

Is There a Glass Ceiling over Europe? Exploring the Gender Pay Gap across the Wage Distribution

Wiji Arulampalam*

Alison L. Booth[†]

Mark L. Bryan[‡]

*University of Warwick and IZA,

[†]Australian National University, University of Essex, CEPR, and IZA,

[‡]Institute for Social and Economic Research, University of Essex,

Is There a Glass Ceiling over Europe? Exploring the Gender Pay Gap across the Wage Distribution

Wiji Arulampalam, Alison L. Booth, and Mark L. Bryan

Abstract

Using harmonized data for the years 1995-2001 from the European Community Household Panel, the authors analyze gender pay gaps by sector across the wage distribution in eleven countries. In estimations that control for the effects of individual characteristics at different points of the distribution, they calculate the part of the gap attributable to differing returns between men and women. The magnitude of the gender pay gap, thus measured, varied substantially across countries and across the public and private sector wage distributions. The gap typically widened toward the top of the wage distribution (the “glass ceiling” effect), and in a few cases it also widened at the bottom (the “sticky floor” effect). The authors suggest that differences in childcare provision and wage setting institutions across EU countries may partly account for the variation in patterns by country and sector.

KEYWORDS: gender pay gap, glass ceiling

IS THERE A GLASS CEILING OVER EUROPE? EXPLORING THE GENDER PAY GAP ACROSS THE WAGE DISTRIBUTION

WIJI ARULAMPALAM, ALISON L. BOOTH, and MARK L. BRYAN*

Using harmonized data for the years 1995–2001 from the European Community Household Panel, the authors analyze gender pay gaps by sector across the wage distribution in eleven countries. In estimations that control for the effects of individual characteristics at different points of the distribution, they calculate the part of the gap attributable to differing returns between men and women. The magnitude of the gender pay gap, thus measured, varied substantially across countries and across the public and private sector wage distributions. The gap typically widened toward the top of the wage distribution (the “glass ceiling” effect), and in a few cases it also widened at the bottom (the “sticky floor” effect). The authors suggest that differences in childcare provision and wage setting institutions across EU countries may partly account for the variation in patterns by country and sector.

Although the mean gender wage gap has been extensively studied in the labor economics literature, only relatively recently has attention shifted to investigating the degree to which the gender gap might vary across the wage distribution and why. Albrecht, Bjorklund, and Vroman (2003), using 1998 data for Sweden, showed that the gender wage gap was increasing throughout the wage distribution and accelerating at the top, and they interpreted this as evidence of a glass ceiling in Sweden. De la Rica, Dolado, and

Llorens (2005) undertook a similar analysis using 1999 data for Spain. They stratified their sample by education group and found that the gender wage gap was expanding over the wage distribution only for the group with college/tertiary education. For less educated groups, the gender wage gap was wider at the bottom than the top. Thus it appears that in Spain there was a glass ceiling for the more educated but not for the less educated. Using a different decomposition methodology in the quantile regressions framework and Spanish data for 1995, del Rio, Gradin, and Canto (2005) obtained results similar to those of de la Rica et al.

The purpose of our paper is to investigate these issues further in order to see if the glass ceiling phenomenon was prevalent across pre-enlargement Europe (the European Union

*Wiji Arulampalam is Professor of Economics, University of Warwick, and Research Fellow, IZA; Alison L. Booth is Professor of Economics, Australian National University and University of Essex, and Research Fellow, CEPR and IZA; and Mark L. Bryan is Chief Research Officer, Institute for Social and Economic Research, University of Essex. The authors gratefully acknowledge the support of the Leverhulme Trust. For helpful comments, they thank Bernd Fitzenberger, Jose Machado, Blaise Melly, Jeremy Smith, and participants in the 7th Labour Econometrics Conference, Auckland, 13–14 August 2004, the 2nd SOLE/EALE World Conference, San Francisco, 2–5 June 2005, and the 2005 Conference of the European Panel Users Network, Colchester, U.K., 30 June–2 July 2005.

Data from the European Community Household Panel Survey 1994–2001 are used with the permission of Eurostat. The data can be purchased from Eurostat. For replication purposes, copies of the computer programs used to generate the results in the paper are available from Mark Bryan (markb@essex.ac.uk).

before the admission of 10 new countries in 2004). Using harmonized data from the European Community Household Panel, we analyze gender pay gaps across the wage distribution for eleven countries using the quantile regression (QR) framework. We investigate the extent to which gender affects the location (conditional mean), scale, and shape of the conditional wage distribution, and whether or not these patterns differ across the public and private sectors.

We first chart the gender pay gap using raw data. We then compare the raw gender gaps with estimates that control for men's and women's attributes using the QR framework. This enables us to see how the gender pay gaps would have differed if women had kept the same distribution of characteristics but had been rewarded like men. Unlike ordinary least squares (OLS), QR methods allow for the possibility that characteristics have different returns at different points of the distribution. Following Albrecht et al. (2003), we interpret a widening gender wage gap at the top of the wage distribution as a *glass ceiling*. We refer to the situation in which the gender pay gap widens at the bottom of the wage distribution as a *sticky floor*.¹ Toward the end of the paper we discuss various hypotheses that could explain our empirical findings.

