Risk Aversion or Discrimination?
Understanding the Gender Pay Gap Using Matched Data

by

Sandra E. Black
Department of Economics
UCLA, IZA and NBER
sblack@econ.ucla.edu

Astrid Kunze
Norwegian School of Economics and Business Administration and IZA
astrid.kunze@nhh.no

Kjell G. Salvanes
Norwegian School of Economics and Business Administration
Statistics Norway and IZA
kjell.salvanes@nhh.no

May 18, 2004

Preliminary and Incomplete

Introduction

The last 20 years has seen great labor market progress of women across many different countries. While a number of papers have attempted to decompose these changes by focusing on differences in human capital between men and women along with changes in the overall structure of wages[1], there is a much smaller literature relating differences in wages to differences in employer characteristics. In particular, how much can be explained by the sorting of workers across establishments? Is it true that women work in industries/occupations/jobs that pay less, and is this changing over time? If yes, what are the possible explanations for this sorting? Is it discrimination or is it gender differences in preferences?

Much of the literature of the sorting of workers across industries, occupations, and establishments has been limited by the absence of large matched employer-employee datasets. The two most prominent studies are the one by Groshen (1991) and Bayard, et al. (2003). However, these two studies come to very different conclusions. Groshen's study of five specific industries suggests that sorting can explain almost the entire wage gap; after controlling for percent females in the establishment, occupation, and establishment-occupation cell, she concludes that no gender wage gap remains. In their more representative study, Bayard et al. (2002) conclude that a substantial gap remains after controlling for occupation, industry, and establishment-occupation cells. The question therefore remains: Are men and women treated equally within occupations in the same firm?

This paper has two goals. The first is to reconcile the findings in the literature using a unique employer-employee matched dataset from Norway that follows the population of individuals over 15 years and has detailed information on occupations and jobs along with actual weekly hours of work as well as weekly pay including bonus payments. Because of the

1 See Blau and Kahn (1999) for more details.

availability of this detailed data, we are better able to examine the sorting of workers and determine which establishment characteristics can help explain differences in pay between men and women. Importantly, our data also contains information on plants' revenues and profits. This allows us an extra dimension with which to explore the relationship between sorting and firm performance.

We find evidence that sorting of workers can explain some, although not all, of the observed gender pay gap between men and women. Given the evidence that sorting of workers does play some role in the determination of gender differences in pay, our second goal will be to seek to understand what causes this sorting. There are two obvious explanations. The first is that women have preferences for these lower-paying jobs because of other features of the job. One possibility is that men and women have different preferences for risk and lower paying firms also provide lower variability in wages. While there is some evidence of this, the literature is quite limited in this area. The second is that firms are engaging in discriminatory behavior and paying equally capable women less than men.

Preliminary evidence suggests that women are more likely than men to work in firms with low profits as well as low variability of profits. We first document and then explore this relationship in an attempt to distinguish whether this represents risk-aversion on the part of women or discrimination on the part of the firm.

Section 2 of our paper discusses the relevant literature on sorting and gender wage differentials. Section 3 discusses our data, while section 4 presents our methodology and results. Section 5 then explores the results a bit further, focusing on the relationship between gender pay and firm performance, both in levels and variability. Section 6 presents concluding remarks and discusses future extensions to this work.

Background

There is a large literature examining gender differences in labor market outcomes, particularly focusing on the role of changing human capital accumulation or differences in the overall wage structure as explanations.[2] Because of data limitations, however, fewer studies have been able to examine the sorting of workers across firms and jobs.

Dolado, Felgueroso, and Jimeno (2003) compare the European Union and the United States in terms of occupational segregation by gender and find that occupational segregation is higher in the EU than in the US and is mostly due to a lower share of women in executive and managerial jobs. Given the broad nature of the paper, covering a number of countries but with data limitations including only 1 year of data, these results provide a good starting point for more detailed work on the topic.

Some of the best known papers of this type include work by Groshen (1991) and more recent work by Bayard, Hellerstein, Neumark and Troske (2003). Groshen uses data on a limited number of sectors and finds evidence that sorting of men and women into industries and jobs can explain much of the gender wage gap. Bayard, et al, however, using a different dataset on matched employer employees covering more industries and occupations, find sorting to be less significant and argue that there is a substantial residual left unexplained, even after extensive controls for industry and occupation are included.

