results in the portion of the raw wage gap explained by personal characteristics to be smaller, suggesting that females are more likely to be employed in generally lower-paying occupations.

Figure 3 displays a similar pattern of male-female wage differences occurring in the public sector. Accounting only for differences in education, the estimated mean wage gap falls to only one percent under the counterfactual assumption that women are paid like men. Figure 3c suggests that if education and experience were the only factors which affected wage rates, the female earnings density would lie slightly above that of males and average female wages would be one percent higher than average male wages. As with the private sector estimates, the inclusion of occupation in the counterfactual calculation lowers the mean wage and shifts the estimated wage density to the left, again promoting the notion that females in the public sector tend to occupy comparatively lower-paying positions than men with otherwise similar human capital.

Figure 3: Estimates of female wage densities given male wages, public sector

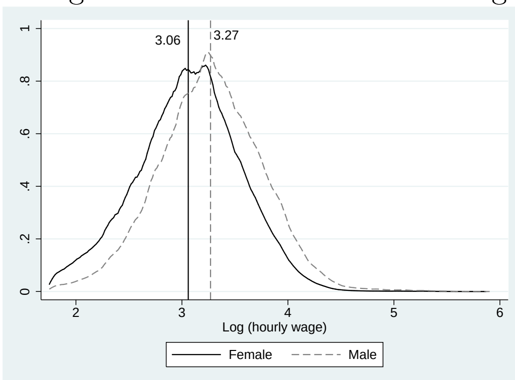

(a) Observed wage densities

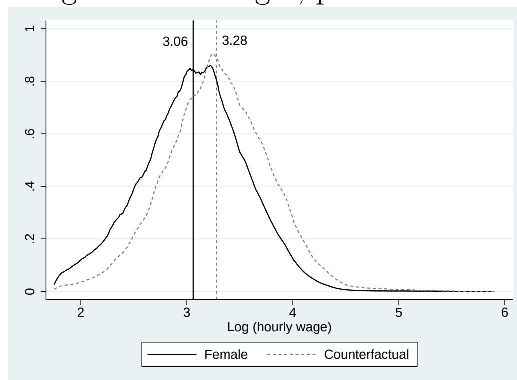

(b) $\theta = f(education)$

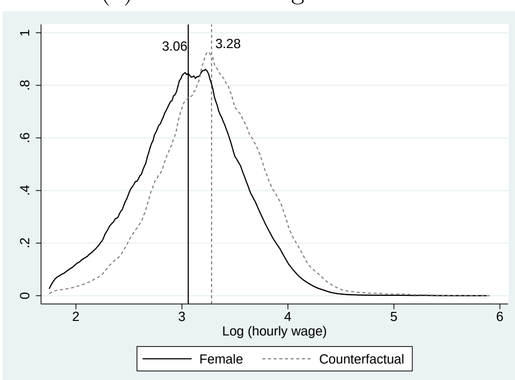

(c) $\theta = f(education, experience)$

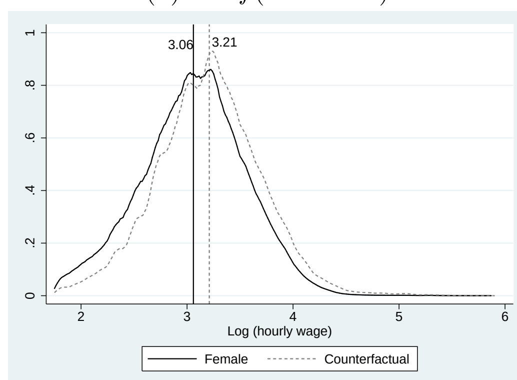

(d) $\theta = f(education, experience, occupation)$

Figure 3: Differences between the female and counterfactual wage densities are interpreted as the effect of differential compensation practices between men and women. Gaps between the male and counterfactual wage densities can be interpreted as a measure of the raw wage gap that can be explained by differences in observable human capital attributes.

The nonparametric reweighting process just demonstrated is elegant in its simplicity. Without even quantitatively measuring treatment effects, it has been shown that when both sexes receive the same conditional wage function, the simulated female earnings density would much more closely mirror the male earnings density. The model also allows me to inspect the potential importance of personal characteristics on the wage gap by including/excluding them from the calculation of individual weights $\theta(X)$. Omitting a variable in the weight estimation process is analogous to excluding that variable from the wage equation in linear Blinder-Oaxaca models.

From this omission process, it appears that occupational preference is the primary human capital determinant of gender wage gaps.

4.4 Estimated treatment effects

Abstracting momentarily from discussion of wage densities as a whole, I shift my attention to measures of the mean effect of gender on earnings. The reweighting approach demonstrated above has many applications in the program evaluation literature, and the following discussion adopts similar terminology. Throughout this section, females will be considered the treatment group and males are the control group. Treatment effects are calculated by measuring the impact on wages while keeping workforce characteristics constant. Let

$$Q(X) = E_f(w|X) - E_m(w|X), \tag{8}$$

thus $Q(X)$ is the wage differential between females and males *conditional* on characteristics X. Also allow $g(X|S)$ to be the distribution of characteristics for sex S.[19] The average treatment effect on the treated is then estimated as

$$\begin{aligned} \tau_{ATT} &= E(w(female)|S=1) - E(w(male)|S=1) \\ &= \int Q(X)g(X|S=1)dX. \end{aligned} \tag{9}$$

The female composition of characteristics is thus used to estimate a weighted mean treatment effect. Interpretation of (8) is as above: If both sexes were given characteristics equal to those of females, τ_{ATT} measures the wage advantage (or penalty) realized by the average member of the female workforce, or the effect of gender on the *average* female's earnings.

This nonparametric static is comparable to the "unexplained" portion of the raw wage gap in the well-known linear version of the Blinder-Oaxaca decomposition in

[19] $S = 1$ denotes the female population while $S = 0$ represents males.