The Data, Variables, and Raw Gender Wage Gap

Our data are from the European Community Household Panel (ECHP), a large-scale survey conducted annually from 1994 to 2001. The ECHP was specifically designed to be harmonized at the input stage: in most countries a standard questionnaire was used, with harmonized definitions and sampling

criteria. Although a standardized questionnaire does not overcome the nuances of interpretation and meaning between different languages, the harmonized format greatly facilitates cross-country comparisons. We include in our analysis the eleven European countries listed in Table 1. We omit Greece and Portugal owing to apparent gaps in the training data and because of the small estimating subsamples with usable information for those two countries. The ECHP data for Britain and Germany were adapted from those countries' existing national household surveys, while data from the full harmonized questionnaire are available for the other countries. Sample sizes are reported in column (5) of Appendix Table A1.

The education, industry, and occupation variables are all coded according to standard, internationally comparable definitions. Education levels are defined according to UNESCO's International Standard Classification of Education (ISCED). ISCED was intended for education policy analysis and was designed to be invariant to differences in national education systems.² The ECHP distinguishes between education completed to the lower secondary stage (ISCED 0–2), upper secondary education (ISCED 3), and post-secondary or tertiary education (ISCED 5–7). The data on industrial sector are categorized according to the European Union's Classification of Economic Activities in the European Community (NACE), and occupation is defined using the International Standard Classification of Occupations (ISCO-88). The Data Appendix lists the occupation and industry groups.

We initially estimated the gender pay gap separately for waves 2 and 8, in order to chart any changes that might have occurred between 1995 and 2001. Since there was little difference between the two sets of estimates, in our main model we estimate the gender gap over the entire sample of waves 2 to 8 inclusive,³ pooling all the waves and

¹Booth, Francesconi, and Frank (2003) first defined a *sticky floor* as the situation arising when otherwise identical men and women might be appointed to the same pay scale or rank, but the women are appointed at the bottom and the men further up the scale. Such a strategy can evade some discrimination laws, since the appointment rank is the same. Here we use the term more generally to describe the situation in which the gender pay gap widens at the bottom of the wage distribution, as will be further explained below.

²For details, see http://www.unesco.org/education/information/nfsunesco/doc/isced_1997.htm.

³We omitted wave 1 for two reasons. First, it does not contain information about whether or not the respon-

also including wave dummies as explanatory variables in addition to the usual set of exogenous variables. For the pooled sample we do not require individuals to be present in all waves or in consecutive waves. We therefore have new entrants across waves, and we lose some individuals through attrition. Thus we have a changing composition of individuals. As reported in the final column of Table A1, respondents were typically present in the panel for four waves. As we discuss below, we account for multiple observations on the same individuals in the calculation of the standard errors.

Because we wish to avoid conflating issues having to do with gender and early educational enrollments, we exclude from our analysis individuals under the age of 22 years, as well as paid apprentices and those on special employment-related training schemes (who account for less than 1% of the sampled age group). Among older workers there may also be differential withdrawal from the labor force, depending, for example, on how early retirement schemes operate. We therefore exclude workers aged 55 years and over. For each country, our estimating subsamples—stratified by gender—comprise full-time and part-time public and private sector employees who were (i) between the ages of 22–54 years inclusive, (ii) working at least 15 hours per week, (iii) not employed in agriculture, and (iv) with valid observations on all the variables used in the wage equations. The 15-hour-per-week cutoff was necessary because of the nature of the ECHP data, where, in the first two waves, we were unable to distinguish individuals regularly working fewer than 15 hours from those out of the labor force. In addition, for those working

fewer than 15 hours, the ECHP across all waves provides no information on firm size, public/private sector, or tenure. Thus our estimating subsamples will under-represent low-hours part-timers.⁴

The dependent variable is the log of the average hourly wage, including overtime payments, in the respondent's main job, deflated to 2001 prices.⁵ The deflators are the European Union's harmonized indices of consumer prices (HICP; see *Eurostat Yearbook 2003*). The ECHP provides a rich set of controls, which are listed in the notes under Table 2. Unfortunately, the ECHP does not collect any information on either union status or union coverage, and so we are unable to control for those variables in our estimation. The data do not contain information on labor market experience, but we do include tenure (6 categories) and a binary indicator denoting whether the individual has had a spell of unemployment since 1989, which should capture some of the variation in workers' labor market attachment.

Throughout, we estimate our models for three subsamples of data: first a combined sample comprising both public and private sector workers, and then two disaggregated subsamples, comprising public and private sector workers, respectively. In the remainder of this section, we discuss the estimates of the raw gender wage gap for all three samples and then briefly present the methodology used to estimate the gender wage gap.

Table 1 reports estimates of the raw gender wage gap by country. The raw gap for the combined sample of both public and private sector workers appears in Panel A, for public sector workers only in Panel B, and for private sector workers only in Panel C.

dent's employment contract was fixed term/casual. If temporary contract coverage varies between men and women, temporary contracts could be an important determinant of the gender wage gap. Second, the deflator used (the EU harmonized index of consumer prices, from Eurostat) is only available from wave 2. Also note that Austria did not join the ECHP until wave 2 and that Finland did not join until wave 3 (following its accession to the EU in 1995). Thus, we have seven waves of data for all countries except Finland, for which we have six waves.

⁴For most countries, low-hours part-timers represent only a tiny fraction of workers. Exceptions are Britain (6.4% of the subsample), Denmark (3.2%), the Netherlands (9.8%), and Ireland (4.0%). In all other countries the proportion of low-hours part-timers is under 3%.