Unfortunately, both of these papers suffer from data limitations. Groshen uses the Bureau of Labor Statistics Industry Occupational Wage Surveys with information on establishments, wages and jobs for one year in six manufacturing industries. Bayard, et al, have more extensive coverage and are not limited to manufacturing sectors, but they are limited in the quality of their match. They take self-reported demographic characteristics from the Census long

2 Blau and Kahn (1997), O'Neill and Polachek (1993)

file and matched the addresses of the place of employment to the Census Bureau SSEL file, which contains information on payroll expenses, employment and whether or not the establishment is part of a multi-establishment firm. Because workers and firms are matched based on industry-cell location, they are only able to match establishments unique to an industry-location cell. The matching process is not random, leaving them with a selected group of workers and firms. Our data are superior in a number of ways, including high quality information on a number of worker and firm characteristics covering the population of firms and workers within these firms.

Datta Gupta and Rothstein (2001) use Danish data to examine how the segregation of women into occupations, industries, establishments, and job cells impacts the gender wage differential. They conclude that occupation has a much larger role than industry or establishment in accounting for the gender gap in full-time private sector wages in Denmark. However, substantial residual wage inequality persists even within job cells.

There are two advantages of our paper over the previous literature. First, we have detailed establishment level and individual level data on the population of workers and firms in the manufacturing sector in Norway over an extensive time period, providing a significant data advantage over studies on the United States. Within this data set, we have very detailed job information, including a job hierarchy measure within occupation that allows us to compare levels not only within occupation but across occupations as well. The second advantage of our paper is our ability to match the employee-firm data to firm performance measures such as profits over an extended period of time. As a result, we can take the next step in this literature and focus on the potential causes of sorting among women and men.

Data

For the empirical analyses we use Norwegian matched employer-employee data. The data set consists of employment and wage data for white collar workers based on data for the period 1986 to 1997. The data are comprised of three different datasets.

The primary data set is a matched employer-employee data set for the population of Norwegian workers and establishments for the period 1986-1997. This data set consists of information on education, annual earnings, family background, family characteristics, and income history for each individual dating back to 1967. Workers have a unique personal identifier with which they are matched to every plant they worked for in the period as well to other labour market outcomes such as employment status and a variety of employment programs. Worker characteristics are collected from a variety of sources and are considered to be of extremely high quality.[3]

For the manufacturing sector, we are then able to match firms to information on profits and other plant level variables used in the analysis. Finally, since we are focusing on the sorting of workers across occupations, jobs and plants, we add in detailed information on the occupation, job and hourly wage data for the population of white collar workers employed by companies that are members of the main employers' association in Norway, the Confederation of Norwegian Business and Industry (NHO).[4] As a result, our analysis is restricted to white collar manufacturing workers and for the years 1986-97.

---

3 See Møen, Salvanes and Sørensen (2003) for a description of this matched employer-employee data set for Norway.

4 The employers' organisation (NHO) has about 16,000 member companies which employ about 450,000 workers (white and blue collar) covering the private sector in Norway. The total labour force in Norway was approximately 2.3 million workers in 2003, with one third is employed in the public sector. The data is based on establishment records for all workers which are members of the organisation. Norwegian law requires all employers to report data on wages and employment to Statistics Norway. Until 1997 NHO collected data for their member plants under this law, and Statistics Norway collected data for all other plants. From 1997 Statistics Norway collects data for all sectors. The data are considered to be very accurate in terms of hourly and monthly wages as well as hours worked, since the wage data was a major input in the collective bargaining process in Norway between the employers' organisation and the unions. See Holden and Salvanes (2003) for an assessment of the wage data from this data source as compared to other sources of earnings data from Norwegian registers.

The wage data we are using in this study are based on the monthly salary of white collar workers reported on September 1. The wage includes the value of fringe benefits, overtime, and bonuses. Hourly wages are constructed from monthly gross earnings divided by the contracted number of working hours per week times 4.33 and are deflated using the consumer price index for Norway with 1990 as the base year.[5]

Our data also contain detailed information on work experience for all workers, which is particularly important when considering gender differences in wages.[6] Work experience is constructed from the pension point history file. This file contains the number of pension points a worker has collected throughout a year and the yearly total income.[7] Importantly, we observe when someone leaves and then re-enters the labour market; this is particularly relevant for women, who may have intermittent labor market experience during periods of child-rearing.[8]

One of the key advantages of the Norwegian data over other data, including that from other Scandinavian countries, is our ability to look at levels of jobs within occupations. That is, we know the hierarchy of jobs within a particular occupation, so we can see how much segregation there is within occupations within plants. In addition, the Norwegian data is structured in such a way that levels within hierarchies are consistent across occupations, so we can look at the more general question of whether, regardless of occupation, women are in lower-level jobs.