Table 3: Estimates of ATT from linear and nonparametric models

Model	Private		Public	
	τ_{ATT}	$\frac{100 \times \tau_{ATT}}{-0.287}$	τ_{ATT}	$\frac{100 \times \tau_{ATT}}{-0.207}$
	Raw differential = -0.287		*Raw differential = -0.207*	
	Controls for education, experience, occupation			
Linear B-O	-0.230	80.1	-0.154	74.4
	(0.005)	-	(0.005)	-
Nonparametric	-0.214	74.6	-0.147	71.0
	(0.004)	-	(0.004)	-
	Controls for education, experience			
Linear B-O	-0.266	92.7	-0.230	111.1
	(0.006)	-	(0.005)	-
Nonparametric	-0.263	91.6	-0.225	108.7
	(0.006)	-	(0.004)	-
	Controls for education			
Linear B-O	-0.266	92.7	-0.227	109.7
	(0.006)	-	(0.004)	-
Nonparametric	-0.266	92.7	-0.227	109.7
	(0.006)	-	(0.004)	-

Table 3: Percent columns refer to the portion of the raw earnings differential attributed to treatment effects. Linear B-O decomposition models include non-interacted dummies for the indicated control variables. Nonparametric treatment effects are estimated as specified above.

the sense that both measure the portion of the gap attributable to differences in skill prices. In the case where the linear model contains full interaction terms for all explanatory variables, both linear and generalized models will produce fully non-parametric estimates of τ_{ATT}. I test the effects of linear parametric specification of the wage equation by estimating τ_{ATT} using two specifications: one with a simple set of dummy variables for each explanatory variable and one fully nonparametric with a full complement of interaction terms. To calculate standard errors, a simple bootstrap was used. Statistics were calculated in 1000 repetitions of sample size N drawn with replacement, then bootstrapped standard errors were derived from the results.

Estimates of the average treatment effect on the treated are presented in table 3. The nonparametric method estimates the average private sector female receives an hourly log wage penalty of 0.214. This unexplained earnings gap is seen in figure 2d, and simply measures the distance between means of the female and counterfactual wage densities.[20] Using the same sample, the linear B-O decomposition estimates the average female faces a wage gap of 0.230. Women in the public workforce receive a smaller wage penalty than those in the private sector, a finding consistent with the previous literature. Again, the parametrically-specified B-O model estimates a larger mean unexplained wage gap than the nonparametric model.

Removing occupation from the vector of explanatory variables raises private sector estimates of τ_{ATT} to 0.266 and 0.263 for the linear and nonparametric specifications, respectively. Interpretation of the nonparametric result is also relatively simple: if men had the same measurable education and experience as women (but still kept the male occupation distribution), the average female would earn an hourly wage 26.3

[20]Mean log wages in figure 2d are rounded to the nearest hundredth, hence the average treatment effect of 0.214 is rounded to 0.21 in the figure.

percent lower than the average male wage. In the public sector, females have greater measurable human capital than males, yet are still estimated to receive a 22.5 percent hourly wage penalty. Removing experience from the set of controlled variables results in little change in estimated treatment effects, although it is worth noting that the fewer factors I control for in the model the larger the difference between estimates of the parametric and nonparametric models.

Why exactly do the estimates from these two specifications differ? The answer lies in the B-O linear functional form specification. Barsky et al. (2002) cites estimation of counterfactuals as the key to bias hidden in linear B-O decompositions. While least squares estimation minimizes mean square error for estimated male or female wage regressions the MSE is not necessarily minimized when estimating counterfactual wages for the opposite group. As a result, these models will either under- or overestimate what true wages would be in the counterfactual scenario, attributing either too much or too little of the mean wage gap to differences in conditional wage functions between genders. In contrast, nonparametric methods avoid specification of a functional form, and thus will not introduce potential bias due to errors in functional form. Thus the nonparametric specification, which Barsky et al. (2002) term the *Generalized Blinder-Oaxaca Decomposition*, will produce more unbiased estimates of the effects of gender on earnings and the female wage gap.

4.5 Differential profile estimates

While a comparison of average treatment effects may shed some light on possible bias arising from the functional form assumptions of past Blinder-Oaxaca studies of the gender wage gap, it may not tell the whole story. Indeed in its most simple form (without use of multiple interaction terms between human capital variables), the linear version of B-O assumes that experience, or age, affects wage rates in a way common to all skill groups. Thus regardless of a female's education or occupational choice, her wages relative to that of men are assumed to follow a given path as she ages.[21] In this section I test the validity of this assumption by comparing experience profile estimates from the nonparametric and linear decomposition models, as well as estimates gleaned from a wage equation using a female dummy variable. First, gender wage differentials were estimated using three models: the two decompositions utilized in the previous section and a dummy variable model.[22] Next, I find the mean of the differential estimates for each education/age/occupation/sector group. I am then able to map differential profiles for every age, for each skill group.

I begin by examining wage differential estimates for highly educated workers in health treatment occupations. For this skill group in the private sector, average male wages are over 50 percent higher than average female wages even though women are much more likely to choose this particular occupation (see appendix table A.1). As can be seen in figure 4, women of all ages are estimated to face differential wages from that of comparable males in the private sector. In the public sector, the youngest cohort of women are estimated to receive a wage *premium* of over 20 percent in the nonparametric model and a wage *penalty* of almost 20 percent in the linear decomposition model. Such differences in wage gap estimates must be attributed to the functional form assumptions underlying the linear B-O method.

Also of note are differences in the differential age profiles between the three models. While the dummy method provides an estimate of the average effect of sex on wages,

[21] Even if some interaction terms are used in the estimation, bias may still persist unless all explanatory variables are fully interacted with each other. In this case, estimates of the wage gap will yield identical results as the nonparametric model.

[22] The dummy variable differential is obtained by regressing hourly log wages on a female dummy variable. As previously, the linear B-O specification includes single dummy variables, one for each education, experience, and occupation category, while the nonparametric specification includes full interaction terms in addition to these. High school dropouts with 0-10 years of experience and who work in management are the base group for all decompositions.