⁵The log wage was calculated from the ECHP variables as $\log(\text{wage}) = \log(\text{PI211MG} * (12/52) / \text{PE005A}) = \log(\text{normal gross monthly earnings from main job including overtime} * (12/52) / \text{weekly hours in main job including overtime})$. No specific information is provided on overtime hours and premia.

Table 1. The Raw Gender Wage Gap in Eleven European Countries, 1995–2001.

Country	Men (Proportion) (1)	Mean (2)	10 th Percentile (3)	25 th Percentile (4)	Median (5)	75 th Percentile (6)	90 th Percentile (7)
A—Pooled							
Austria	0.580	0.234	0.268	0.236	0.210	0.194	0.201
Belgium	0.538	0.100	0.100	0.100	0.074	0.080	0.150
Britain	0.505	0.246	0.238	0.248	0.234	0.248	0.252
Denmark	0.508	0.132	0.107	0.091	0.099	0.155	0.248
Finland	0.496	0.184	0.112	0.119	0.170	0.243	0.266
France	0.541	0.142	0.136	0.127	0.113	0.122	0.139
Germany	0.570	0.204	0.230	0.189	0.178	0.206	0.227
Ireland	0.551	0.201	0.253	0.233	0.209	0.166	0.129
Italy	0.593	0.063	0.089	0.072	0.054	0.037	0.028
Netherlands	0.597	0.183	0.151	0.146	0.155	0.196	0.232
Spain	0.622	0.138	0.154	0.127	0.114	0.071	0.044
B—Public Sector							
Austria	0.519	0.135	0.153	0.122	0.087	0.093	0.115
Belgium	0.478	0.073	0.061	0.058	0.033	0.065	0.136
Britain	0.349	0.212	0.213	0.185	0.216	0.197	0.217
Denmark	0.336	0.114	0.128	0.085	0.105	0.119	0.175
Finland	0.350	0.259	0.164	0.196	0.260	0.316	0.307
France	0.451	0.116	0.112	0.095	0.110	0.128	0.139
Germany	0.432	0.128	0.105	0.098	0.146	0.157	0.157
Ireland	0.521	0.110	0.133	0.140	0.079	0.040	0.093
Italy	0.513	0.006	-0.002	0.010	0.001	-0.021	0.046
Netherlands	0.481	0.200	0.144	0.187	0.191	0.196	0.232
Spain	0.527	0.054	0.083	0.068	0.058	-0.005	0.065
C—Private Sector							
Austria	0.606	0.292	0.286	0.286	0.275	0.273	0.266
Belgium	0.577	0.137	0.121	0.135	0.120	0.140	0.199
Britain	0.564	0.306	0.269	0.304	0.309	0.326	0.311
Denmark	0.636	0.134	0.104	0.115	0.091	0.167	0.240
Finland	0.600	0.167	0.121	0.135	0.146	0.199	0.242
France	0.588	0.202	0.180	0.156	0.170	0.201	0.228
Germany	0.624	0.262	0.294	0.253	0.231	0.247	0.282
Ireland	0.566	0.273	0.264	0.267	0.263	0.267	0.313
Italy	0.638	0.153	0.145	0.111	0.130	0.146	0.194
Netherlands	0.643	0.208	0.177	0.176	0.176	0.217	0.278
Spain	0.654	0.230	0.252	0.206	0.205	0.244	0.207

Notes: The raw wage gap is measured as the difference between the log male and log female hourly wage. The log wage was calculated from the ECHP variables as $\log(\text{wage}) = \log(\text{PI211MG} * (12/52) / \text{PE005A}) = \log(\text{normal gross monthly earnings from main job including overtime} * (12/52) / \text{hours in main job including overtime})$. It was then deflated to 2001 prices using harmonized indices of consumer prices (HICP) from the *Eurostat Yearbook 2003*.

Except for the coefficients in *italics*, all coefficients are significantly different from zero at the 5% level.

The data for Finland cover 1996–2001 only.

Column (1) of Table 1 indicates the male percentage of the various subsamples by country. Men formed the majority of the work force in all countries except Finland (Panel A). Britain and Denmark had a very similar gender composition, at 50.5% and 50.8%, respectively. The male proportion

was highest in Spain (62.2% of employees), closely followed by the Netherlands (59.7%) and Italy (59.3%). The mean raw gender gap is presented in column (2). Inspection of columns (3)–(7) reveals that in three countries—Ireland, Italy, and Spain—the raw gender gap is decreasing as we move from

the 10th to the 90th percentiles. In contrast, there is a striking increase of 15–16 percentage points as we shift from the 10th to the 90th percentiles in Finland and from the 50th to the 90th percentiles in Denmark. In Britain, while the raw gender wage gap was rather large (its mean was 24.6%), it was relatively constant across the distribution.

This simple comparison suggests the presence of considerable heterogeneity across our EU countries. It also shows that measuring the gender pay gap at the mean of each distribution (that is, comparing an “average” woman with an “average” man) can produce a misleadingly simple picture of how men’s and women’s wages differ. This mean gap can hide larger or smaller gaps between high-paid men and women, or between low-paid men and women.⁶

Next consider the raw gaps for the public sector, presented in Panel B. We were interested in stratifying our sample by sector because institutions in the public sector typically differ greatly from those in the private sector. In the public sector, organizations are largely non-profit and thus isolated from the rigors of the market economy. Thus, in principle, they could more easily follow “tastes for discrimination” in their wage-setting behavior. However, they are also subject to government objectives and policies. The European Union countries have adopted strong regulations in favor of equal opportunities, and it is likely that these are more stringently enforced in the public sector than in the private sector.

