

THE GENDER WAGE GAP: EXTENT, TRENDS, AND SOURCES*

Francine D. Blau

Cornell University, NBER, IZA, DIW and CESifo

Lawrence M. Kahn

Cornell University, IZA, CESifo, and NCER (Australia)

February 2015

*We are indebted to Jason Cook and Ankita Patnaik for excellent research assistance.

1. Introduction

The gender wage gap has now been intensively investigated for a number of decades, but also remains an area of active and innovative research. In this article, we provide new empirical estimates delineating the extent of and trends in the gender wage gap and their sources. We then survey the literature to identify what has been learned about the sources of the gap, both those that can be readily included in conventional analyses and those which cannot. Our primary focus will be on the United States, although we also place the United States in a comparative perspective, particularly as such comparisons help to further our understanding of the sources of the gender wage gap. The focus on the United States is in part designed to make our task more manageable, as there has been an explosion of research on this topic across many countries. Nonetheless, we believe much of what we have learned for the United States is applicable to other countries, particularly other economically advanced nations. In our comprehensive review of the literature, we particularly emphasize areas where there has been exciting new research, including research on gender differences in psychological attributes/noncognitive skills and mathematics test scores, and on the reversal of the gender education gap. In some areas, we will rely to some extent on recent reviews to establish baseline findings so that we may focus our own discussion in areas where we can add value by reviewing more recent work or focusing on areas where there is some controversy.

The long-term trend has been a substantial reduction in the gender wage gap, both in the United States and in other economically advanced nations (Blau and Kahn 2008). However, the shorter term picture in the United States has been more mixed. The period of strongest wage convergence between men and women was the 1980s, and progress has been slower and more uneven since then. Moreover, a number of other related trends appear to have plateaued or

slowed since the 1990s, including increases in female labor force participation rates and reductions in occupational segregation by sex.

The plan of the paper is as follows. In Section 2, we begin by documenting the changes in the gender gap that have occurred in the United States since the 1950s based on published data. We then provide new analyses for the 1980 to 2010 period that include decompositions of the changes in the gender wage gap into portions associated with key characteristics such as schooling, experience, industry, occupation and union status. We also examine how women fared relative to men at various points in the wage distribution. Our decompositions show the importance of these measured factors in accounting for the levels and changes in the gender pay gap. We also find that an unexplained gap remains and, moreover, that it has been stable subsequent to a dramatic narrowing over the 1980s.

In the remaining sections we probe what is known about the various factors that contribute to the gap, including the extent of and trends in these factors. Some of the variables we consider are measured in our data set and included in our analysis in Section 2, as well as other similar type analyses. Other factors are not included and presumably help to provide insight into the sources of the unexplained gap. However, it is important to point out that the effects of factors that are not explicitly included in traditional regression analyses may actually be taken into account to some extent by measured variables. For example, women have been found to be more risk averse than men on average which could lower their relative wages. However, to the extent that this factor operates through gender differences in occupational sorting, e.g., if it results in women avoiding occupations with greater variance in earnings, regression analyses that control for occupation will adjust for this factor.

Our consideration of explanatory factors begins with a discussion of supply side factors in Sections 3-6, focusing on labor force participation, education and mathematics test scores, labor force experience, training and turnover, gender roles and identity, and psychological attributes/noncognitive skills. We follow this, in Section 8, with a consideration of gender differences in occupations, industries, and firms, while in Sections 9 and 10, we discuss discrimination and public policy, including the role of the government's antidiscrimination effort and family policy. In Section 11 we turn to a consideration of the role of wage setting institutions and overall demand factors. Finally, Section 12 presents conclusions.

2. Overview of the US Gender Wage Gap

In this section, we use published data, information from the Michigan Panel Study of Income Dynamics (PSID), and the March Current Population Survey (CPS) to establish the facts on the levels and trends in the US gender wage gap and on their sources (in a descriptive sense). Accounting for the sources of the level and changes in the gender pay gap will provide guidance for understanding recent research studying gender and the labor market.

Figure 1 shows the long-run trends in the gender pay gap over the 1955-2014 period based on two published series: usual weekly earnings of full-time workers and annual earnings of full-time, year-round workers.¹ After many years with a stable female/male earnings ratio of roughly 60%, women's relative wages began to rise sharply in the 1980s, with a continued, slower and more uneven rate of increase thereafter. By 2013, women earned about 78% of what men did on an annual basis and about 82% on a weekly basis.

¹ The data cover all workers 16 years of age and older from 1979 onward, and those at least 14 years old for years preceding 1979. For additional information on sources see Blau, Ferber, and Winkler (2014), p. 148.

To better understand the sources of the gender wage gap, we use data from the PSID, which is the only data source that has information on actual labor market experience (a crucial variable in gender analyses) for the full age range of the population. We focus on men and women age 25-64 who were full-time, non-farm, wage and salary workers and who worked at least 26 weeks during the preceding year. The focus on full-time workers and those with substantial labor force attachment over the year is designed to identify female and male workers with fairly similar levels of labor market commitment. The sample is also restricted to family heads and spouses/cohabitators because the PSID only supplies the crucial work history information for these individuals. Due to this and other limitations in coverage by the PSID, described in the Data Appendix, we present some additional data on the gender pay gap using the fully nationally representative March CPS.² The empirical results in this section are of interest in and of themselves and also serve to set the stage for the literature review to follow by providing a frame of reference for how each of the measured factors discussed relates to the overall gap. Our data cover the 1980-2010 period, in which, as Figure 1 shows, women have made major gains in relative wages.

Table 1 shows the evolution of the female-male ratio of average hourly earnings at the mean and also at the 10th, 50th, and 90th percentiles over the 1980-2010 period.³ The Table shows the gender wage ratio for four years—1980, 1989, 1998, and 2010—based on both PSID and CPS data. Because in both data sets earnings refer to the previous year, we use, for example, the 1981 wave to measure wages in 1980. The overall pattern is very similar across these samples, and it also largely matches that in the published data shown in Figure 1, increasing

² Additional information on the details of our data preparation and analysis is available in the online Data Appendix.

³ Entries are calculated as $\exp(D)$, where D is the female log wage at the mean, or at the indicated percentile, minus the corresponding male log wage.

one's confidence in the PSID.⁴ Specifically, the mean female/male wage ratio was 62-64% in 1980, rising to 72-74% by 1989, with a further increase to 79-82% by 2010. This pattern, along with that in Figure 1, suggests that women made their largest gains in relative wages during 1980s, a result we have studied in some detail in prior work (Blau and Kahn 2006), where we explicitly compared the 1980s and the 1990s.

Table 1 also shows gender wage ratios at the bottom (10th percentile), middle (50th percentile) and top (90th percentile) of the wage distribution. Within each percentile, the time pattern is similar to that for the overall mean: the gender wage ratio rose over the period, with the largest gain during the 1980s. However, a closer examination shows that women gained least, in a relative sense, at the top. In both the PSID and CPS in 1980, women at the top had a slightly higher pay ratio than those in the middle and a slightly lower pay ratio than those at the bottom. Yet by 2010, in both data sets, women's relative pay at the top was less than that at the middle and bottom of the distribution, often considerably so. Specifically, in 2010, women's relative pay at the top was 8-9 percentage points less than that at the middle or bottom in the PSID, while it lagged the middle and bottom by 6-11 percentage points in the CPS. Later in this section, we will consider the role of measured factors in accounting for the slower reduction at the top and in following sections we will attempt to shed additional light by reviewing the literature on the labor market for highly skilled workers.

At the same time that the gender pay gap has been narrowing, women have been increasing their relative labor market qualifications and commitment to work. Tables 2 and 3

⁴ The unemployment rate was 7.1% in 1980, 5.3% in 1989, 4.5% in 1998, and 9.6% in 2010. The high level of unemployment in 2010 may raise concerns about the representativeness of that year for studying the gender pay gap. Reassuringly, however, we found similar results when we ended our PSID sample in 2006, before the Great Recession began; moreover, as noted, the overall pattern of the gender pay gap shown in Table 1 matches that based on annual data shown in Figure 1.

show the extent of such changes among our PSID sample of full-time workers. Table 2 focuses on the prime human capital determinants of men's and women's wages, education and actual full-time experience. In the case of education, there was a dramatic reversal of the gender gap. In 1981, women had lower average levels of schooling than men and were less likely to have exactly a bachelor's or advanced degree. Over the period, women narrowed the education gap with men and, by 2011, women had higher average levels of schooling and were more likely to have an advanced degree than men.⁵ While men had a slightly higher incidence of having exactly a bachelor's degree, women were more likely to have at least a bachelor's degree (i.e. the sum of the Bachelor's Degree Only and Advanced Degree categories).⁶

In the case of labor market experience, the story is one of substantial narrowing of the gender experience gap. In 1981, men had nearly 7 more years of full-time labor market experience on average than women. By 2011, however, the gap had fallen markedly to only 1.4 years, with the fastest rate of increase in women's relative experience occurring during the 1980s. Thus, on these two basic measures of human capital—schooling and actual labor market experience—women made important gains during the 1981-2011 period, reversing the education gap and greatly reducing the experience gap.

Table 3 further explores trends in the determinants of wages by showing gender differences in the incidence of high-level jobs as well as collective bargaining coverage. Rising employment in managerial or professional jobs may be an indicator of increasing human capital or work commitment, even controlling for levels of schooling and actual labor market

⁵ Tables 2 and 3 refer to 1981, 1990, 1999, and 2011 rather than 1980, etc., as shown in Table 1, because earnings refer to the previous year, while other variables are measured as of the survey date.

⁶ CPS data also show that, in 1981, men had higher levels of schooling and incidence of bachelor's or advanced degrees than women; by 2011, women in the CPS had higher levels of schooling than men, as in the PSID. Moreover, in the CPS, in 2011, women not only had a higher incidence of advanced degrees than, but also a slightly higher incidence of exactly a college degree than men.

experience. For example, such jobs may entail higher levels of responsibility and pressure than other jobs, and only those with the appropriate training and commitment may be qualified to take them. Increases in women's relative representation in such jobs may then be a further indicator of their rising human capital and labor market commitment. However, women's representation in such jobs may also be affected by employer discrimination in entry or promotions. Women's improvements may therefore also reflect reductions in discrimination. Both interpretations are plausible. First, it seems likely that women's increasing levels of schooling and, as discussed below, increasing representation in lucrative fields of study, as well as their rising experience levels would be expected to lead to their greater representation in high-level positions. Second, given women's increasing qualifications and commitment to the labor market, employer incentives for statistical discrimination (this concept is discussed further below) have likely been reduced.

Under either interpretation, studying these differences can yield insights into the sources of changes in the gender pay gap. Table 3 shows remarkable increases in women's relative representation in such high-level jobs. The male advantage in managerial jobs fell from 12 percentage points in 1981 to just two percentage points in 2011. Moreover, while women were more likely to work in professional jobs throughout the period than men were, their advantage grew from five percentage points in 1981 to nine percentage points in 2011. However, many women in professional jobs remain employed in traditionally female occupations such as nursing or K-12 teaching that are generally less lucrative than traditionally male professions. We therefore also show in Table 3 gender differences in the incidence of employment in "male" professional jobs, which we define as professional jobs other than nursing or K-12 and other non-college teaching positions, most of which were predominantly male at the start of our

period. While men were four percentage points more likely than women to be in such jobs in 1981, by 2011, the gender gap had been virtually eliminated. At the same time women were making these occupational gains, they were greatly reducing their concentration in administrative support and clerical jobs.⁷

In addition to these occupational changes, one notable feature of the post-1980 labor market is the steady reduction in the portion of the economy covered by collective bargaining. Table 3 shows that this reduction hit men much harder than women. Specifically, men's collective coverage fell from 34% in 1981 to 17% in 2011, while women's coverage only declined from 21% to 19%.⁸ As is the case with women's gains in education, full time labor market experience, and employment in high-level occupations, we expect the elimination of the gender gap in collective bargaining coverage to contribute to a reduction in the gender pay gap.

How have these improvements in women's relative labor market qualifications and employment location affected the gender pay gap? We study this issue by decomposing changes in the gender pay gap over the 1980-2010 period using log wage regressions. We proceed in two stages. First, we estimate wage models that only control for education, experience, race/ethnicity, region, and metropolitan area residence. We term this the "human capital specification," since other than basic controls, we include only human capital variables—education and experience. Second, we augment this model with a series of industry, occupation and union coverage dummy variables. We term this equation the "full specification." Because

⁷ We obtained very similar results using the March CPS for studying the gender gaps in managerial, professional, and "male" professional employment.

⁸ While the PSID data show women as now having slightly higher collective bargaining coverage than men, US Bureau of Labor Statistics data, although also showing a narrowing in the gender gap, show men continuing to retain a small edge. Specifically, in 1983, among those 16 years and older, 27.7% of men were covered by collective bargaining, compared to 18.0% of women; by 2011, men's coverage had decreased to 13.5%, while women's declined to 12.5% (<http://data.bls.gov/pdq/SurveyOutputServlet> , accessed August 18, 2014).

these latter variables may have an ambiguous interpretation—i.e., they may represent human capital, other labor market skills, and commitment, on the one hand, or employer discrimination, on the other hand—we present both versions. Note that we do not control for marital status or number of children, since these are likely to be endogenous with respect to women’s labor force decisions. Our decompositions can be viewed as reduced forms with respect to family formation decisions.

We measure education by controlling for years of schooling, plus dummy variables for having exactly a bachelor’s degree and an advanced degree. We include measures of both full-time and part-time labor market experience and their squares. Race and ethnicity are controlled for using four mutually-exclusive categories: white non-Hispanic (the excluded category), black non-Hispanic, other non-Hispanic, and Hispanic. We control for three of the four Census regions as well as including a dummy variable for residence in a metropolitan area. In the full specification, we additionally control for a series of 15 industry (including government employment) and 20 occupation dummy variables and a collective bargaining coverage dummy variable. The construction of these categories took account of changes in the PSID’s coding scheme over the period and is described in the online Data Appendix.

2.1 Explaining the Gender Wage Gap at the Mean

Figure 2 shows female to male log wage ratio, unadjusted for covariates (i.e. reproduced from Table 1), adjusted for the covariates in the human capital specification, and adjusted for the covariates in the full specification. The adjusted ratios are defined as $\exp(\text{res})$, where res is the mean female residual from the male log wage equation.⁹ While the adjusted differential (or

⁹ Results were broadly similar when we computed the male residuals from female equations, although the unexplained residual was somewhat higher in absolute value for men. We focus on the female residuals from the male equations since it seems most relevant to ask how women would fare if they were treated like men in the labor market. The US labor force aged over the 1980-2010 period, and it is well known that gender pay gaps increase

unexplained gap) is sometimes taken as an estimate of discrimination, it may be a biased one since it may include the effects of omitted factors. We consider this issue in greater detail Section 9, while our discussion of research on selection, unmeasured attributes such as competitiveness or risk aversion, and possible glass ceilings will shed light on some possible sources of the pay gap that cannot be explained by measured characteristics.

The results for the unadjusted ratios in Figure 2 mirror the trends from the published data, showing a large increase in the female-to-male wage ratio over the 1980s, with continued but smaller gains in subsequent decades. Over the 1980-2010 period as a whole, the unadjusted ratio increased substantially from 62.1 to 79.3 percent. The adjusted ratios also rose considerably over this period, from 71.1 to 82.1 percent in the human capital specification and from 79.4 to 91.6 percent in the full specification. However all of these gains occurred in the 1980s. This means that, while a reduction in the residual or unexplained gap played an important role in the narrowing of the gender wage gap over the 1980s, it has not been a factor since then (see, also Blau and Kahn 2006). Figure 2 also indicates that the difference between the human-capital adjusted ratio and the unadjusted ratio fell dramatically over the 1980-2010 period, reflecting women's increasing human capital levels relative to men's. By 2010, the human capital variables (and the other variables included in this specification) explained very little of the gender wage gap: the unadjusted ratio was 79% compared to the adjusted ratio of 82%. As Goldin (2014) has commented, "As women have increased their productivity enhancing characteristics and as they 'look' more like men, the human capital part of the wage difference

with age. To investigate whether aging has influenced our picture of the trends in the gender wage gap, we re-weighted our data using 1980 age weights using a quartic in age in a procedure based on DiNardo, Fortin and Lemieux (1996). Men and women in our wage samples were in fact 3-4 years older in 2010 than 1980. When we repeated our analyses using 1980 age weights, we found that the overall female to male wage ratio would have been 80.7% in 2010, compared to its actual value of 79.3% as shown in Table 1 and Figure 2, a slight increase as expected. However, the adjusted ratios were very similar to those shown in Figure 2.

has been squeezed out.” As the results in Table 2 show, this represents to some extent countervailing factors: women are now *better* educated than men but they continue to lag (slightly) in actual labor market experience. In the full specification, the adjusted ratio (91.6 percent) remained considerably higher than in the human capital specification (82.1 percent), suggesting a continued substantial role for occupation and industry in explaining the gender wage gap (recall that union differences have now been virtually eliminated).

Table 4 provides further detail on the contribution of particular labor market characteristics to the reduction in the gender wage gap. Specifically, it shows the fraction of the total gender wage gap in 1980 and 2010 accounted for by gender differences in each group of variables for both the human capital and full specifications. The entries are the male-female differences in each variable multiplied by the corresponding male coefficients from the current year wage regression. In Panel A, one sees the contribution of traditional human capital variables—education and experience—not controlling for industry, occupation or union status. This specification in effect allows human capital to affect these intervening variables and thus gives the reduced form effect of education and experience in explaining the gender wage gap. In 1980, the male advantage in education raised the gender wage gap somewhat (0.0219), while the male experience gap contributed substantially (0.114 log points) to the gap. By 2010, due to the education reversal, women’s *higher* level of education slightly *raised* their relative wage (0.019 log points), while the much smaller (compared to 1980) male advantage in labor market experience accounted for only a small portion (0.037 log points) of the gender wage gap. Together, human capital factors (education and experience) accounted for 27% of the gender wage gap in 1980 compared to only 8% in 2010. Moreover, the decline in the unexplained gap—from 0.341 log points in 1980 to .197 log points in 2010 also contributed substantially to

the narrowing of the gender gap over the period, although, as we have seen, the decrease in the unexplained gap occurred only during the 1980s. Unexplained factors accounted for a larger *share* of the gender gap in the human capital specification in 2010 (85%) than in 1980 (71%).

Table 4, Panel B, shows the decomposition of the gender pay gap using the full specification. Interestingly, the effects of education and experience are quite similar to that in Panel A, implying that the impact of these measures of human capital operates primarily within industries, occupations and union coverage status. In 1980, gender gaps in union coverage, industry, and occupation together accounted for 0.126 log points, or 26% of the gender pay gap. In all, gender differences in covariates are able to account for 52% of the gender pay gap. A major shift over the period is that industry and occupation accounted for a considerably larger share of the reduced 2010 gender wage gap (51 percent vs. 26% in 1980), while the effect of unionization virtually vanished. As in the case of the human capital specification, a marked decline in the unexplained gap (from 0.231 log points in 1980 to 0.088 log points in 2010) contributed to the narrowing of the gender wage gap, and, again, this decrease occurred over the 1980s. In this case, however, unexplained factors account for a smaller *share* of the gender gap in 2010 (38%) than in 1980 (49%). The growing importance of location in the labor market (i.e., industry and occupation) in accounting for the gender gap suggests that future research on explanations might fruitfully focus on gender differences in employment distributions and their causes. This meshes well with increased attention to the role of firms as firm-worker matched data increasingly become available.

One puzzling finding in Table 4 is that, despite the occupational improvements of women shown in Table 3, gender differences in occupation accounted for a larger pay gap than in 2010 than in 1980 (0.076 vs. 0.051 log points). However, while women upgraded their occupations

during this period, the wage consequences of gender differences in occupations became larger as well. When we evaluate the impact of changes in the gender gaps in covariates on the change in the gender wage gap *using a constant set of male wage coefficients* (for 1980 or 2010), women's relative occupational upgrading is now found to lower the gender pay gap by 0.03-0.04 log points (results not shown). This highlights that shifts in labor market prices can affect women's progress in narrowing the gender wage gap. The role of wage structure in affecting changes over time in relative wages of women, as well as differences across countries in the magnitude of the gender wage gap, is considered in Section 10.

2.2 Explaining the Gender Wage Gap Across the Wage Distribution

As we saw in Table 1, as of 2010, (i) there was a relatively large gender gap at the top of the distribution and (ii) the wage gap fell more slowly over the 1980-2010 period at the top than at other portions of the distribution. These two patterns suggest the notion of a “glass ceiling” in which women face barriers in entering the top levels of the labor market and which we discuss in more detail in Section 8. To provide some further evidence on this phenomenon, we decompose the gender pay gap at specific percentiles of the distribution into portions due to covariates and portions due to wage coefficients. The latter component corresponds to the unexplained gap and, while as noted above, it is sometimes taken to be a measure of discrimination, it may be a biased estimate.

To study the unexplained gap across the distribution, we use a method developed by Chernozhukov, Fernández-Val and Melly (2013) which decomposes unconditional intergroup gaps (in our case, male-female gaps) at a given percentile into a portion due to the distribution of personal characteristics and a portion due to different wage functions conditional on personal characteristics. This latter portion corresponds to the unexplained gap. As discussed by the

authors, the method involves computing the distribution of personal characteristics and the conditional wage distribution by gender. For example, let log wages be denoted by Y , y be a specific value of log wages, m represent males, f represent females, and X be a vector of covariates affecting wages. Then,

$$(1) F_{Y[m,m]}(y) = \int F_{Ym|Xm}(y|x) dF_{Xm}(x)$$

$$(2) F_{Y[f,f]}(y) = \int F_{Yf|Xf}(y|x) dF_{Xf}(x)$$

$$(3) F_{Y[m,f]}(y) = \int F_{Ym|Xm}(y|x) dF_{Xf}(x)$$

where $F_{Y[m,m]}$ refers to the unconditional distribution of log wages with the male wage function and the male personal characteristics, with a corresponding definition for $F_{Y[f,f]}$; $F_{Y[m,f]}$ is the hypothetical wage distribution that would face women if they were rewarded according to the male wage function; $F_{Ym|Xm}$ refers to the conditional distribution of male wages given their characteristics; and F_{Xm} refers to the distribution of male personal characteristics, with corresponding definitions for $F_{Yf|Xf}$ and F_{Xf} .

To decompose the differences between the unconditional male and female wage distributions, we note that:

$$(4) F_{Y[m,m]} - F_{Y[f,f]} = \{F_{Y[m,m]} - F_{Y[m,f]}\} + \{F_{Y[m,f]} - F_{Y[f,f]}\}$$

The first term in brackets in equation (4) shows the effect of differing distributions of personal characteristics, while the second term shows the wage function effect.

To implement the decomposition, Chernozhukov, Fernández-Val and Melly (2013) suggest computing the empirical distribution of the X variables and using quantile regressions for the conditional wage distribution. We follow that procedure and estimate 100 quantile regressions. In addition, we compute the standard errors using bootstrapping with 100

repetitions. In Table 5, we present the decomposition results for the 10th, 50th and 90th percentiles.¹⁰ At each percentile, women's covariates improve relative to men's over the period in both the human capital and full specifications. In the human capital specification, women's improved qualifications lead to a 0.09 to 0.10 log point reduction in the gender gap across the distribution. In the full specification, women's improved covariates lead to a 0.11 log point reduction in the gender pay gap at the 10th and the 50th percentiles but only a 0.06 log point improvement at the 90th percentile. This latter finding suggests that among the highly skilled, women's occupations, industries and union status did not improve as much relative to men's at the top as among those with middle or low skill levels.

The coefficient effects shown in Table 5 imply much larger reductions in the unexplained gap at the 10th and 50th percentiles than at the 90th percentile. In the human capital specification, the unexplained gap fell by 0.18 to 0.20 log points at the 10th and 50th percentiles, but only by 0.06 log points at the 90th percentile; in the full specification, the corresponding reductions in the unexplained gap were 0.16 to 0.18 log points at the 10th and 50th percentiles but only 0.06 log points at the 90th percentiles. By 2010, the unexplained gap was larger at the 90th percentile than at the 10th or 50th percentile in both specifications; in contrast, in 1980, the unexplained gap was smaller at the 90th than at the 50th, although still larger than at the 10th percentile.¹¹

¹⁰ The decomposition allows us to recover the unconditional distribution of wages by adding the effects of the covariates and wage coefficients, and the results closely match the actual percentiles. The Chernozhukov, Fernández-Val and Melly (2013) approach is similar in principle to the method of unconditional quantile regressions suggested by Firpo, Fortin and Lemieux (2009).

¹¹ While the unexplained gap in the full specification for 2010 appears very low at the 10th percentile, we are reluctant to place a strong interpretation on this in light of its relatively large standard error. Taking the coefficient at face value suggests a larger role for differences by occupation, industry, and unionism in accounting for gender wage gaps at the bottom than at the other percentiles (especially the 90th) where gender difference in wages within occupation, industry, and union status appear to play a relatively larger role.

These coefficient effects suggest the possibility of a glass ceiling among highly skilled women, although they could also result from unmeasured factors leading highly skilled men to earn particularly high relative wages. We discuss research on discrimination in Section 9. However, we note that such a result—either a relatively large unexplained gender gap at the top or more slowly falling gender pay gaps at the top than elsewhere in the distribution is a common finding in the recent literature on the gender gap that uses quantile regression methods to study these issues. For example, in earlier work (Blau and Kahn 2006), we used PSID data and found that the unexplained gender pay gap in 1998 at the 90th percentile was larger than at lower percentiles and that it had fallen less since 1979. Similarly, using PSID data Kassenboehmer and Simming (2014) found that the unexplained gender gap fell by less at the 90th percentile than at lower regions of the distribution over the 1993-5 to 2004-8 period and that, in 2004-8, there was a somewhat larger unexplained gap at the 90th than at the 50th percentile. Moreover, European research also typically finds a larger unexplained gap at the top than the middle of the distribution. For example, Arulampalam, Booth and Bryan (2007) used European Community Household Panel microdata on 11 countries for 1995-2001 and found such “glass ceiling effects” for private sector workers in ten of 11 countries studied, and Albrecht, Björklund and Vroman (2003) found similar effects using Swedish data for the 1990s. Thus, a relative high or relatively slowly falling gender pay at the top appears to be a common result in research studying the gender gap across the distribution.

2.3 Summary

Our overview of the US gender wage gap shows a substantially decreased but persistent wage gap between men and women. Decompositions of the gender pay gap indicate the importance of education and experience, as well as industry, occupation and union status in

accounting for the levels and changes in the gender pay gap. They also highlight the diminished role of human capital factors in accounting for the gender wage gap over time—due both to the reversal of the education gap between men and women and the narrowing of the gender gap in experience. Gender differences in occupation and industry have remained important in explaining the gender wage gap, despite occupational upgrading of women relative to men. However, the role of unions in accounting for gender differences in wages has virtually disappeared. While a decrease in the unexplained gap played a role in narrowing the gender wage gap in the 1980s, an unexplained gender wage gap remains and has been roughly stable since the 1980s decline. We also found that gender wage gap is currently larger at the top of the wage distribution and has decreased more slowly at the top than at other points in the distribution. This remains the case even after accounting for measured characteristics. We now turn to a discussion of the underlying factors affecting the observed sources of the gender pay gap, as well as in factors that may be included in the unobserved gap in accounting exercises like this one. We also probe for insights on why the gap is larger at the top.

3. Labor Force Participation

3.1 Trends

Labor force participation is a crucial factor in understanding developments in women's wages. This is the case both because the receipt of wages is conditional on labor force participation and employment, and also because women's labor force attachment is a key factor influencing the gender wage gap. U.S. women's labor force participation rates increased dramatically in the five decades following World War II and this increase, driven by rising participation rates of married women, underlies what Goldin (2006) has termed the "quiet revolution" in gender roles that underlies women's progress in narrowing the gender wage gap

and other dimensions of labor market outcomes. For that reason, we briefly summarize the trends in female labor force participation in the United States.