Figure 4: Gender wage gaps, health diagnosis and treatment occupations

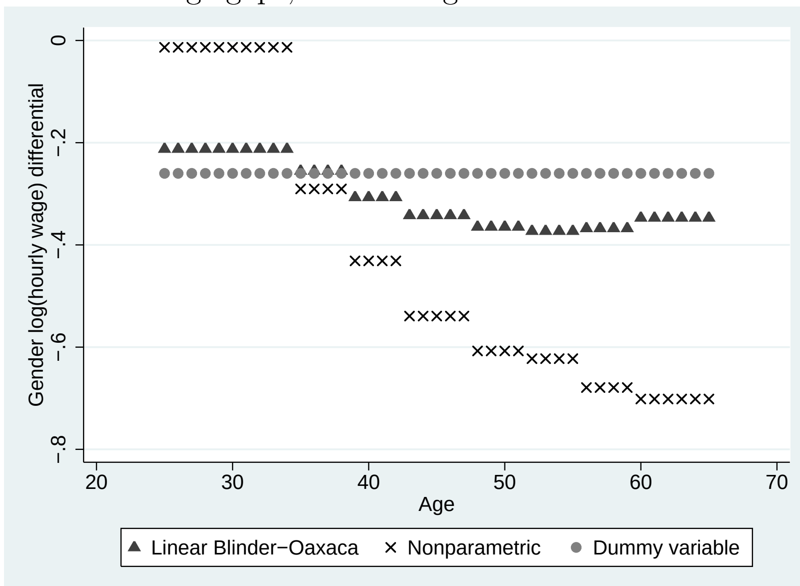

(a) Private

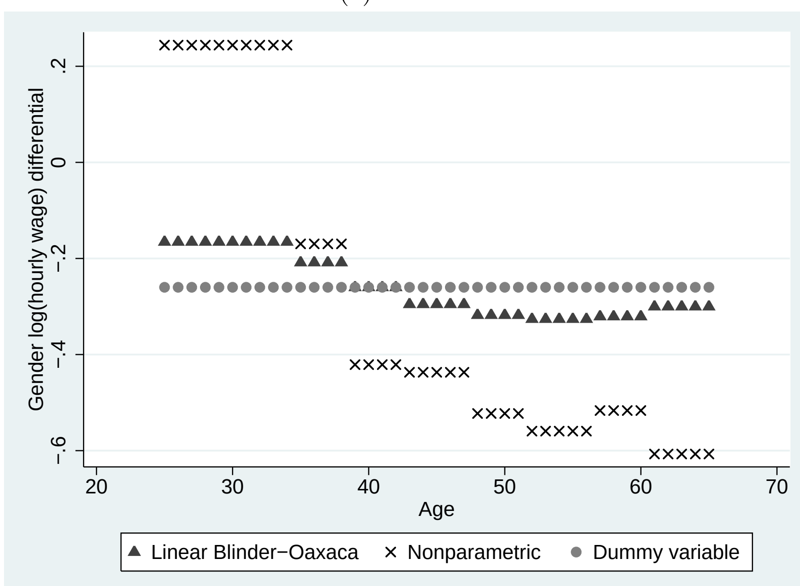

(b) Public

Figure 4: Gender wage gaps are estimated separately by sector. Dummy variable estimates are obtained by regressing log wages on a female dummy variable separately for each sector.

interpretation of this statistic at any single level of experience (or age) is likely to be invalid. For young cohorts, the dummy variable approach tends to overestimate

Figure 5: Gender wage gaps, education occupations

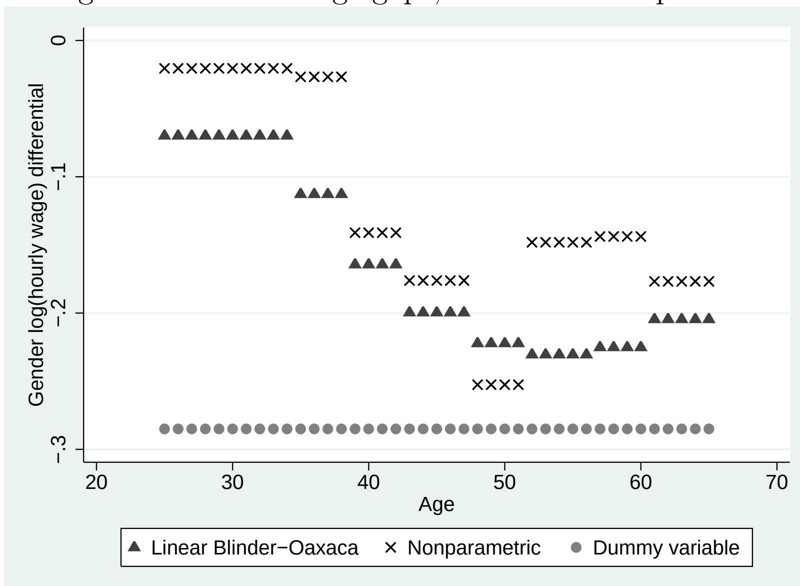

(a) Private

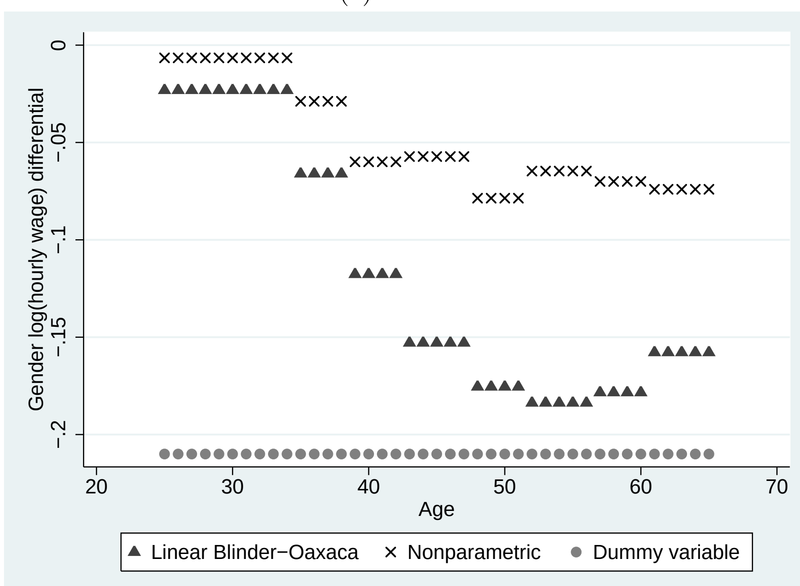

(b) Public

Figure 5: Gender wage gaps are estimated separately by sector. Dummy variable estimates are obtained by regressing log wages on a female dummy variable separately for each sector.

the true wage gap while underestimating the wage gap faced by more experienced women. Likewise, linear B-O estimates confront the same issue when compared to nonparametric estimates.