⁶Overall wage inequality differs substantially across countries. In our data, the country with the most compressed raw log hourly wage distribution (public and private sectors combined) is Denmark, followed by Italy, the Netherlands, Finland and Belgium, and Austria. The country with the most unequal wage distribution is Ireland, followed by Spain, Britain, and France and Germany. The 90th-10th percentile differentials of the raw log wage distributions are as follows: Austria, 0.94 log points; Belgium, 0.90; Britain, 1.20; Denmark, 0.72; Finland, 0.90; France, 1.13; Germany, 1.01; Ireland, 1.32; Italy, 0.88; Netherlands, 0.89; and Spain, 1.30. Although calculated from our samples of prime-aged workers, these figures are reasonably close to the 90-10 log wage differentials reported by Blau and Kahn (1996), whose sample included four of the countries considered here, and in OECD (1996). Both of these studies used datasets different from ours.

Simple OLS pooling tests reject joint equality of the public-private sector coefficients in every country, confirming that this is a valid separation.

As column (1) in Panel B shows, the public sector had a majority female work force in seven of our eleven countries. Only in Austria, Ireland, Italy, and Spain were men in the majority in the public sector, and even in these countries, the majority was slim (the highest proportion of men was 52.7% in Spain). In the private sector, on the other hand (Panel C), men predominated across all countries, and in six countries they accounted for over 60% of the private sector work force.

The raw wage gap measured at the mean (column 2) was generally higher in the private sector than in the public sector. While the raw average gender wage gap in the public sector was in excess of 20% in Britain, Finland, and the Netherlands, in Belgium, Italy, and Spain it was under 10%, and indeed in Italy it is found to have been insignificantly different from zero. In contrast, in the private sector, the raw average gender gap exceeded 13% in all countries, and in Britain and Austria it was close to 30%. In France, Germany, Ireland, the Netherlands, and Spain the gap was around or over 20%.

How does the raw gender wage gap vary across the unconditional distribution? In the public sector, Italy is the only country where the raw gender gap is found to have been statistically insignificant in all parts of the distribution except at the top (see column 7), where it was still very much smaller (at about 5%) than in other countries. In Finland and the Netherlands, the raw gap increases monotonically as we move up the unconditional wage distributions, and in Belgium, Denmark, and Germany, the gap is also higher toward the top of the distribution. In Ireland and Spain, the gap moves in the opposite direction. In Britain, a remarkably level raw gap of about 20% occurs across the distribution.

Measures of the raw gap exhibit similar patterns in the private sector (Table 1, Panel C). The gender wage gap varies little over the distribution in Britain. The gap increases moving up the wage distribution in Finland and Netherlands, and is also higher toward

the top in Belgium, Denmark, France, and Ireland. In contrast to what is found in the public sector in Italy, the wage gap is now significantly different from zero and is U-shaped. We find a similar pattern in Germany.

In summary, we find that in both the public and the private sectors there was a tendency in some countries for the gender wage gap to be higher at the top of the wage distribution than in the middle region, hinting at a possible “glass-ceiling” effect. However, the gender wage gap was wider at the bottom end, too, for public sector workers in five countries (Austria, Britain, Denmark, France, and Spain) and for private sector workers in four countries (France, Germany, Italy, and Spain). This hints at a “sticky floor” effect in some countries. But these are only raw gender gaps. In order to find out how much of the observed raw wage gap can be explained by the differences in the returns to various characteristics, we next turn to the quantile regression results.

Wage Gap Estimates from Quantile Regressions

The Econometric Model

Instead of looking at the effects of gender and other covariates on the conditional mean of the log wage distribution, we look at the effects of gender and other covariates on different *quantiles* of the log wage distribution.⁷ The effects of covariates on the location, scale, and shape of the conditional wage distribution can be easily estimated using a quantile regression (QR) framework. Since the QR framework allows the characteristics to have different returns at different quantiles, at each point of the distribution it can control more fully for differences between men’s and women’s wages that are attributable to their characteristics.

Following Buchinsky (1998), we specify the θ^{th} ($0 < \theta < 1$)⁸ conditional quantile of the log wage (w) distribution for the i^{th} individual ($i = 1, \dots, N$) in wave t ($t = 1, \dots, T$) as

$$(1) \quad \text{Quant}_{\theta}(w_{it} | \mathbf{x}_{it}) = \alpha(\theta) + \mathbf{x}_{it}' \boldsymbol{\beta}(\theta),$$

implying

$$(2) \quad w_{it} = \alpha(\theta) + \mathbf{x}_{it}' \boldsymbol{\beta}(\theta) + \varepsilon_{\theta it},$$

with $\text{Quant}_{\theta}(\varepsilon_{\theta it} | \mathbf{x}_{it}) = 0$.

For each sector, we estimate this model for men and women separately. Note that if the underlying model were truly a location model—in the sense that the changes in explanatory variables caused only a change in the location of the distribution of w and not in the shape of the distribution—then all the slope coefficients would be the same for all θ .⁹ We use Stata 9 to estimate the coefficients of our QR model.