The sharp increase in female participation rates is illustrated in Figure 3, which shows the rate rising from 31.8 percent in 1947 to 57.2 percent in 2013. The *gender gap* in participation rates was further reduced by the steady decline in male participation rates over this period. As may be seen in Figure 3, the growth in female participation rates began to slow and then plateau in the 1990s. Female participation rates have fallen in the wake of the Great Recession, mirroring a similar pattern among men.

There is a voluminous literature on the sources of rising female labor force participation rates dating at least from Mincer's (1962) insightful analysis of the early post-World War II increase. Consistent with Mincer's original analysis, numerous studies have continued to find that rising real wages for women have played a major role in explaining the rise in married women's labor force participation. The substitution effect due to increases in female wages more than outweighed the negative income effect due to increases in their husbands' incomes during periods of rising male wages.¹² Moreover, during the 1970s and 1980s, husbands' real incomes stagnated overall and declined for less educated men. While this factor contributed to increases in women's labor force participation during this period, consistent with Mincer's initial insight, it accounted for relatively little of the increase, with rising wages continuing to play the more important role (Juhn and Murphy 1997; Blau and Kahn 2007). Indeed, the married women with the largest increase in market hours since 1950 were those with high-wage husbands (see Juhn and Murphy 1997 and McGrattan and Rogerson 2008), likely drawn in by widening wage

¹² See, e.g., Blau and Kahn (2007) and references therein. For an excellent discussion of longer term factors, see Goldin (2006).

inequality and rising returns to skill (e.g., Autor, Katz and Kearney 2008). Rising returns to skill likely also underlie the much larger increases in labor force participation rates for highly educated women relative to their less educated counterparts (Blau 1998, Blau, Ferber, and Winkler, Figure 6-6).

A number of other factors apart from rising wages and increasing educational attainment have also been found to be important in explaining women's increasing labor force participation. These include the greater availability of market substitutes for home work and improvements in household technology (e.g., Greenwood, Seshadri, and Yorukoglu 2005), the development and dissemination of the birth control pill (Goldin and Katz 2002; Bailey 2006; Bailey, Hershbein, and Miller 2012), and demand shifts that favored occupations like clerical work where women were well represented (Goldin 1990; Oppenheimer 1976). At the same time, however, studies focused on conventional economic variables (wages, nonlabor or husband's income, education, and demographic variables) for periods of rapid increase in female participation rates (i.e., prior to the 1990s) generally find that measured variables, including the key wage and income variables, cannot fully explain the observed increases.¹³ This suggests an important role for shifts in preferences and other unmeasured factors. Cotter, Hermsen, and Vanneman (2011) and Fortin (forthcoming) provide some evidence on attitudes, although establishing causation in this relationship is challenging, since people may adjust their attitudes in light of their labor force behavior and outcomes as well as vice versa.

A final point to note is that, between 1980 and 2000, female own wage and income elasticities declined substantially in magnitude (Blau and Kahn 2007; Heim 2007). This is of significance in that it has brought female elasticities closer to male elasticities, and, though a

¹³ See Blau and Kahn (2007) and references therein.

gender difference remains, may be interpreted as an indicator that women are coming to more closely approximate men in terms of the role that market work plays in their lives (Goldin 2006; Blau and Kahn 2007).

3.2 Selection and the Gender Wage Gap

Changes over time in female participation rates raise the issue of selection bias (Heckman 1979; Gronau 1974), since data on wages are available only for a self-selected group of labor force participants. As noted above, inclusion in the wage sample requires employment and, depending on the study, there may be additional requirements, for example, being a wage and salary worker (i.e., not self-employed), working full-time, working full-year or a minimum number of weeks in a year, etc. Selection bias is likely to be a more serious issue for women's than men's wages because the closer the wage sample is to 100 percent of the underlying population, the smaller the selection bias.¹⁴

In considering wage differences between men and women, the focus would ideally be on wage offers rather than observed wages; selection bias arises because the latter are influenced by individuals' decisions about whether or not to participate in the wage and salary sector. Self-selection into the wage sample may take place on either measured or unmeasured factors and both may affect trends in observed wages. Our decompositions in Section 2 standardized for shifts in measured factors; however selection on unmeasured factors can bias the estimated coefficients in wage regressions and potentially result in misleading estimates of levels and trends in adjusted gender wage gaps. If inclusion in the wage sample is selective of those with higher (lower) wage offers, the mean of observed wages will be higher (lower) than the mean of

¹⁴ See Mulligan and Rubinstein's (2008) discussion of the identification-at-infinity method of correcting for selection and associated references.

wage offers. And, moreover, there are plausible scenarios under which the magnitude and even the sign of the selection bias may change over time. For example, intuitively we would expect changes in labor force participation rates to change the extent of selection bias and, as we have seen, not only have female participation rates increased over time, the pace of the increase has varied, with rapid rises prior to 1990, followed by slower growth and eventual plateauing thereafter. Moreover, as Mulligan and Rubinstein (2008) point out, selection patterns may change over time even in the absence of changing participation rates with, for example, changes in skill prices as they emphasize.¹⁵ Or, as another example, Blau and Kahn (2006) point to changes in public policies, specifically welfare and the Earned Income Tax Credit (EITC), as affecting selection in the 1990s. Thus, the direction of any potential selection bias on either wage levels or trends is unclear a priori.

Does selection produce misleading estimates of levels and trends in gender wage gap and is the effect sizable? The evidence on this is mixed. Blau and Beller (1988) examine the impact of selection bias on the trends in the gender earnings gap over the 1970s (1971-81), using a standard Heckman 2-step selectivity bias correction for both the male and female wage equations. The first stage was identified by the inclusion of the individual's nonlabor income, a dummy for whether s/he was age 62 or over (and hence entitled to early social security benefits), and the number of family members who were aged 18-64. Demographic variables such as marital status and number of children that are sometimes used to identify the selection correction were included in the wage equation as well as the selection equation.

¹⁵ Similarly in his comparison of black and white wages for women, Neal (2004) points out that selection may operate differently even for two groups that have roughly similar participation rates.

Blau and Beller found that, while published data on the median earnings of year-round, full-time workers showed little change in the gender pay gap during the 1970s, expanding the sample to include all workers (i.e., part-year and part-time) and using a regression approach to standardize for weeks and hours worked increased the estimate of earnings gains in an OLS context. When they corrected for selectivity bias, they found that wage offers resulted in substantially higher estimates of wage gains for white women relative to white men than did observed wages. Although the effect of the selectivity bias correction was to lower the estimated increase in the earnings ratio for blacks, the coefficients on the selectivity variables were not significant.

Blau and Kahn (2006) examined the gender wage gap over the 1979-98 period, using wage data from the PSID for 1979, 1989 and 1998. They adjusted for selection in several stages. They first progressively expanded their wage sample, first by adding part-time workers to their base sample of full-time workers; then, for those still lacking wage observations, by using the longitudinal nature of their data set to recover real (PCE deflated) wages for the most recent year available in a four year window. For the remaining individuals, in the spirit of Neal and Johnson's (1996) and Neal's (2004) analyses of black-white wage differentials, they estimated median regressions and included some additional individuals by making assumptions about whether they placed above or below the median of real wage offers. Specifically, they assumed that individuals with at least a college degree and at least eight years of actual full-time labor market experience had above the median wage offers for their gender, and that those with less than a high school degree and less than eight years of actual full-time labor market experience had below-median wage offers for their gender.

For each year, Blau and Kahn (2006) find that selection bias is positive, i.e., that the raw and human capital adjusted gender gap in wage offers is larger than the corresponding gaps for observed wages. However, their results suggest that the direction of the selectivity effect on earnings *growth* differed between the 1980s and 1990s. Specifically, they found that in the 1980s, convergence was slower after correcting for selection; however, in the 1990s, convergence was faster after the correction. They argue that the results for the 1980s are consistent with evidence that employment gains for married women were largest for wives of higher-wage men who themselves are likely to be more skilled (on both measured and unmeasured characteristics), while the pattern for the 1990s may reflect the large entry of relatively low-skilled, female single-family heads during this decade (as we have seen increases in married women's participation rates had slowed), which has been linked to changes in welfare policies and the expansion of the EITC (e.g., Meyer and Rosenbaum 2001). For the 1979-98 period as a whole, the results presented in Blau and Kahn (2006) suggest the selectivity adjustment had a nontrivial but small impact on the trends in either the unadjusted or adjusted differential.

Results obtained by Mulligan and Rubinstein (2008) suggest a much more significant role for selection in accounting for the convergence in observed wages between 1975-79 and 1995-99. Using data from the Current Population Survey and focusing on workers employed full time and full year, they implement two approaches to adjusting for selecting bias: a Heckman 2-step estimator and an identification at infinity estimator. Their Heckman 2-step estimator is identified by inclusion of number of children aged 0-6 interacted with marital status in the first stage. They also estimate results using the identification at infinity estimator which entails estimating some of the wage equation parameters on a sample that is selected based on observed characteristics

such that nearly all of the sample is predicted to be employed. In most cases they find virtually no evidence of closing of the gender wage gap once selection has been accounted for. Mulligan and Rubinstein (2008) explain their findings in terms of rising wage inequality that has increased the returns to skill. In response, women with less human capital may drop out of the workforce, while those with more human capital may enter. While it is possible to control for some indicators of human capital in the CPS (i.e., formal education), it is also quite possible that some indicators are unmeasured, giving rise to a change in the composition of the female workforce based on unmeasured characteristics and hence an important role for the selectivity bias adjustment. Consistent with this story, they find that selection of women into the full-time, full-year workforce was actually negative in the 1970s and shifted to positive in the 1990s.

Finally, Jacobsen, Khamis and Yuksel (2014) estimate wage equations for each year in the 1964-2013 period using March CPS data in order to construct a measure of lifetime earnings. Following Mulligan and Rubinstein, they use a Heckman 2-step estimator; they account for selection by excluding marital status from the wage equation, but report that their results were similar when they employed the same identification restriction as Mulligan and Rubinstein (the interaction of marital status and number of children aged 0-6) for periods when the children variable was available. They find increasingly positive selection into employment toward the end of their sample period, like Mulligan and Rubinstein (2008). However, in contrast to Mulligan and Rubinstein's (2008) wage results, they find that the gender gap in lifetime earnings closed in the 1980s although it then stopped converging. These findings for lifetime earnings are broadly similar to the unadjusted wage trends reviewed in Section 2.

Possible selection bias in measuring the gender wage gap is an important and complex issue. Thus, it may not be surprising that efforts to address it have not yet achieved a consensus.

Some differences arise because each of the reviewed studies not only focuses on a different data set or time period, but each uses a different approach to correcting for selection or implements it differently—including different definitions of the wage sample and different specifications of estimating equations. The PSID (used by Blau and Kahn 2006 and our data source in Section 2) permits a control for actual labor market experience, which will perforce be an unmeasured factor in a study based on the CPS (e.g., Blau and Beller 1988; Mulligan and Rubinstein 2008, and Jacobsen, Khamis and Yuksel 2014), which does not contain this information. More fundamentally, available approaches to correcting for selection bias each have their own strengths and weaknesses. One issue raised by estimation of the Heckman 2-step estimator is that an exclusion restriction (i.e., a variable that affects labor supply but does not affect wages) is needed (or at least desirable). The studies employing this approach reviewed here based identification on variables that could be argued to directly affect wages (such as nonlabor income in the case of Blau and Beller 1988 or marriage and children in the cases of Mulligan and Rubinstein 2008 and Jacobsen, Khamis and Yuksel 2014). Moreover, while it doesn't require exclusion restrictions, the identification at infinity method used by Mulligan and Rubinstein (2008) raises some concern because the experience of the groups identified as high employment probability may not be representative of the larger male and female wage samples.¹⁶ Finally, while the approach used by Blau and Kahn (2006) of adding observations above and below the median based on high-education, high-experience or low-education, low-experience does not raise identification issues, it does require the assumption that the wage offers for the identified

¹⁶ For example, Mulligan and Rubinstein (2008) find that when the sample is restricted to those with characteristics that predict a .8 or higher probability of being employed, it includes .8% (1970s), .5% (1980s), and 1.2% (1990s) of the white female full time, full year observations. This amounts to roughly 300 female observations per five year CPS cross section. Moreover, Mulligan and Rubinstein (2008) find that the conclusion of no true wage convergence is affected by the specification of the model used to predict employment probabilities.

groups are above median or below median, *conditional* on their measured human capital levels. This is an assumption that may reasonably be questioned, particularly at the high end.¹⁷

Thus, we see the issue of selection bias as an area where continued research, and perhaps new methodologies are needed to resolve the debate,¹⁸ though we note that with the substantial upgrading of women's education, experience levels, and occupations that we documented in Section 2, it seems highly likely to us that unadjusted gaps, at least, have declined.

4. Education and Mathematics Test Scores

4.1 Education

Education is an area in which there has been a reversal of the gender differential, as our analysis of the PSID in Section 2 showed. In the United States, traditionally, men were more likely than women to go to college and beyond. So, for example, in 1971, women received 43 percent of associate and bachelor's degrees, 40 percent of master's degrees, 14 percent of Ph.D.s, and 6 percent of first professional degrees (awarded in post-college professional training programs, including medicine, law, dentistry, pharmacy, veterinary medicine, and theology). By 1980, women had caught up to men in college graduation and subsequently they have surpassed them. As of 2011, women earned 57 percent of bachelor's degrees and 62 percent of associate degrees. There have been comparable gains at the post-graduate level—with women receiving 61 percent of master's degrees, 51 percent of Ph.D.'s and 49 percent of first professional degrees

¹⁷ Blau and Kahn argue that this assumption is more likely to be valid for the group with low-education, low-experience group (placed below the median). When they repeat the analyses adding only the low-education, low-experience group, their results are virtually identical.

¹⁸ Blundell, Gosling, Ichimura and Meghir (2007) propose a method that uses bounds, tightened by restrictions based on economic theory, to estimate changes in the distributions of wages that allow for the nonrandom selection into work. Applying this method to the United Kingdom, they find evidence of increases in the relative wages of women. Olivetti and Petrongolo (2008) explore the role of selection in employment in accounting for international differences in the gender wage gap using alternative assumptions on the position of imputed wages of the nonemployed in the spirit of Neal (2004) and Blau and Kahn (2006).

(Blau, Ferber and Winkler 2014, Chapter 8).¹⁹ The broad outlines of these trends prevail across the economically advanced nations and many developing countries as well (Goldin, Katz and Kuziemko 2006, and Becker, Hubbard, and Murphy 2010).

In addition, the type of education women receive has changed toward more mathematics- and career-oriented programs. Substantial gender differences in college majors remain, but college majors are considerably less gender segregated than they were in the 1960s (Blau, Ferber, and Winkler 2014, Chapter 8). Much of these gains were achieved by the 1980s, however, with less progress since then (England and Li 2006; Bronson 2013). Significantly, women continue to lag in the STEM (science, technology, engineering and mathematics) fields, particularly in mathematically-intensive fields (Ceci, Ginther, Kahn, and Williams 2014).

As relatively more highly educated female cohorts have replaced earlier ones, women have now become more highly educated than men in the overall population (Blau, Ferber, and Winkler 2014, Chapter 8). The female advantage is particularly evident in the labor force (see Section 2), which is still more highly selected on education for women than for men. The reasons why women have overtaken men in education are not fully understood but seem to take in both pecuniary and nonpecuniary factors. The edge men traditionally enjoyed in college and beyond could be rationalized a human capital investment framework. Women's shorter expected worklife reduced their gains to investing in large amounts of formal schooling, although other factors, including familial attitudes, social gender norms, and discrimination by educational institutions could be factors as well. From the human capital perspective, women's rising labor force attachment might be expected to raise the returns to their investment in higher education

¹⁹ These figures are based on published data from the Department of Education for 1960-61 and 2010-11. In 2011, women also received 46 percent of master's degrees in business, which are not included in this tabulation of first professional degrees in the Department of Education statistics.

and thus to narrow the educational gender gap. Working in the same direction, reductions in occupational segregation associated with the increased entry of college women into higher-paying, formerly male managerial and professional jobs likely provided a further economic incentive for women to invest in college.²⁰ These employment gains likely reflect, at least in part, the government's antidiscrimination in employment effort spearheaded by the enforcement of Title VII of the Civil Rights Act and the implementation of Affirmative Action for government contractors (evidence on this is discussed in Section 10).

A number of additional factors appear to have contributed to the increase in women's educational attainment. First is the development of "the pill" and its growing availability to young, unmarried women beginning in the late 1960s and early 1970s. The availability of the pill was associated with and facilitated a delay in marriage and childbearing, which in turn enabled women to pursue professional training after college (Goldin and Katz 2002 and Bailey 2006). Second, passage and enforcement of Title IX of the Civil Rights Act, which banned discrimination in educational institutions, likely led to changes in admission and other practices that facilitated and encouraged women's increased participation in higher education. Third, and associated with these and other developments, social norms and views on gender appropriate education investments most likely also changed. Finally, as Goldin, Katz, and Kuziemko (2006) show, girls were well positioned to increase their college attendance in terms of their high school grade point averages and class rank, which surpassed those of boys even during the era in which boys' college-going exceeded girls'. Moreover, while girls' high school preparation and test

²⁰ Decreases in occupational segregation were especially pronounced among the college educated (Blau, Brummund, and Liu 2013a and b). It is unclear whether the college wage premium is higher for women than men. Earlier work by Dougherty (2005) and others suggested that returns measured in this way were higher for women, but Hubbard (2011) presents evidence that this is not the case when topcoding in the major data set used in these studies, the Current Population Survey, is corrected.

scores in science and mathematics initially lagged those of boys', these gaps were reduced as girls' expectations of attending college increased.

While the above considerations may help to explain why women have caught up to men in education, or at least why they have *reduced* the gender education gap (since women's expected labor force attachment is still less than men's), women's surpassing men is more puzzling—especially since, as noted earlier, this is an international phenomenon. A number of possible explanations for this have been offered, and all may play a role to some extent.

First, a college education not only increases own income but also results in family-related income gains due to assortative mating. Such gains are likely to be larger for women than men, since, in the majority of couples, men are still the higher earners. Moreover, college-educated women have lower divorce rates and a lower incidence of out-of-wedlock births, making them less likely than their less-educated counterparts to become lower income, single family heads. To the extent this association is causal, this factor would also increase the returns to college more for women than men. DiPrete and Buchmann (2006) find that such family-related income gains (adjusted for family size) increased more for women than for men, suggesting that this may be part of the reason for the increase in women's college-going. Further, in the event of a divorce, Bronson (2013) argues that college provides insurance value and presents evidence that this consideration helps to explain the growth in women's college attendance.

Second, there appear to be gender differences in noncognitive skills—for summaries and discussions, see Goldin, Katz and Kuziemko (2006) and Becker, Hubbard, and Murphy (2010)—that suggest girls have lower nonpecuniary costs of investing in college than boys. For one thing, as noted earlier, girls have traditionally excelled relative to boys in secondary school academic performance and this was the case even when they were less likely than boys to go to college.

This suggests that girls find school less difficult or unpleasant than boys. There is evidence, for example, that boys spend much less time doing homework than girls (Porterfield and Winkler 2007). In addition, boys have a much higher incidence of school disciplinary and behavior problems, ranging from minor infractions to school suspensions and participation in criminal activity. Boys are also two to three times more likely to be diagnosed with attention deficit hyperactivity disorder (ADHD). The reasons for these gender differences have not been fully determined but one factor may be the later maturation of boys (Goldin, Katz and Kuziemko 2006). To the extent that females have lower total (pecuniary plus nonpecuniary) costs of investing in education on average than males, they will have a larger response to given increases in the benefits of college.

Becker, Hubbard, and Murphy (2010) also put great emphasis on gender differences in the *distribution* of noncognitive skills. They present evidence that the variance in noncognitive (or what they call nontraditional) skills is smaller for women than men, suggesting that under some circumstances the elasticity of supply to college will be higher for women than men. This depends on the location of the relevant portion of the distribution of costs. If, as appears likely, the relevant portion is close to the mean of costs, the density of individuals that can respond to an increase in benefits is larger for a lower-variability distribution that peaks around the mean—as is the case for women. If women have a higher elasticity of supply to college, then even for equal changes in the benefits, women can overtake men in college attainment.

6.3 Mathematics Test Scores

The gender differential in mathematics scores is potentially linked to gender differences in wages and occupations. Traditionally, U.S. boys and young men have had higher average mathematics test scores than girls and young women, as well as higher representation at top

performance levels. As noted above, the gender difference has narrowed as high school curricula of boys and girls have gotten more similar. Indeed, some evidence indicates that boys no longer have higher average math test scores during their high school years than girls.²¹ However, there is continuing evidence of a gender difference at top performance levels, with males outnumbering females at the very high ranges of science and math tests, and females outnumbering males at the very high ranges of reading and language tests (e.g., Pope and Sydnor 2010). The male advantage at the upper end of math test scores has been cited as a factor in the underrepresentation of women in STEM fields, although this contention has been the focus of considerable debate.²² Of particular interest, a significant strand of recent research focuses on the social determinants of these differences and implicitly asks whether gender differences in math performance may be influenced by educational policy and other environmental factors.

Evidence that social influences matter comes from a variety of sources. For example, Pope and Sydnor (2010) examined geographic variation in gender patterns in the standardized test scores of public school students in the U.S. In all states, they found the pattern noted above: average math scores of boys and girls were roughly equal, but boys were disproportionately represented at the top in math and science. They also found that girls were disproportionately represented at the top in reading. Importantly, however, the *extent* of the gender difference

²¹ Hyde, Lindberg, Linn, Ellis, and Williams (2008); this study used data from state assessments of cognitive performance. However, Fryer and Levitt (2010) continue to find a gender gap at the high school level using the Early Childhood Longitudinal Study Kindergarten Cohort, which is a sample of children entering kindergarten in 1998. Both studies are critical of SAT data since the pool of students taking the test is not representative of the full population and selection into the test may differ by gender.

²² For example, while Hyde et al. (2008) note the slightly greater variance of male test scores in their data, they argue that gender differences along this dimension are “are insufficient to explain lopsided gender patterns in participation in some STEM fields.” In their extensive review, Ceci, Ginther, Kahn and Williams (2014) are also skeptical that math differences can account for the gender underrepresentation in math intensive fields. For an early study delineating the relationship between mathematical ability and field choice and its relationship to male-female differences in earning and occupations, see Paglin and Rufolo (1990).

among the top performers varied considerably across states and regions. This finding alone suggests environmental factors play a role. Moreover, Pope and Sydnor (2010) found that in states where boys were particularly heavily represented among top scorers in math and science, girls were particularly heavily represented among top scorers in reading. This finding suggests that sex-stereotyping is going on in some states, with boys and girls being differentially treated in some ways that contribute to the gender imbalance in test scores.

International comparisons may also be instructive. Using data from the Programme for International Student Assessment (PISA), Guiso, Monte, Sapienza, and Zingales (2008) find that although girls' average math scores are below boys' in many countries, as in the case of US states, the gender difference in scores varies considerably across countries, again suggest the importance of environmental factors. Moreover, in some countries, such as Sweden and Norway, there is no gender gap in math scores, while, in Iceland, girls' mean scores are higher than boys'. Similarly, there is considerable variation in the proportion of girls compared to boys scoring at the highest levels and, for example in Iceland, more girls scored at the highest levels than boys. The importance of social factors is further suggested by Guiso et al.'s (2008) finding that girls' math scores (at the mean as well as their representation at the highest levels) were positively related to indicators of country-level gender equity (as measured by various indicators, including the World Economic Forum's Gender Gap Index).²³ Further international evidence in support of this conclusion is provided by Nollenberger, Rodríguez-Planas, and Sevilla's (2014) study of the effect of source country characteristics on the mathematical performance of girls

²³ However, Fryer and Levitt (2010) report that the finding that the gender math gap is linked to measures of gender equality is sensitive to the inclusion of Muslim countries where, although women have very low status, there is little or no gender gap in math. Hoffman, Gneezy, and List (2011) provide intriguing experimental evidence in support of the role of environment with their finding that the gender gap in spatial abilities disappears when one compares a matrilineal to an adjacent patrilineal society in Northeast India.

relative to boys among second-generation immigrants (i.e., native born children of immigrants). Also using PISA data on multiple countries, they find that the higher the degree of gender equality (as measured by gender-equity indexes) in the *parents'* source country, the higher the math performance of girls relative to boys in the destination country.

Gender differences in psychological attributes may also contribute to gender differences in math scores. Below we survey evidence on gender differences in attitudes toward competition; and find that there is some evidence not only that women shy away from competition relative to men, but that they respond differently (and less favorably than men) to competitive pressure. As pointed out by Niederle and Vesterlund (2010) the competitive pressures associated with test taking may thus “magnify and potentially distort underlying gender differences in skills,” particularly in an area like mathematics where there is a strong belief that men are better at it (p. 140).

Evidence on “stereotype threat,” further suggests that men’s and women’s math test scores may be influenced by gender stereotypes and their salience in the test-taking situation. Stereotype threat is the notion that the negative stereotype that women have weaker math ability can undermine their performance. Spencer, Steele, and Quinn (1999) test for stereotype threat in explaining gender differences in performance on a difficult mathematics test among university men and women matched for equivalent math backgrounds and interest (there was no gender difference in performance on an easy test). They had two treatments: (i) half of the participants were told that the test had shown gender differences in the past, while (ii) the other half were told that the test had been shown to be gender fair—that men and women performed equally well on the test. In the latter case, men and women performed equally. But when subjects were told that the same test had shown gender differences in the past, women scored lower on the test than

men.

Finally, recent work by Lavy and Sand (2015), using Israeli data, suggests that girls may face teacher discrimination in the evaluation of their math skills and that teacher bias has long-term impacts. They identify biased primary school teachers by differences between their average grading of boys' and girls' performance on a (nonblind) teacher-graded in-class exam compared to the performance of the same boys and girls on a blind (anonymously-graded) national exam. Overall, they find teachers over-assess boys relative to girls, with girls outscoring boys in the national exam and boys outscoring girls in the classroom exam. They use random assignment of teachers to examine the longer-term consequences of the extent of primary school teacher bias, and find that it enhances boys' scores on later standardized national exams in middle school and high school and their enrollment in advanced high school math courses, while girls are negatively affected.

5. Labor Force Experience, Training and Turnover

In this section we focus on the empirical literature that illuminates the importance for the gender wage gap of work-related indicators of human capital, including experience, tenure, training and turnover. Dating from the seminal work of Mincer and Polachek (1974), gender differences in experience and labor force attachment have been seen as central to the understanding of the gender wage gap. Assuming a traditional division of labor by gender in the family, women will anticipate shorter and more discontinuous work lives as a consequence of their family responsibilities; they will thus have lower incentives to invest in on-the-job training than men. Their resulting smaller human capital investments and reduced labor market experience will lower their relative earnings. The Mincer and Polachek analysis also highlighted the role of human capital depreciation during workforce interruptions as further lowering the

wages of women upon their return to market work. Women are also expected to choose occupations for which human capital investments are less important and in which the skill depreciation that occurs during time spent out of the labor force is minimized (Polachek 1981).