To create this structure, each worker is assigned to an occupational group and a level within the occupational group. The groups include technical white collar workers; foremen;

5 Wages are trimmed such that nominal wages below 50 Nkr per hour and above 300 Nkr are excluded.
6 See O'Neill and Polachek (1993) for a discussion of the importance of experience in studying the gender wage gap.
7 Note that workers collect pension points as well when they receive benefits through unemployment or parental leave, sickness absence.
8 In 1986 10 percent of women in our data work part time and in 1995 30 percent are reported in part time work. Exclusion of part timers leads to exclusion of approx. 150000 person year spell.

administration; retail sales, storage, and a miscellaneous group for those workers who don't fall within the first five groups. Within occupational groups, each worker is assigned a level from zero (the highest) to 7. These codes are created and assigned by the Norwegian Employer's Organization to be comparable across firms and industries and are created to use for wage negotiations; as a result, the data are considered very reliable.

Our final sample includes full-time white collar manufacturing workers from 1986-1997; this is comprised of approximately 28,500 male observations per year and approximately 7,000 female observations per year.[9]

Summary Statistics

Table 1 presents summary statistics for individual characteristics by gender and year for our sample. Men are somewhat older and have slightly more years of education. Our measure of actual years of experience shows that at the mean men have worked approximately 4 years more than women; this difference, nevertheless, has decreased by 0.5 years across the 10 period of observation. Because we have selected full-time workers, hours of work do not vary among men and women. Comparison of the mean hourly wages reported in 1990 prices and in NOK show a substantial gender wage gap in the manufacturing sector among white collar workers. In 1987 the female to male hourly wage ratio was 0.73, and in 1997 it was 0.79. This is slightly higher than for the population in Norway which had a wage ratio of 0.72 in 1987 and 0.77 in 1997. While these ratios may seem low, they are in fact higher than comparable measures for representative data sets for the U.S. where a ratio of 0.65 was reported for 1987, and around 0.73

9 The employer-employee matched date file source we use in the following includes 1,223,036 year person observations of which 71 percent are men. For our analysis sample, we exclude from this data source workers for which we do not have firm characteristics, and specifically profit information. This is we exclude those not in the manufacturing sector and who are not white collar workers. This leads to the exclusion of 710,956 person year observations. Furthermore, we exclude workers possibly still in education, age less than 23, and (early) retired people, age older than 60. This results in the exclusion of another 36,707 observations. Finally, in this version of the paper we focus on full-time workers which leads to the exclusion of another 46,833 observations.

for 1997 (Blau and Kahn (2000)). Similar to the international evidence for the 80s, we observe a downward trend in the gender wage gap.

Looking at the wage gap across the wage distribution at various percentiles, we find that the raw gap is actually larger at the upper end of the distribution. Figure 1 plots logarithmic wage differentials between men and women at the 10^{th} percentile, the median and the 90^{th} percentile of their corresponding wage distributions. Differentials vary by almost 20 percentage points although the evidence suggests a significant improvement in women's relative status over time. The improvement is particular pronounced among the upper tail of the distribution, consistent with earlier work suggesting the existence of a glass ceiling in other Scandinavian countries. (Albrecht et al 2003) In our empirical analysis, we will be able to directly examine such effects since we have information not only on the occupation people work in but also at what level within that occupation the individual is located.

As a first measure of the degree of sex segregation in the labour market, we compute the Duncan and Duncan index grouping individuals by education, industry, occupation, and level within occupation. (See Table 2.) To show the trend over time, we report the index for three years.[10] Note that an index close to 1 means complete segregation in the labour market; thus, men and women work in completely different fields. An index close to zero means an equal distribution. The index itself can be interpreted as how many men and women would have to change work place cells in order to achieve an equal distribution.

These indices suggest a large degree of sex segregation in the labour market along almost all the dimensions we examine. On a more positive note, the movement over time is towards a more equal distribution in almost all dimensions (except across industries). In line with the

10 We calculate the Duncan and Duncan index as $DD = 0.5 * \sum_{j=1}^{J} |(f_j / F) - (m_j / M)|$ where j=1,...J work place cells, and f_j (m_j) is the number of women (men) in work place cell j. F(M) is the total number of female (male) workers.*

evidence from other countries, this is what one would expect given that women are catching up to men in terms of education and other human capital measures and there is evidence of decreasing employment discrimination and an expansion of white collar jobs over this time period. We find little evidence of sorting across industries, although this is not surprising given that we are focusing only on the manufacturing sector. However, there is significant sorting across education and occupation cells and largest across hierarchies within occupations. This is quite consistent with the fact that women are under-represented in managerial levels and over-represented in lower-level jobs.

Methodology and Results: Evidence of Sorting?