We do not use the conventional method of estimating gender pay gaps in a linear regression framework that was first introduced by Blinder (1973) and Oxaca (1973). Instead, we use a different method to calculate the gender gap at the θ^{th} quantile due to differing returns adjusted for characteristics as detailed below. That is the gap measured as the difference in pay that women would have faced at the θ^{th} quantile if their labor market characteristics had been rewarded as men’s were.¹⁰

Estimation Strategy and the Decomposition Method

As noted above, we initially estimate the models separately by gender and country,

⁹Quantile regression models are more general than simple linear regression models allowing for heteroskedastic errors, since they allow for more general dependence of the distribution of w (the dependent variable) on the \mathbf{x}' s instead of just the mean and the variance of the conditional mean alone.

¹⁰We also calculated the gap using men’s characteristics as the reference, which answers the question of what the gender wage gap would be if women had men’s characteristics but were still rewarded as women (or equivalently if men were rewarded as women). The results were qualitatively similar, the main difference being that for some countries the estimates of the gap were somewhat smaller in the higher parts of the wage distribution. For example, the estimated gap at the 90th percentile (using pooled data) was appreciably smaller in six countries (typically by about 6 percentage points). The practical effect of this is to reduce the glass ceiling effect and increase the sticky floor effect. To some extent, the choice of base is arbitrary and these changes do not modify our conclusions.

⁷The linear conditional quantile regression model was first introduced by Koenker and Bassett (1978). For a recent survey of these models, see Buchinsky (1998).

⁸ $\theta = 0.5$ refers to the Median.

Table 2. Estimated Wage Gap [percentage of the raw gap explained by different returns]: Pooled Model without Industry and Occupational Dummies.

Country	OLS	10%	25%	50%	75%	90%
Austria	0.268 [114]	0.230 [86]	0.213 [90]	0.222 [106]	0.261 [135]	0.304 [151]
Belgium	0.146 [146]	0.102 [102]	0.118 [118]	0.133 [179]	0.165 [206]	0.214 [143]
Britain	0.233 [95]	0.196 [83]	0.204 [82]	0.223 [95]	0.250 [101]	0.286 [114]
Denmark	0.104 [79]	0.050 [47]	0.061 [68]	0.091 [92]	0.150 [97]	0.213 [86]
Finland	0.238 [129]	0.154 [138]	0.179 [151]	0.230 [135]	0.284 [117]	0.304 [114]
France	0.215 [151]	0.189 [139]	0.177 [139]	0.183 [162]	0.215 [176]	0.279 [201]
Germany	0.148 [73]	0.124 [54]	0.126 [67]	0.132 [74]	0.150 [73]	0.197 [87]
Ireland	0.248 [123]	0.236 [93]	0.260 [112]	0.263 [126]	0.250 [151]	0.232 [179]
Italy	0.146 [232]	0.146 [164]	0.127 [176]	0.115 [214]	0.137 [367]	0.183 [649]
Netherlands	0.131 [72]	0.047 [31]	0.072 [49]	0.107 [69]	0.173 [89]	0.242 [104]
Spain	0.185 [134]	0.236 [114]	0.213 [136]	0.180 [135]	0.147 [164]	0.124 [227]

Notes: The model includes dummies for whether training was received in the last year, age, education, tenure, marital status, health status, any experience of unemployment since 1989, part-time status, fixed term and casual contracts, region (where possible), sector, and year. Dummies were also included for cases in which there was a very large number of missing values. See the Appendix for further details.

All wage gaps were statistically significant at the 1% level.

and thus have 22 sets of estimations for each specification (Table 2 samples). Then we disaggregate by sector, and estimate the models separately by gender, sector, and country, yielding 44 sets of estimations for each specification (Table 4 subsamples). In the interest of space, and given the focus of our paper, we do not provide the complete set of estimates for each country. Instead we move straight to the calculations of the gender wage gap obtained from the QR model. Full details of all estimated effects for each country can be obtained from the authors on request. The calculations enable us to evaluate whether a glass ceiling, a sticky floor, or both are present.

All models include the full set of other controls listed beneath Tables 2 and 4 (including wave dummies). Table 2 shows the results from the wage gap calculations obtained from the pooled model, excluding controls for occupation and industry but including a dummy variable for private sector. Table 4 gives the results disaggregated by sector, without and with occupation and industry dummy variables.

We now provide more details on our calculations of the wage gaps. First we estimated the quantile regressions for each gender (and by sector where necessary). Then we calculated the predicted wage at different parts of the wage distributions by gender

(and sector). The wage gap in which we are interested measures the effect of different returns to men and women when *women's* characteristics are used in the counterfactual calculations. A positive wage gap implies that the returns to women's characteristics are lower than those of men, and a negative gap implies the reverse. Instead of using average characteristics of the female sample to calculate the counterfactuals, we follow the bootstrap procedure suggested by Machado and Mata (2005) and use the distribution of women's characteristics to calculate the decompositions directly at particular quantiles of interest.¹¹ The procedure involves estimating marginal density of wages that are consistent with the estimated conditional densities given by (2) and the hypothesized distribution of characteristics. In practical terms, this was carried out as follows:

Step 1: Generate a random sample of size $n = 5,000$ from a uniform distribution $U[0,1]$: $\theta_1, \dots, \theta_n$. This will yield a series of numbers telling us which quantiles are to be estimated.