We may obtain further insights by distinguishing between general training (which is transferable across firms) and firm-specific training (which imparts skills which are unique to a particular enterprise).²⁴ Women will especially avoid jobs requiring large investments in firm-specific skills because the returns to such investments are reaped only as long as one remains with a particular employer. At the same time, employers are expected to be reluctant to hire women for such jobs because they bear some of the costs of firm-specific training. (Since general training is transferable, a simple model predicts that employees will bear the costs and reap the returns to such training, although modern analyses of general training find that firms share the costs and benefits here as well (Acemoglu and Pischke 1999)). Such employer behavior would be consistent with models of statistical discrimination where, given employer uncertainty about worker productivity, firms may discriminate against groups like women or minorities based on real or perceived average differences in productivity (Phelps 1972; Aigner and Cain 1977). As Altonji and Blank (1999) point out, such discrimination is plausible given evidence that firms face uncertainty about the productivity of their workers.²⁵

Recent work by Goldin (2014) continues to highlight the role of work force interruptions in lowering women's wages but outlines a different mechanism for this effect. Goldin (2014) analyzes the impact of interruptions in the context of a broader analysis of the impact of temporal

²⁴ See Becker (1993, 1st ed. 1964) for this distinction and Blau, Ferber, and Winkler (2014) for a graphical development of its application to gender differences in on-the-job training investments.

²⁵ See, for example Farber and Gibbons 1996; and Altonji and Pierret 1997; or, more recently, Kahn (2013) and Kahn and Lange (2014).

flexibility (or the lack thereof) in impacting the gender wage gap. In particular, she focuses on the disproportionate rewards in some occupations/firms for working long hours and particular hours.²⁶ Her main focus is on hours of work, but as she notes, interruptions can also be viewed in this context. She argues that the explanation for a high wage penalty for temporal flexibility can best be understood through the lens of personnel economics rather than human capital theory. In particular, she sees such pay differences as arising because of differences across workplaces in the value of long hours rather than of differences across individuals in amounts of human capital. The result is a classic compensating differential equilibrium à la Rosen (1986). Workers place different values on temporal flexibility (with women placing a higher value than men) and firms or sectors confront different cost to providing it—workers sort across workplaces accordingly.

Goldin points to (and presents empirical support for) the importance occupational characteristics that make providing flexibility extremely costly in some sectors and relatively inexpensive in others. So, the wage penalty for flexibility is likely to be high in jobs that require meeting deadlines (time pressure), being in contact with others to perform the job, maintaining and establishing interpersonal relationships, adhering to preset schedules, and doing work for which other workers are not close substitutes. As an example, there may be a high penalty to shorter hours or workforce interruptions for lawyers at a large, high-powered firm, not because of the smaller amount of human capital acquired by those working fewer hours or the depreciation of their human capital stock during time out of work, but rather due to interruptions

²⁶ In related work, Cha and Weeden (2014) examine the role of an increase in the prevalence of long (50 or more) work hours and the rising returns to long hours in slowing convergence in the gender wage gap during the 1979-2009 period. They find that this factor worked to increase the gender wage gap by about 10 percent of the total change over this period—mainly due to the rising return to long hours (the gender gap in the incidence of long hours was relatively constant). This factor was particularly important in managerial and professional occupations.

in servicing clients and the inability to smoothly hand over work to other employees. We shall return to her findings below in the context of our discussion of individual occupations.

The Goldin analysis is interesting in itself and also highlights that findings that show returns to long hours and labor market experience and penalties to workforce interruptions are susceptible to other interpretations than human capital. In addition to the factors that Goldin highlights affecting the costs of providing flexibility, others include signaling²⁷—longer hours and workforce continuity may signal greater willingness to work hard, as well as greater motivation and commitment—and discrimination. Related to the signaling argument, discrimination may be due to statistical discrimination against the “type” of worker who puts a high premium on temporal flexibility.

5.1 Gender Differences in Labor Force Experience and Work Hours

As we have seen in Section 2, and as borne out in a wide literature, there is considerable evidence that overall gender differences in labor market experience account for a significant portion of the gender wage gap and that decreases in the gender experience gap help to account for the corresponding decline in the gender wage gap that we have observed in recent decades (e.g., Blau and Kahn 1997, Blau and Kahn 2006, O’Neill and Polachek 1993; Gayle and Golan 2012).²⁸ This was illustrated by our findings in Section 2, where we saw that as women’s labor force attachment has increased, gender differences in labor market experience have been greatly reduced, although not eliminated. For example, Table 2 shows that while in 1981 women full-

²⁷ See, for example, Landers, Rebitzer, and Taylor (1996); see Goldin (2014) for additional references.

²⁸ Bailey, Hershbein, and Miller (2012) explore the role of access to the pill in altering women’s human capital investments (labor market experience and education) and hence lowering the gender wage gap. Weinberger and Kuhn (2010) examine the extent to which the decline in the gender wage gap was associated with changes across cohorts in the relative rate of wage growth after labor market entry (slopes), versus changes in relative earnings levels at labor market entry (levels). They find that the former (plausibly associated with post-school investments including experience) accounts for about 1/3 of the decline, with the remainder associated with changes across cohorts (i.e., each entry cohort faring better than its predecessor).

time workers had 6.8 years less full time labor market experience than men, this difference had declined to 1.4 years by 2011.²⁹ Thus, as shown in Table 4, gender differences in experience and labor force attachment now account for a considerably smaller share of the gender wage gap than previously: in 1980 the gender difference in experience accounted for a gap of 0.114 log points (24 percent of the gap); this had declined to 0.037 log points (16 percent of the gap) in 2010. (These results are based on the human capital specification so as to measure the full (direct plus indirect) effect of experience.)

As we have seen, Mincer and Polachek (1974) also pointed out that skills of women may be expected to depreciate during periods of time out of the labor force, further contributing to the gender wage gap. They also expected women's wages to rebound to some extent after reentry due to a burst of investment upon their return to employment. Mincer and Ofek (1982), for example, present evidence consistent with this expected pattern. In perhaps the most detailed examination of these issues, Light and Ureta (1995) endeavor for younger workers to fully characterize past employment experience, measuring the fraction of time worked during each year. They find that the *timing* of experience accounts for as much as 12% of the unadjusted gender wage gap. It would be interesting to know how findings from this earlier literature, regarding timing of gaps as well as their importance, would hold up for more recent cohorts of women who are considerably more attached to the labor force. (Mincer and Ofek (1982) employ data from the National Longitudinal Survey of mature women observed over the 1966-1974

²⁹ Published government data on tenure (length of time with a particular employer) also indicate a precipitous drop in the gender gap. In 1966, men's median tenure was 2.4 years more than women's; by 2012, the gender gap had fallen to only 0.1 years. And the share of long-term workers, those with tenure of 10 or more years, was only slightly higher for men (35 percent) than for women (33 percent). See, U.S. Department of Labor, Bureau of Labor Statistics, "Job Tenure of Workers, January 1966," *Special Labor Force Report* No. 77 (1967); and U.S. Department of Labor, Bureau of Labor Statistics, "Employee Tenure in 2012," *News Release* (September 18, 2012), available at <http://www.bls.gov/news.release/pdf/tenure.pdf> (accessed December 1, 2012). Note median tenure data are for workers 16 and over; the share of long tenure is for workers 25 and over.

period and Light and Ureta (1995) use data from the National Longitudinal Surveys of young women and men over the 1966-1981 (women) or 1968-1984 (men) periods.)³⁰ Consistent with a possibly diminished role of workforce interruptions, Blau and Kahn (2013b) find that, although coefficients on variables measuring time out of the labor force are generally negative (although not always significant), estimates of the unexplained gender wage gap are not sensitive to their inclusion, actually, not only in 1999, but in 1990 and 1980 as well.³¹ Their data from the Panel Study of Income Dynamics did not permit them to look at the timing of interruptions, but Spivey (2005), using data from the National Longitudinal Survey of Youth 1979, found that timing of experience can explain only a negligible portion of the gender wage gap among workers observed over the 1979-2000 period.³²

The foregoing results that suggest a relatively small role for experience in explaining the gender wage gap currently are for the labor market as a whole. In contrast, recent influential work has highlighted the particular importance of labor force experience, attachment, and hours worked in some occupations, including business and professions like law, where work histories and current hours seem to be a particularly important determinant of gender wage differences. This work is of particular interest in that the findings for these occupations are particularly applicable to the upper end of the wage distribution where, as we have seen, the gender wage gap has declined more slowly than at other regions. Also of interest are findings from Goldin (2014)

³⁰ Light and Ureta (1995) note that they observe men as well as women with nonwork spells and advocate their detailed model of work histories as the preferred approach to analyzing the contribution of experience to wages regardless of sex.

³¹ Data are for workers aged 18-65 in the indicated year.

³² Respondents were 14-22 in 1979. Spivey provides a useful review of the literature on the wage effects of workforce interruptions.

that point to the high penalty for flexibility in some high wage occupations in causing a gender wage gap within them.

We begin with Noonan, Corcoran, and Courant's (2005) study of lawyers³³ and Bertrand, Goldin, and Katz's (2010) study of MBAs. Noonan, Corcoran, and Courant (2005) focused on two cohorts of graduates of the University of Michigan Law School 15 years after graduation; the first cohort was surveyed between 1987 and 1993 and the second between 1994 and 2000. The results for the two cohorts were quite similar. The gap in pay between women and men was found to be relatively small at the outset of their careers, but 15 years later, men earned over 50 percent more. A considerable portion of this difference reflected choices that workers themselves made, including the greater propensity of women lawyers to currently work shorter hours and to have worked part time in the past or to have taken some time out after child birth. Also important was job setting (type and size of employer). Taken together, these factors and other measured variables explained 69 percent of the gender gap in earnings (not including job setting) and 77-80 percent (including job setting).

Bertrand, Goldin, and Katz (2010) examined earnings of MBAs who graduated between 1990 and 2006 from the Booth School of Business of the University of Chicago (they were surveyed in 2006-2007). Like the study of lawyers, the researchers reported a relatively small gender differential at the outset of the career. However, averaged across the full set of MBA graduates (individuals who had been out for 1 to 16 years), men earned 0.29 log points (33 percent) more than women. By 10-16 years post-degree, men earned 0.60 log points (82 percent) more. The study found that the gender gap could largely be explained by labor supply factors like weekly hours and actual post-MBA work experience, which are in turn related to career-

³³ See also Goldin's (2014) reexamination of these data that arrives at broadly consistent findings.

family tradeoffs. A small portion of the wage difference was also due to men's academic advantage in the MBA program (higher GPA and more finance courses taken). For the sample as a whole, these and other measured factors accounted for 78 percent of the gender difference in earnings.³⁴

This research suggests substantial penalties for shorter hours, lesser experience and workforce interruptions among JDs and MBAs. With respect to hours, it should be noted that both of these decompositions focus on annual earnings, leaving open whether the importance of current hours reflects simply a proportional reduction in earnings or an additional hourly wage penalty for shorter hours. Moreover, these results could be seen in the context of the human capital model, and the particular importance of human capital in these occupations. Goldin (2014), however, views these results as more consistent with her analysis of the high penalties to *flexibility* in these occupations discussed above, including a convex return to current hours.³⁵ More generally, for college graduates in the 95 highest earnings occupations, she examined the relationship between an index of occupational characteristics associated with high costs of flexibility and unexplained gender wage gaps. The index of flexibility was found to be positively related to (i.e., to increase) the (adjusted) gender log wage gap.³⁶ She also computed the elasticity of annual earnings with respect to weekly hours in each occupation and found that

³⁴ This was calculated by evaluating the regression coefficient on the female dummy in Table 3, specification 6; see Bertrand, Goldin, and Katz (2010), p. 239. This was the most highly specified regression that did not include a control for "reason for choosing job," which we felt inappropriate to include in a regression estimating the unexplained gap potentially due to discrimination.

³⁵ Goldin (2014) notes that about two thirds of the total penalty from job interruptions among those in the Chicago MBA sample who were 10 to 16 years out is due to taking *any* time out. Cumulative time not working is only about one year for these women, which would seem a relatively modest interruption to elicit large penalties in a human capital context.

³⁶ The differential is expressed as female minus male log earnings and is thus negative. Goldin controls for a quartic in age, race, hours and weeks worked, and education. The American Community Survey data which she employs do not permit her to control for actual labor market experience.

occupations with higher elasticities had more negative log earnings gender gaps. She found that business occupations and law had high values on the inflexibility index as well as high elasticities of annual earnings with respect to weekly hours. In contrast, technology and science jobs scored much lower on the inflexibility measure and had smaller elasticities. Her findings are consistent with a high cost for flexibility in law and business and a lower cost in science and technology—with gender wage gaps lining up accordingly. The latter finding is surprising in a human capital context where it might be expected that human capital acquisition and depreciation of skills would be particularly important in science and technology jobs. As a further contrast to business and law, she provides a case study of pharmacists (see also Goldin and Katz 2012) in which industry developments and technological factors have greatly reduced the costs of flexibility and the gender pay gap has fallen accordingly.

5.2 Gender Differences in Training and Quitting

Considerable empirical evidence supports the prediction of the human capital model that women will receive less on-the-job training than men (e.g., Altonji and Spletzer 1991; Barron and Black 1993). This finding is consistent with employer and worker decisions based on a lower expected probability of women remaining with the firm or in the workforce. A study by Royalty (1996) is particularly illuminating in that she explicitly examined the role of women's higher (predicted) probability of turnover in explaining the gender training difference. While Royalty supports the expectation that expected turnover helps to account for the gender difference in training, interestingly, she finds that a major portion of the training gap remains unexplained even after this and other determinants of training are taken into account. This finding, which is analogous to an unexplained gap in an analysis of the gender wage differential is consistent with a role for discrimination, although, as in that case it may also be due to omitted factors.

As in the case of experience, it would be interesting to see this literature updated to account for the impact of rising women's labor force attachment on the findings. This is especially the case in that younger cohorts of women now have higher educational attainment than men, and more educated workers are believed to get more on-the-job training than less educated workers as implied by their steeper experience-earnings profiles. Using data from the NLSY1979 through 2006, Kosteas (2013) found that, consistent with expectations based on the human capital model, women with more traditional gender role attitudes (as measured in 1979) were less likely to invest in training. This suggests that the gender training difference remains relevant.

Since gender differences in quit behavior can differentially impact the wages and occupations of men and women, it is important to ascertain the extent and sources of such differences. In general, while some evidence suggests that women workers may have higher quit rates on average than men, most of this difference has been found to be due to the types of jobs they are in and the worker's personal characteristics.³⁷ That is, *all else equal*, women are no more likely to quit than their male counterparts. Indeed, it is unclear that even the average gender difference in quitting still prevails. Using data on young workers from the 1987 wave of the National Longitudinal Survey of Youth (NLSY) 1979, Royalty (1998) finds the average probability of staying on the job is not significantly different for men and women.

However, consistent with women placing a greater priority on family responsibilities to the detriment of their labor market outcomes, evidence indicates that the *reasons* for quitting tend to differ between men and women (Sicherman 1996, Royalty 1998, and Keith and

³⁷ This finding dates back to the first detailed work on this topic by Viscusi (1980) and Blau and Kahn (1981) and is reflected in the findings of more recent studies, e.g., Sicherman (1996); and Royalty (1998).

McWilliams 1995). In particular, women are more likely to quit their jobs for family-related reasons or to exit to nonemployment. These types of quits are likely to negatively affect subsequent earnings as workers may trade-off wages for other job characteristics or lose seniority and firm-specific training if they exit the labor force. Men are more likely to quit for job-related reasons; such quits tend to positively affect subsequent earnings as workers are bid away to higher-paying opportunities.³⁸ Interestingly, the study by Royalty (1998) noted above found for young workers using 1987 data that this gender difference in the pattern of quits is concentrated among workers with a high school education or less, while little gender difference was found for those who had attended college. This suggests that the workforce attachment of the latter group may be more nearly equal to their male counterparts. More broadly, it would be of interest to see analyses of both quitting and the reasons for quitting updated to see whether the outlines of the earlier findings still hold. In light of the declining gender differences in labor force attachment, it is reasonable to expect that gender differences in quit behavior have further diminished.

6. The Impact of Gender Roles and Motherhood

Traditional gender roles and women's greater responsibility for nonmarket work may negatively affect women's labor market outcomes beyond their impact of labor force attachment per se. In this section we first consider the motherhood wage penalty, which has gotten considerable attention in the literature. We then review other ways in which traditional gender roles can reduce women's relative wages.

6.1 The Motherhood Wage Penalty

Considerable empirical evidence indicates a negative relationship between children and

³⁸ These wage effects are examined in Keith and McWilliams (1995).

women's wages, commonly known as the motherhood wage penalty.³⁹ While the observed empirical association could be causal, it could also be due to selection. The selection argument is plausible in that women with lower wage offers will have lower costs of children. However, there are also a number of reasons for expecting a causal effect, beyond an impact on work experience and the incidence of part-time work. First, particularly in the era before parental leave was mandated, but even to some extent today, the birth of a child may cause a woman to break her tie to her current employer, either to withdraw from the labor force entirely or to switch to a more "child-friendly" job. To the extent this occurs, she forgoes the returns to any firm-specific training she might have received as well as returns, if any, to having made a particularly good job match. Second, anticipation of this possibility could deter both women and their employers from making large investments in the firm-specific training of women of childbearing age. Third, motherhood may reduce women's productivity in a variety of ways not readily captured in wage analyses including, for example, less effort expended at work, constraints on work schedules and travel, and reluctance to be promoted to a more demanding job.

A final possibility is that mothers may face discrimination and there is persuasive experimental evidence from Correll, Benard, and Paik (2007) that this is the case. In this study, the authors first conducted a laboratory experiment in which they asked student evaluators to assess résumés of equally-qualified same-sex (female or male) job applicants who differed only as to parental status. Parental status was indicated by including information on the application such as PTA coordinator for a parent and fundraiser for a neighborhood association for a

³⁹ For a recent review of the literature and comparative findings across economically advanced countries see, Sigle-Rushton and Waldfogel (2007). Early influential treatments include, Fuchs (1988), Korenman and Neumark (1992), and Waldfogel (1998).

nonparent. Mothers were perceived by evaluators as less competent and less committed to paid work and the evaluators recommended lower starting salaries for them. In contrast, the evaluators did not penalize men for being fathers. Indeed, they perceived fathers to be more committed and recommended higher starting salaries for them. Correll, Benard, and Paik (2007) further confirmed their lab findings using a field experiment in which they sent résumés and cover letters from fictional, equally-qualified, same-sex applicants to employers advertising for job openings. They found that prospective employers called mothers back only about half as often as nonmothers, while fathers were not disadvantaged in the hiring process, although, in contrast to the lab experiment, fathers were not advantaged relative to nonfathers.

Why might employers discriminate against mothers relative to nonmothers? One possibility is that they perceive mothers as a group to be less productive, on average, due to the types of factors we have discussed. This could result in statistical discrimination against mothers based on employers' perceptions of average differences in productivity between mothers and nonmothers.

There has also been some research focusing on the impact of family status on men's wages, with most of the attention on the observed positive association between marriage and male earnings controlling for measured characteristics. Here again the question arises as to whether this relationship is causal and, if so, why. The possibility that it reflects selection is intuitively plausible in that, even today, men tend to be the primary wage earners in most families. This gives women a considerable incentive to select spouses with higher earnings potential. There are, however, also reasons for expecting that the relationship between marriage and male wages is causal. Specialization in the family à la Becker (1981, enlarged edition 1991) allows married men to focus on the market while their wives have primary responsibility for nonmarket production. Related to this, traditional notions of gender roles which view the husband as the

primary earner may increase married men's effort and motivation and hence their wages. It is also possible that employers discriminate in favor of married men—this is hinted at by the findings on parental status discussed above. Overall, as in the case of the motherhood wage penalty, the empirical evidence suggests that some portion of the observed relationship is causal.⁴⁰

6.2 Housework

Women's generally greater nonmarket responsibilities could impact labor market outcomes in a number of ways. Becker's (1985) theoretical analysis focused on the longer hours that married women and mothers tend to spend in these activities which could reduce the effort that they put into their market jobs, controlling for hours, and thus decrease their hourly wages compared to men. Indeed, it has been found that additional hours spent in housework are associated with lower wages, all else equal, although results are stronger for married women than married men (see, e.g., Hersch and Stratton 1997 and 2002). An interesting result in Hersch and Stratton (2002) links the strength of the negative effects to the type of housework that women are typically more likely to perform—routine tasks like meal preparation, cleaning, shopping, and laundry—that are more likely to be engaged in on a daily basis and that, the authors argue, are more likely to interfere with market productivity. One concern about these findings is that, while the Hersch and Stratton studies pay careful attention to endogeneity by estimating instrumental variable and fixed effect models, such evidence might, nonetheless, be contaminated by reverse causation.⁴¹

⁴⁰ For useful reviews of the literature, see, Ribar (2004) and Rodgers and Stratton (2010). For an early influential study, see Korenman and Neumark (1992). There is also some evidence that fatherhood increases male earnings, particularly when the mother experiences a workforce interruption (Lundberg and Rose 2000).

⁴¹ However, in both Hersch and Stratton (1997) and (2002), the authors are unable to reject the hypothesis that housework is exogenously determined with respect to wages.

6.3 Gender and the Location of the Family

Early theoretical and empirical work by Frank (1978), Mincer (1978), and Sandell (1977) called attention to the determinants of the location of the family which may be seen as potentially influencing the relative wages of married women. To the extent that families place priority on the husband's, rather than on the wife's, career in determining the location of the family, her earnings are likely to be decreased. She may be a "tied mover," relocating when it is not advantageous for her to leave a job where she has accumulated firm-specific training or which is a particularly good match. Alternatively, she may be a "tied stayer," unwilling to relocate despite better opportunities elsewhere. As Mincer pointed out, the tendency of women to be tied-movers need not merely reflect adherence to traditional gender roles. It is economically rational for the family to place greater emphasis on the employment and earnings prospects of the larger earner (generally the husband) whose gains to migration are likely to outweigh the losses of the spouse who is a tied mover—for early evidence see Sandell (1977). Cooke, Boyle, and Couch (2009) present recent evidence that this is indeed still the case on average, i.e., that migration is associated with a significant increase in total family earnings, despite declines in women's earnings.

Anticipation of a lesser ability to determine the geographic location of the family may also lead women to select occupations in which jobs are likely to be readily obtained in any labor market, thus constraining their occupational choices to geographically flexible jobs. As Benson (forthcoming) points out, even as women have entered higher-level, traditionally-male occupations in recent years, their entry into the more geographically-dispersed occupations (e.g., physicians, accountants, pharmacists, and managers) has been considerably greater than the more geographically-clustered (e.g., specialized engineers and physical scientists).

Some recent work has elaborated on how location decisions are likely to be affected as some couples, particularly college-educated “power couples,” may depart from these patterns and try to accommodate both careers by making a joint location decision. Costa and Kahn (2000) report that college-educated couples became increasingly located in large metropolitan areas over the 1970-1990 period. They argue that this is because large metropolitan areas offer more potential job matches for such couples. They point out that the share of dual career households increased among the college educated over this period and note Goldin’s (1997) evidence that the career-orientation of college-educated women also increased. They also note that, if returns to education are higher in larger cities, power couples have a greater income loss of locating outside of them than do other dual career couples. Costa and Kahn show that the concentration of power couples in larger metropolitan areas is greater than for other household types and exceeds what would be predicted for observationally identical single individuals, thus supporting the colocation argument.

On the other hand, Compton and Pollak (2007), using longitudinal data, do not find that power couples (again, in which both spouses have college degrees) are more likely to migrate to larger cities than other couples. Rather, their findings are more consistent with the aggregate results on location determination summarized above that it is the education (and presumably the earning power) of the husband that principally affects the couple’s propensity to migrate to a large metropolitan area, suggesting that, even within the group of power couples relocations may still adversely affect women’s wages relative to men’s.⁴² This is plausible in that, even among power couples, it is likely that the husband is the higher earner and more likely to be in an

⁴² They suggest that the location trends delineated by Costa and Kahn are due to higher rates of power couple formation in larger metropolitan areas. They also note that the trend of increasing concentration of power couples in larger metropolitan areas did not continue between 1990 and 2000.

occupation that is geographically clustered.

6.4 Gender Roles and Gender Identity

Recent work by Bertrand, Kamenica, and Pan (2013) points to possible far-reaching effects of adherence to traditional gender roles on the relative outcomes of men and women. They draw on Akerlof and Kranton's (2010) development of the concept and implications of identity, defined as a sense of belonging to a social category, combined with a view about how people who belong to that category should behave. Departures from these norms are perceived as generating costs and hence people seek to avoid them.

Bertrand, Kamenica, and Pan probe the consequences of the view that a wife should not earn more than her husband and find it to impact a number of outcomes. For example, they find that, within marriage markets, as the probability that a randomly chosen woman would outearn a randomly chosen man increases, marriage rates decline. Similarly, couples in which the wife outearns her husband have lower rates of marital satisfaction and are more likely to divorce. Of particular relevance to the issues under consideration here, they find that, in couples in which the wife's potential income is likely to exceed her husband's (based on predicted income), the wife is less likely to be in the labor force and, if she does work, her income is lower than predicted. Such a selection pattern would lower the observed relative wages of employed married women. Also of interest, given the inverse relationship between housework and wages, they find, based on time use surveys, that the gender gap in nonmarket work is *increased* if the wife earns more than her husband. This finding is particularly surprising given that Beckerian notions of comparative advantage would lead us to expect the opposite (Becker 1981, enlarged edition

1991),⁴³ assuming that relatively higher earning women do not generally have even higher relative values of nonmarket time. A possible interpretation of this pattern is that these high earning wives are attempting to compensate for violating the gender norm of earning more than their husbands. As we have seen, greater housework time is expected to negatively affect wages.

The findings from Bertrand, Kamenica, and Pan suggest that additional explorations of gender norms by economists would be fruitful in understanding the gender wage gap and other gender differences in outcomes. This paper may also be seen as part of recent efforts, discussed below, to understand the potential contribution of gender differences in attitudes and preferences to gender differences in outcomes. However, while the findings in this paper are striking, it is possible that the strength of this norm may be diminishing. First, the share of married couple families in which the wife outearns her husband has been growing steadily, as married women's labor force participation and education levels have increased and the male-female wage gap has declined. For example, this share increased by over 80 percent between 1988 and 2012, both among families in which both members of the couple had earnings (from 15.9 to 29.0 percent) and among married couples overall (from 8.2 to 15.4 percent).⁴⁴ Moreover, there is evidence that in the bulk (60 percent) of couples in which the wife outearns her husband, this disparity is relatively permanent—that is, it persists over a three-year period (Winkler, McBride, and Andrews 2005). Second, attitudes seem to be becoming more permissive along this dimension. A 2013 attitude survey found that only 28 percent of adults agreed that “It’s generally better for a

⁴³ That is, the division of labor in the family should be determined by the comparative advantage of each spouse in market vs. nonmarket activity.

⁴⁴ U.S. Census Bureau, Historical Income Tables-Families, “Table F-22. Married-Couple Families with Wives' Earnings Greater Than Husbands' Earnings: 1988 to 2012,” <http://www.census.gov/hhes/www/income/data/historical/families/>, accessed June 26, 2014. Similarly, Wang, Parker, and Taylor (2013) report that, among married couples with children under age 18, the share of families in which the mother earned more than the father increased from 3.8 percent in 1960 to 22.5 percent in 2011.

marriage if the husband earns more than his wife” compared to 40 percent in 1997. College graduates had especially permissive views, with only 18 percent supporting this view (Wang, Parker, and Taylor 2013). While an adherence to traditional gender roles strong enough to influence behavior in the manner delineated by Bertrand, Kamenica, and Pan (2013) need not be conscious and overt, it is nonetheless of interest that such views, as expressed, are becoming more permissive. Moreover, this has been occurring at the same time the share of couples where the wife outearns her husband has been increasing; this points to the likelihood that couples are acting on their more permissive views and also to the possibility that behavior (the increasing incidence of such families) influences norms and attitudes as well as the reverse.