The methodology we apply is straightforward and similar to that employed by Groshen (1991) and Bayard et al (2003), although the earlier studies were limited in the quantity and quality of their data and often used measures of the percentage female in a particular cell as opposed to the more flexible cell indicators. We investigate the influence of worker sorting on the gender wage gap by estimating a standard log wage equation including an indicator variable if the individual is female and adding controls for both worker and firm characteristics to understand how much of the differential can be explained by these characteristics. Table 3a presents the results for 1987 and Table 3b presents the results from 1997 data. Comparing Column 1 in the two tables, we can see that the raw gender wage gap among white-collar manufacturing workers declined between 1987 and 1997. However, a significant fraction of this improvement can be explained by changing demographics of women; including indicators for level of education (66 levels/types of education), experience, and experience squared reduce the gap to approximately 12.3 percent in 1987 and 13.5 percent in 1997. This is quite consistent with the evidence from the United States which finds a substantial role for changing demographic characteristics in recent improvements in women's labor market status. (Blau and Kahn 1997).

Including indicators for 5-digit industry and occupation groups (7 different administrative groups) exacerbate the gender wage gap, suggesting that women are sorted into higher paid industries and occupations. (See Columns 3 and 4.) Columns 5 and 6 demonstrate that level in the hierarchy within occupations and sorting across plants does appear to explain a significant portion of the remaining gender wage gap; women appear to sort into lower-wage plants and especially into lower level within occupations.

Columns 1-4 suggest our results are consistent with those of Bayard et al. (2003) using data from the United States. Although the sorting of workers is important in explaining the gender wage gap, a large proportion – a little less than half - of the gender gap is still attributable to the gender of the worker. However, when we do include the very detailed information about where in the hierarchy within occupation cells a worker is placed, more than two thirds of the gender gap disappears, indicating that sorting across hierarchy levels of occupation is very important. Once we control for the within-occupation hierarchy, our results begin to resemble the results in Groshen (1991). She reports that most of the wage gap is due to job-cell segregation and very little of the within job-cell wage gap persists; up to 6 percent. Our findings suggest that sorting of workers can explain much more of the gender wage gap than suggested in Bayard et al. (2003), although not as much as the findings reported in Groshen (1991).

Explanations

Given that we find evidence of sorting of workers across occupations, plants, and industries, we next explore possible causes of this sorting. In particular, how can firm characteristics affect the sorting of workers? Do women have different preferences for establishment characteristics? One of the main relationships we investigate is whether there is a correlation between gender and wage volatility. Do women prefer lower-risk establishments? There are a number of papers that have attempted to uncover evidence of different preferences for risk among men and women.[11]

Barber and Odean (2001) test theories from the psychology literature that suggest gender differences in ”overconfidence” using common stock investments of men and women and conclude that men do trade excessively, and significantly more than women. In addition, they provide evidence that women appear to invest in slightly less-risky assets, providing some evidence that women may in fact be more risk averse. Sunden and Surette (1998) show that gender significantly affects how individuals choose to allocate assets in defined-contribution plans, with women investing in lower-risk assets.

Most recently, DeLeire and Levy (2001) use data from the Current Population Survey and the Bureau of Labor Statistics to understand the relationship between occupational choice and risk. They estimate conditional logit models of occupation choice as a function of the risk of work-related death and other job characteristics and find evidence that women choose safer jobs than men. Within gender, single mothers and fathers are most averse to fatal risk, presumably because they have the most to lose. They conclude that men and women’s different preferences for risk can explain about one-quarter of the fact that men and women choose different occupations. We focus on another form of risk--income uncertainty—and how that affects sorting of workers among the population of Norway.

It may also be that women prefer a less-competitive environment, and high variability in profits is a proxy for the types of competition within a firm. Recent evidence by Gneezy,

11 Garen (1988) examines the relationship between job risk and wages and finds evidence that risk is an important determinant of wages, suggesting that unobserved heterogeneity in the returns to risk is important and that OLS underestimates the wage premia for fatality and injury risk. However, they do not distinguish between effects by gender.

Niederle, and Rustichini (2003) provides some evidence that women do not perform as well as men in a competitive environment. Using experimental data from a computer lab, Gneezy et al show that when women and men are not competing for a particular "reward", performance is not statistically different between the two groups. However, when they are competing, men's performance improves while women's does not change and men therefore do better (statistically significant). While it is difficult to extrapolate these results to the work environment, they may suggest that women may seek to avoid "tournament" style jobs/firms, and hence may sort into firms with less variability in wages/profits.

In order to understand differential sorting based on risk among women and men, we turn to the relationship between establishment performance such as profitability and the variance of profitability over time and the sorting of women into these plants. Table 4 provides summary statistics of our plant level data. We use two measures of variability of profits. The profit variable is defined as revenues minus labour costs minus costs of material and the user price of capital.[12] We generate the standard deviation of real profits (1990NK) within a plant across the entire period in our sample, that is 1987 until 1997. A limitation of this variable is that it does not capture the tails in the distribution of profits well. Therefore, as an alternative measure, we also examine the difference in the 90th percentile in the distribution of real profits within a plant and the 10th percentile. The percentage of females within a plant is the ratio of number of females in all years s within plant j and all workers within plant j over the full time period.