Next, I examine gender wage differentials for workers in education occupations, an industry in which both men and women are highly represented. Figure 5 compares the gender wage differential age profile for each model. Again, a simple comparison of mean earnings in each sector (i.e. the dummy variable approach) is likely to yield misleading information for this particular skill group. The estimated progression of relative wages also differs between linear B-O and nonparametric models. For public sector educators, the gap between differential estimates from these two specifications is more than 0.1 for ages 43-65. In general, the nonparametric technique yields smaller estimates of the gender wage gap. Because this model assumes nothing about the relationship between wages and skills, it should be viewed as the most accurate as is possible given the quality of the data.

Examination of differential experience profiles for other skill groups reveals similar evidence and generally supports the hypothesis that parametric specification is likely to lead to errors of functional form, although no formal hypothesis test is offered here. For some skill groups, the linear specification does very well at matching the nonparametric differential estimate. For other groups, the two models produce very different estimates.

Two points should be gleaned from this exercise. First, nonparametric estimates of the female wage gap document significant heterogeneity in the effects of sex and age on wage rates. Not only does the simplistic dummy variable approach assume homogeneous effects, but also does a generally poor job of reflecting the actual differential faced by large portions of the female workforce. Secondly, the linear Blinder-Oaxaca decomposition estimated here projects the gender wage differential to "progress" equally throughout a worker's lifetime regardless of other observable characteristics. As the nonparametric estimates show, this is likely a poor assumption. Thus even if the linear B-O model represents a significant upgrade over the dummy variable approach, it will also likely produce biased estimates of the mean wage gap.[23]

5 Conclusion

In terms of empirical techniques used to evaluate the male/female wage differential, there has been little progress made in recent years. The overwhelming majority of recent literature on the subject utilizes various forms of the linear Blinder-Oaxaca decomposition, focusing on the mean effects of gender on wage outcomes. In this aspect, the generalized Blinder-Oaxaca decomposition proposed in this paper offers a significant upgrade of past models because it allows researchers to estimate entire counterfactual wage densities rather than focusing solely on expected wages in the counterfactual setting. The potential benefits to this approach are clear and extend to a wide variety of issues. For example, suppose government officials wished to assess the impact of a policy on the wage rates of public employees. Given an adequate number of observations in both the treatment and control groups, one could easily estimate what the wage density of the affected workforce would be in the absence of the policy, allowing focus on the group of public employees as a whole instead of focusing on the policy's impact on the *average* worker.

Another important update of the current literature is the absence of potential specification error in nonparametric estimates of the gender wage gap. When all covariates are accounted for, the generalized B-O model predicts that being female lowers the wage of average women in the private sector by approximately 21.4 percent; the mean gender effect for public sector women is -14.7 percent. The linear B-O model predicts larger unexplained wage differentials than the generalized decomposition. Because the same sample and covariates are used, we can attribute the difference in

[23]That is unless all explanatory variables are interacted with each other.

estimates to specification error of the wage equation in the linear model. Thus when the precise relationship between wages and personal characteristics is unknown, which is likely the case in *most* studies, nonparametric estimation allows researchers to leave the functional form of this relationship unspecified, eliminating possible biased results due to misspecification.

While the potential benefits of the generalized approach are great, it would be remiss to leave the costs unaddressed. The first issue is that of result interpretation. The aforementioned nonparametric results may be interpreted as an estimate of wage discrimination *if and only if* all wage-determining variables are observed and controlled for by the researcher. If this is not the case, we can attribute the measured wage gaps to everything other than those factors directly observed. Difficulty in interpretation is also confounded by the second issue with nonparametric estimation: data sufficiency. More precise interpretation typically requires researchers to control for a greater number of personal characteristics. However, the addition of explanatory variables creates problems of common support among the two sample populations, and requires many more observations. Thus there is an inherent tradeoff between precise interpretation of results and increased data requirements. By choosing to focus on only three explanatory variables, I hope to eliminate common support issues as much as possible while being cautious when interpreting results. In the current study, I believe the benefits to the nonparametric approach outweigh the costs, although researchers with limited datasets should consider these carefully before attempting nonparametric estimation.

Results of the nonparametric analysis imply that although differences in education, experience, and occupation play a role in wage outcome differences between men and women, other factors are surely at play. In particular, there is evidence to suggest that differential conditional wage functions account for at least part of the noticeable difference between wage densities of males and females.[24] Both qualitative and quantitative results propose that of the human capital variables examined in the paper, occupation plays the largest role in observed wage gaps. This result supports previous research which suggests females generally receive lower pay because they tend to choose relatively low-paying occupations.

The purpose of this paper has been to investigate differences in male and female wage rates and their effect on observed earnings densities using a generalized Blinder-Oaxaca decomposition. Clearly, estimation shows that given male wage rates, the wage density of females would shift significantly to the right, and much more closely resemble that of males. The model's pliability leaves many avenues of this issue open for discovery, and suggests that improved methods of differential estimation, free of specification error, are available to researchers with large datasets.

[24] I make this claim under the assumption that unobservable personal characteristics do not account for all of the measured treatment effects.

References

Baker, Michael; Benjamin, Dwayne; Desaulniers, Andree, and Grant, Mary. The distribution of the male/female earnings differential, 1970-1990. *Canadian Journal of Economics*, 1995.

Barón, Juan and Cobb-Clark, Deborah. Occupational segregation and the gender wage gap in private- and public-sector employment: A distributional analysis. 2010.

Barsky, Robert; Bound, John; Charles, Kerwin, and Luptin, Joseph. Accounting for the black-white wealth gap: A nonparametric approach. *Journal of the American Statistical Association*, 2002.