7. Psychological Attributes/Noncognitive Skills

Labor economists have become increasingly interested in the effect of noncognitive or “soft” skills—including psychological attributes, preferences, and personality—on labor market outcomes and behavior (Heckman and Kautz 2012). This trend has been driven by a number of factors but perhaps most important is that, although considerable evidence supports the importance of traditional economic variables in explaining labor market behavior and outcomes, there is almost always a sizeable component of any behavior or outcome that is not explained by these variables, leading researchers to reach out beyond the confines of traditional economic models for explanations. With respect to gender, intriguing findings suggest a number of psychological attributes that differ between women and men. For example, women have been found to be less willing than men to negotiate and compete and to be more risk averse (for reviews, see, Bertrand 2011; Croson and Gneezy, 2009). Gender differences in these characteristics have been proposed as an explanation for women’s lower wages and lower representation in high-level jobs.

In considering research on gender differences in psychological attributes or noncognitive skills, some cautions must be borne in mind. First, even if men and women do differ on average, it is not possible at this point to know the role of nature versus nurture. We do not attempt to address this fundamental issue here, however, we consider it important that research suggests social factors play a part and have highlighted such findings. Moreover, whatever their origin (nature or nurture), gender differences may still be malleable—so, for example, women may be encouraged to negotiate and given tips on improving their negotiating skills. Second, gender differences in all psychological attributes do not necessarily favor men. For example, there is some evidence that women have better interpersonal or “people” skills than men (Borghans, ter Weel, and Weinberg 2014). Another area where differences favor women is that, as we saw in our discussion of education, the greater behavioral problems of boys appear to contribute to their lower rate of college going. Also, it should be noted that a particular psychological attribute—like men’s willingness to compete or lower risk aversion—may be an advantage in some settings but a disadvantage in others.⁴⁵ In addition, as we shall see below, the same trait may be rewarded differently for men and women, or indeed even be penalized for women when it is rewarded for men.

Finally, much of the evidence on gender differences in psychological attributes has been gleaned from laboratory experiments and there are reasonable concerns about generalizing the results of such experiments outside the lab. And, while confirmation of lab results in the field is suggestive, even in this case, there may be questions about how well the experiment represents what would occur in a real world setting (Harrison and List 2004, and Pager 2007). Moreover,

⁴⁵ For example, Eckel and Füllbrunn (2015) provide experimental evidence from a financial asset market that female traders are less likely to produce speculative price bubbles.

importantly, findings from laboratory or field experiments generally cannot be easily translated into accounting for a particular portion of the gender wage gap. Studies based on survey questions in data sets that include information on respondents' attitudes and preferences along with other characteristics and labor market outcomes are more promising in this regard but elicit their own sets of concerns about endogeneity and precisely what it is (i.e., what particular trait or traits) one is really measuring.

Capitalizing on two excellent recent reviews (Bertrand 2011; Croson and Gneezy, 2009), we discuss this work selectively. And, in light of the above cautions, we particularly focus on research that contributes to our understanding of the applicability and broader significance of the findings from lab experiments, as well as on research that sheds light on the role of social factors in producing the observed gender differences.

7.1 Evidence from Survey Data

While, as noted above, much of the research on the role of psychological traits comes from experimental data, a few studies have used survey data to examine the impact of gender differences in traits and attitudes on the gender wage gap in a wage regression context. For example, Fortin (2008) examines the role of a number of noncognitive traits on the wages of young workers and finds that women place greater emphasis on people/family and men on money/work. She finds that gender differences on these and other traits but particularly the importance placed on money/work have a modest but significant role in accounting for the gender wage gap. Importantly, these factors are measured prior to labor market exposure so as to be more exogenous to the actual rewards offered each group in the labor market. Manning and Swaffield (2008) consider the impact of a large number of psychological factors, also measured prior to labor market entry, for young men and women in the United Kingdom. They also

conclude that such psychological factors play a modest role, accounting for at most 4.5 log points out of an almost 25 log point gap that opened up between young men and women ten years after labor market entry. Finally, as noted above, not all gender differences in noncognitive factors analyzed in survey data favor men in their relationship to wages. Borghans, ter Weel, and Weinberg (2014) present evidence that there is a female advantage in “people skills” and find evidence of a people skills premium in wages.

Just as the contribution of gender differences in the traditional human capital variables may change over time, thus helping to account for the decline in the gender wage gap, so may the impact of noncognitive traits. In this regard, it is interesting that Fortin (2008) finds some convergence in gender differences in noncognitive skills, with the role played by noncognitive factors shrinking in a more recent compared to an earlier cohort. Further, Borghans, terWeel, and Weinberg (2014) find evidence of a growing importance of interpersonal interactions (in part due to increased computer use) in affecting wages that can help explain rising female relative wages

While such findings are informative in elucidating some of the possible omitted factors that lie behind the unexplained gap in traditional wage regressions, the coefficients on noncognitive skills in a wage equation cannot necessarily be given a causal interpretation. Both wages and attitudes, for example, may be determined by the same exogenous factor(s). And, as in the case of the traditional productivity proxies discussed above, there may be important feedback effects from differential treatment in the labor market (and the anticipation of such differential treatment) to noncognitive traits. So, for example, gender differences in the importance placed on money may influence wages through negotiating behavior or effort, but the source of women's lower emphasis on money could be, at least in part, anticipation of lower income due to

labor market discrimination. Finally, in analyses based on survey data, there is likely to be some ambiguity as to precisely what one is measuring—might responses indicating say leadership skills be picking up the effect of some unmeasured factor like self-confidence? For these reasons, just as research on labor market discrimination (discussed below) has tended to move towards experimental evidence, at least in confirming findings based on statistical analyses of survey data, there has been a parallel development in studying the impact of psychological characteristics. We move to a consideration primarily of experimental evidence in the next sections.

7.2 Negotiation

Researchers have found that men's and women's average propensity to negotiate differs, with women being much less likely to do so (Babcock and Laschever 2003; see also reviews in Bertrand 2011; and Croson and Gneezy, 2009). Women's lower propensity to negotiate over salaries, raises, or promotions, could reduce their pay relative to men's. The gender difference could reflect social factors, including women being socialized to feel that they are being pushy or overbearing (unfeminine) if they negotiate—i.e., pursue their own goals in the face of conflict with others (Babcock and Laschever 2003). Consistent with the notion that the female gender role is seen as incongruent with negotiating, a meta-analysis by Mazei et al (forthcoming) found that gender differences in negotiating outcomes were reduced when negotiators negotiated on behalf of *another individual*. Moreover, women may have likely learned that their negotiating can trigger a negative response from others. For example, in a series of laboratory experiments, Bowles, Babcock, and Lai (2007) asked study participants to evaluate managers based on a transcript or a video of a job placement interview. They found that participants were disinclined to work with female managers who negotiated for higher compensation but that negotiating had

little effect on their evaluation of male managers.

Results from a field experiment by Leibbrandt and List (forthcoming) confirm the gender differences in negotiating behavior obtained in the lab studies but suggest that such differences may be sensitive to the cues given. In examining the response of applicants to job advertisements, they found that men were more likely to negotiate than women when there was no *explicit* statement that wages were negotiable. However, when it was explicitly stated that wages were negotiable, the gender difference disappeared and even reversed.

While it may be possible to enhance women's negotiating skills and reduce the gender difference in negotiating, it is also important to realize that there are limitations to what may be achieved by doing so. Negotiation is a form of bargaining and as such the outcome is influenced by the alternatives available to the individual. To the extent that women face discrimination in the labor market that lowers their wages relative to men's, their expected outcome from the bargaining process will be smaller than for their male counterparts. Moreover, to the extent that, as we have seen, women who negotiate elicit negative responses compared to men, the gender difference in the prospective result from negotiating is further widened.

7.3 Competition

There is evidence from laboratory experiments that, on average, men are more competitively inclined than women (Bertrand 2011; Croson and Gneezy 2009). In Niederle and Vesterlund's (2007) influential study, for example, subjects were given a task (adding up sets of two-digit numbers) for which there was no average gender difference in performance. Subjects received feedback on their own performance but not on their performance relative to others. When subsequently given a choice between a noncompetitive compensation scheme (a piece rate—pay according to the number of problems correctly solved) and a competitive

compensation scheme (a tournament where only the highest scorer out of a group of 4 was compensated), men overwhelmingly (73 percent) selected the tournament while only a minority (35 percent) of the women did so. Low performing men actually competed more than high performing women. Interestingly, while high-scoring women lost out financially by shying away from competition, low-performing men competed too much from a payoff-maximizing perspective. The gender difference in attitudes towards competition could be a disadvantage for women in the labor market, potentially lowering their relative pay and leading them to avoid certain occupations or business settings, although these findings also suggest that men may sometimes compete more than is optimal.

An interesting recent study suggests that differences in attitudes toward competition observed in the lab do translate into differences in career choices. Buser, Niederle, and Oosterbeek (2012) collected data on the competitiveness of high school students in the Netherlands through in-class experiments and then tracked their subsequent education choices across four study profiles at age 15. While boys and girls had very similar levels of academic ability, boys were substantially more likely than girls to choose the more prestigious profiles. The authors found that up to 23 percent of the gender difference in profile choice could be attributed to gender differences in competitiveness, as assessed by the in-class experiments.

Some evidence that women shy away from competitive environments is also indicated by a recent large-scale field experiment. Flory, Leibbrandt and List (forthcoming) randomly assigned job-seekers into viewing advertisements with different compensation schemes. Consistent with the results of lab experiments, the more heavily the compensation package tilted towards rewarding the individual's performance relative to a coworker's performance, the more the applicant pool shifted to being more male dominated. However, there was little or no gender

difference when compensation was *only slightly* (rather than heavily) based on performance relative to a coworker's or when the job was to be compensated based on team (rather than individual) relative performance. Moreover, the sex-type of the job mattered. The occupation under study was administrative assistant. A male-oriented ad described tasks focused around sports. The "female" ad was similar in other respects but the focus was general—the authors deemed this a female-type job because, nationally, administrative assistant is a predominantly female occupation (79 percent female in 2001). Strikingly, there were no gender differences in propensity to apply under any of the compensation schemes for the female treatment—the gender differences described above were only obtained for the male-type job. While it would have been interesting to see results for a completely neutral occupation, these findings suggest a strong interaction between the gender role or identity of the task and men's and women's propensity to compete. Moreover, while individual responses to compensation schemes were not correlated with readily observable characteristics like education and experience, a blind analysis of the quality of interview questionnaire responses suggested that the highly competitive regime disproportionately attracted low-ability males. As the authors note, this is consistent with Niederle and Vesterlund's (2007) finding that "males compete too much" in terms of maximizing monetary payoffs.

While much of this evidence does indeed suggest that men are, on average, more attracted to competitive environments than women are, what are the effects of this difference on the gender pay gap? Manning and Saidi (2010) note that the kind of laboratory evidence we have discussed implies that women will be less attracted than men to jobs in which pay is contingent on individual performance. Using the British Workplace Employment Relations Survey for 1998 and 2004, they find that women were indeed less likely to have jobs with pay

for performance than men. However, this gender difference accounted for only a very small portion of the British pay gap overall and among managerial workers. Thus, the impact of gender differences in competitiveness on the gender pay gap based on this evidence appears to be very limited.

Finally, also of interest is a study that compared the results of lab experiments testing for gender differences in preferences for competition in two different cultures (Gneezy, Leonard and List 2009). The findings of this study strongly suggest that men's and women's attitudes towards competition are influenced by broader social factors. The authors found that, consistent with the previous results in developed countries, men opted to compete at roughly twice the rate of women in a traditional patriarchal society (the Maasai of Tanzania). However, in a matrilineal/matrilocal society where inheritance and residence are determined by the female lineage (the Khasi of India), women chose the competitive environment more often than men.

There is also some evidence that competition increases the *relative* performance of men compared to women when both participate in the activity, although the evidence on this is more mixed (Croson and Gneezy 2009). On the one hand, Gneezy, Niederle and Rustichini (2003), for example, found no significant difference in performance by gender under piece rates for a maze solving task on the computer. However, when pay was competitive, men's performance was increased significantly and women's stayed the same, yielding a gender difference. On the other hand, Niederle and Vesterlund's (2007) study discussed above found that the performance of both men and women improved similarly under the tournament and that there was still no gender difference in performance.⁴⁶

Some particularly compelling evidence on the impact of competition on performance is

⁴⁶ See also the review of recent studies in Örs, Palomino, and Peyrache (2013).

presented in a recent study by Örs, Palomino, and Peyrache (2013). The authors examined gender differences in performance for the same group of subjects on real-world academic achievement examinations that differed in their levels of competition. They found that men performed better than women on the highly competitive entrance exam for admission to the Master of Science in Management at the École des Hautes Études Commerciales (HEC) in Paris even though, for the same cohort, women performed significantly better than men on the national baccalauréat exam two years prior, which the authors characterize as “noncompetitive.” Moreover, among the subset admitted to HEC, women outperformed the same males in first year grades in nonmathmatically-oriented classes (where grades are based on relative performance only in a very loose sense).

7.4 Risk Aversion

Based on the laboratory experiments they review, Croson and Gneezy (2009) report that women are, on average, more risk averse than men.⁴⁷ All else equal, occupations with more variable earnings are expected to pay a compensating wage differential to induce workers to accept the higher levels of risk. To the extent women are more likely to avoid such jobs, women’s greater risk aversion could lower their earnings relative to men (Bertrand 2011). Risk aversion could also plausibly affect job performance in particular occupations, such as money managers.

Interestingly, Croson and Gneezy (2009) report that, while women are found to be more risk averse among persons drawn from the general population or among university students, studies that focus on managers and professionals have found little or no evidence of gender

⁴⁷ A review and analysis by Nelson (forthcoming) finds the results to be more mixed, with some studies reporting higher female average risk taking and many cases in which the male advantage lacked statistical significance.

differences in financial risk preferences. For example, one study of mutual fund managers found that funds managed by men and women did not differ in risk or performance. Similarly, male and female managers and entrepreneurs displayed similar risk propensities. It is not possible to know whether such findings are due to the type of selection we have just discussed (with more risk-taking people of both sexes choosing to enter or remain in particular fields) or learning (people who initially differ in their risk propensities may learn from their professional environment). In either case, however, these findings suggest that while women's relative aversion to risk may lower their relative earnings due to occupational sorting, this factor probably does not help to explain *within* occupational earnings differences (or at least not within the occupations studied). Further, to the extent these findings are due to learning, it suggests that these preferences can be shaped by environment.

7.5 Gender Differences in Agreeableness

There is also evidence of gender differences in what are known in psychology as the “big five” personality traits. These traits are openness to experience, conscientiousness, extroversion, agreeableness, and neuroticism.⁴⁸ One of the most consistent gender differences has been found for agreeableness, with women being found to be more agreeable than men (Bertrand 2011). Agreeableness refers to being more trusting, straightforward, altruistic (warm), compliant, modest, and sympathetic.

Perhaps not surprisingly given labor market realities, a study that examined the return to agreeableness in a regression context found that men earned a premium for being *disagreeable*. However, this attribute was *not* found to be related to women's wages. Thus, the gender difference in agreeableness contributed to the gender earnings gap both because men were

⁴⁸ For a definition of each trait, see Mueller and Plug (2006).

considerably more disagreeable than women, but also because only men were rewarded for this trait (Mueller and Plug 2006).⁴⁹ Judge (2012) found broadly similar results regarding the effect of agreeableness. Although, in this case, both men and women were rewarded for being disagreeable, the reward was *much larger* for men than women. Taken together, the results of both studies show a gender gap in the rewards to this personality trait and hint at a double bind for women. As in the case of negotiation, women face potential penalties for not engaging in this behavior but, if they do, may elicit negative or less positive responses than men.

8. Occupations, Industries, and Firms

In this section we consider empirical evidence on the extent and dimensions of employment segregation by sex. The results in Section 2 indicate that, while the share of the gender wage gap due to human capital (education and experience) has declined noticeably, the share accounted for by locational factors like occupation and industry has increased from 27% of the 1980 gap to 49% of the much smaller 2010 gap. Moreover, occupational upgrading by women contributed to the narrowing of the gap over this period, though much of this effect was offset by adverse (to women) movements in returns to occupations. The firm dimension, not accessible in data sets like the PSID and CPS that were used above, has also been shown to be important. Finally, gender differences in representation across the hierarchies within occupations, as particularly emphasized in discussions of the glass ceiling, constitute another dimension of employment differences that is also generally not captured by these data sets, at

⁴⁹ One must of course be cautious in basing conclusions about the contribution of gender differences in individual regression coefficients (on either categorical or continuous variables) to the gender pay gap in isolation from the other coefficients in the model (Oaxaca and Ransom 1999). Additional results were that men were also rewarded for emotional stability and openness to experience, while women were rewarded for conscientiousness and openness to experience.

least directly. Indirectly, some light on this may be shed by quantile regression analyses focusing at the top, as illustrated by our estimates in Section 2.

Of these dimensions of employment differences, occupational differences between men and women have received the most attention. Gender differences in occupations have been and continue to be striking, although they have declined significantly since 1970. In terms of general outlines, in 1970, women were considerably more concentrated than men in administrative support and service occupations, and a bit more highly represented in professional jobs overall, and particularly in predominantly female professions like teaching and nursing. Men were considerably more likely to be in managerial jobs and much more concentrated than women in blue collar occupations, including relatively high-paying craft and skilled positions. They were also considerably more likely than women to be in predominantly male professions like law, medicine, and engineering. Since 1970, women have reduced (but not eliminated) their overrepresentation in administrative support and service jobs and made significant inroads into management and male professions. There has been little change in gender differences in representation in blue collar jobs.⁵⁰

Beyond broad occupations, the Census provides information on some 500+ detailed occupational classifications. The widely used segregation index developed by Duncan and Duncan (1955) provides a useful summary measure of this data. This index may be interpreted as the percentage of females (or males) who would have to change jobs for the occupational distribution of women and men to be the same, with a value of 0 indicating no segregation and a value of 100 indicating complete segregation. Early work suggested little change in the extent of

⁵⁰ This discussion is based on Blau, Ferber and Winkler (2014), Chapter 7; 1970 occupational data were converted into Census occupational categories for 2000 using a crosswalk developed in Blau, Brummund and Liu (2013a).

occupational segregation prior to 1970 (Gross 1968, Jacobs 1989). Starting in 1970, there was considerable progress in reducing the extent of occupational segregation (Beller 1982, Bianchi and Rytina 1986). A useful summary of the 1970-2009 trends is provided by Blau, Brummund, and Liu (2013 a and b) based on a comparable set of Census occupational categories for 2000 .⁵¹ They report that the index was 64.5 in 1970 and fell to 51.0 by 2009, a sizable decline from an extremely high initial level. However, the index declined at a diminished pace over the decades, falling by 6.1 points over the 1970s and 4.3 points over the 1980s, but only 2.1 points over the 1990s and just 1.1 points (on a decadal basis) over the 2000s. Blau, Brummund and Liu (2013a and b) also report that trends differed across educational groups. More substantial progress was made by highly educated women, who succeeded in moving into formerly male managerial and professional occupations. Gains were smaller for less-educated women, reflecting the lack of progress in integrating male blue-collar occupations.

While the overall decline in the segregation was substantial, the 51 percent figure for the index in 2009 indicates that occupational differences between men and women remain large. A sizable literature indicates that female occupations pay less than male occupations for workers with similar measured characteristics (e.g., Levanon, England, and Allison 2009).⁵² Our estimates in Section 2 imply that occupational differences can explain (in an accounting sense) one third of the gender wage gap in 2010. This estimate includes controls for actual labor market

⁵¹ Their findings are similar to earlier studies for overlapping periods, where available.

⁵² Early studies highlighting the empirical importance of occupational and in some cases industry differences in explaining the gender wage gap include, Fuchs 1971; Blinder 1973; Oaxaca 1973; and Sawhill 1973. For examples of early studies examining the effect of percent female in the occupation on earnings, see Sorensen (1990) and Macpherson and Hirsch (1995); there is also a wide literature in sociology examining this issue, see Levanon, England, and Allison (2009) for a review. More recently some research suggests that “care work”—occupations in which “concern for the well-being of others is likely to affect the quality of services provided”—may pay less *ceteris paribus* (for a review see Folbre 2012, quotation is from p. 66). Women are disproportionately represented in such jobs.

experience but is based on only 21 occupations. Nonetheless, it is very similar to Goldin's (2014) estimate for a number of samples (based on education and labor force attachment) based on the American Community Survey (2009-2011) using the full set of three-digit occupations, but with no control for actual experience (which is not available in the ACS). Our results in Table 4 also indicate that occupation is the largest single factor accounting for the gender pay gap, with the second being industry (15 categories) at 18 percent. Taken together occupation and industry differences account for over one half of the gender wage gap. There has been less focus on industry differences in explaining the gender wage gap.

Another related dimension of employment differences between men and women that has also gotten less attention, perhaps in part due to data limitations, is gender differences in the distribution of employment by firm. An early study by Blau (1977) presented evidence of high levels of employment segregation of men and women by firm *within* narrowly-defined occupational categories and their relationship to gender wage differentials within occupations. She developed a model in which employer tastes for discrimination against women à la Becker (1971, orig. pub. 1957) are widespread, but the ability to exercise them is constrained by the firm's position in the wage hierarchy, which is determined by a variety of institutional and market forces and could not easily be altered to accommodate employer discriminatory preferences (comparable to the notion of firm effects today). Consistent with this model, she found women were concentrated in firms that paid lower wages to *both* men and women across all occupations, and conversely men tended to be employed at the firms which paid higher wages to both sexes. Subsequent work confirmed the continued importance of differences in the distribution of employment across firms in accounting for overall gender wage differences (Groschen 1991; Bayard, Hellerstein, Neumark, and Troske 2003).

With the growing availability of matched firm-worker data, the firm dimension is becoming an increasingly active area of research for understanding the gender wage gap. For example, recent work has considered the role of monopsony in explaining the gender wage gap. A number of studies (discussed in greater detail below) find, consistent with a role for monopsony, that women have lower labor supply elasticities to the firm than men. One of these studies, Webber (forthcoming), uses matched firm-worker data and reports that women's lower labor supply elasticities are primarily due to cross-firm, rather than within firm, differences in elasticities, suggesting a reason why firms that disproportionately employ women tend to be lower paying overall. As another example, a recent study by Card, Cardoso, and Kline (2014), using Portuguese firm-worker data, investigates the relative importance of sorting across firms (i.e., women's greater likelihood of working at low wage firms) and within firm bargaining (with women receiving less of the premium men receive in working for high-wage firms) in explaining the gender wage gap. They find evidence that both factors play a role.

Finally, not only do men and women tend to work in different occupations, they also tend to be employed at different levels of the hierarchy within occupations. This is the case in a number of arenas, ranging from business to academia to unions. So, for example, recent data on Fortune 500 companies indicate that, although women are nearly half of managers, they comprise only 14.3 percent of executive officers, and 3.8 percent (19) of CEOs, and hold just 16.6 percent of board seats.⁵³ Or, in law, women are less likely than men to be employed as partners in large firms (over 50)—as was true for 26 percent of male compared to 14 percent of

⁵³ Data on executive officers are for 2012 and for CEOs and boards of directors for 2011. See, *2012 Catalyst Census: Fortune 500 Women Executive Officers and Top Earners*, available at <http://www.catalyst.org/> (accessed December 11, 2012); Catalyst, *Catalyst Pyramid: U.S. Women in Business* (New York: Catalyst, 2012), available at www.catalyst.org (accessed December 21, 2012); and “*2012 Catalyst Census: Fortune 500 Women Board Directors*,” available at <http://www.catalyst.org/> (accessed December 11, 2012).

female 1979-1985 graduates of the University of Michigan Law School fifteen years after graduation (Noonan, Corcoran, Courant 2005). Similarly, in 2012 only 15 percent of AFL-CIO executive council members were women.⁵⁴ And, as a final example, in academia, the female share decreases as we move up the ranks—from assistant professors (61 percent) to associate professors (50 percent) to full professors (28 percent) (Blau, Ferber, and Winkler, Chapter 7).

In all these cases, it is difficult to determine whether the scarcity of women at the top is simply due to the fact that women are relative newcomers and it takes time to move up through the ranks (the “pipeline” argument) or whether it represents particular barriers to women’s advancement (i.e., a “glass ceiling”). Moreover, a lower representation of women at higher levels could be due to discrimination or subtle barriers facing women but could also reflect greater work-family conflicts for women that reduce their productivity and/or interest in high level positions.

Nonetheless, there are indicators that at least some of the gender difference reflects discrimination. For example, a number of studies (e.g., Blau and DeVaro 2007, Cobb-Clark 2001, and McCue 1996), find that women are less likely to be promoted, all else equal, although some do not (e.g., Hersch and Viscusi 1996). For academics, some studies find lower probabilities of promotion for women, even after accounting for indicators of qualifications like number of publications, although results differ by field and gender differences appear to have diminished in recent years (Ginther and Kahn, forthcoming and Ceci, Ginther, Kahn, and Williams 2014).⁵⁵

The possibility of discrimination is further suggested by studies in both the corporate world (Bell

⁵⁴ “About AFL-CIO,” at www.aflcio.org, accessed March 2012. The American Federation of Labor-Congress of Industrial Organizations (AFL-CIO) is the largest union federation in the United States.

⁵⁵ Focusing on the most recent research, evidence of a *ceteris paribus* female shortfall in promotions is found for economics and the life sciences, but not for other social sciences and natural sciences.

2005, Shin 2012, Kurtulus and Tomaskovic-Devey 2012) and academia (Ehrenberg, Jakubson, Martin, Main, and Eisenberg 2012) which find that women at the lower ranks fare better (in terms of representation or wages) when women are more highly represented at the higher ranks.⁵⁶

Whatever the sources of these differences, research suggests they can have substantial consequences for gender wage differences. For example, our own data analyses in Section 2 indicated that gender wage gaps at higher levels of the wage distribution were larger and fell more slowly than at lower levels. And, as we noted, this result appears in line with other research both in the United States and abroad. As another example, Bertrand and Hallock (2001) of gender differences in pay among the five highest-paid executives in S&P 1500 firms found that the 2.5 percent of executives in their sample who were women earned 45 percent less than their male counterparts. Female executives were younger and thus had less seniority, and this contributed to the gender difference. However, three-quarters of the gender pay gap was due to the fact that women managed smaller companies and were less likely to be the CEO, chair, or president of their company.

9. Evidence on Labor Market Discrimination

To the extent that gender differences in outcomes are not fully accounted for by productivity differences due to gender differences in human capital and other supply-side sources, models of labor market discrimination offer an explanation. Theoretical work in this area was initiated by Becker's (1971, orig. pub. 1957) model of racial discrimination. Becker

⁵⁶ A recent paper by Bertrand, Black, Jensen, and Lleras-Muney (2014) examining the effects of corporate board quotas for women in Norway found that the reform increased the representation of women on corporate boards and reduced their pay gap relative to male board members. They also found evidence suggestive of a growing representation of female employees at the very top of the firms' income distribution (top 5 highest earners), but not elsewhere in the firms' income distribution (i.e., they found no evidence of "trickle down" below the top 5 earners).

conceptualized discrimination as a taste and analyzed three cases: those in which the discriminatory tastes were held by employers, co-workers, and customers or clients. Under certain circumstances, such discrimination will cause a wage differential between men and women. Discriminatory employers will only hire women at a sufficient wage discount that compensates them for the disutility of employing women. Discriminatory male workers will demand a wage premium to work with women thus raising men's relative wages, and the reluctance of discriminatory customers or clients to buy goods or services provided by women will make women less productive in terms of revenue brought in, thus depressing women's relative wages.

Becker (1971, orig. pub. 1957) and others (e.g., Arrow 1973) have pointed out that competitive forces should reduce or eliminate employer discrimination in the long run because the least discriminatory firms, which hire more lower-priced female labor, would have lower costs of production and should drive the more discriminatory firms out of business. One answer, suggested initially by Becker himself, is that discrimination will be located in sectors of the economy that are not competitive.