Comparison of the two measures of spread in profitability shows that the distribution does appear to have thick tails which are not picked up by the standard deviation measure. The

12 Because some firms may fail to report profits for all years, we drop observations where the level of profits=0 and the standard deviation of profits is equal to zero.

percentage of females in white collar manufacturing jobs is 28.8 percent, which is lower than in the labor market overall.[13]

In order to treat measures of variability in profits as a proxy for "risk", we must first establish that variability in profits is correlated with variability in wages. Table 5 provides some evidence of this, showing the relationship between profit variability (within plant over time, from 1987-1997) and wage variability (within plant over time). We can see that there is a positive relationship, even controlling for the percentage female. Although we are interested in the sorting of women based on income risk, or the variability of wages, we use profits as a proxy for wages. We are interested in the variability induced by the inherent riskiness of the firm; individual wages incorporate both establishment level risk as well as individual level risk. Variability in profits provides a good proxy for the first type of variability. There is an extensive literature relating profits to rent-sharing among workers. (See Black and Strahan (2001) and Bertrand and Mullainathan (1999) for evidence.)

We next look at the relationship between the average level of profits in a firm and the level of wages, including controls for the percentage of women in the firm. Table 6 presents these results, and we can see that the percentage of women in a firm is associated with lower levels of profits. However, it is not clear what drives this result, as this could be the result of sorting and not a causal relationship.[14]

In order to push this relationship one step further, we next investigate whether women work in plants with less variability of profits, holding other characteristics constant. We assume that variability of profits within a plant over time captures wage risk for a worker. Workers in

13 Labour force participation rates in Norway differ by approximately 10% among men and women.

14 This negative correlation between percentage women and establishment outcomes has been observed in the data but there is little evidence to support any time of causal relationship. See Black and Lynch (2001) for one such instance.

plants with very low variability in profits have little hope for obtaining part of the rents, and, hence, they may experience relatively stable wages. Workers in plants with high variability in profits have, by contrast, large possibilities to take part in gains, but also losses, leading to relatively large dispersion in wage growth across individuals. Put differently, jobs in plants with low variability in profits may carry a lower wage risk and plants with high profit variability a high wage risk.

The basic empirical relationship we are interested in is:

$$prof_j = \alpha + \beta \ percentage\ females_j + u_j ,$$

where j indexes plants. The dependent variable is a measure for the variability of profits within a plant. The key parameter is β, showing the relationship between the percentage of females within a plant and variability of profits within that same plant. We would expect a negative coefficient if in fact women are less risk loving and therefore prefer to work in plants with low profit variability and hence high wage stability.

Tables 7 and 8 present our estimation results. Simple regressions show that there is a significant negative relationship between the variation in profits within a plant and the percentage of female workers within the plant. This coefficient remains significant even after controlling for heterogeneity with respect to industry and occupation, though it is reduced by approximately 20 percent with these controls.[15] A much more important variable is the firm size variable. This variable is strongly negatively correlated with the percentage of female variable and, therefore, the coefficient is reduced by 20 percent. Nevertheless, this does not change the main finding and we still find a significant and negative relationship.

15 The F-test suggests that the coefficients are jointly significant at the 5% level.

Conclusion

Using a unique matched employer-employee dataset from Norway, we investigate the effect of sorting across industry, occupation, job, and level within job on the relative wages of women. While the evidence using data from the United States is mixed, we find evidence that worker sorting can explain a substantial portion of the gender wage gap in Norway. In addition, the trends in female relative wages and human capital accumulation in Norway are similar to the patterns in other European countries as well as the United States, suggesting that the results we find here may provide insight into women's progress in other countries as well.

The second section of the paper turns to understanding possible explanations for sorting across firms by gender. The existing literature provides some evidence that women and men may have different preferences for risk. In this paper, we examine the relationship between worker sorting and the variability of profits as a proxy for risk and find that women appear to sort into establishments with lower average profits (and pay) but also lower variability of profits. These results provide intriguing suggestive evidence that there may be a relationship between women's preferences for risk and their choice of plant even within sectors. Future versions of this paper hope to explore this relationship more fully, examining differences across plant sizes as well as possible differences among women with different education or family structures. We also hope to rule out other explanations, such as heterogeneity of plants along other dimensions.

To date, our results are based on examining white-collar workers in the manufacturing sector. Future work will also incorporate information on the service sector and the public sector as well as information on blue-collar workers. We suspect, however, that including these other sectors will only serve to reinforce our findings.