Bertrand, Marianne; Goldin, Claudia, and Katz, Lawrence. Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, 2010.

Blau, Francine and Kahn, Lawrence. Swimming upstream: Trends in the gender wage differential in the 1980's. *Journal of Labor Economics*, 1997.

Blinder, Alan. Wage discrimination: reduced form and structural estimates. *Journal of Human Resources*, 1973.

Boudarbat, Brahim and Connolly, Marie. The gender wage gap among recent post-secondary graduates in canada: a distributional approach. *Canadian Journal of Economics*, 2013.

Butcher, Kristen and DiNardo, John. The immigrant and native-born wage distributions: Evidence from united states censuses. *Industrial and Labor Relations Review*, 2002.

DiNardo, John; Fortin, Nicole, and Lemieux, Thomas. Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. *Econometrica*, 1996.

Filer, Randall. Male-female wage differences: The importance of compensating differentials. *Industrial and Labor Relations Review*, 1985.

Firpo, Sergio; Fortin, Nicole, and Lemieux, Thomas. Unconditional quantile regressions. *Econometrica*, 2009.

Fortin, Nicole. The gender wage gap among young adults in the united states: The importance of money versus people. *Journal of Human Resources*, 2008.

Fortin, Nicole; Lemieux, Thomas, and Firpo, Sergio. Decomposition methods in economics. *Handbook of Labor Economics*, 2011.

Galego, Aurora and Pereira, Joao. Evidence on gender wage discrimination in portugal: Parametric and semi-parametric approaches. *The Review of Income and Wealth*, 2010.

Gregory, Robert and Borland, Jeff. Recent developments in public sector labor markets. *Handbook of Labor Economics*, 1999.

Gronau, Reuben. Sex-related wage differentials and women's interrupted labor careers-the chicken or the egg. *Journal of Labor Economics*, 1988.

Groshen, Erica. The structure of the female/male wage differential: Is it who you are, what you do, or where you work? *Journal of Human Resources*, 1991.

Hellerstein, Judith; Neumark, David, and Troske, Kenneth. Market forces and sex discrimination. *Journal of Human Resources*, 2002.

Lazear, Edward and Rosen, Sherwin. Male-female wage differentials in job ladders. *Journal of Labor Economics*, 1990.

Oaxaca, Ronald. Male-female wage differentials in urban labour markets. *International Economic Review*, 1973.

Rapaport, Carol. Apparent wage discrimination when wages are determined by nondiscriminatory contracts. *American Economic Review*, 1995.

Rosenbaum, Paul. Model-based direct adjustment. *Journal of the American Statistical Association*, 1987.

Sheather, S.J. and Jones, M.C. A reliable data-based bandwidth selection method for kernel density estimation. *Journal of the Royal Statictical Association B*, 1991.

Wood, Robert; Corcoran, Mary, and Courant, Paul. Pay differences among the highly paid: The male-female earnings gap in lawyers' salaries. *Journal of Labor Economics*, 1993.

Appendix A

Table A.1: Mean female wage gaps in the private sector, by occupation

Occupation	Pct. males	Pct. females	$\bar{w}_f^{occ} - \bar{w}_m^{occ}$
Managers	15.07	11.50	-0.348
Accountants and auditors	1.68	3.15	-0.378
Analysts, HR, and labor relations	2.08	2.97	-0.393
Inspectors and management support	1.10	1.46	-0.232
Engineers and architects	3.96	0.66	-0.162
Mathematical and computer scientists	2.47	1.44	-0.133
Natural scientists	0.55	0.37	-0.171
Health diagnosis and treatment	1.47	6.10	-0.511
Teachers (primary and secondary)	0.44	1.83	-0.351
College and university instructors	0.20	0.31	-0.203
Librarians, archivist and curators	0.02	0.05	-0.084
Social scientists and planners	0.24	0.32	-0.275
Lawyers and judges	1.03	0.64	-0.219
Artists and entertainers	1.71	1.71	-0.194
Health technicians and specialists	0.42	2.88	-0.145
Science technicians	1.44	0.56	-0.181
Software developers	2.13	0.73	-0.137
Other operators and technicians	0.61	1.29	-0.204
Service sales representatives	1.70	2.08	-0.483
Commodities sales representatives	7.49	6.98	-0.476
Office supervisors	1.09	3.18	-0.226
Receptionists and typists	0.31	3.01	-0.257
Secretaries	0.17	7.68	-0.158
Records and financial clerks	0.50	5.65	-0.181
Communication, mail, and message dist.	3.12	2.93	-0.201
Adjusters and investigators	1.48	5.16	-0.235
Miscellaneous administrative support	0.52	3.85	-0.314
Police, fire, and protective services	0.63	0.30	-0.080
Food preparation and services	2.07	3.61	-0.193
Health service occupations	0.25	4.10	-0.148
Building and personal services	2.10	2.54	-0.232
Farming, forestry, and fishing occupations	1.60	0.37	-0.131
Automobile and electric mechanics	8.12	0.44	-0.052
Construction workers	7.81	0.23	-0.177
Extractive occupations	0.06	0.00	-0.018

Table A.1: Mean female wage gaps in the private sector, by occupation

Occupation	Pct. males	Pct. females	$\bar{w}_f^{occ} - \bar{w}_m^{occ}$
Precision production workers	4.67	2.25	-0.364
Machine operators	4.61	2.71	-0.316
Welders and hand-assembly occupations	3.21	2.11	-0.270
Transportation and mat. moving occs.	7.76	0.77	-0.215
Helpers and other laborers n.e.c.	4.14	2.07	-0.214
Total	100.00	100.00	-0.287

Table 4: The two middle columns calculate the percentage of private sector employees for each sex and occupation. The rightmost column displays the difference between average female log earnings and average male log earnings for each occupational grouping.