While Becker emphasized monopolistic elements in the product market, a related approach targets monopsonistic power on the part of the employer in the labor market (e.g., Madden 1973; Black 1995).⁵⁷ Monopsony could help to explain how discriminatory gender wage differences arise and persist if employers wield greater monopsony power over men than women workers. For this explanation of the wage gap to hold, women's supply of labor to the *firm* must be less wage elastic than men's, giving employers greater monopsony power over

⁵⁷ See Manning (2003) for a systematic development of the “new monopsony” literature and its application to the gender wage gap among other issues.

women than men. This might seem counter-intuitive at first, in that there is clear evidence that women have a *larger* own-wage elasticity of labor supply to the *labor market* than men, although in the United States the gender difference has been decreasing as women's elasticities have declined in magnitude since 1980 (Blau and Kahn, 2007; Heim, 2007). However, a variety of factors could still potentially result in women having a smaller responsiveness to wage changes at the firm level. Perhaps the most intriguing possibility is discrimination itself. Black (1995) develops a model in which search costs give employers a degree of monopsony power. If there is discrimination against women, women will face higher search costs than men, increasing employers' monopsony power over them.

In addition, models of statistical discrimination (Phelps 1972) were later developed, in part to explain the persistence of discrimination in the long run in the face of competitive forces. Such models assume uncertainty and imperfect information; thus differences between groups in the expected value of productivity or in the reliability with which productivity may be predicted may result in differences in the treatment of members of each group. As a consequence, firms may pay women less, exclude them from jobs requiring substantial firm-specific training, or deny them promotions (for promotions, see Lazear and Rosen 1990).

It has been argued that such statistical discrimination (making decisions on the basis of the average characteristics of the group) is consistent with profit maximization and can thus persist in the face of competitive forces. However, Aigner and Cain (1977) contend that such models are no more convincing in explaining the persistence of discrimination than models based on tastes. To the extent that employers' views are correct, the lower expected productivity of women will reduce their wages but women as a group will be paid their expected productivity. This does not constitute labor market discrimination as economists define it, i.e., pay differences

that are not accounted for by productivity differences. Moreover, they argue that when employer beliefs regarding average differences are erroneous, discrimination clearly exists but discrimination based on such misperceptions is even less likely to persist in the long run than discrimination based on tastes. Finally, although they acknowledge that less reliable predictions of a group's productivity combined with risk aversion by employers could produce a discriminatory differential, a perfectly elastic supply of risk neutral entrepreneurs would be expected to erode discriminatory differentials based on this factor.

However, such models may have greater usefulness in explaining the persistence of gender differentials in pay than Aigner and Cain's analysis suggest. Consider first the situation where employer perceptions are correct, is it appropriate to consider this a form of 'discrimination' in any sense? From a normative perspective, the answer may be yes, to the extent that basing employment decisions on a characteristic like sex—a characteristic that the individual cannot change—could be viewed as inequitable. Indeed, the practice of judging an individual on the basis of group characteristics rather than upon his or her own merits seems the very essence of stereotyping or discrimination. Such behavior is certainly not legal, for example, under antidiscrimination laws and regulations.

Now consider the situation where employer perceptions are incorrect. If statistical discrimination is accompanied by feedback effects, this may be a credible source of persistent discriminatory pay differences (Arrow 1973; Lundberg and Startz 1983). For example, if employers incorrectly expect that women are more likely to quit their jobs, they may respond by giving women less firm-specific training or assigning them to dead-end jobs. Faced with fewer incentives to remain on the job, women may respond by exhibiting exactly the higher turnover that employers expect.

Further insight on the persistence of discrimination is suggested by what Bertrand, Chugh and Mullainathan (2005) have termed implicit discrimination. This is based on findings from social psychologists that discriminatory attitudes and stereotyping may be unconscious (e.g., Fiske 1998)—suggesting that they would not be easily eliminated. Indeed, as gender discrimination has become less socially acceptable, it has likely become less overt and more subtle, as well as unconscious. Finally, as our discussion of statistical discrimination above suggests, discrimination can adversely affect women's human capital investments and labor force attachment by lowering the market rewards to this behavior—i.e., through feedback effects (e.g., Weiss and Gronau 1981).

Models based on tastes for discrimination are also consistent with employment segregation, but do not necessarily predict it will occur. If wages are flexible, it is possible that discrimination will result in lower pay for women, but produce little or no segregation. However, if discriminatory tastes against women in traditionally male pursuits are both strong and prevalent, women may tend to be excluded from these areas. If such segregation does occur, it may or may not be associated with gender pay differentials. In the presence of sufficient employment opportunities in the female sector, equally qualified women may earn no less than men. The relationship between occupational segregation and earnings differentials in an otherwise competitive setting is clarified in Bergmann's (1974) overcrowding model. If potentially equally qualified men and women are segregated by occupation, the wages in male and female jobs will be determined separately by the supply and demand for labor in each sector. Workers in male jobs will enjoy a relative wage advantage if the supply of labor is more abundant relative to demand for female than for male occupations.

9.1 Statistical Evidence on the Unexplained Gap

One strand of empirical research on labor market discrimination uses regression methods and versions of the Blinder-Oaxaca decomposition (Blinder 1973, Oaxaca 1973) to calculate unexplained female wage shortfalls (i.e., a wage gap not accounted for by gender differences in measured characteristics) as estimates of the extent of discrimination. For example, in Section 2, we presented results on the unexplained gap for a number of years based on PSID data. We found an unexplained gender wage gap in each year, although the magnitude of the gap had declined over time. The finding of such an unexplained gap is fairly standard in the literature (for reviews, see e.g., Altonji and Blank 1999, Stanley and Jarrell 1998, and Hersch 2006). Such an unexplained or residual wage gap is often taken as an estimate of labor market discrimination. However, as is well known, such estimates are suggestive, but not conclusive. Discrimination is overstated if men have higher levels of unmeasured productivity (or poorer working conditions). On the other hand, if women are better endowed with unmeasured characteristics on average, as may be the case with some variables, like people skills as we have seen, regression methods would understate discrimination. The unexplained gap will also understate discrimination if some of the explanatory variables such as experience, occupation, industry or union status have themselves been influenced by discrimination—either directly through the discriminatory actions of employers, co-workers or customers or indirectly through women’s responses to lower returns, i.e., through feedback effects. For these reasons, the literature has moved in the direction of research designs that use various strategies to overcome the problems of traditional statistical studies. For example, some studies use samples of men and women such as lawyers or MBAs in which samples are more homogeneous and the controls for qualifications are much more detailed than in commonly-used databases such as the CPS or the PSID. Presumably

omitted-variable biases are less severe in such homogeneous samples.⁵⁸ In addition, experimental research, such as audit studies, tests for discrimination under circumstances where, by construction, men and women have identical qualifications. Finally, we will briefly consider the small number of studies that have tested other predictions of Becker's (1971, orig. pub. 1957) discrimination model to see whether or not the results are consistent with discrimination.

9.2 Evidence on Subgroups

As discussed above, studies applying the same statistical techniques as labor-market wide studies, but focusing on more homogeneous groups of workers like lawyers and MBAs may provide more convincing evidence of labor market discrimination. In addition, given their data sources, they are able to control for detailed characteristics (e.g., grade point averages while in school), not available in broader studies. We have already considered such studies above and found that they also provide deeper insights into the supply-side sources of gender differentials, particular the important role of hours worked and workforce interruptions in demanding professions. Here we focus on their implications for estimates of discrimination.

One qualification that must be made however in interpreting the results of such studies for this purpose is that, when we focus on specific occupations, we introduce an additional element of selection, beyond selection into the labor force discussed above. The direction of such selection is unclear a priori, however it seems reasonable to us that, when we focus on high-level, traditionally male oriented professions, women may be a positively selected group relative to men. If this is the case, then studies of such occupational subgroups will understate the extent of discrimination.

⁵⁸ While we believe the findings of such studies are instructive for studying discrimination, we do not mean to imply that the Bertrand, Goldin, and Katz (2010) study which we reference was designed for this purpose.

The studies of lawyers (Noonan, Corcoran, and Courant 2005) and MBAs (Bertrand, Goldin, and Katz 2010) referenced earlier find that, even if one accounts for variables related to family status, like work force interruption and fewer hours worked, unexplained gender earnings differences remain which are potentially due to discrimination, although they are of course susceptible to other explanations. In the law study, men earned 11 percent more, controlling for an extensive list of worker qualifications and other factors, including grades while in law school, detailed work history data, and information about type and size of employer. In the MBA study, men earned nearly 7 percent more even accounting for work force interruptions, fewer hours worked, and gender differences in business school GPAs and finance courses taken.

There has also been research analyzing gender differences in the most mathematically-intensive academic fields (geoscience, engineering, economics, mathematics/computer science and the physical sciences). The findings of this literature have recently been reviewed by Ceci, Ginther, Kahn and Williams (2014). These results are mixed, with some studies finding little gender salary gap in these fields once experience and productivity are controlled for, while others find that a male salary premium persists even after controlling for these factors.

9.3 Experimental Evidence

Given the problems with traditional statistical studies, researchers have been interested in uncovering alternative sources of evidence on discrimination. As noted above, one approach that provides particularly persuasive evidence of discrimination is experiments, either naturally occurring labor market events that may be seen and analyzed as if they were experiments or actual experiments in which the researcher manipulates the treatment so as to test for discrimination, either in the laboratory or in the field. The advantage of experimental studies is that they offer estimates of the role of discrimination that are potentially less contaminated by

unmeasured factors. This is a rapidly growing research approach and we illustrate the findings by a selection of studies that impart the flavor and show the breadth of these findings.

The first study we consider is Goldin and Rouse's (2000) investigation of the impact of the natural experiment created when symphony orchestras began to adopt "blind" auditions for musicians in which a screen is used to conceal the identity of the candidate. They found that the adoption of the screen substantially increased the probability that a woman would advance out of preliminary rounds and be the winner in the final round. The switch to blind auditions was found to explain one quarter of the increase in the percentage female in the top five symphony orchestras in the United States, from less than 5 percent of all musicians in 1970 to 25 percent in 1996.

A second study, Neumark (1996), was a field experiment or hiring audit. Male and female pseudo-job seekers were given similar résumés and sent to apply for jobs waiting on tables at the same set of 65 Philadelphia restaurants. The results provided statistically significant evidence of discrimination against women in high-priced restaurants (where earnings of workers are generally higher). In these restaurants, a female applicant's probability of getting an interview was 40 percentage points lower than a male's and her probability of getting an offer was 50 percentage points lower.

A third experimental study, a field experiment by Moss-Racusin, Dovidio, Brescoll, Graham, and Handelsman (2012) sheds light on possible bias in academic science. Science faculty from the fields of biology, chemistry, and physics at six large, research-intensive universities (three public and three private) were asked to provide feedback on the application materials of (fictitious) senior undergraduate students who they were told ultimately intended to go to graduate school and had recently applied for a science laboratory manager position. The

faculty participants believed they were evaluating a real student who would subsequently receive the faculty participants' ratings as feedback to help his or her career development. Participants were randomly assigned to receive an application either from a female student (Jennifer) or a male student (John); all other information was otherwise set up to be identical. Faculty participants rated the male applicant as significantly more competent and suitable for the position than the (identical) female applicant. Participants also set a starting salary for male applicants that was almost \$4,000 higher than the salary offered to female applicants, and offered more career mentoring to the male applicants. Female faculty were equally likely to exhibit bias against the female students as male faculty.

A fourth study, by Reuben, Sapienza, and Zingales (2014), implemented a laboratory experiment where some subjects (employers) hired other subjects (applicants) to perform an arithmetic task that, on average, men and women perform equally well. Their findings are consistent with negative stereotyping of women in math-related areas. They found that when employers had no information about applicants other than appearance (which makes sex clear), both male and female employers were twice as likely to hire a man as a woman. The discrimination (sex differential) was similar when applicants self-reported their expected performance, largely because men tended to overestimate future performance (women also slightly underestimated theirs)—and employers did not correct for this. Gender discrimination in hiring was reduced, but not eliminated (i.e., women were still under-hired), when employers were provided with full information about applicants' previous performance on the task. One very interesting feature of this study is that subjects were given the Implicit Association Test (IAT), a computer-based behavioral assessment designed to measure implicit or unconscious

gender stereotyping or bias.⁵⁹ They found that that IAT scores were correlated with the initial bias in sex-related beliefs (when employers only knew the sex of the applicant) and with a bias in updating expectations when performance information was self-reported (i.e., not sufficiently correcting for male overestimation). While, as we have noted, discrimination against women persisted even when information about applicants' previous performance was available, the extent of such discrimination was not correlated with IAT score.

Finally, we point to the results of the study by Correll, Bernard, and Paik (2007) summarized above that suggests that women, but not men, face discrimination based on their parental status. Using both laboratory and field experiments, the authors found that the participants had less favorable views regarding the résumés of equally-qualified mothers relative to those of nonmothers, while fathers were not disadvantaged relative to nonfathers. Such a finding suggests discrimination against women based on parental status.

9.4 Evidence on Predictions of Discrimination Models

As we have seen, Becker (1971, orig. pub. 1957) and others (e.g., Arrow 1973) have pointed out that competitive forces should reduce or eliminate employer discrimination in the long run because the least discriminatory firms, which hire more lower-priced female labor, would have lower costs of production and should drive the more discriminatory firms out of business. For this reason, Becker suggested that discrimination would be more severe in firms or sectors that are shielded to some extent from competitive pressures. Consistent with this reasoning, Hellerstein, Neumark and Troske (2002) found that, among plants with high levels of product market power (and hence the ability to discriminate), those employing relatively more

⁵⁹ See, Greenwald, McGhee, and Schwartz (1998). Subjects took the version measuring the association between sex and science-related abilities.

women were more profitable. Similarly, Black and Strahan (2001) reported that, with the deregulation of the banking industry beginning in the mid-1970s, the gender wage gap in banking declined. (Deregulation was viewed as increasing competitiveness within the industry.) And Black and Brainerd (2004) found that increasing vulnerability to international trade (i.e., increased competitive pressure) reduced apparent gender wage discrimination in concentrated industries, again as predicted by the Becker model. In a similar vein, Heyman, Svalerty, and Vlachos' (2013) study based on Swedish worker-firm matched data found evidence that a firm takeover was associated with a reduction in the gender wage gap. They interpret takeovers as a manifestation of competitive pressure.

There is also some evidence consistent with statistical discrimination against women, based on employers' difficulty in distinguishing more from less career oriented women. So, for example, Gayle and Golan (2012) propose a model in which workers have private information on their costs of participating in the labor force. They show that this asymmetric information is quantitatively important in explaining of the gender pay gap. Similarly, Thomas (2015) proposes a model which shows that if there is asymmetric information about worker's future labor participation, the imposition of mandated maternity leave policies can increase the gender gap in promotion. This is because such policies make it more difficult for employers to distinguish between more and less family-oriented women, since they disproportionately raise post-birth employment by the former. Consistent with the model, she presents evidence that the Family and Medical Leave Act of 1993 increased women's probability of remaining employed but lowered their probability of promotion and that information asymmetry played a role in producing this result.

Finally, as we discussed above, greater monopsony power of employers over women than

men workers provides a possible mechanism for the existence and persistence of a discriminatory gap. This requires greater elasticity of labor supply to the firm for men than women. Evidence on gender differences in labor supply elasticities at the firm level for the United States is mixed. On the one hand, using aggregate labor market data, Viscusi (1980), Blau and Kahn (1981), and Light and Ureta (1992) all find that women's quit rates are at least as wage responsive as men's; Manning (2003) too finds no evidence of lower female separation elasticities in data for the United States and the United Kingdom. On the other hand, Ransom and Oaxaca (2010) report some evidence consistent with the monopsony model as an explanation for gender wage differentials at a chain of grocery stores, as do Ransom and Simms (2010) for schoolteachers in Missouri. Moreover, using economy-wide linked employer-employee data, Webber (forthcoming) finds evidence of lower labor supply elasticities for women. Internationally, Barth and Dale-Olsen (2009) and Hirsch, Schank, and Schnabel (2010) find evidence using matched employer-employee data that men's turnover is more wage-elastic than women's in Norway and Germany, respectively.

10. Evidence on the Impact of Policy

Women's relative skills and the degree of employer discrimination can be affected by equal employment opportunity laws and regulations, as well as by government policies directed at the difficulties of combining work and family. In this section, we briefly consider what is known about these types of policies and their impacts, focusing primarily on the United States.

The United States was a world leader in implementing equal employment opportunity policy as the first economically advanced nation to pass and implement antidiscrimination laws and regulations (Blau and Kahn 1996b). The centerpiece of the government's antidiscrimination

activities is Title VII of the Civil Rights Act of 1964 which broadly bans discrimination by sex (as well as by race, religion and national origin) in virtually all aspects of the employment relationship, including hiring and firing, training, promotion, wages, and fringe benefits and covers all businesses employing 15 or more workers. Title IX, an important amendment to the Civil Rights Act passed in 1972, prohibits sex discrimination in most educational institutions. In addition, the Equal Pay Act of 1963 mandates equal pay for men and women who do substantially equal work in the same establishment. Further, under some circumstances, affirmative action, or “pro-active steps ... to erase differences between women and men, minorities and nonminorities, etc.” (Holzer and Neumark, 2000, p. 484), is also required, primarily for government contractors under an Executive Order promulgated in 1965 and amended to include women in 1967. Affirmative action has also been voluntarily adopted by many employers.

In thinking about the impact of the government’s antidiscrimination enforcement effort, one question that arises is whether the time path of the increase in women’s relative earnings (see Figure 1) appears compatible with an effect of these laws and regulations. This question arises because we see no indication of a notable improvement in women’s relative earnings in the immediate post-1964 period that might be attributable to the effects of the government’s antidiscrimination effort; the gender pay ratio remained basically flat through the late 1970s or early 1980s, after which it began to increase. In contrast, blacks experienced considerable increases in their relative earnings in the decade following the passage of the civil rights laws that many scholars attribute, at least in part, to the impact of the antidiscrimination laws (e.g., Donohue and Heckman 1991).

Nonetheless, there is some evidence from a variety of detailed, micro-level studies of a positive effect of government equal employment opportunity policies on women's earnings and occupations. Beller (1979, 1982) used enforcement activity as an indicator of the strength of government sanctions under Title VII and found evidence of improvements over the 1967-1974 period in women's relative earnings (Beller 1979) and their probability of being employed in a predominantly male occupation (Beller 1982). Carrington, McCue and Pierce (2000) took firm size as an indicator of coverage and enforcement and found that, over the 1963-87 period, the relative employment of women and blacks by larger employers increased. Kurtulus (2012) found that the share of women and minorities in high-paying skilled occupations grew more over the 1973–2003 period at federal contractors than other employers. These gains took place primarily prior to or in the early years of the Reagan Administration and after 1991; a pattern that matches what is known about climate of enforcement of affirmative action and antidiscrimination laws more broadly, including a winding down of the enforcement effort during the Reagan years. Kurtulus' (2012) findings are consistent with an earlier study by Leonard (1990), which found faster employment growth for black and white females at contractor establishments over the 1974-80 period. Finally, Holzer and Neumark (1999 and 2000 b) measured affirmative action by employer self-reports (this could include both mandated and voluntary programs) and found cross-sectional evidence that affirmative action reallocates women and minorities to the affirmative action sector by increasing both their applications and employment. This is likely to raise their relative wages since the authors find that such firms are higher paying and, in addition, have smaller race and sex differences in wages (see also Holzer and Neumark 2000a for a review).

We find these results of female gains due to equal employment policy not entirely implausible, despite the time pattern of aggregate female relative earnings gains, for at least two reasons. First, we note that some improvements in women's status do indeed date to the 1970s—chiefly, the growth in women's enrollments in professional schools and the beginning of a substantial decline in occupational segregation. The educational shifts may reflect, at least in part, the impact of Title IX, but also a response to perceived increases in labor market opportunities that improved the incentives for women to train for nontraditional jobs. Since occupational segregation by sex was considerably more pronounced than by race (Fuchs 1988 and Jacobsen 1994), such occupational shifts may have been more necessary for women than for blacks to reap the gains from the government's antidiscrimination efforts, resulting in a greater lag in the impact of the government's equal employment opportunity policies on women's relative earnings. Second, these laws and regulations were rolled out during a period of extremely high growth in female labor supply; the negative wage effects of this expansion in labor supply could have camouflaged an otherwise positive effect of the government's efforts. On the other hand, it is puzzling that the largest female relative wage gains and the strongest evidence of a decline in the unexplained gender wage gap were during the 1980s (see Section 2 and Section 11), when the government's antidiscrimination effort was noticeably scaled back.

Turning to work-family policy, we focus on parental leave, although we note that there are a wide range of other possible policies, including child care that might be considered. The Family and Medical Leave Act (FMLA) of 1993 mandates that eligible workers be allowed to take up to 12 weeks of unpaid leave for birth or adoption, acquiring a foster child, illness of a

child, spouse, or parent, or their own illness.⁶⁰ (Firms may voluntarily provide longer and/or paid leave.) Workers are entitled to their jobs upon returning from the leave. While such policies may encourage firm-specific investments, thus raising women's relative wages (since parental leave is much more likely to be taken by women than men), they may also encourage labor force withdrawal for longer periods of time than otherwise, reducing women's accumulation of experience. Mandated leaves, particularly of long duration, may also diminish women's opportunities by increasing employer costs of hiring women and hence providing incentives to discriminate against them. Mandated leaves might also reduce women's relative wages to finance the benefit (e.g., Gruber 1994). Thus, the effect of parental leaves on the gender wage gap is theoretically ambiguous. Empirical evidence for the United States suggests that the effect of the FMLA has been modest; it has been found to have a small positive effect on employment and no effect on wages (Baum 2003 and Waldfogel 1999). Results are broadly similar for California's introduction of 6 weeks of paid leave (with a replacement ratio of 55 percent) in 2004. Employment probabilities in the post-leave period were increased; the effect on wages was not statistically significant (see, Baum and Ruhm 2013). A recent study by Thomas (2015) does however suggest that FMLA increased the gender gap in promotion.

Since provision of parental leave in the United States is considerably less generous (in both duration and payment) than in other economically advanced countries, international comparisons may shed light on potential effects of more generous leave policies. In a study of 9 Western industrialized countries, Ruhm (1998) found that female earnings were unaffected by rights to short parental leaves, while longer leaves (more than 5 or 6 months) lead to reductions

⁶⁰ In addition, the Pregnancy Discrimination Act of 1978 (an amendment to Title VII of the Civil Rights Act) prohibits employers from discriminating against workers on the basis of pregnancy.

in women's relative wages. These findings are consistent with results from Blau and Kahn (2013a), which found that the greater expansion of family-friendly policies in other economically advanced countries than in the United States between 1990 and 2010 increased female labor force participation in these countries relative to United States, but was associated with a lower likelihood of women having full-time jobs or working as managers or professionals. (The mean duration of leave in these other countries was 57 weeks in 2010, up from 37 weeks in 1990.) Taken together, these results suggest that a number of offsetting factors may be at work, with a little impact on wages for shorter leaves and a negative effect dominating for long periods of mandated parental leave.

11. Wage Structure, Demand and Institutions

Much research on the gender pay gap focuses on gender differences in qualifications, including experience, or treatment by firms (discrimination). In addition, however, men and women work in a world economy in which labor market prices, such as the returns to education or experience, are affected by larger forces of supply and demand as well as labor market institutions. We now consider research that studies the impact of these larger economic forces on the gender pay gap.

A useful starting point is a key insight of Juhn, Murphy and Pierce (1991), a study of black-white wage differentials, that the overall wage structure can affect the relative wages of specific groups. By "wage structure," we mean the returns that the labor market offers for various skills and for employment in various industries or occupations. For example, countries with strong unions that raise the wages of less-skilled workers tend to have a relatively compressed wage structure, while, in the United States, wages are determined in a more decentralized manner, resulting in a more dispersed wage structure. The wage structure can also

change over time as rewards to skills and premiums for employment in high-wage occupations and industries increase or decrease.

Both the human capital and discrimination explanations of the gender pay gap suggest a potentially important role for wage structure in determining how women fare relative to men across countries or over time. We illustrate by some examples focused on the temporal dimension. For example, despite important recent gains, women still have less experience than men, on average. If the labor market return to experience rises over time, women will be increasingly disadvantaged by their lesser amount of experience. In addition, both the human capital and discrimination models suggest reasons why women are likely to be employed in different occupations and perhaps in different industries than men. This implies that an increase in the returns to employment in “male” occupations or industries will also place women at an increasing disadvantage. In fact, the patterns of rising overall wage inequality in the labor market, particularly in the 1980s, resulted from precisely such increases in the market rewards to skill and to employment in high-wage male sectors (Blau and Kahn 1997). This means that women as a group were essentially “swimming upstream” in a labor market growing increasingly unfavorable for workers with below-average skills—in this case, below-average experience—and for workers employed in disproportionately female occupations and industries. Yet the 1980s were precisely the time period in which women made the largest gains.

11.1 U.S. Evidence on the Impact of Wage Structure on the Gender Wage Gap

How were U.S. women able to swim upstream and narrow the gender wage gap in the face of economy-wide forces working against them? Blau and Kahn (1997 and 2006) found that this was the outcome of two broad sets of countervailing factors. On the one hand, working to decrease the gender wage gap, women increased their qualifications relative to men and, in the

1980s, the unexplained gender gap also narrowed substantially. On the other hand, working to widen the gender wage gap, particularly during the 1980s, were changes in wage structure (or returns to characteristics) that favored men over women. Of particular importance were a rise in the return to experience and increases in returns to employment in occupations and industries where men are more highly represented. The sizable increase in the supply of women over the 1980s is another factor that likely worked to widen the gender wage gap as well. The decrease in the gender wage gap occurred because the factors favorably affecting women's wages were large enough to more than offset the impact of unfavorable shifts in returns and increasing female labor supply.

However, the matter may be more complicated than a simple decomposition of the trends would suggest. While rising demand for skill did shift labor market prices in a way that worked against women on net in the 1980s, the underlying labor market demand shifts that widened overall wage inequality appear to have favored women relative to men in certain ways. Thus these demand shifts likely also contributed to a decrease in the unexplained gender gap identified in Blau and Kahn (1997 and 2006) and Section 2. Overall, manufacturing employment declined, particularly in the 1980s. In addition, some evidence indicates that technological change produced within-industry demand shifts that favored white-collar relative to blue-collar workers in general. Given that men have tended to hold a disproportionate share of manufacturing and blue-collar jobs, these shifts would be expected to benefit women relative to men (Berman, Bound and Griliches 1994; Blau and Kahn 1997 and 2006). Further, evidence suggests that increased computer use favors women's wages compared to men (Krueger 1993; Weinberg 2000; Welch 2000; Autor, Levy and Murnane 2003; Beaudry and Lewis 2014). This may reflect women's greater comparative advantage in cognitive relative to manual or motor skills ("brains"

versus “brawn” to borrow Welch’s (2000) terminology).⁶¹ Moreover, Borghans, ter Weel and Weinberg (2014) present evidence that interpersonal interactions have become more important with the spread of computers. Since women’s interpersonal skills tend to exceed men’s, on average, this factor worked to increase women’s wages relative to men’s (Borghans, ter Weel and Weinberg 2014).

Finally, our earlier work and the data analysis shown in Tables 1-5 show that the gender pay gap closed much more slowly after 1990 than during the 1980s. Some evidence for the importance of demand shifts in causing this slowdown comes from Blau and Kahn (2006), who find that demand shifts related to industries and occupations favoring women were smaller in the 1990s than in the 1980s. Moreover, Borghans, ter Weel and Weinberg (2014) find that the growth in the demand for interpersonal skills was faster in the 1980s than in the 1990s. In both of these studies, the slowdown in demand shifts favorable to women coincided with the slowdown in gender wage convergence.