References

Albrecht, James, Anders Bjorklund, and Susan Vroman. "Is There a Glass Ceiling in Sweden?" Journal of Labor Economics, Vol 21, Number 1, January 2003.

Barber, Brad M. And Terrance Odean. "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment." Quarterly Journal of Economics, February 2001.

Bayard, Kimberly, Judith Hellerstein, David Neumark, and Kennth Troske. "New Evidence on Sex Segregation and Sex Differences in Wages from Matched Employee-Employer Data." Journal of Labor Economics, 2003, Vol 21, No. 4.

Bertrand, Marianne and Mullainathan, Sendhil. "Is There Discretion in Wage Setting? A Test Using Takeover Legislation." *Rand Journal of Economics*, Autumn 1999, *30*(3), pp. 535-54.

Black, Sandra E. and Lisa M. Lynch. "How to Compete? The Impact of Workplace Practices and Information Technology on Productivity." Review of Economics and Statistics, August 2001.

Black, Sandra E. and Philip Strahan. "Division of Spoils? Rent Sharing and Discrimination in a Regulated Industry." American Economic Review, September 2001.

Blau, F. and L. Kahn. "Gender Differences in Pay." Journal of Economic Perspectives. Fall, 2000. pp. 75-99.

Blau, F. and L. Kahn, "Swimming Upstream: Trends in the Gender Wage Differential in the 1980s." *Journal of Labor Economics*, 15 (1997), 1-42.

Blau, F. and L. Kahn, The Gender Earnings Gap: Learning from International Comparison," *American Economic Review*, May 1992, 533-38.

Datta Gupta, Mabanita, and Donna S. Rothstein. "The Impact of Workers and Establishment Level Characteristics on Male-Female Wage Differentials: Evidence from Danish Matched Employee-Employer Data." Centre for Labour Market and Social Research Working Paper 01-09, August 2001.

DeLeire, Thomas and Helen Levy. "Gender, Occupation Choice and the Risk of Death at Work." NBER Working Paper 8574, November 2001.

Dolado, J.J., F. Felgueroso, and J.F. Jimeno. "Female Employment and Occupational Changes in the 1990s: How is the EU Performing Relative to the U.S.?" European Economic Review 45 (2001) 875-889.

Garen, John, "Compensating Wage Differentials and the Endogeneity of Job Riskiness." Review of Economics and Statistics, February 1988. Vol 70, No 1.

Gneezy, U; M. Niederle, and A. Rustichini, "Performance in Competitive Environments: Gender Differences." Quarterly Journal of Economics, August 2003.

Groshen, Erica L. "The Structure of the Female/Male Wage Differential: Is it Who You Are, What You Do, or Where You Work?" Journal of Human Resources, Vol 26, Number 3, Summer 1991 pages 457-472.

Groshen, Erica L. "Sources of Intra-Industry Wage Disperation: How Much Do Employers Matter?" Quarterly Journal of Economics, August 1991.

O'Neill, J., and S. Polachek, "Why the Gender Gap Narrowed in the 1980s," *Journal of Labor Economics*, January 1993, part 1, 205-228.

Sunden, Annika E. and Brian J. Surette. "Gender Differences in the Allocation of Assets in Retirement Savings Plans." American Economic Review, May 1998, p.207.

Table 1
Summary Statistics

	Men		Women	
	1987	**1997**	**1987**	**1997**
Age	42	42	38	40
	(9)	(9)	(10)	(9)
Education	11.9	12	11.0	11.4
	(2.6)	(3.4)	(2.0)	(3)
Experience	12.4	17.8	7.9	13.9
	(3.9)	(6.6)	(3.3)	(5.4)
Hours worked per week	37.9	37.3	37.8	37.2
	(.28)	(.06)	(.32)	(.08)
Wage	124.1	135.8	91.02	107.6
	(33.9)	(34.4)	(19.99)	(27.2)
N	23,500	29,628	5,210	9,119

Note: Calculated using a Norwegian sample of full-time workers, aged 23-60 and in manufacturing white collar workers, 1987-1997.

Table 2
Gender Segregation - Duncan and Duncan Index

	Education Type	2-Digit Industry	Occupation	Job	Level
1987	0.52	0.23	0.64	0.67	0.36
1992	0.50	0.20	0.57	0.62	0.37
1997	0.49	0.22	0.51	0.58	0.32
Number of cells	220-260**	133	22	140-190**	7

Note: We calculate the Duncan and Duncan index as $DD = 0.5 * \sum_{j=1}^{J} |(f_j / F) - (m_j / M)|$ where j=1,...J work place cells, and f_j (m_j) is the number of women (men) in work place cell j. F(M) is the total number of female (male) workers.*We drop spells if hierarchy code (jobcode)<1000. **Number of cells vary across years.