Table A.2: Mean female wage gaps in the public sector, by occupation

Occupation	Pct. males	Pct. females	$\bar{w}_f^{occ} - \bar{w}_m^{occ}$
Managers	11.09	9.11	-0.198
Accountants and auditors	1.36	2.03	-0.263
Analysts, HR, and labor relations	0.86	1.80	-0.196
Inspectors and management support	1.94	1.33	-0.108
Engineers and architects	2.80	0.36	-0.087
Mathematical and computer scientists	2.41	1.40	-0.076
Natural scientists	2.04	0.76	-0.105
Health diagnosis and treatment	1.45	4.64	-0.336
Teachers (primary and secondary)	12.55	31.72	-0.153
College and university instructors	4.28	2.86	-0.212
Librarians, archivist and curators	0.38	1.28	-0.131
Social scientists and planners	0.78	0.67	-0.142
Lawyers and judges	2.04	1.25	-0.139
Artists and entertainers	0.72	0.58	-0.085
Health technicians and specialists	1.04	1.30	-0.148
Science technicians	1.31	0.45	-0.143
Software developers	0.83	0.44	-0.062
Other operators and technicians	0.53	0.79	-0.447
Service sales representatives	0.27	0.26	-0.216
Commodities sales representatives	0.28	0.42	-0.297
Office supervisors	1.71	2.16	-0.262
Receptionists and typists	0.20	1.37	-0.206
Secretaries	0.21	8.60	-0.196
Records and financial clerks	0.51	3.60	-0.202
Communication, mail, and message dist.	4.56	3.65	-0.117
Adjusters and investigators	0.53	1.48	-0.203
Miscellaneous administrative support	0.88	4.27	-0.248
Police, fire, and protective services	18.64	2.98	-0.153
Food preparation and services	0.27	1.57	-0.379
Health service occupations	0.44	1.72	-0.199
Building and personal services	4.94	2.62	-0.170
Farming, forestry, and fishing occupations	1.56	0.22	0.010
Automobile and electric mechanics	3.96	0.20	-0.050
Construction workers	3.17	0.08	-0.041
Extractive occupations	0.01	0.00	-0.283

Table A.2: Mean female wage gaps in the public sector, by occupation

Occupation	Pct. males	Pct. females	$\bar{w}_f^{occ} - \bar{w}_m^{occ}$
Precision production workers	1.85	0.31	-0.261
Machine operators	0.51	0.13	-0.293
Welders and hand-assembly occupations	0.51	0.14	-0.153
Transportation and mat. moving occs.	3.76	1.09	-0.286
Helpers and other laborers n.e.c.	2.82	0.35	-0.329
Total	100.00	100.00	-0.207

Table 5: The two middle columns calculate the percentage of private sector employees for each sex and occupation. The rightmost column displays the difference between average female log earnings and average male log earnings for each occupational grouping.

Figure A.1: Male and Female Earnings Densities, Both Sectors

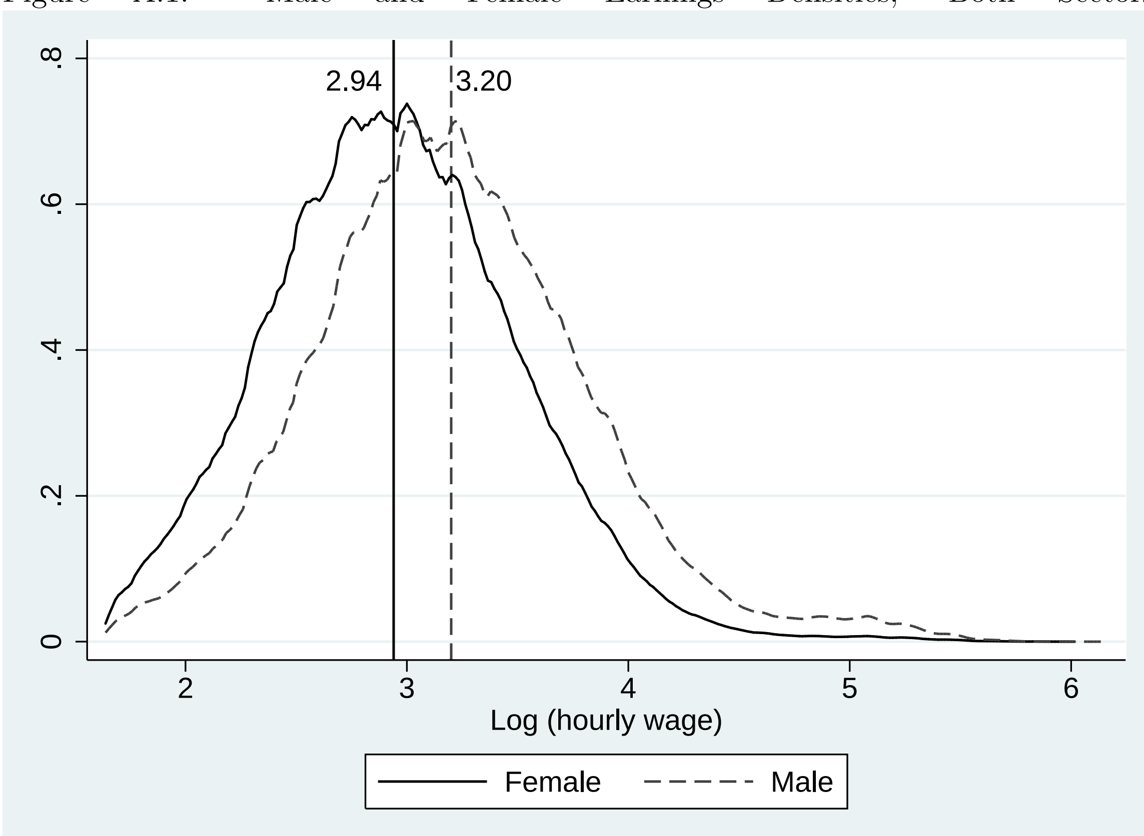

Appendix B

The purpose of this appendix is to provide a clear explanation of the empirical process, from sample creation to differential estimation. The following pages give precise instructions to obtaining my estimates, including variable categorization, sample exclusions, and coding (if deemed necessary). Anyone attempting to recreate my study may contact me at emakela@g.clemson.edu if they run into insurmountable issues.