11. 2 International Comparative Evidence on the Impact of Wage Structure on the Gender Wage Gap

As mentioned earlier, many other countries have far more centralized wage-setting institutions than the United States, resulting in a far higher degree of wage compression. Centralized collective bargaining tends to reduce wage differentials through the negotiation of relatively high wage floors, which raise the relative wages of those near the bottom of the distribution, including women (Blau and Kahn, 1996a). In countries such as those of Continental Europe, unions cover a much larger portion of the labor market than in the United States, and

⁶¹ Bacolod and Blum (2010) present evidence that there has been an increase in the labor market return to cognitive skills and a corresponding decrease in the return to motor skills. This has likely benefited women relative to men since women tend to be more highly represented in occupations where cognitive skills are important while men are more likely to be in jobs that emphasize motor skills.

wage-setting is much more centralized. For example, among a group of four Scandinavian countries (Denmark, Finland, Norway and Sweden) and the four largest Continental European countries (France, Germany, Italy and Spain), as of 2010, 62 to 91 percent of workers were covered by collective bargaining, with a median of 82 percent; in contrast, only 13 percent of U.S. workers were covered by unions (unpublished data used in OECD 2012). The gender pay gap was tended to be smaller in these European countries than in the United States. Specifically, as of 2012, the gender pay gap in these countries ranged from 6 to 19 percent, with a median of 11 percentage points; however, the OECD reported a 2012 U.S. gender pay gap of 19 percent (OECD 2014, p. 288).

The impact of wage compression on the gender pay gap has been studied formally in research that examines international differences as well as within-country changes in the gender pay gap. Much of this research uses methods developed by Juhn, Murphy and Pierce (1991)—JMP—to study changes over time in black-white differentials in the United States. These methods were first used to study international differences in the gender pay gap by Blau and Kahn (1992 and 1996b). In the context of the gender pay gap, these methods decompose differences or changes in the gender pay gap into portions that are gender-specific (i.e., gender differences in measured characteristics and the size of the unexplained gap adjusted for residual inequality) and portions that are related to the overall wage structure (i.e., differences in the returns to measured characteristics and the prices of unmeasured skills). For example, Blau and Kahn (1996b) compared the US gender wage gap in the late 1980s with that in nine other countries (Australia, Austria, West Germany, Hungary, Italy, Norway, Sweden, Switzerland and the United Kingdom). They addressed a puzzle. U.S. women compared favorably with women in the other countries on several measures of skills relative to men, and, moreover, as we noted

earlier, the United States had a longer commitment to policies of equal pay and equal employment opportunity for women than the other countries. Yet the gender pay gap in the United States ranked above average for their sample of countries. The fact that US women had better measured characteristics and placed higher in the distribution of male residuals than women in the other countries, suggested that gender-specific factors could not explain the higher US gender wage gap. However, wage structure (prices) had large effects raising women's relative wages in the other countries compared with the United States; indeed, wage structure was more than sufficient to account for the full amount of the difference between the US gender wage gap and that in other countries. Blau and Kahn (1992; 1996b) interpreted this pattern as reflecting the impact of international differences in labor market institutions. Kidd and Shannon (1996) also found an important role for wage structure in explaining the substantially lower gender wage gap in Australia than in Canada for 1989–90. Their results are consistent with Blau and Kahn's (1992; 1996b) in that Australia has much higher coverage by collective bargaining than Canada.

While the studies surveyed so far all find that overall wage compression lowers the gender pay gap, a recent study using JMP methods finds only a small role for wage structure in explaining differences in the 2002 gender pay gap across nine countries: Italy, Spain, Portugal, the Netherlands, Norway, the Czech Republic, Latvia, Slovakia, and Lithuania (Simón 2012). The author also finds that across this set of countries, collective bargaining coverage is positively associated with gender pay gap, and that the overall 90-10 wage gap is negatively associated with the gender pay gap (although both correlations were statistically insignificant). These atypical results appear related to the particular countries included in the study and in particular the high representation in his sample of former Communist countries. The gender pay gap in these countries has behaved differently from that in Western countries (generally rising in

contrast to the falling gender pay gap in Western countries), and below, we discuss research on changes in the gender pay in countries making a transition from Communism to a market-oriented system.⁶²

In some instances, countries' labor market institutions have changed dramatically and this provides an opportunity to study the impact of relatively exogenous changes in the overall wage structure on the gender pay gap. For example, in Sweden between 1968 and 1974, the trade union movement engineered a major compression of wages. Using the JMP decomposition, Edin and Richardson (2002) found that this change in wage structure contributed to a reduction of the gender wage gap during this period. Datta Gupta, Oaxaca, and Smith (2006) used similar methods to study changes in the Danish gender wage gap between 1983 and 1995. This was a period of increased decentralization of the wage determination process, a development that would be expected to lead to a more dispersed wage structure and therefore a rising gender wage gap. The authors indeed found that the gender wage gap in Denmark increased during this period and that most of the increase can be accounted for by the widening overall wage structure.

One of the most dramatic changes in the world over the last 25 years has been the fall of Communism. In former Soviet Bloc countries and in China, highly centralized wage-setting institutions with considerable wage compression were replaced with market-oriented, decentralized wage setting. These changes in institutions may be expected to widen the gender pay gap. Several studies of changes in the gender pay gap in such countries, using JMP

⁶² See also European Commission (2006, p. 31) which found that, while there was no overall relationship between the aggregate 90/10 wage ratio and the gender pay gap in the EU in 2002, which included formerly Communist countries, for the pre-enlargement sample of EU countries (the EU-15), there was indeed a positive relationship between wage inequality and the gender pay gap. Moreover, in contrast to Simón, Kahn (2014) found for a larger sample (30-31 OECD countries), a significantly positive correlation between the male 50-10 wage gap and the gender pay gap and a significantly negative association between collective bargaining coverage and the gender pay gap.

decomposition methods, do indeed find such results. For example, Brainerd (2000) studied the Czech Republic, Hungary, Poland, Russia, the Slovak Republic, and Ukraine and found in each case that after the fall of Communism, the wage structure became more dispersed and raised the gender pay gap. Moreover, Orazem and Vodopivec (2000) found similar results for Slovenia after the fall of Communism there, although there was little effect of the changing wage structure on the gender pay gap in Estonia. Finally, Zhang, Han, Liu and Zhao (2008) studied the gender gap in China over the 1988-2004 period, one in which the labor market became much less centralized as China's economy became much more market oriented. The authors used JMP-style methods to find that the resulting spread in the wage structure raised the gender pay gap considerably.

As discussed in detail by Blau and Kahn (2006), the JMP decomposition assumes that male wage regression coefficients and male residuals are relevant indicators of the (measured and unmeasured) prices facing women in the labor market. Some support for this assumption is provided by the fact that wage coefficients and residual distributions have changed similarly for men and women over time in the United States and that men's and women's wage coefficients and residual distributions are similar to each other within countries at a point in time (Blau and Kahn, 2002). But it is possible to directly test whether male wage compression leads to a smaller gender wage gap, and Blau and Kahn (2003) have done so by compiling a microdata-set for 22 countries over the 1985–94 period. Looking across countries, they find that the gender wage gap is positively affected by a country's male skill prices (that is, the level of male wage inequality adjusted for measured characteristics), as well as by the relative net supply of women (that is, supply net of demand). A likely interpretation is that more compressed male wages are an indicator of smaller wage differentials in general, as suggested above in our discussion of

centralized wage-setting institutions. Bolstering this interpretation is the authors' further finding that, other things equal, greater coverage by collective bargaining reduces the gender wage gap. It thus appears that high wage floors negotiated by unions do serve to lower the gender wage gap.

If firms take labor costs as given, then high union-negotiated wage floors should lower female relative employment. And this is precisely what Bertola, Blau and Kahn (2007) find in a study of relative employment in 17 countries over the 1960–96 period. Specifically, they find that greater coverage by highly centralized unions lowers female employment and raises female unemployment compared with men's. Moreover, an additional institutional feature of many European labor markets is employment protection legislation, which makes it very expensive to fire workers. These systems are believed to reduce worker discharges but also to discourage future hiring due to anticipated firing costs (Bertola 1992). The benefits of such systems are likely to accrue to incumbent workers, while new entrants may pay a cost in reduced availability of permanent jobs. Since women are more likely to be new entrants than men are, these systems of employment protection are not likely to have gender-neutral effects.

As a result of employment protection of permanent jobs, many European countries such as Spain, Finland and Sweden have a relatively high incidence of temporary jobs, which provide an alternative. Temporary jobs are typically lower-paying and of course less secure than permanent jobs (Boeri 2011; Kahn forthcoming). More stringent permanent employment protection legislation is expected to lead to a higher incidence of temporary jobs because it is much less costly to fire workers from them (Blanchard and Landier 2002), and women are expected to be disproportionately affected. Consistent with this expectation, Kahn (2007) found, using microdata from several countries, that more stringent permanent employment protection

legislation led to higher concentration of women in temporary jobs. This discussion of wage floors and employment protection suggests that such labor market interventions may protect employed workers in permanent jobs, many of whom may be women; however, it is also women who disproportionately pay the costs these institutions entail.

12. Conclusion

We have shown that the gender pay gap in the United States fell dramatically from 1980 to 1989, with further smaller reductions through 2010. Using PSID microdata, we documented the improvements over the 1980-2010 period in women's education, experience and occupational representation, as well as the elimination of the female shortfall in union coverage, and showed that they played an important role in the reduction in the gender pay gap. Particularly notable is that, by 2010, conventional human capital variables (education and labor market experience) taken together explained little of the gender wage gap in the aggregate. This is due to the reversal of the gender difference in education, as well as the substantial reduction in the gender experience gap. On the other hand, gender differences in location in the labor market—distribution by occupation and industry—continued to be important in explaining the gap in 2010. A decrease in the unexplained gap over the 1980s contributed to the robust convergence in the gender wage gap over that decade, with the unexplained gap falling sharply from 21-29% in 1980 to 8-18% by 1989. However, the unexplained gap did not fall further, remaining in this range over the succeeding 20 years. We also found that both the raw and the unexplained gender pay gap declined much more slowly at the top of the wage distribution than at the middle or the bottom. By 2010, the raw and unexplained female shortfalls in wages, which had been fairly similar across the wage distribution in 1980, were larger for the highly skilled than for others, suggesting that developments in the labor market for executives and highly

skilled workers especially favored men. Research on these labor markets suggests a particularly important role for work force interruptions and shorter hours in explaining gender wage gaps high skilled occupations than for the workforce as a whole, although the interpretation of these findings in a human capital framework has been challenged. Goldin (2014), for example, argues that they more likely represent wage penalties for temporal flexibility.

We then provided a discussion of the causes of women's improvements in measured skills, emphasizing the remarkable reversal of the gender gap in college attendance as well as women's increasing commitment to the paid labor force. In light of the persistent unexplained pay gap, we then discussed recent research on gender differences in factors that standard data sets cannot measure, or which have not been the focus of conventional wage gap studies. We considered the ways in which conventional gender roles and gender identity as well as the presence of children, can contribute to the gender wage gap. We also examined evidence on gender difference in mathematics test scores and noncognitive skills such as gender differences in attitudes toward competition, negotiation, and risk aversion. While male advantages in some of these factors may help to explain not only some of the unexplained gender wage gap but also gender differences in occupations and fields of study, women may have advantages in some noncognitive skills, like interpersonal skills. Unfortunately data on these factors are generally not available in a form that permits an accounting of their role in explaining the aggregate gender wage gap. Thus it is not possible at this point to attribute a specific portion of the gender wage gap to these factors. Whether or not these currently unmeasured factors would help to explain the gender pay gap, discrimination is still a possible factor, and we cited some recent research based on experimental evidence that strongly suggests that discrimination cannot be discounted as contributing to the persistent gender wage gap. Indeed, we noted some experimental evidence

that discrimination against mothers may help to account for the motherhood wage penalty as well. Finally, we reviewed research that finds that, given men's and women's differing skill levels and locations in the economy (by occupation, industry, and firm), overall labor market prices can have a significant effect on the gender wage gap. In particular, the more compressed wage structures in many Continental European countries have served to lower the gender pay gap there relative to the United States, with its lack of wage-setting institutions that would reduce gender wage differentials.

In our view, there are many unanswered questions that we believe could be fruitfully studied by those interested in the gender pay gap. First, while there are gender differences in some psychological attributes/noncognitive skills, more work is needed to confirm these differences outside the laboratory setting where much of the research has been focused, although we have reviewed some recent studies that have done so. As mentioned earlier, there is also relatively little research that would enable us to determine the quantitative importance of these differences for the gender wage gap. Moreover, we found evidence that these gender differences can themselves be affected by social context and thus might not be independent causes of the gender pay gap in the first place. Further probing of these questions would be very useful. Second, field experiments arguably provide credible exogenous variation in the economic environment facing workers as well as real-world settings, and these will likely continue to provide insights into gender differences in preferences, behavior, and labor market outcomes. Third, the continued importance of gender differences in employment by industry and occupation, as well as by firm, suggest the fruitfulness of research aimed at better understanding the sources of these gender difference as well as their consequences. The growing availability of matched firm-worker data should facilitate such research. Finally, some innovative policies have

been developed regarding issues of work and family, including parental leave entitlements that incentivize fathers' leave taking (Dahl, Løken, and Mogstad 2013; Patnaik 2015). The long run impact of these policies on gender and the labor market as well as the division of labor within the family is an important research topic.

Data Appendix

The analysis in Tables 1-7 is based on microdata taken from the indicated waves of the PSID and the March CPS. The PSID is the only data source which has information on actual labor market experience for the full age range of the population. However, because the PSID only supplies this work history information for family heads and spouses/cohabitators, it does not cover adults who are living with relatives, such as grown children living at their parents' house. In addition, the PSID's base sample began with roughly 5,000 families from 1968, when immigrants were a much smaller portion of the population. This means that the current PSID sample, which consists of these original families plus splitoffs, undercounts immigrants today. For these reasons, we also show data from the CPS which are more representative of the whole U.S. population.

We focus on men and women age 25-64 who were full time, non-farm, wage and salary workers and who worked at least 26 weeks during the preceding year. This age group has, for the most part, left school, allowing us to abstract from issues of combining work and school attendance. Limiting the top of the age range to 64 to some degree abstracts from normal retirement issues (patterns were very similar when we limited the sample to ages 25-54). In addition, by limiting our sample to those who worked full time and had at least 26 weeks of work in the prior year, we are focusing on those with a relatively strong labor market commitment. This sample restriction leads to a relatively homogeneous sample with respect to this commitment, allowing us to reach more accurate conclusions about the prices women and men face in the labor market. We exclude the self-employed and those in agriculture on the grounds that it is difficult to separate labor income from capital income or income in kind for these groups. Our basic dependent variable is the log of average hourly earnings, which we compute in the PSID by dividing annual labor earnings by annual hours worked and then take the log; when we use the CPS, we divide annual wage and salary earnings by annual hours worked.

For early years of the PSID, separate values for wage and salary income and self-employment/farm income are not available for wives. In earlier work (Blau and Kahn 2004) we showed that this omission didn't have an important effect on average hourly earnings among household heads, a group for which we had data on wage and salary earnings. While the PSID does not topcode earnings, the CPS does. To adjust for this in the CPS, we multiplied the topcoded value by 1.45. (In each year, less than 2% of the sample was topcoded.) In both data sets, we exclude those earning less than \$2/hr in 2010 dollars, using the Personal Consumption Expenditures deflator (taken from www.bea.gov). This cutoff equals 28-38% of the real Federal minimum wage across our sample period (see <http://www.bea.gov/iTable/iTable.cfm?ReqID=9&step=1#reqid=9&step=3&isuri=1&904=1980&903=4&906=a&905=2014&910=x&911=0>, accessed August 19, 2014 and <http://www.dol.gov/whd/minwage/chart.htm>, accessed August 19, 2014). We experimented with other cutoffs, including a flat \$3/ hour in 2010 dollars, as well as using 50% of each year's real minimum wage as a cutoff. The results were very similar to those presented here.

References

- Acemoglu, Daron and Jörn-Steffen Pischke. 1999. "The Structure of Wages and Investment in General Training." *Journal of Political Economy* 107 (3): 539-573.
- Aigner, Dennis J. and Glen G. Cain. 1977. "Statistical Theories of Discrimination in Labor Markets." *Industrial & Labor Relations Review* 30 (2): 175-87.
- Akerlof, George A. and Rachel E. Kranton. 2010. *Identity Economics: How Our Identities Affect Our Work, Wages, and Well-being*. Princeton: Princeton University Press.
- Albrecht, James, Anders Björklund, and Susan Vroman. 2003. "Is There a Glass Ceiling in Sweden?" *Journal of Labor Economics* 21 (1): 145-177.
- Altonji, Joseph G., Rebecca M. Blank. 1999. "Race and Gender in the Labor Market." In *Handbook of Labor Economics*, Volume 3c, edited by Orley C. Ashenfelter and David Card, 3143-3259. Amsterdam: North-Holland.
- Altonji, Joseph G. and Charles R. Pierret. 2001. "Employer Learning And Statistical Discrimination," *Quarterly Journal of Economics* 116 (1): 313-350.
- Altonji, Joseph G. and James R. Spletzer. 1991. "Worker Characteristics, Job Characteristics, and the Receipt of On-the-Job Training." *Industrial & Labor Relations Review* 45 (1): 58-79.
- Arrow, Kenneth. "The Theory of Discrimination." 1973. In *Discrimination in Labor Markets*, edited by Orley Ashenfelter and Albert Rees, 3-33. Princeton, NJ: Princeton University Press.
- Arulampalam, Wiji, Alison L. Booth, and Mark L. Bryan. 2007. "Is There a Glass Ceiling Over Europe? Exploring the Gender Pay Gap Across the Wage Distribution." *Industrial & Labor Relations Review* 60 (2): 163-186.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality in the 1990s: Revising the Revisionists." *Review of Economics and Statistics* 90 (2): 300-23.
- Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118 (4): 1279-1333.
- Babcock, Linda and Sara Laschever. 2003. *Women Don't Ask: Negotiation and the Gender Divide*. Princeton: Princeton University Press.
- Bacolod, Marigee and Bernardo S. Blum. 2010. "Two Sides of the Same Coin: U.S. 'Residual Inequality' and the Gender Gap." *Journal of Human Resources* 45 (1): 197-242.
- Bailey, Martha J. 2006. "More Power to the Pill: The Impact of Contraceptive Freedom on Women's Lifecycle Labor Supply." *Quarterly Journal of Economics* 121 (1): 289-320.
- Bailey, Martha J., Brad Hershbein, and Amalia R. Miller. 2012. "The Opt-In Revolution? Contraception and the Gender Gap in Wages." *American Economic Journal: Applied Economics* 4 (3): 225-254.

- Baum, Charles L. 2003. "The Effects of State Maternity Leave Legislation and the 1993 Family and Medical Leave Act on Employment and Wages." *Labour Economics* 10 (5): 573–96.
- Baum, Charles L. and Christopher Ruhm. 2013. "The Effects of Paid Family Leave in California on Labor Market Outcomes." NBER Working Paper 19741.
- Barron, John and Dan A. Black. 1993. "Gender Differences in Training, Capital and Wages." *Journal of Human Resources* 28 (2): 342–64.
- Barth, Erling and Harald Dale-Olsen. 2009. "Monopsonistic Discrimination, Worker Turnover, and the Gender Wage Gap." *Labour Economics* 16 (5): 589–597.
- Bayard, Kimberly, Judith Hellerstein, David Neumark, and Kenneth Troske. 2003. "New Evidence on Sex Segregation and Sex Difference in Wages from Matched Employee-Employer Data." *Journal of Labor Economics* 21 (4): 887–923.
- Beaudry, Paul and Ethan Lewis. 2014. "Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-City Evidence from 1980–2000." *American Economic Journal: Applied Economics* 6 (2): 178–194.
- Becker, Gary S. 1971 (orig. pub. 1957). *The Economics of Discrimination*, Second edition. Chicago: University of Chicago Press.
- Becker, Gary S. 1981, enlarged edition 1991. *A Treatise on the Family*, Cambridge, MA: Harvard University Press.
- Becker, Gary S. 1985. "The Allocation of Effort, Specific Human Capital, and the Differences Between Men and Women in Earnings and Occupations." *Journal of Labor Economics* 3 (1, pt. 2): S33–S58.
- Becker, Gary S. 1993, 1st ed. 1964. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, 3rd ed., Chicago: University of Chicago Press.
- Becker, Gary S., William H. J. Hubbard, and Kevin M. Murphy. 2010. "Explaining the Worldwide Boom in Higher Education of Women." *Journal of Human Capital* 4 (3): 203–41.
- Bell, Linda. 2005. "Women-Led Firms and the Gender Gap in Top Executive Pay." Institute for the Study of Labor Working (IZA) Working Paper No. 1689.
- Beller, Andrea H. 1979. "The Impact of Equal Employment Opportunity Laws on the Male/Female Earnings Differential." In *Women in the Labor Market*, edited by Cynthia B. Lloyd, Emily Andrews, and Curtis L. Gilroy, 304–330. New York: Columbia University Press.
- Beller, Andrea H. 1982. "Occupational Segregation by Sex: Determinants and Changes." *Journal of Human Resources* 17 (3): 317–92.
- Benson, Alan. Forthcoming. "Re-thinking the Two-Body Problem: The Segregation of Women into Geographically-Flexible Occupations." *Demography*.
- Bergmann, Barbara R. 1974. "Occupational Segregation, Wages and Profits When Employers Discriminate by Race or Sex." *Eastern Economic Journal* 1 (2): 103–110.

- Berman, Eli, John Bound and Zvi Griliches. 1994. "Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturers." *Quarterly Journal of Economics* 109 (2): 367-297.
- Bertola, Giuseppe, 1992. "Labor Turnover Costs and Average Labor Demand." *Journal of Labor Economics* 10(4): 389-411.
- Bertola, Giuseppe, Francine D. Blau, and Lawrence M. Kahn. 2007. "Labor Market Institutions and Demographic Employment Patterns." *Journal of Population Economics* 20 (4): 833-867.
- Bertrand, Marianne. 2011. "New Perspectives on Gender." In *Handbook of Labor Economics*, Volume 4b, edited by Orley C. Ashenfelter and David Card, 1545-1592. Amsterdam: Elsevier, Ltd.
- Bertrand, Marianne, Sandra E. Black, Sissel Jensen, and Adriana Lleras-Muney. 2014. "Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labor Market Outcomes in Norway." National Bureau of Economic Research Working Paper No. 20256.
- Bertrand, Marianne, Dolly Chugh and Sendhil Mullainathan. 2005. "Implicit Discrimination." *American Economic Review* 95(2): 94-98.
- Bertrand, Marianne, Claudia Goldin, and Lawrence F. Katz. 2010. "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors." *American Economic Journal: Applied Economics* 2 (3): 228–55,
- Bertrand, Marianne and Kevin F. Hallock. 2001. "The Gender Gap in Top Corporate Jobs." *Industrial & Labor Relations Review* 55 (1): 3–21.
- Bertrand, Marianne, Emir Kamenica, and Jessica Pan. 2013. "Gender Identity and Relative Income within Households." mimeo, University of Chicago.
- Bianchi, Suzanne M. and Nancy Rytina. 1986. "The Decline of Occupational Sex Segregation during the 1970s: Census and CPS Comparisons." *Demography* 23(1): 79–86.
- Black, Dan A. 1995. "Discrimination in an Equilibrium Search Model." *Journal of Labor Economics* 13 (2): 309–34.
- Black, Sandra E. and Elizabeth Brainerd. 2004. "Importing Equality? The Effects of Globalization on Gender Discrimination." *Industrial & Labor Relations Review* 57 (4): 540–59.
- Black, Sandra E. and Philip E. Strahan. 2001. "The Division of Spoils: Rent-Sharing and Discrimination in a Regulated Labor Market." *American Economic Review* 91 (4): 814–31.
- Blanchard, Olivier and Augustin Landier. 2002. "The Perverse Effects of Partial Labour Market Reform: fixed--Term Contracts in France." *Economic Journal* 112(480): F214-F244
- Blau, Francine D. 1977. *Equal Pay in the Office*. Lexington, MA: Lexington Books.
- Blau, Francine D. 1998. "Trends in the Well-Being of American Women, 1970–1995." *Journal of Economic Literature* 36 (1): 112-165.

- Blau, Francine D. and Andrea H. Beller. 1988. "Trends in Earnings Differentials by Gender, 1971-1981." *Industrial & Labor Relations Review* 41 (4): 513-529.
- Blau, Francine D., Peter Brummund, and Albert Yung-Hsu Liu. 2013a. "Segregation by Gender 1970-2009: Adjusting for the Impact of Changes in the Occupational Coding System." *Demography* 50 (2): 471-92.
- Blau, Francine D., Peter Brummund, and Albert Yung-Hsu Liu. 2013b. "Erratum to: Segregation by Gender 1970-2009: Adjusting for the Impact of Changes in the Occupational Coding System." *Demography* 50 (2): 493-494.
- Blau, Francine D. and Jed DeVaro. 2007. "New Evidence on Gender Differences in Promotion Rates: An Empirical Analysis of a Sample of New Hires." *Industrial Relations* 46 (3): 511-550.
- Blau, Francine D., Marianne A. Ferber, and Anne W. Winkler. 2014. *The Economics of Women, Men, and Work*, 7th ed. Upper Saddle River, NJ: Prentice Hall/Pearson.
- Blau, Francine D. and Lawrence M. Kahn. 1981. "Race and Sex Differences in Quits by Young Workers." *Industrial and Labor Relations Review* 34 (4): 563-77.
- Blau, Francine D. and Lawrence M. Kahn. 1992. "The Gender Earnings Gap: Learning from International Comparisons." *American Economic Review* 82 (2): 533-538.
- Blau, Francine D. and Lawrence M. Kahn. 1996a. "International Differences in Male Wage Inequality: Institutions versus Market Forces." *Journal of Political Economy* 104 (4): 791-837.
- Blau, Francine D. and Lawrence M. Kahn. 1996b. "Wage Structure and Gender Earnings Differentials: An International Comparison." *Economica* 63 (supplement): 29-62.
- Blau, Francine D. and Lawrence M. Kahn. 1997. "Swimming Upstream: Trends in the Gender Wage differential in the 1980s." *Journal of Labor Economics* 15(1, pt. 1): 1-42.
- Blau, Francine D. and Lawrence M. Kahn. 2002. *At Home and Abroad: U.S. Labor Market Performance in Institutional Perspective*. New York: Russell Sage.
- Blau, Francine D. and Lawrence M. Kahn. 2003. "Understanding International Differences in the Gender Pay Gap." *Journal of Labor Economics* 21 (1): 106-144.
- Blau, Francine D. and Lawrence M. Kahn. 2004. "The U.S. Gender Pay Gap in the 1990s: Slowing Convergence." National Bureau of Economic Research Working Paper 10853. Cambridge, Mass.
- Blau, Francine D. and Lawrence M. Kahn. 2006. "The U.S. Gender Pay Gap in the 1990s: Slowing Convergence." *Industrial & Labor Relations Review* 60 (1): 45-66.
- Blau, Francine D. and Lawrence M. Kahn. 2007. "Changes in the Labor Supply Behavior of Married Women: 1980-2000." *Journal of Labor Economics* 25 (3): 393-438.
- Blau, Francine D. and Lawrence M. Kahn. 2008. "Women's Work and Wages," In *The New Palgrave Dictionary of Economics*, 2nd ed., edited by Steven N. Durlauf and Lawrence E. Blume, 762-772. London: Palgrave Macmillan.
- Blau, Francine D. and Lawrence M. Kahn. 2013a. "Female Labor Supply: Why is the US Falling Behind?" *American Economic Review* 103 (3): 251-256.