Step 2: For each θ from step 1, estimate the coef-

¹¹A similar procedure was used by Albrecht et al. (2003) and de la Rica et al. (2005). Machado and Mata (2005, Section 2.4) provided a detailed discussion of various methodologies that have been used to calculate counterfactual densities.

ficients $\beta_m(\theta)$ and $\beta_f(\theta)$ in equation (2) using the male dataset and female dataset, respectively.

Step 3: Randomly draw 5,000 women (with replacement) and use their characteristics to predict the wages using the estimated coefficients ($\beta_m(\theta)$ and $\beta_f(\theta)$) from step 2, generating two sets of predicted wages covering the whole distribution. Note that because the block-bootstrapping procedure is used to account for clustering at the individual level, there are more than 5,000 observations in practice (see footnote 14). This enables us to calculate (i) the marginal distribution of women's wages and (ii) the marginal distribution of men's wages that would obtain if their characteristics were distributed as women's are.

Step 4: Using the distributions calculated in step 3, we estimate the wage gap as the difference between the predicted wage at each quantile using the newly generated wage distribution for women and the counterfactual distribution for men.¹²

Although we realize that individuals might self-select into a particular sector/industry/occupation, limitations of the dataset prevent us from addressing the issue of self-selection. However, in order to see how the results change, we first present the results from the estimation that pools the sectors, then separately estimate the wage gaps for each sector with and without the industry and occupational controls. We discuss the latter estimates toward the end of the next subsection.

Estimates for the Combined Sample (Public and Private Sectors)

The wage gap estimates obtained following the method just described are reported in Table 2, together with the percentage of the raw gap that is explained by different returns.¹³ To facilitate comparison with the

usual procedure, in the first column of Table 2 we report the gender wage gap estimated from ordinary least squares (OLS) using average female characteristics.

The first interesting point to note from this table is that all the estimated gender wage gaps from the model are positive. Thus, in all countries, even assuming that men and women have identical distributions of characteristics, there is a gender pay gap across the wage distribution (due to differing returns). Notice also that these estimates are all significantly different from zero at the 1% level.¹⁴

To facilitate comparison, we summarize in Table 3 the results from the QR model estimates. We first define a glass ceiling as existing if the 90th percentile wage gap is higher than the estimated wage gaps in other parts of the wage distribution by at least two percentage points. The sticky floor phenomenon is defined to exist if the 10th percentile wage gap is higher than the 25th percentile wage gap by at least two percentage points. The results are summarized in columns (1) and (5). Column (1) of Table 3 shows that there was a glass ceiling in nine countries (Ireland and Spain are the exceptions). Alternative definitions of a glass ceiling—see columns (2) and (3)—produce similar conclusions. Note also that the estimated wage gap is found to increase in Belgium, Britain, Denmark, Finland, Germany, and the Netherlands all along the wage distribution. In contrast, the estimated wage gap decreases as one moves up the wage distribution in Spain.

There is also some evidence of sticky floors, but in only two countries (Italy and Spain) using the 10-50 difference and in only one country (Spain) using the 10-25 difference. Here women at the bottom (the 10th percentile) are found to have been more disadvan-

¹²An alternative procedure to the above would be to make use of the raw female wage distribution in the comparison. Although the simulated female wage distribution using the female characteristics and returns is very similar to the raw female wage distribution, we have chosen to use the former because the comparisons are made using wage distributions that have been simulated using the same procedure.

¹³The controls included in the equations are listed under Table 2. In addition to human capital variables and job characteristics, they include year dummies to allow for cyclical effects on the gender wage gap at

each quantile. Unfortunately, the ECHP contains no information about individual union membership or coverage by collective bargaining.

¹⁴A bootstrap sample of size 200 was used for the calculation of the standard errors. The sampling procedure used in the calculation also used the block-sampling method to account for clustering at the individual level because of the panel nature of the dataset (Fitzenberger 1998).

Table 3. Summary of Table 2 Quantile Regression Results—Pooled Model.

Country	Glass Ceiling Measured by: ^a			Sticky Floor Measured by: ^b		Estimated Profile of Wage Gap along Distribution	Estimated Range of Wage Gap (%)
	90-All Gaps	90-75 Difference	90-50 Difference	10-50 Difference	10-25 Difference		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Austria	✓	✓	✓				23–30
Belgium	✓	✓	✓			Increasing	10–21
Britain	✓	✓	✓			Increasing	20–29
Denmark	✓	✓	✓			Increasing	5–21
Finland	✓	✓	✓			Increasing	15–30
France	✓	✓	✓				18–28
Germany	✓	✓	✓			Increasing	12–20
Ireland							23–26
Italy	✓	✓	✓	✓			12–18
Netherlands	✓	✓	✓			Increasing	5–24
Spain				✓	✓	Decreasing	12–24

Notes:

^aA glass ceiling is defined to exist if the 90th percentile wage gap exceeds the reference gap by at least 2 points.

^bA sticky floor is defined to exist if the 10th percentile wage gap exceeds the reference wage gap by at least 2 points.

tagged than those at the 25th percentile. In general, the wage gap at the mean is found to provide a very incomplete picture of the differing returns faced by women and men at various points of the wage distribution.

The proportion of the observed raw wage gap that is explained by the differences in returns to characteristics is shown in square brackets for each country in Table 2. A value greater than 100% implies that women had characteristics that compensated them for any “discrimination”—defined here as different returns to the same characteristics—that they faced in the labor market. For example, in five countries—Belgium, Finland, France, Italy, and Spain—women typically had better characteristics than men. The same was also true for Irish women whose earnings placed them in the top parts of the distribution. Next we turn to the estimates obtained from the disaggregated subsamples, where public and private workers are examined separately.