Creating the sample

I begin with the full sample from the American Community Survey (henceforth ACS) from the years 2000-2011. Data was downloaded from IPUMS USA thanks to ipu. Variable codes referenced in this appendix are those from IPUMS, and in some instances may differ from the coding structure in the raw ACS data files.

After importing the sample, I first remove all observations reported as living in group "institutions" (332558 observations deleted) or "other group quarters" (236691 observations deleted). Then, individuals over the age of 65 (3612641 observations deleted) and under the age of 25 (8058616 observations deleted) are removed. To remove any possible simultaneity in the effects of work and schooling on work productivity, I remove all observations currently attending school (762624 observations deleted). Those who report as having less than 9 years of education are also dropped from the sample (549852 observations deleted).

Since the study focuses on the work income, those observations who were not employed and currently working are removed from the sample, along with those employed by the armed forces (3563403 observations deleted). I then remove workers who report as self-employed (1062678 observations deleted), unpaid family workers (17634 observations deleted), and those who work for non-profit organizations

(724216 observations deleted). To conform with previous studies, I focus solely on full-time, full-year (FTFY) workers, thus I then remove individuals who work fewer than 40 weeks per year (570795 observations deleted) or fewer than 30 hours in a typical workweek (392305 observations deleted). Since veterans typically differ from non-veterans in terms of unmeasured human capital and would thus likely violate the model's statistical assumptions, I remove all veterans from the sample (716000 observations deleted). Also excluded are 3603 remaining observations who report as being employed outside the United States.

I then wish to remove those individuals who report wage income below the legal threshold, as these observations are either miscoded or misreported in light of all other restrictions.[25] The following code is used to convert earnings to 2011 dollars, then calculate hourly wages for each observation:

- replace incwage=incwage*cpi11
- gen ww=43.5*(wkswork2==4)+48.5*(wkswork2==5)+51*(wkswork2==6)
- gen totalhours=uhrswork*ww
- gen wage=incwage/totalhours

All observations reporting hourly wage income less than $5.75 (the lowest real minimum wage in the sample period) are then removed from the sample (118583 observations deleted).

In the spirit of abiding by the unconfoundedness assumption required to interpret my results, the next sample restriction seemed prudent. It is my opinion that there are likely to be systematic differences in unmeasured human capital stocks among different racial groups. To eliminate the possibility of this affecting my results, I construct my sample using only white observations, dropping all individuals who

[25] In real terms, the lowest minimum wage in the sample time frame was in 2006 when the real minimum wage (in 2011 dollars) was $5.75.

report their race as non-white (956787 observations deleted). For similar reasons, I also make a few sample restrictions based on occupation. These omissions are listed here:

- Social, recreation, and religious workers: drop if occ1990>173 & occ1990<177 (27378 observations deleted)
- Supervisors and proprietors of sales jobs: drop if occ1990==243 (170155 observations deleted)
- Farm operators and managers: drop if occ1990>472 & occ1990<477 (4875 observations deleted)

The final sample restrictions are made to adhere to the model's common support requirement. Because the wage effects of gender are clearly not equal in the public and private sectors, I conduct the analysis separately for each, meaning that I require common support for both sectors rather than the sample as a whole. For each sector, I calculate the ratio of workers in each education/experience/occupation group.[26] If there exists a male or female skill group in either sector with no observations in the comparable skill group of the opposite sex, all individuals in this skill group are dropped from the sample. The (unsimplified) code for this restriction is as follows:

- bys sex public: egen count_ss=count(logwage)
- bys sex public ed exp occ: egen count_sseeo=count(logwage)
- gen w_eeo=count_sseeo/count_ss
- gen weeo_mg=w_eeo if male==1 & public==1
- gen weeo_fg=w_eeo if male==0 & public==1
- gen weeo_mp=w_eeo if male==1 & public==0
- gen weeo_fp=w_eeo if male==0 & public==0
- bys ed exp occ: egen mg_eeo=mean(weeo_mg)
- bys ed exp occ: egen fg_eeo=mean(weeo_fg)
- bys ed exp occ: egen mp_eeo=mean(weeo_mp)

[26] I will henceforth refer to each education/experience/occupation grouping a "skill group".

- bys ed exp occ: egen fp_eeo=mean(weeo_fp)
- gen malew_eeo_g=mg_eeo/fg_eeo
- gen femw_eeo_g=fg_eeo/mg_eeo
- gen malew_eeo_p=mp_eeo/fp_eeo
- gen femw_eeo_p=fp_eeo/mp_eeo
- drop if malew_eeo_g==. (10006 observations deleted)
- drop if femw_eeo_g==.
- drop if malew_eeo_p==. (327 observations deleted)
- drop if femw_eeo_p==.

As is shown above, 10,006 government workers are dropped from the sample for common support purposes, along with 327 private sector employees.

In sum, I begin with the entire ACS sample from the years 2000-2011, which includes a total of 26,075,611 observations. After all excluded observations are removed, the final sample consists of 3,328,908 private sector observations (1,829,855 male and 1,499,053 female) and 854,976 public sector observations (349,928 male and 505,048 female).

Variable creation

The paper focuses on three primary explanatory variables. Education is clearly defined in the ACS. School dropouts are defined as all observations reporting fewer than 12 years of education, or as having attended grade 12 but not having attained a diploma. High school graduates include those individuals with a GED or alternative credentials. The "some college" category includes all workers who report attending at some college, but have not attained a bachelor's degree.[27] The advanced degree educational category includes all observations who report as having attended college for 5 years or more, or have been granted a master's, doctoral, or professional degree.

[27] This category includes workers who have been granted a 2 year associates degree of any kind.