- Blau, Francine D. and Lawrence M. Kahn. 2013b. "The Feasibility and Importance of Adding Measures of Actual Experience to Cross-Sectional Data Collection." *Journal of Labor Economics* 3 (2, part 2): S17-S58.
- Bliner, Alan. 1973. "Wage Discrimination: Reduced Form and Structural Estimates." *Journal of Human Resources* 8 (4): 436-455.
- Blundell, Richard, Amanda Gosling, Hidehiko Ichimura, and Costas Meghir. .2007 "Changes in the Distribution of Male and Female Wages Accounting for Employment Composition Using Bounds." *Econometrica* 75 (2): 323-363.
- Boeri, Tito. 2011. "Institutional Reforms and Dualism in European Labor Markets." In *Handbook of Labor Economics*, Volume 4b, edited by Orley C. Ashenfelter and David Card, 1173-1236. Amsterdam: Elsevier, Ltd.
- Borghans, Lex, Bas terWeel, and Bruce Weinberg. 2014. "People Skills and the Labor-Market Outcomes of Underrepresented Groups." *ILRRReview* 67 (2): 287-334.
- Bowles, Hannah R., Linda Babcock and Lei Lai. 2007. "Social Incentives for Sex Differences in the Propensity to Initiate Negotiation: Sometimes It Does Hurt to Ask." *Organizational Behavior and Human Decision Processes* 103 (1): 84-103.
- Brainerd, Elizabeth. 2000. "Women in Transition: Changes in Gender Wage Differentials in Eastern Europe and the Former Soviet Union." *Industrial & Labor Relations Review* 54 (1): 138-162.
- Bronson, Mary Ann. 2013. "Degrees are Forever: Marriage, Educational Investment, and Lifecycle Labor Decisions of Men and Women." Unpublished manuscript UCLA.
- Buser, Thomas, Muriel Niederle, and Hessel Oosterbeek. 2012. "Gender, Competitiveness and Career Choices." NBER Working Paper No. 18576.
- Card, David, Ana Rute Cardoso, and Patrick Kline. 2014. "Bargaining, Sorting and the Gender Wage Gap: The Role of Firms in the Relative Pay of Women." Working Paper UC Berkeley.
- Carrington, William J., Kristin McCue and Brooks Pierce. 2000. "Using Establishment Size to Measure the Impact of Title VII and Affirmative Action." *Journal of Human Resources* 35 (3): 503-523.
- Cha, Youngjoo and Kim A. Weeden. 2014. "Overwork and the Slow Convergence in the Gender Gap in Wages." *American Sociological Review* 79 (3): 457-484.
- Ceci, Stephen J., Donna K. Ginther, Shulamit Kahn, and Wendy M. Williams. 2014. "Women in Academic Science: A Changing Landscape." *Psychological Science* 15 (3):75-141.
- Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly. 2013. "Inference on Counterfactual Distributions." *Econometrica* 81 (6): 2205-2268.
- Cobb-Clark, Deborah A. 2001. "Getting Ahead: The Determinants of and Payoffs to Internal Promotion for Young U.S. Men and Women." In *Research in Labor Economics*, vol. 20, edited by Solomon W. Polachek, 339-372. Amsterdam: Elsevier Science, JAI.

- Compton, Janice and Robert A. Pollak. 2007. "Why Are Power Couples Increasingly Concentrated in Large Metropolitan Areas?" *Journal of Labor Economics* 25 (3): 475–512.
- Cooke, Thomas J., Paul Boyle, and Kenneth Couch. 2009. "A Longitudinal Analysis of Family Migration and the Gender Gap in Earnings in the United States and Great Britain." *Demography* 46 (1): 147-167.
- Correll, Shelley J., Stephen Benard, and In Paik. 2007. "Getting a Job: Is There a Motherhood Penalty?" *American Journal of Sociology* 112 (5): 1297–1338.
- Costa, Dora L. and Matthew E. Kahn. 2000. "Power Couples: Changes in the Locational Choice of the College Educated, 1940-1990." *Quarterly Journal of Economics* 115 (4): 1287-1315.
- Cotter, David, Joan M. Hermsen, and Reeve Vanneman. 2011. "The End of the Gender Revolution? Gender Role Attitudes from 1977 to 2008." *American Journal of Sociology* 117, (4): 259-289.
- Croson, Rachel and Uri Gneezy. 2009. "Gender Differences in Preferences." *Journal of Economic Literature* 47 (2): 448-474.
- Dahl, Gordon, Katrine Løken, and Magne Mogstad. 2013. "Peer Effects in Program Participation." *American Economic Review* 104(7): 2049-2074.
- Datta Gupta, Nabanita, Ronald L. Oaxaca, and Nina Smith. 2006. "Swimming Upstream, Floating Downstream: Comparing Women's Relative Wage Progress in the United States and Denmark." *Industrial & Labor Relations Review* 59(2): 243-266.
- DiNardo, John, Nicole M. Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach." *Econometrica* 64(5): 1001-1044.
- DiPrete, Thomas A. and Claudia Buchmann. 2006. "Gender-Specific Trends in the Value of Education and the Emerging Gender Gap in College Completion." *Demography* 43 (1): 1–24.
- Donohue III, John J. and James Heckman. 1991. "Continuous Versus Episodic Change: The Impact of Civil Rights Policy on the Economic Status of Blacks." *Journal of Economic Literature* 29 (4): 1603-1643.
- Dougherty, Christopher. 2005. "Why Are the Returns to Schooling Higher for Women than for Men?" *Journal of Human Resources* 40 (4): 969–988.
- Duncan, Otis Dudley and Beverly Duncan. 1955. "A Methodological Analysis of Segregation Indexes." *American Sociological Review* 20 (2): 210-17.
- Eckel, Catherine C. and Sascha C. Füllbrunn. 2015. "Thar SHE Blows? Gender, Competition, and Bubbles in Experimental Asset Markets." *American Economic Review* 105 (2): 906–920.
- Edin, Per-Anders and Katarina Richardson. 2002. "Swimming with the Tide: Solidarity Wage Policy and the Gender Earnings Gap." *Scandinavian Journal of Economics* 104 (1): 49-67.

- Ehrenberg, Ronald G., George H. Jakubson, Mirinda L. Martin, Joyce B. Main, and Thomas Eisenberg. 2012. "Diversifying the Faculty Across Gender Lines: Do Trustees and Administrators Matter?" *Economics of Education Review* 31 (1): 9-18.
- England, Paula and S Li. 2006. "Desegregation Stalled: The Changing Gender Composition of College Majors, 1971-2002." *Gender and Society* 20 (5): 657-677.
- European Commission (2006). *The Gender Pay Gap—Origins and Policy Responses: A Comparative Review of 30 European Countries*. Luxembourg, Office for Official Publications of the European Communities.
- Farber, Henry S. and Robert Gibbons. 1996. "Learning and Wage Dynamics." *Quarterly Journal of Economics* 111 (4):1007-1047.
- Firpo, Sergio, Nicole Fortin, and Thomas Lemieux. 2009. "Unconditional Quantile Regressions." *Econometrica* 77 (3): 953-973.
- Fiske, Susan T. "Stereotyping, Prejudice, and Discrimination." 1998. In *Handbook of Social Psychology*, edited by D. T. Gilbert, S. T. Fiske, and G. Lindzey, 357–411. New York: McGraw-Hill.
- Flory, Jeffrey A., Andreas Leibbrandt, and John A. List. forthcoming. "Do Competitive Work Places Deter Female Workers? A Large-Scale Natural Field Experiment on Gender Differences in Job-Entry Decisions." *Review of Economic Studies*.
- Folbre, Nancy, ed. 2012. *For Love and Money: Care Provision in the United States*. New York: Russell Sage Foundation.
- Fortin, Nicole M. 2008. "The Gender Wage Gap Among Young Adults in the United States; The Importance of Money Versus People." *Journal of Human Resources* 43 (4): 884–918.
- Fortin, Nicole M. Forthcoming. "Gender Role Attitudes and Women's Labor Market Participation: Opting Out, AIDS, and the Persistent Appeal of Housewifery," *Annals of Economics and Statistics*, special issue on Economics of Gender.
- Frank, Robert, H. 1978. "Why Women Earn Less: The Theory and Estimation of Differential Overqualification." *American Economic Review* 68 (3): 360-73.
- Fryer, Roland G. and Steven D. Levitt. 2010. "An Empirical Analysis of the Gender Gap in Mathematics," *American Economic Journal: Applied Economics* 2 (2): 210–240.
- Fuchs, Victor R. 1971. "Differences in Hourly Earnings Between Men and Women." *Monthly Labor Review* 94 (5): 9-15.
- Fuchs, Victor R. 1988. *Women's Quest for Economic Equality*. Cambridge: Harvard University Press.
- Gayle, George-Levi and Limor Golan. 2012. "Estimating a Dynamic Adverse-Selection Model: Labour-Force Experience and the Changing Gender Earnings Gap 1968–1997." *Review of Economic Studies* 79 (1): 227-267.
- Ginther, Donna K. and Shulamit Kahn. Forthcoming. "Women's Careers in Academic Social Science: Progress, Pitfalls and Plateaus." In *The Economics of Economists*, edited by Alessandro Lanteri and Jack Vromen. Cambridge: Cambridge University Press.

- Gneezy, Uri, Kenneth L. Leonard, and John A. List. 2009. "Gender Differences in Competition: Evidence from a Matrilineal and a Patriarchal Society." *Econometrica* 77 (5): 1637–64.
- Gneezy, Uri, Muriel Niederle and Aldo Rustichini. 2003. "Performance in Competitive Environments: Gender Differences." *Quarterly Review of Economics*, 118 (3): 1049-74.
- Goldin, Claudia. 1990. *Understanding the Gender Gap: An Economic History of American Women*. New York: Oxford University Press.
- Goldin, Claudia. 1997. "Career and Family: College Women Look to the Past." In *Gender and Family Issues in the Workplace*, edited by Francine D. Blau and Ronald G. Ehrenberg, 20-64. New York: Russell Sage.
- Goldin, Claudia. 2006. "The Quiet Revolution That Transformed Women's Employment, Education, and Family." *American Economic Review* 96 (2): 1-21.
- Goldin, Claudia. 2014. "A Grand Convergence: Its Last Chapter." *American Economic Review* 104 (4): 1091-1119.
- Goldin, Claudia and Lawrence F. Katz. 2002. "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions." *Journal of Political Economy* 110 (4): 730–70.
- Goldin, Claudia and Lawrence M. Katz. 2012. "The Most Egalitarian of All Professions: Pharmacy and the Evolution of a Family-Friendly Occupations." National Bureau of Economic Research Working Paper 18410.
- Goldin, Claudia, Lawrence F. Katz, and Ilyana Kuziemko. 2006. "The Homecoming of American College Women: The Reversal of the College Gender Gap." *Journal of Economic Perspectives* 20 (4): 133–56.
- Goldin, Claudia and Cecilia Rouse. 2000. "Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians." *American Economic Review* 90 (4): 715–41.
- Greenwald Anthony G., Debbie E. McGhee, Jordan L.K. Schwartz. 1998. "Measuring Individual Differences in Implicit Cognition: The Implicit Association Test." *Journal of Personality and Social Psychology* 74(6): 1464–1480.
- Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu. 2005. "Engines of Liberation." 2005. *Review of Economic Studies* 72 (1): 109–133.
- Gronau, Reuben. 1974. "Wage Comparisons-a Selectivity Bias." *Journal of Political Economy* 82 (6): 1119-43.
- Groshen, Erica L. 1991. "The Structure of the Female/Male Wage Differential: Is It Who You Are, What You Do, or Where You Work?" *Journal of Human Resources* 26 (3): 457–72.
- Gross, Edward. 1968. "Plus ca change...? The sexual structure of occupations over time." *Social Problems* 16 (2): 198–208.
- Gruber, Jonathan. 1994. "The Incidence of Mandated Maternity Benefits." *American Economic Review* 84 (3): 622-641.
- Guiso, Luigi, Ferdinando Monte, Paola Sapienza, and Luigi Zingales. 2008. "Education Forum: Culture, Gender, and Math." *Science* 320: 1164–1165.

- Harrison, Glenn W. and John A. List. 2004. "Field Experiments." *Journal of Economic Literature* 52 (4): 1009-55.
- Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47 (1): 153-61.
- Heckman, James J. and Tim Kautz, T. 2012. "Hard Evidence on Soft Skills." Adam Smith Lecture. *Labour Economics*, 19 (4): 451-464.
- Heim, Bradley T. 2007. The Incredible Shrinking Elasticities: Married Female Labor Supply, 1979-2003. *Journal of Human Resources* 42 (4): 881-918.
- Hellerstein, Judith K., David Neumark, and Kenneth Troske. 2002. "Market Forces and Sex Discrimination." *Journal of Human Resources* 37 (2): 353-80.
- Hersch, Joni. 2006. "Sex Discrimination in the Labor Market." *Foundations and Trends in Microeconomics* 2 (4): 281-361.
- Hersch, Joni and Leslie S. Stratton. 1997. "Housework, Fixed Effects and Wages of Married Workers." *Journal of Human Resources* 32 (2): 285-307.
- Hersch, Joni and Leslie S. Stratton. 2002. "Housework and Wages." *Journal of Human Resources* 37 (1): 217-229.
- Hersch, Joni and W. Kip Viscusi. 1996. "Gender Differences in Promotions and Wages." *Industrial Relations* 35 (4): 461-72.
- Heyman, Fredrik, Helena Svaleryd, and Jonas Vlachos. 2013. "Competition, Takeovers, and Gender Discrimination." *ILRReview* 66 (2): 409-432.
- Hirsch, Boris, Thorsten Schank, and Claus Schnabel. 2010. "Differences in Labor Supply to Monopsonistic Firms and the Gender Pay Gap: An Empirical Analysis Using Linked Employer-Employee Data from Germany." *Journal of Labor Economics* 28 (2): 291-330.
- Hoffman, Moshe, Uri Gneezy, and John A. List. 2011. "Nurture Affects Gender Differences in Spatial Abilities." *Proceedings of the National Academy of Science* 108 (36): 14786-14788.
- Holzer, Harry and David Neumark. 1999. "Are Affirmative Action Hires Less Qualified? Evidence from Employer-Employee Data," *Journal of Labor Economics* 17 (3): 534-69.
- Holzer, Harry J. and David Neumark. 2000 a. "Assessing Affirmative Action." *Journal of Economic Literature* 38 (3): 483-568.
- Holzer, Harry J. and David Neumark. 2000 b. "What Does Affirmative Action Do?" *Industrial and Labor Relations Review* 53 (2): 240-271.
- Hubbard, William H. J. 2011. "The Phantom Gender Difference in the College Wage Premium," *Journal of Human Resources* 46 (3): 568-86.
- Hyde, Janet S, Sara M. Lindberg, Marcia C. Linn, Amy B. Ellis, and Caroline C. Williams. 2008. "Gender Similarities Characterize Math Performance." *Science* 321: 494-495.
- Jacobs, Jerry. A. 1989. "Long-term Trends in Occupational Segregation by Sex." *American Journal of Sociology* 95 (1): 160-173. Jacobsen, Joyce P. 1994. "Trends in Workforce Sex Segregation, 1960-1990." *Social Science Quarterly* 75 (1): 204-211.

- Jacobsen, Joyce., Melanie Khamis, and Mutlu Yuksel. 2014. "Convergences in Men's and Women's Life Patterns: Lifetime Work, Lifetime Earnings, and Human Capital Investment." IZA Discussion Paper No. 8425.
- Judge, Timothy A. 2012. "Do Nice Guys—and Gals—Really Finish Last?" *Journal of Personality and Social Psychology* 102 (2): 390-407.
- Juhn, Chinhui and Kevin M. Murphy. 1997. "Wage Inequality and Family Labor Supply." *Journal of Labor Economics* 15 (1, pt. 1): 72-97.
- Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce. 1991. "Accounting for the Slowdown in Black-White Wage Convergence." In *Workers and Their Wages*, edited by Marvin Koster, 107-143. Washington, DC: AEI Press.
- Kahn, Lawrence M. 2007. "The Impact of Employment Protection Mandates on Demographic Temporary Employment Patterns: International Microeconomic Evidence." *Economic Journal*, 117 (521): F333-F356.
- Kahn, Lawrence M. 2014. "Wage Compression and the Gender Pay Gap." Working Paper, Cornell University.
- Kahn, Lawrence M. Forthcoming. "The Structure of the Permanent Job Wage Premium: Evidence from Europe." *Industrial Relations*.
- Kahn, Lisa B. 2013. "Asymmetric Information between Employers." *American Economic Journal: Applied* 5 (4): 165-205.
- Kahn, Lisa B. and Fabian Lange. 2014. "Employer Learning, Productivity and the Earnings Distribution: Evidence from Performance Measures." *Review of Economic Studies* 81(4): 1575-1613.
- Kassenboehmer, Sonja C. and Mathias G. Sinning. 2014. "Distributional Changes in the Gender Wage Gap." *ILRRReview* 67 (2): 335-361.
- Keith, Kristen and Abigail McWilliams. 1995. "The Wage Effects of Cumulative Job Mobility." *Industrial & Labor Relations Review* 49 (1): 121-37.
- Kidd, Michael P. and Michael Shannon. 1996. "The Gender Wage Gap: A Comparison of Australia and Canada." *Industrial & Labor Relations Review* 49(4): 729-746,
- Korenman, Sanders and David Neumark. 1991. "Does Marriage Really Make Men More Productive?" *Journal of Human Resources* 26 (2): 282-307.
- Korenman, Sanders and David Neumark. 1992. "Marriage, Motherhood, and Wages." *Journal of Human Resources* 27(2): 233 – 55.
- Kosteas, Vasilios D. 2013. "Gender Role Attitudes, Labor Supply, and Human Capital Formation." *Industrial Relations* 52 (4): 915-940.
- Krueger, Alan B. 1993. "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984-1989." *Quarterly Journal of Economics* 108 (1): 33-60.
- Kurtulus, Fidan Ana. 2012. "Affirmative Action and the Occupational Advancement of Minorities and Women During 1973-2003." *Industrial Relations* 51 (2): 213-246.

- Kurtulus, Fidan Ana and Donald Tomaskovic-Devey. 2012. "Do Female Top Managers Help Women to Advance? A Panel Study Using EEO-1 Records." *The Annals of the American Academy of Political and Social Science* 639 (1): 173-197.
- Landers, Renée M., James B. Rebitzer, and Lowell J. Taylor. 1996. "Rat Race Redux: Adverse Selection in the Determination of Work Hours in Law Firms." *American Economic Review* 86 (3): 329-348.
- Lavy, Victor and Edith Sand. 2015. "On the Origins of Gender Human Capital Gaps: Short and Long Term Consequences of Teachers' Stereotypical Biases." NBER Working Paper 20909.
- Lazear, Edward P. and Sherwin Rosen. 1990. "Male-Female Wage Differentials in Job Ladders." *Journal of Labor Economics* 8 (2): S106-S123.
- Leibbrandt, Andreas and John A. List. Forthcoming. "Do Women Avoid Salary Negotiations? Evidence from a Large-Scale Natural Field Experiment." *Management Science*.
- Leonard, Jonathan. 1990. "The Impact of Affirmative Action Regulation and Equal Opportunity Law on Employment." *Journal of Economic Perspectives* 4 (4): 47-63.
- Levanon, Asaf, Paula England, and Paul Allison. 2009. "Occupational Feminization and Pay: Assessing Causal Dynamics Using 1950-2000 U.S. Census Data." *Social Forces* 88 (2): 865-981.
- Light, Audrey and Manuelita Ureta. 1995. "Early-Career work Experience and Gender Wage Differentials." *Journal of Labor Economics* 13 (1): 121-154.
- Lundberg, Shelly J. and Rose, Elaina. 2000. "Parenthood and the Earnings of Married Men and Women." *Labour Economics* 7 (6S): 689-710.
- Lundberg, Shelly J. and Richard Startz. 1983. "Private Discrimination and Social Intervention in Competitive Labor Markets." *American Economic Review* 73 (3): 340-47.
- Macpherson, David A. and Barry T. Hirsch. 1995. "Wages and Gender Composition, Why Do Women's Jobs Pay Less?" *Journal of Labor Economics* 13 (3): 426-71.
- Madden, Janice F. 1973. *The Economics of Sex Discrimination*, Lexington, MA: Lexington Books.
- Manning, Alan. 2003. *Monopsony in Motion: Imperfect Competition in Labor Markets*, NJ: Princeton University Press.
- Manning, Alan and Farzad Saidi. 2010. "Understanding the Gender Pay Gap: What's Competition Got to Do With It?" *Industrial & Labor Relations Review* 63 (4): 681-698.
- Manning, Alan and Joanna Swaffield. 2008. "The Gender Gap in Early-Career Growth." *Economic Journal* 118 (530): 983-1024.
- Mazei, Jens, Joachim Hüffmeier, Philipp Alexander Freund, Alice F. Stuhlmacher, Lena Bilke, and Guido Hertel. Forthcoming. "A Meta-Analysis on Gender Differences in Negotiation Outcomes and Their Moderators." *Psychological Bulletin*.
- McCue, Kristin. 1996. "Promotions and Wage Growth." *Journal of Labor Economics* 14 (2): 175-209.

- McGrattan, Ellen R. and Richard Rogerson. 2008. "Changes in the Distribution of Family Hours Worked Since 1950." In *Frontiers of Family Economics* Volume 1 edited by Peter Rupert, 115-138. Bingley, UK: Emerald Group Publishing Ltd.
- Meyer, Bruce D. and Dan T. Rosenbaum. 2001. "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers." *Quarterly Journal of Economics* 116 (3): 1063–1114.
- Mincer, Jacob. 1962. "Labor Force Participation of Married Women." In *Aspects of Labor Economics*, Universities National Bureau of Economic Research Conference Studies, no. 14, edited by H. Gregg Lewis, 63–97. Princeton, NJ: Princeton University Press.
- Mincer, Jacob. 1978. "Family Migration Decisions." *Journal of Political Economy* 86 (5): 749–73.
- Mincer, Jacob and Haim Ofek. 1982. "Interrupted Work Careers: Depreciation and Restoration of Human Capital," *Journal of Human Resources* 17 (1): 3-24.
- Mincer, Jacob and Solomon W. Polachek. 1974. "Family Investments in Human Capital: Earnings of Women." *Journal of Political Economy* 82 (2, pt. 2): 76–108.
- Moss-Racusin, Corinne A., John F. Dovidio, Victoria L. Brescoll, Mark J. Graham, and Jo Handelsman. 2012. "Science Faculty's Subtle Gender Biases Favor Male Students." *Proceedings of the National Academy of Science (PNAS)* 109 (41): 16474–16479.
- Mueller, Gerrit and Erik J.S. Plug. 2006. "Earnings Effects of Personality." *Industrial & Labor Relations Review* 60 (1): 3–22.
- Mulligan, Casey B. and Yona Rubinstein. 2008. "Selection, Investment, and Women's Relative Wages." *Quarterly Journal of Economics* 123 (3): 1061–110.
- Neal, Derek. 2004. "The Measured Black-White Wage Gap Among Women Is Too Small." *Journal of Political Economy* 112 (1, pt. 2): S1–S28.
- Neal, Derek and William R. Johnson. 1996. "The Role of Premarket Factors in Black-White Wage Differences." *Journal of Political Economy* 104 (5): 869-895.
- Nelson, Julie A. Forthcoming. "Are women really More Risk-Averse Than Men? A Re-Analysis of the Literature Using Expanded Methods." *Journal of Economic Surveys*.
- Neumark, David, with the assistance of Roy J. Blank and Kyle D. Van Nort. 1996. *Quarterly Journal of Economics* 111 (3): 915–42.
- Niederle, Muriel and Lise Vesterlund. 2007. "Do Women Shy Away from Competition?" *Quarterly Journal of Economics* 122 (3): 1067–101.
- Niederle, Muriel and Lise Vesterlund. 2010. "Explaining the Gender Gap in Math Test Scores: The Role of Competition." *Journal of Economic Perspectives* 24 (2): 129–44.
- Nollenberger, Natalia, Núria Rodríguez-Planas, and Almudena Sevilla. 2014. "The Math Gender Gap: The Role of Culture." IZA Working Paper 8379.
- Noonan, Mary C., Mary E. Corcoran, and Paul Courant. 2005. "Pay Differences Among the Highly Trained: Cohort Differences in the Sex Gap in Lawyers' Earnings." *Social Forces* 84 (2): 853–72.

- Oaxaca, Ronald. 1973. "Male-Female Wage Differences in Urban Labor Markets." *International Economic Review* 14 (3): 693–709.
- Oaxaca, Ronald L. and Michael Ransom. 1999. "Identification in Detailed Wage Decompositions." *Review of Economics and Statistics* 81(1): 154–157.
- OECD. 2012. *OECD Employment Outlook 2012*. Paris: OECD.
- OECD. 2014. *OECD Employment Outlook 2014*. Paris: OECD.
- Olivetti, Claudia and Barbara Petrongolo. 2008. "Unequal Pay or Unequal Employment? A Cross-Country Analysis of Gender Gaps." *Journal of Labor Economics* 26 (4): 621–54
- O'Neill, June and Solomon W. Polachek. 1993. "Why the Gender Gap in Wages Narrowed in the 1980s." *Journal of Labor Economics* 11 (1, pt. 1): 205–228.
- Oppenheimer, Valerie. 1976 (originally published 1970). *The Female Labor Force in the United States: Demographic and Economic Factors Governing Its Growth and Changing Composition*. Westport, CT: Greenwood Press.
- Orazem, Peter F. and Milan Vodopivec. 2000. "Male-Female Differences in Labor Market Outcomes During the Early Transition to Market: The Cases of Estonia and Slovenia." *Journal of Population Economics* 13 (2): 283–303.
- Örs, Evren, Frederic Palomino and Eloic Peyrache. 2013. "Performance Gender-Gap: Does Competition Matter?" *Journal of Labor Economics* 31(3): 443–99.
- Pager, Devah. 2007. "The Use of Field Experiments for Studies of Employment Discrimination: Contributions, Critiques, and Directions for the Future." *The Annals of the American Academy of Political and Social Science* (609): 103–33.
- Paglin, Morton and Anthony M. Rufolo. 1990. "Heterogeneous Human Capital, Occupational Choice, and Male–Female Earnings Differences." *Journal of Labor Economics* 8 (1, pt. 1): 123–44.
- Patnaik, Ankita. 2015. "Reserving Time for Daddy: The Short and Long-Run Consequences of Fathers' Quotas." Unpublished Working Paper, Cornell University.
- Phelps, Edmund S. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review* 62 (4): 659–661.
- Polachek, Solomon W. 1978. "Sex Differences in College Major." *Industrial & Labor Relations Review* 31 (4): 498–508.
- Polachek, Solomon W. 1981. "Occupational Self-Selection: A Human Capital Approach to Sex Differences in Occupational Structure." *Review of Economics and Statistics* 63 (1): 60–69.
- Pope, Devin G. and Justin R. Sydnor. 2010. "Geographic Variation in the Gender Differences in Test Scores." *Journal of Economic Perspectives* 24(2): 95–108
- Porterfield, Shirley L. and Anne E. Winkler. 2007. "Teen Time Use and Parental Education: Evidence from the CPS, MTF, and ATUS," *Monthly Labor Review* 130 (5): 37–56.
- Ransom, Michael R. and Ronald L. Oaxaca. 2010. "New Market Power Models and Sex Differences in Pay." *Journal of Labor Economics* 28 (2): 267–89.