Figure 1

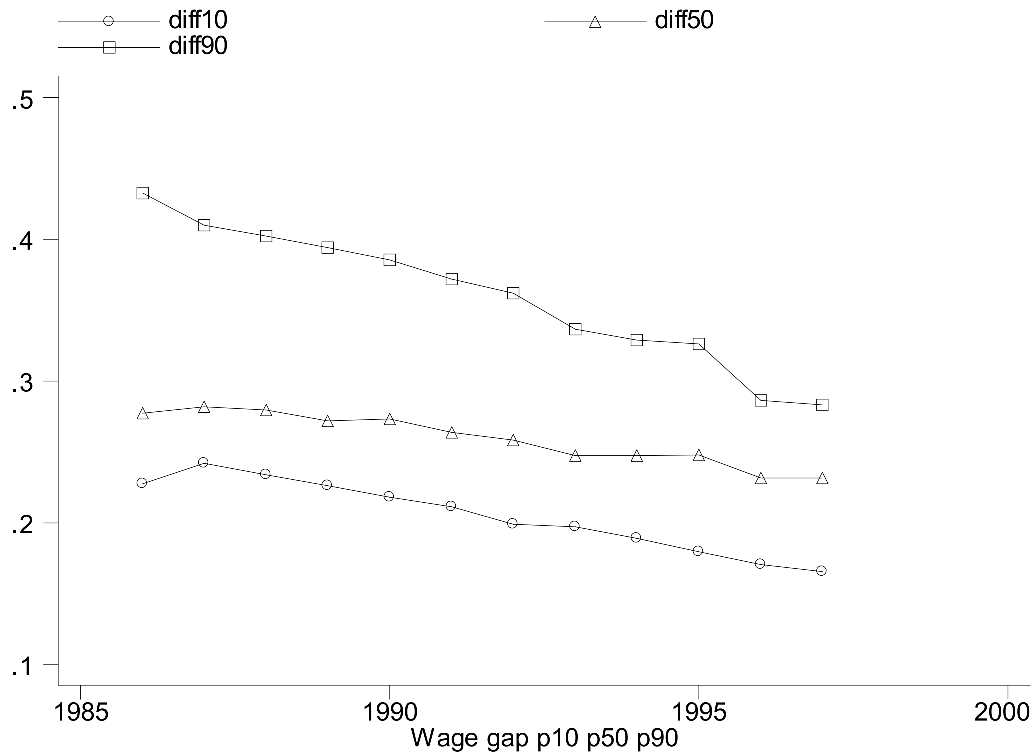
diff10
diff50
diff90
.5
.4
.3
.2
.1
1985
1990
1995
2000
Wage gap p10 p50 p90

Table 3a
Wage Equations: 1987 Data

Dependent Variable: log(wage)	(1)	(2)	(3)	(4)	(5)	(6)
Female	-.289 (.003)	-.123 (.003)	-.131 (.003)	-.151 (.003)	-.097 (.002)	-.082 (.002)
Demographic Controls	No	Yes	Yes	Yes	Yes	Yes
5-digit Industry Controls	No	No	Yes	Yes	Yes	Yes
Occupation Controls	No	No	No	Yes	Yes	Yes
Level in occupation hierarchy	No	No	No	No	Yes	Yes
Plant Controls	No	No	No	No	No	Yes
Adjusted R-Square	.1786	.5878	.6250	.6443	.7703	.8259
N	36,778	36,717	36,717	36,717	36,717	36,717

Table 3b
Wage Equations: 1997 Data

Dependent Variable: log(wage)	(1)	(2)	(3)	(4)	(5)	(6)
Female	-.224 (.003)	-.135 (.002)	-.143 (.002)	-.160 (.002)	-.076 (.002)	-.070 (.002)
Demographic Controls	No	Yes	Yes	Yes	Yes	Yes
5-digit Industry Controls	No	No	Yes	Yes	Yes	Yes
Occupation Controls	No	No	No	Yes	Yes	Yes
Level in occupation hierarchy	No	No	No	No	Yes	Yes
Plant Controls	No	No	No	No	No	Yes
Adjusted R-Square	.1437	.4888	.5219	.5364	.7692	.8110
N	48,003	47,829	47,829	47,735	47,735	46,940

Table 4: Summary Statistics at the Plant Level

Variable	Mean	Std. Dev.
Mean of Profits within plants	13,919	83,216
Std of Profits within plants	9,019	27,646
90^{th}-10^{th} Difference in Profits within plants	22,883	74,288
Percent Female within plant	.23	.20
Percent Married within plant	.65	.20
Firm Size	97	171
Age of Workers	41.9	4.6
Work Experience	14.3	2.6
Number of plants	2,251	

Note: Norwegian sample of 23-60 year old, full-time white collar workers in manufacturing.