Experience categories account not only for potential workforce experience, but also for estimated work intensity. Potential experience is calculated as:

- gen potexperience=age-18
- replace potexperience=age-20 if ed==2
- replace potexperience=age-21 if ed==3
- replace potexperience=age-22 if ed==4
- replace potexperience=age-25 if ed==5

where ed=1,2,3,4,5 corresponds to education categories HS dropout, HS graduate, some college, bachelor's degree, and advanced degree, respectively. Because workers of different skill groups might accrue human capital at varying rates, I use March CPS samples (years 2000-2013) to calculate the average total hours worked for workers based on their birth cohort, education, and sex:

- gen birthyr=year-age
- drop if birthyr>1986
- gen cohort=1*(birthyr<1940)+2*(birthyr>1939 & birthyr<1945)
  +3*(birthyr>1944 & birthyr<1950)+4*(birthyr>1949 & birthyr<1955)
  +5*(birthyr>1954 & birthyr<1960)+6*(birthyr>1959 & birthyr<1965)
  +7*(birthyr>1964 & birthyr<1970)+8*(birthyr>1969 & birthyr<1975)
  +9*(birthyr>1974 & birthyr<1980)+10*(birthyr>1979)
- bys cohort ed sex: egen wksw=mean(wkswork1)
- bys cohort ed sex: egen hrsw=mean(uhrswork)
- gen totalhrs=wksw*hrsw
- replace totalhrs=totalhrs/2000

To construct experience categories, I merge the compressed CPS dataset with the ACS microdata by cohort, education, and sex. Experience is then calculated as the product of potential experience and the work intensity variable totalhrs. This variable is then converted into a categorical variable based partially on the even distribution of observations in each category:

- gen experience=potexperience*totalhrs
- gen exp=1*(experience<10)+2*(experience>10 & experience<15)
  +3*(experience>15 & experience<20)+4*(experience>20 & experience<25)
  +5*(experience>25 & experience<30)+6*(experience>30 & experience<35)
  +7*(experience>35 & experience<40)+8*(experience>40)

Worker occupation was classified as closely as possible following the 1990 census occupation classification system. Using occupational categories which are too wide and inclusive would risk invalidating differential estimates because this would increase the risk of systematic differences in job characteristics between men and women. However, using occupation categories which are too narrowly defined would result in greater difficulty achieving common support and would likely mean the elimination of many more observations from the sample. The following table contains the precise occupational codes used to create the occupation variable in the paper. Excluded occupational categories are listed above.

Estimation

Once the sample is created, estimation is relatively straightforward. The code sample referenced above can be used to calculate the reweighting variables θ in the paper's empirical model.[28] More precisely, to estimate the counterfactual wage density in figure 2d, one needs only to estimate the wage density for males in the private sector with observations reweighted by the function femw_eeo_p (seen just above). Estimation of counterfactual wage densities 2b and 2c is done in a similar fashion, simply excluding experience and occupation variables in the weight calculations performed in the previous code sample.

The process is almost the same when calculating nonparametric average treat-

[28]When all explanatory variables are accounted for, femw_eeo_g is equivalent to $\theta_f(X)$ in (7) if we wish to measure the counterfactual wage density for females in government work. Likewise $\theta_f(X)$ is equal to femw_eeo_p if we wish to analyze gender wage differences in the private sector.

Categorization of occupational codes	
Occupation	Code
Managers	occ1990<23
Accountants and auditors	occ1990¿22 & occ1990<25
Analysts, HR, and labor relations	occ1990>24 & occ1990<28
Inspectors and management support	occ1990>27 & occ1990<40
Engineers and architects	occ1990>42 & occ1990<60
Mathematical and computer scientists	occ1990>60 & occ1990<69
Natural scientists	occ1990>68 & occ1990<84
Health diagnosis and treatment	occ1990>83 & occ1990<110
Teachers (primary and secondary)	(occ1990>153 & occ1990<164) *(ind1990!=850)
College and university instructors	(occ1990>153 & occ1990<164) *(ind1990==850)
Librarians, archivist and curators	occ1990>163 & occ1990<166
Social scientists and planners	occ1990>165 & occ1990<174
Lawyers and judges	occ1990>177 & occ1990<180
Artists and entertainers	occ1990>180 & occ1990<201
Health technicians and specialists	occ1990>201 & occ1990<210
Science technicians	occ1990>210 & occ1990<226
Software developers	occ1990==229
Other operators and technicians	occ1990>225 & occ1990<240 & occ1990!=229
Service sales representatives	occ1990>243 & occ1990<257
Commodities sales representatives	occ1990>257 & occ1990<300
Office supervisors and equip. operators	occ1990>302 & occ1990<309
Receptionists and typists	occ1990>313 & occ1990<325
Secretaries	occ1990==313
Records and financial clerks	occ1990>325 & occ1990<345
Communication, mail, and message dist.	occ1990>344 & occ1990<374
Adjusters and investigators	occ1990>374 & occ1990<379
Miscellaneous administrative support	occ1990>378 & occ1990<390
Police, fire, and protective services	occ1990>410 & occ1990<428
Food preparation and services	occ1990>428 & occ1990<445
Health service occupations	occ1990>444 & occ1990<448
Building and personal services	occ1990>447 & occ1990<470
Farming, forestry, and fishing occupations	occ1990>478 & occ1990<500
Automobile and electric mechanics	occ1990>500 & occ1990<550
Construction workers	occ1990>550 & occ1990<600
Extractive occupations	occ1990>600 & occ1990<620
Precision production workers	occ1990>620 & occ1990<700

Machine operators	occ1990>700 & occ1990<780
Welders and hand-assembly occupations	occ1990>780 & occ1990<800
Transportation and mat. moving occs.	occ1990>800 & occ1990<860
Helpers and other laborers n.e.c.	occ1990>860 & occ1990<900 + occ1990==405

ment effects. When evaluating τ_{ATT}, I am simply recording the difference between the average female wage and the average male wage, where the male population is reweighted such that it has the characteristics of the female population with which it is being compared. For example, to calculate τ_{ATT} for private sector workers, one would use the following code:

- sum logwage if male==0 & public==0
- global i=r(mean)
- sum logwage if male==1 & public==0 [aw=femw_eeo_p]
- global j=r(mean)
- global ATT=$i-$j
- di $ATT

Linear treatment effect estimates are obtained using the *oaxaca* package, where τ_{ATT} is equivalent to the "unexplained" portion of the measured mean wage gap. Linear wage equations are estimated using only dummy variables for education, experience, and occupation, and do not include interaction terms.

The information provided in this appendix should be adequate to properly reconstruct the dataset and produce the empirical results in the paper. If anyone attempting this task is having difficulty, please feel free to contact me.