The wage gap estimates obtained from the separate public and private sector subsamples are reported in Table 4, together with the percentage of the raw gap that is attributable to differing returns.¹⁵ To facilitate

comparison with the usual procedure, we report in the first column the gender wage gap estimated from ordinary least squares (OLS) using average female characteristics. The findings we focus on, shown on the left-hand side of Table 4, are the results from estimations without industrial and occupational controls.

As for the pooled model, Table 4 shows that all the estimated gender wage gaps are positive, in both the public and the private sectors. Thus, in all countries, even if women had had the same returns as men, they still would have received lower pay across the wage distribution. All the gaps are significantly different from zero at the 1% level. With the exception of Finland and the Netherlands, the estimated wage gaps are also generally higher in the private sector than in the public sector. In Finland the OLS public sector gender gap was 25%, versus 21% in the private sector; and the gender gap in the public sector also exceeded that in the private sector over the whole distribution. In the Netherlands, the OLS gender wage gap was 13–14% in both sectors, with a similar increasing gap across the wage distribution.

We next discuss in more detail how the wage gap varies across the wage distribution.

¹⁵See footnote 13.

Table 4. Estimated Wage Gap [and percentage of the raw gap explained by different returns].

Country	OLS	10%	25%	50%	75%	90%
<i>Industry and Occupation Excluded</i>						
Public Sector						
Austria	0.227 [168]	0.153 [100]	0.140 [115]	0.190 [219]	0.239 [258]	0.289 [250]
Belgium	0.122 [167]	0.065 [105]	0.072 [124]	0.099 [304]	0.141 [218]	0.209 [154]
Britain	0.176 [83]	0.109 [51]	0.138 [75]	0.182 [84]	0.192 [98]	0.241 [112]
Denmark	0.089 [78]	0.058 [45]	0.063 [71]	0.084 [80]	0.122 [102]	0.181 [104]
Finland	0.255 [98]	0.158 [96]	0.192 [98]	0.247 [95]	0.298 [94]	0.313 [102]
France	0.172 [148]	0.145 [129]	0.130 [138]	0.146 [133]	0.189 [148]	0.273 [197]
Germany	0.099 [78]	0.058 [55]	0.072 [73]	0.102 [70]	0.137 [87]	0.167 [106]
Ireland	0.177 [161]	0.167 [125]	0.153 [109]	0.161 [205]	0.163 [407]	0.186 [199]
Italy	0.086 [1431]	0.031 [-1823]	0.037 [361]	0.072 [5490]	0.125 [-593]	0.169 [365]
Netherlands	0.142 [71]	0.049 [34]	0.088 [47]	0.131 [69]	0.183 [93]	0.235 [101]
Spain	0.077 [142]	0.102 [122]	0.096 [140]	0.082 [142]	0.040 [-844]	0.062 [96]
Private Sector						
Austria	0.251 [86]	0.212 [74]	0.207 [72]	0.215 [78]	0.233 [85]	0.269 [101]
Belgium	0.144 [105]	0.090 [75]	0.120 [88]	0.144 [120]	0.174 [124]	0.218 [110]
Britain	0.247 [81]	0.201 [75]	0.224 [74]	0.246 [80]	0.272 [83]	0.302 [97]
Denmark	0.118 [88]	0.045 [44]	0.081 [71]	0.110 [122]	0.163 [98]	0.209 [87]
Finland	0.211 [126]	0.134 [111]	0.165 [122]	0.207 [142]	0.250 [126]	0.284 [117]
France	0.234 [116]	0.197 [110]	0.174 [112]	0.189 [111]	0.236 [117]	0.294 [129]
Germany	0.162 [62]	0.139 [47]	0.142 [56]	0.146 [63]	0.159 [64]	0.200 [71]
Ireland	0.230 [84]	0.185 [70]	0.215 [81]	0.240 [91]	0.256 [96]	0.269 [86]
Italy	0.172 [112]	0.156 [108]	0.138 [124]	0.146 [112]	0.169 [116]	0.205 [106]
Netherlands	0.127 [61]	0.029 [16]	0.068 [38]	0.107 [61]	0.172 [79]	0.249 [90]
Spain	0.211 [92]	0.214 [85]	0.211 [102]	0.207 [101]	0.202 [83]	0.205 [99]

Notes: All models include dummies for whether training was received in the last year, age, education, tenure, marital status, health status, any experience of unemployment since 1989, part-time status, fixed term and casual contracts, private sector firm size, region (where possible), and year. Dummies were also included for cases in which

To facilitate comparisons, we summarize in Table 5 the results from the QR model estimates reported in Table 4 using the same definitions as before (see Table 3). We focus on the models excluding the industry and occupation dummies.

First, consider the public sector estimates, shown in the top panel of Table 4. In nine of the eleven countries—Austria, Belgium, Britain, Denmark, Finland, France, Germany, Italy, and the Netherlands—the gender wage gap was higher at the 90th percentile of the wage distribution than at any of the lower percentiles we examined, pointing to a widespread glass ceiling effect across Europe. Following our precise criterion for the existence of a glass ceiling—that the 90th percentile wage gap must exceed the estimated wage gaps in other parts of the wage distribution by at least two percentage points—Finland drops out of the set. However, it rejoins this

group of countries when the glass ceiling is defined by the 90-50 differential (column 3). Across the countries, the highest wage gap at the 90th percentile is found in Finland, where it increases monotonically from about 16% at the 10th percentile to about 31% at the 90th percentile. The estimated wage gap is also found to increase monotonically along the wage distribution in six other countries—Belgium, Britain, Denmark, Germany, Italy, and the Netherlands—although not as dramatically as in Finland.