Table 5: Relationship between variability of profits and wages within firms

	(1)	(2)	(3)	(4)
Standard deviation wage	574**	587**	513**	530
	(55)	(55)	(57)	(60.7)
	[57]	[65]	[52]	[56]
Percent Female		-9400**	-8333**	-9967
		(2818)	(3034)	(3251)
		[1609]	[1835]	[1996]
Industry Controls	No	Yes	Yes	Yes
Occupation Controls	No	No	No	Yes
Adjusted R-Square	.0448	.049	.0747	.079
N	2,251	2,251	2,251	2,251

Note: All regressions include a constant. ** significant at 5 percent significance level. *significant at 10 percent level. Bootstrapped standard errors in brackets.

Table 6: Estimation of relation between profit within plants and percentage of female workers - dependent variable profit within plant 1986-1997

	(1)	(2)	(3)	(4)	(5)
Percent Female	-7,199**	-12,500 **	-9,864**	-11,630*	-11,006*
	(8,677)	(9,386)	(9,912)	(9349)	(9424)
Firm size				168**	168**
				(10)	(10)
Percent Married					4,482
					(8,428)
Industry Controls	No	Yes	Yes	Yes	Yes
Occupation Controls	No	No	Yes	Yes	Yes
Adjusted R-Square	.0003	.0079	0.01	0.12	0.12
N	2,251	2,251	2,251	2,251	2,251

Note: All regressions include a constant. ** significant at 5 percent significance level. *significant at 10 percent level.

Table 7: Estimation of relation between variability in profits within plants and percentage of female workers - dependent variable standard deviation of profits within plant 1986-1997

Variable	(1)	(2)	(3)	(4)	(5)
Percent Female	-7294**	-5001**	-3975 **	-4992**	-4418**
	(2879)	(3064)	(3230)	(2612)	(2632)
	[1336]	[1597]	[1483]	[1317]	[1472]
Firm size				97**	97**
				(2.81)	(34.47)
				[16.7]	[17.83]
Percent Married					4122**
					(2354)
					[1362]
Industry Controls	No	Yes	Yes	Yes	Yes
Occupation Controls	No	No	Yes	Yes	Yes
Constant	10750	9522	5511	4498	1629
	(897)	(1563)	(6519)	(5272)	(5518)
Adjusted R-square	0.0028	0.045	0.054	0.3816	0.3825
N	2,251	2,251	2,251	2,251	2,251

Note: All regressions include a constant. Standard errors are in round parentheses. Bootstrap standard errors with 200 repetitions are reported in squared parentheses. ** significant at 5 percent significance level. Industry Controls are indicators for the two-digit industry group. Occupation controls are dummy variables for 22 groups.

Table 8: Estimation of relation between variability in profits within plants and percentage of female workers - dependent variable 90th percentile-10th percentile-difference in profits within a plant 1986-1997

	(1)	(2)	(3)	(4)
Percent Female	-17949**	-12352**	-9694**	-12344**
	(7738)	(8248)	(8701)	(7154)
	[3478]	[4180]	[3948]	[3438]
Firm Size				252
				(7.71)
				[49.82]
Industry Controls	No	Yes	Yes	Yes
Occupation Controls	No	No	Yes	Yes
Adjusted R-square	0.0024	0.042	0.049	0.3579
N	2,251	2,251	2,251	2,251

Note: All regressions include a constant. Standard errors are in round parentheses. Bootstrap standard errors with 200 repetitions are reported in squared parentheses. ** significant at 5 percent significance level.

Appendix Table 1
Industry Level Summary Statistics
Manufacturing

	Percent of All Workers	Percent Females*	Average Firm Size
Food	19.8	29.1	76
Textiles, Apparel, and Leather Products	4.7	29.7	59
Wood and Wood Products, including Furniture	12.2	23.1	58
Paper and Paper Products, including Printing and Publishing	14.6	37.6	113
Chemicals and Chemical Petroleum, Coal, Rubber and Plastic Products	9.2	22.2	75
Mineral Products	5.6	16.6	111
Basic Metals	6.1	16.0	270
Fabricated Metal Products	27.2	15.4	94
Other Manufacturing Industries	0.6	41.3	51

Note: * Calculated as the average percentage of females within plant. Percentage of married does not differ across these industry sectors.

Appendix Table 2
Levels of Employment

	Percentage of All Workers	Percentage Female*
Level 1 (Lowest)	2.2	73.4
Level 2	27.0	33.2
Level 3	61.2	21.0
Level 4	8.2	18.3
Level 5	0.9	9.5
Level 6 (Highest)	0.5	2.2

Note: * Calculated as the average percentage of females within plant. Percentage of married does not differ across these industry sectors.