- Ransom, Michael R. and David P. Sims. 2010. "Estimating the Firm's Labor Supply Curve in a 'New Monopsony' Framework: Schoolteachers in Missouri." *Journal of Labor Economics* 28 (2): 331-355.
- Reuben, Ernesto, Paola Sapienza, and Luigi Zingales. 2014. "How Stereotypes Impair Women's Careers in Science." *PNAS* 111 (12): 4403-4408.
- Ribar, David C. 2004. "What Do Social Scientists Know About the Benefits of Marriage? A Review of Quantitative Methodologies." IZA Discussion Paper No. 998.
- Rodgers, William M. III, and Leslie S. Stratton. 2010. "The Male Marital Wage Differential: Race, Training, and Fixed Effects." *Economic Inquiry* 48 (3): 722-742.
- Rosen, Sherwin. 1986. "The Theory of Equalizing Differences." In *The Handbook of Labor Economics*, Vol. 1, edited by Orley Ashenfelter and Richard Layard, 641-92. Amsterdam: Elsevier-North Holland.
- Royalty, Anne Beeson. 1996. "The Effects of Job Turnover on the Training of Men and Women." *Industrial & Labor Relations Review* 49 (3): 506-521.
- Royalty, Anne Beeson. 1998. "Job-to-Job and Job-to-Nonemployment Turnover by Gender and Education Level." *Journal of Labor Economics* 16 (2): 392-443.
- Ruhm, Christopher J. 1998. "The Economic Consequences of Parental Leave Mandates: Lessons from Europe." *Quarterly Journal of Economics* 113 (1): 285-317.
- Ruhm, Christopher J. 2011. "Policies to Assist Parents with Young Children." *The Future of Children* 21 (2): 38-68.
- Sandell, Steven H. 1977. "Women and the Economics of Family Migration." *The Review of Economics and Statistics* 59 (4): 406-414.
- Sawhill, Isabel V. 1973. "The Economics of Discrimination against Women: Some New Findings." *Journal of Human Resources* 8 (3): 383-396.
- Shin, Taekjin. 2012. "The Gender Gap in Executive Compensation: the Role of Female Directors and Chief Executive Officers." *The Annals of the American Academy of Political and Social Science* 639 (1): 258-278.
- Sicherman, Nachum. 1996. "Gender Differences in Departures from a Large Firm." *Industrial & Labor Relations Review* 49 (3): 484-505.
- Sigle-Rushton, Wendy and Jane Waldfogel. 2007. "Motherhood and Women's Earnings in Anglo-American, Continental European, and Nordic Countries." *Feminist Economics* 13. (2): 55-91.
- Simón, Hipólito. 2012. "The Gender Gap in Earnings: An International Comparison with European Matched Employer-Employee Data." *Applied Economics* 44 (15): 1985-1999.
- Spencer, Steven J., Claude M. Steele, and Diane M. Quinn. 1999. "Stereotype Threat and Women's Math Performance." *Journal of Experimental Social Psychology* 35 (1): 4-28.
- Spivey, Christy. 2005. "Time off at What Price? The Effects of Career Interruptions on Earnings." *Industrial & Labor Relations Review* 59 (1): 119-40.

- Sorensen, Elaine. 1990. "The Crowding Hypothesis and Comparable Worth Issue." *Journal of Human Resources* 25 (1): 55–89.
- Stanley, T. D. and Stephen B. Jarrell. 1998. "Gender Wage Discrimination Bias? A Meta-Regression Analysis." *Journal of Human Resources* 33 (4): 947–73.
- Stratton, Leslie S. 2001. "Why Does More Housework Lower Women's Wages? Testing Hypotheses Involving Job Effort and Hours Flexibility." *Social Science Quarterly* 82 (1): 67–76.
- Thomas, Mallika. 2015. "The Impact of Mandated Maternity Benefits on the Gender Differential in Promotions: Examining the role of Adverse Selection." Unpublished manuscript, University of Chicago.
- Viscusi, W. Kip. 1980. "Sex Differences in Worker Quitting." *Review of Economics and Statistics* 62 (3): 388–98.
- Waldfogel, Jane. 1998. "Understanding the 'Family Gap' in Pay for Women with Children." *Journal of Economic Perspectives* 12 (1): 157–70.
- Waldfogel, Jane. 1999. "The Impact of the Family and Medical Leave Act." *Journal of Policy Analysis and Management* 18 (2): 281–302.
- Wang, Wendy, Kim Parker, and Paul Taylor. 2013. *Breadwinner Moms: Mothers are the Sole or Primary Provider in Four-in-Ten Households with Children; Public Conflicted about the Growing Trend*. Washington, DC: Pew Research Center available at www.pewresearch.org accessed June 27, 2014.
- Webber, Douglas, A. Forthcoming. "Firm-Level Monopsony and the Gender Pay Gap." *Industrial Relations*.
- Weinberg, Bruce. 2000. "Computer Use and the Demand for Female Workers." *Industrial & Labor Relations Review* 53 (2): 290–308.
- Weinberger, Catherine J. and Peter Kuhn. 2010. "Changing Levels or Changing Slopes? The Narrowing of the U.S. Gender Earnings Gap, 1959–1999." *Industrial & Labor Relations Review* 63 (3): 384–406.
- Weiss, Yoram and Reuben Gronau. 1981. "Expected Interruptions in Labour Force Participation and Sex-Related Differences in Earnings Growth." *Review of Economic Studies* 48 (4): 607–19.
- Welch, Finis. 2000. "Growth in Women's Relative Wages and in Inequality Among Men: One Phenomenon or Two?" *American Economic Review* 90 (2): 444–49.
- Winkler, Anne E., Timothy D. McBride, and Courtney Andrews. 2005. "Wives Who Outearn Their Husbands: A Transitory or Persistent Phenomenon for Couples?" *Demography* 42 (3): 523–535.
- Zhang, Junsen, Jun Han, Pak-Wai Liu, and Yaohui Zhao. 2008. "Trends in the Gender Earnings Differential in Urban China." *Industrial & Labor Relations Review* 61 (2): 224–243.

**Figure 1: Gender Earnings Ratios of Full-Time Workers
1955-2014**

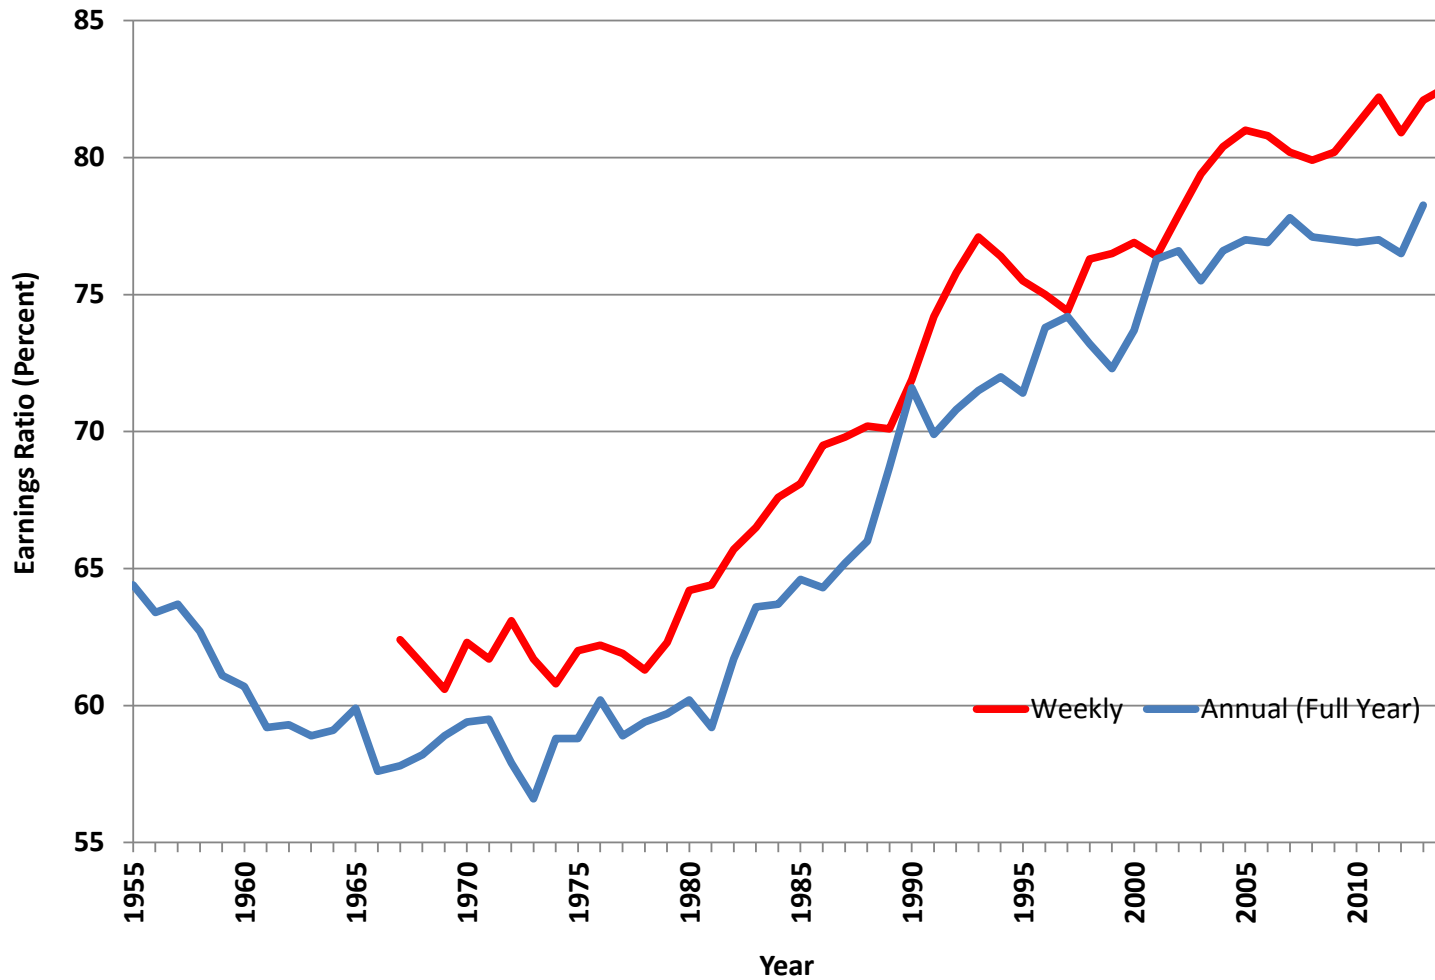
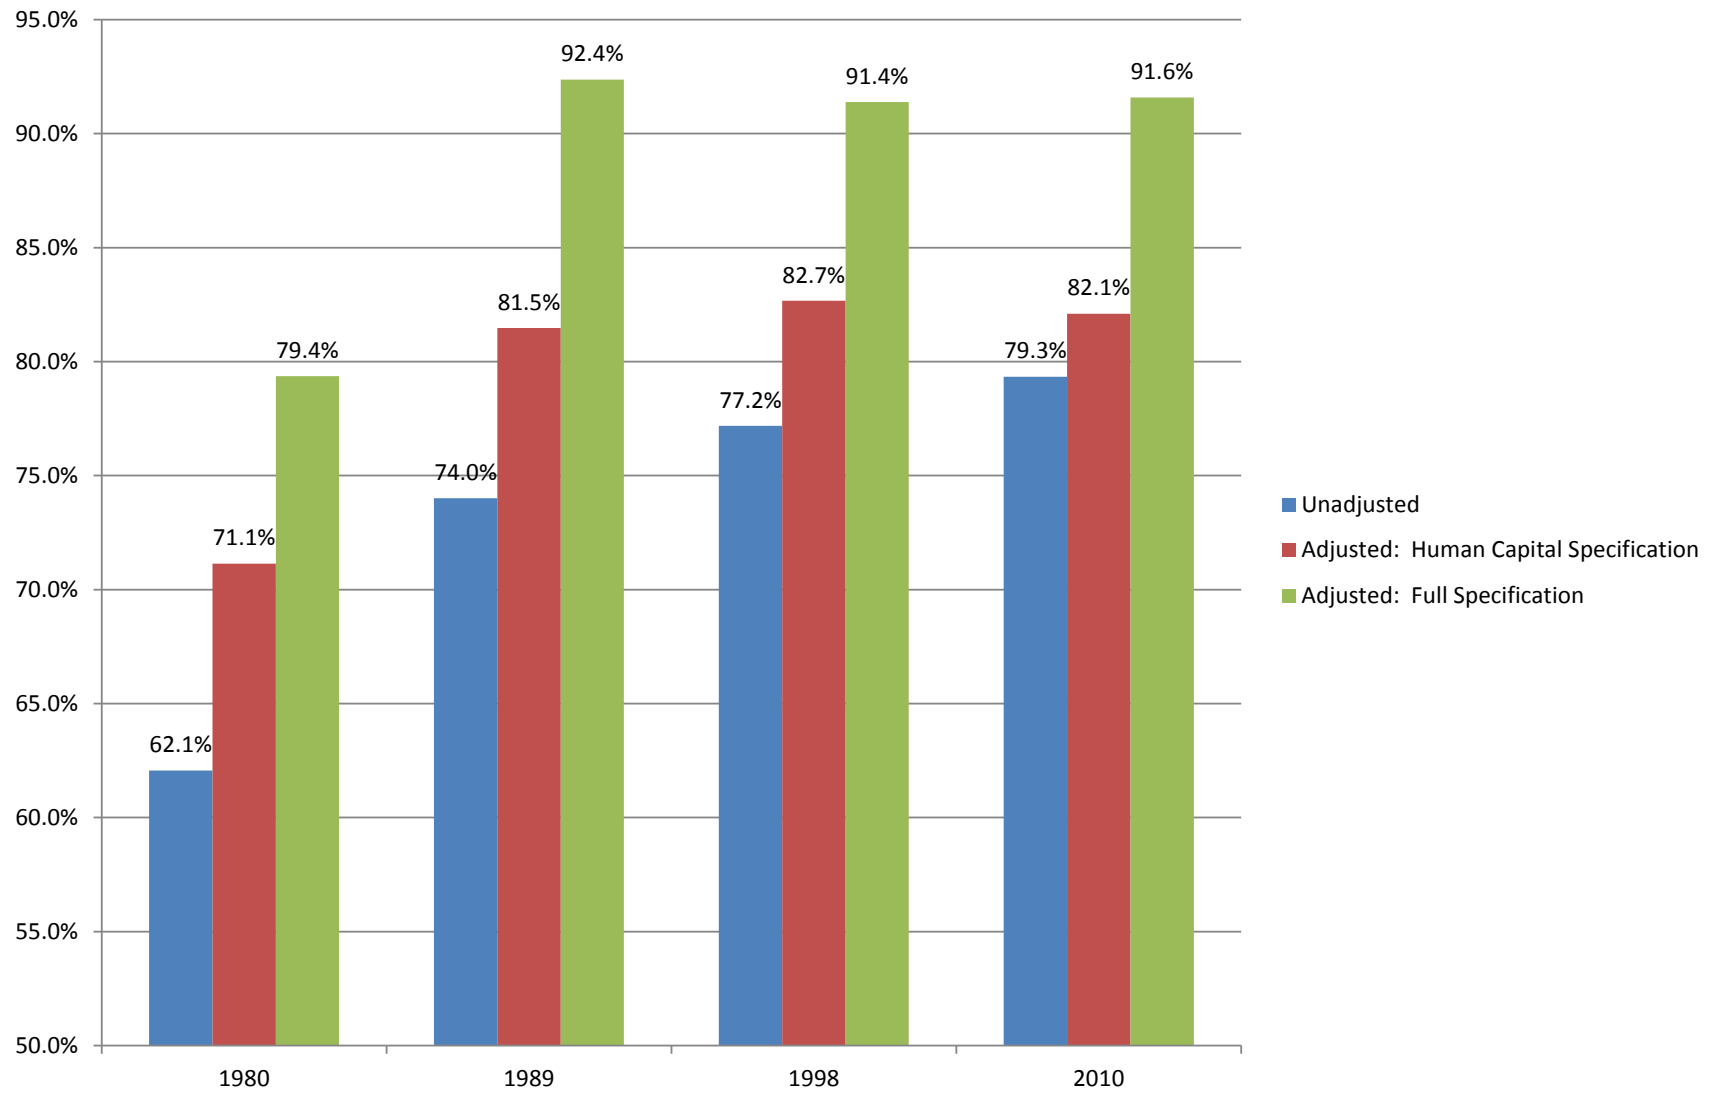


Figure 2: Female to Male Log Wage Ratio, Unadjusted and Adjusted for Covariates (PSID)



**Figure 3: Trends in Female and Male Labor Force Participation Rates, 1947-2014
(age 16 and over)**

Source: Labor Force Statistics from the Current Population Survey available at www.bls.gov; and *Employment & Earnings*, various issues.

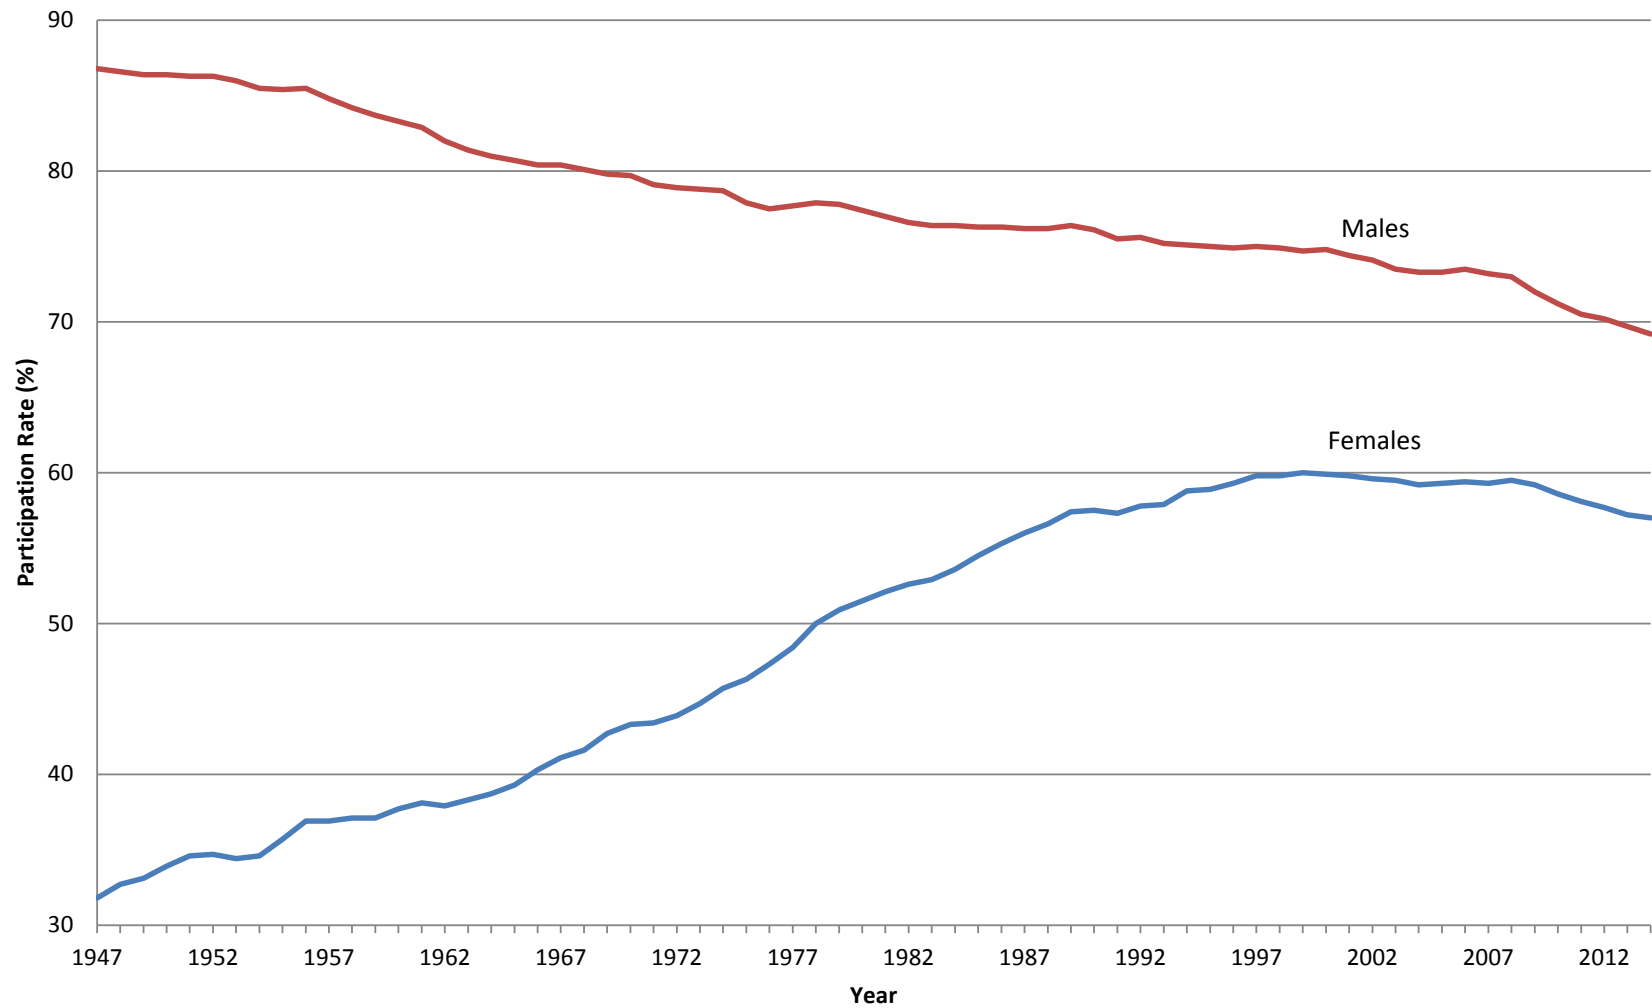


Table 1: Unadjusted Female/Male Log Hourly Wage Ratios, Full Time Workers

Year	Mean	10th Percentile	50th Percentile	90th Percentile
Panel Study of Income Dynamic (PSID)				
1980	62.1%	64.8%	60.1%	62.4%
1989	74.0%	76.3%	72.4%	74.6%
1998	77.2%	80.3%	79.8%	73.8%
2010	79.3%	81.5%	82.4%	73.9%
March Current Populations Survey (CPS)				
1980	63.5%	68.7%	61.9%	64.3%
1989	72.4%	78.1%	72.2%	71.4%
1998	77.1%	81.3%	76.2%	76.1%
2010	82.3%	87.6%	82.2%	76.6%

Notes: Sample includes nonfarm wage and salary workers age 25-64 with at least 26 weeks of employment. Entries are $\exp(D)$, where D is the female mean log wage, 10th, 50th or 90th percentile log wage minus the corresponding male log wage.

Table 2: Schooling and Actual Full Time Work Experience by Gender, PSID

Year	Men	Women	Men-Women
Years of Schooling			
1981	13.3	13.2	0.2
1990	13.8	13.7	0.0
1999	14.2	14.3	-0.1
2011	14.3	14.5	-0.2
Bachelor's Degree Only			
1981	18.1%	15.3%	2.7%
1990	20.0%	17.6%	2.3%
1999	23.4%	22.2%	1.2%
2011	26.2%	24.7%	1.5%
Advanced Degree			
1981	10.0%	7.4%	2.5%
1990	10.3%	8.7%	1.6%
1999	11.7%	10.8%	0.9%
2011	12.9%	15.7%	-2.8%
Years of Full Time Experience			
1981	20.3	13.5	6.8
1990	19.2	14.7	4.5
1999	19.8	15.9	3.8
2011	17.8	16.4	1.4

Notes: Sample includes full time nonfarm wage and salary workers age 25-64 with at least 26 weeks of employment.

Table 3: Incidence of Managerial or Professional Jobs and Collective Bargaining Coverage by Gender, PSID

Year	Men	Women	Men-Women
Managerial Jobs			
1981	21.5%	9.2%	12.3%
1990	21.1%	10.9%	10.2%
1999	21.8%	15.3%	6.5%
2011	18.3%	16.2%	2.2%
Professional Jobs			
1981	17.0%	21.8%	-4.8%
1990	19.4%	26.1%	-6.6%
1999	20.4%	26.9%	-6.4%
2011	21.7%	31.1%	-9.4%
"Male" Professional Jobs			
1981	14.6%	10.1%	4.5%
1990	17.3%	14.1%	3.2%
1999	17.6%	13.2%	4.4%
2011	18.6%	17.8%	0.8%
Collective Bargaining Coverage			
1981	34.5%	21.1%	13.3%
1990	25.4%	19.4%	6.1%
1999	21.5%	18.2%	3.3%
2011	17.4%	18.9%	-1.5%

Notes: Sample includes full time nonfarm wage and salary workers age 25-64 with at least 26 weeks of employment. "Male" Professional jobs are professional jobs excluding nurses and K-12 and other non-college teachers.

Table 4: Decomposition of Gender Wage Gap, 1980 and 2010 (PSID)

Variables	1980		2010	
	Effect of Gender Gap in Explanatory Variables		Effect of Gender Gap in Explanatory Variables	
	Log Points	Percent of Gender Gap Explained	Log Points	Percent of Gender Gap Explained
A. Human Capital Specification				
Education Variables	0.0129	2.7%	-0.0185	-7.9%
Experience Variables	0.1141	23.9%	0.0370	15.9%
Region Variables	0.0019	0.4%	0.0003	0.1%
Race Variables	0.0076	1.6%	0.0153	6.6%
Total Explained	0.1365	28.6%	0.0342	14.8%
Total Unexplained Gap	0.3405	71.4%	0.1972	85.2%
Total Pay Gap	0.4770	100.0%	0.2314	100.0%
B. Full Specification				
Education Variables	0.0123	2.6%	-0.0137	-5.9%
Experience Variables	0.1005	21.1%	0.0325	14.1%
Region Variables	0.0001	0.0%	0.0008	0.3%
Race Variables	0.0067	1.4%	0.0099	4.3%
Unionization	0.0298	6.2%	-0.0030	-1.3%
Industry Variables	0.0457	9.6%	0.0407	17.6%
Occupation Variables	0.0509	10.7%	0.0762	32.9%
Total Explained	0.2459	51.5%	0.1434	62.0%
Total Unexplained Gap	0.2312	48.5%	0.0880	38.0%
Total Pay Gap	0.4770	100.0%	0.2314	100.0%

Notes: Sample includes full time nonfarm wage and salary workers age 25-64 with at least 26 weeks of employment. Entries are the male-female differential in the indicated variables multiplied by the current year male log wage coefficients for the corresponding variables. The total unexplained gap is the mean female residual from the male log wage equation.

Table 5: Decomposition of the Gender Log Wage Gap by Unconditional Distribution Percentile (PSID)

Percentile	1980		2010	
	<u>Specification</u>		<u>Specification</u>	
	Human Capital	Full	Human Capital	Full
A. Effect of Covariates				
10th percentile	0.1767 (0.0234)	0.2729 (0.0374)	0.0721 (0.0249)	0.1648 (0.0453)
50th percentile	0.1215 (0.0167)	0.2381 (0.0279)	0.0237 (0.0151)	0.1274 (0.0235)
90th percentile	0.1139 (0.0188)	0.2281 (0.0260)	0.0265 (0.0203)	0.1246 (0.0329)
B. Effect of Wage Coefficients				
10th percentile	0.2958 (0.0429)	0.1886 (0.0487)	0.1134 (0.0359)	0.0319 (0.0511)
50th percentile	0.3876 (0.0220)	0.2598 (0.0275)	0.1836 (0.0231)	0.0835 (0.0255)
90th percentile	0.3316 (0.0269)	0.2336 (0.0285)	0.2749 (0.0341)	0.1790 (0.0357)
C. Sum of Covariate and Wage Coefficient Effects				
10th percentile	0.4725 (0.0367)	0.4615 (0.0353)	0.1855 (0.0266)	0.1967 (0.0314)
50th percentile	0.5091 (0.0226)	0.4979 (0.0232)	0.2073 (0.0236)	0.2109 (0.0211)
90th percentile	0.4455 (0.0314)	0.4617 (0.0311)	0.3014 (0.0346)	0.3036 (0.0342)

Notes: Sample includes full time nonfarm wage and salary workers age 25-64 with at least 26 weeks of employment. Entries are based on the decomposition of the unconditional gender log wage gap at each indicated percentile, based on methods in Chernozhukov, Fernández-Val and Melly (2013). Standard error are in parentheses and are computed by bootstrapping with 100 repetitions.