We find no evidence of sticky floors, as previously defined, in the public sector: that is, the results indicate that the wage gap at the 10th percentile did not exceed the wage gap at the 25th percentile by 2+ percentage points in the public sector in any of the eleven countries. In general, Belgium, Denmark, Germany, Italy, and Spain were countries with relatively low wage gaps. Once again,

Table 4 (continued).

Country	OLS	10%	25%	50%	75%	90%
<i>Industry and Occupation Included</i>						
Public Sector						
Austria	0.227 [168]	0.191 [125]	0.163 [133]	0.191 [220]	0.221 [239]	0.266 [230]
Belgium	0.090 [124]	0.046 [74]	0.050 [85]	0.070 [215]	0.109 [169]	0.169 [125]
Britain	0.134 [63]	0.091 [43]	0.116 [63]	0.135 [63]	0.144 [73]	0.205 [95]
Denmark	0.070 [62]	0.058 [45]	0.051 [58]	0.059 [56]	0.086 [72]	0.136 [78]
Finland	0.216 [83]	0.115 [70]	0.140 [71]	0.203 [78]	0.269 [85]	0.319 [104]
France	0.096 [83]	0.092 [82]	0.077 [81]	0.078 [71]	0.108 [85]	0.167 [121]
Germany	0.122 [95]	0.111 [105]	0.110 [112]	0.118 [81]	0.140 [89]	0.147 [93]
Ireland	0.184 [167]	0.186 [139]	0.169 [120]	0.177 [225]	0.165 [413]	0.181 [194]
Italy	0.097 [1611]	0.041 [-2425]	0.047 [463]	0.081 [6182]	0.138 [-653]	0.169 [365]
Netherlands	0.121 [61]	0.039 [27]	0.070 [38]	0.112 [59]	0.160 [81]	0.218 [94]
Spain	0.083 [154]	0.090 [109]	0.079 [116]	0.095 [166]	0.069 [-1445]	0.076 [118]
Private Sector						
Austria	0.214 [73]	0.182 [64]	0.177 [62]	0.188 [68]	0.210 [77]	0.247 [93]
Belgium	0.132 [96]	0.100 [83]	0.121 [90]	0.131 [109]	0.148 [106]	0.185 [93]
Britain	0.190 [62]	0.155 [58]	0.172 [56]	0.188 [61]	0.213 [65]	0.227 [73]
Denmark	0.088 [66]	0.032 [31]	0.065 [56]	0.088 [97]	0.123 [74]	0.161 [67]
Finland	0.151 [91]	0.068 [56]	0.112 [83]	0.154 [105]	0.188 [95]	0.205 [85]
France	0.163 [81]	0.146 [81]	0.126 [81]	0.132 [77]	0.152 [76]	0.190 [83]
Germany	0.143 [55]	0.088 [30]	0.109 [43]	0.137 [59]	0.166 [67]	0.213 [75]
Ireland	0.163 [60]	0.081 [31]	0.143 [54]	0.184 [70]	0.195 [73]	0.206 [66]
Italy	0.173 [113]	0.148 [102]	0.135 [122]	0.152 [117]	0.179 [123]	0.220 [114]
Netherlands	0.131 [63]	0.059 [33]	0.091 [51]	0.123 [70]	0.168 [77]	0.222 [80]
Spain	0.181 [78]	0.173 [69]	0.178 [86]	0.184 [90]	0.189 [78]	0.176 [85]

there was a very large number of missing values. See the Appendix for further details.

All estimated wage gaps are statistically significant at the 1% level.

the wage gap at the mean is found to provide a very incomplete picture of the differing returns faced by women and men at various points of the wage distribution.

Second, consider the private sector estimates, shown in the bottom panel of Table 4 and also summarized in the bottom panel of Table 5. In contrast to the public sector, the private sector exhibits very large wage gaps. Furthermore, a glass ceiling effect is found in almost all of the eleven countries. The only exceptions are Spain, which had a fixed gap of about 20–21% along the wage distribution, and Ireland, which had a glass ceiling only under the column (3) definition. The only evidence of sticky floors is in France, where the gap at the 10th percentile was just over two percentage points higher than at the 25th percentile.

We next turn to the proportion of the observed raw wage gap that is explained by

the differences in returns to characteristics, shown in square brackets for each country in Table 4. Since the wage distributions are simulated using a single distribution of (women's) characteristics, differences over the quantiles represent interactions between the characteristics and returns. In the public sector in six countries—Austria, Belgium, France, Ireland, Italy, and Spain—women typically had better characteristics than men, and at the top of the distribution this was true for women in all countries except Spain.¹⁶ In the private sector, Finland, France,

¹⁶The very large values in, for example, the public sector in Italy merit comment. They arise because there is essentially no gap in the raw data (Table 1, panel B), but a large gap when characteristics are controlled for, resulting in a high percentage figure due to differing returns.

Table 5. Summary of Table 4 Quantile Regression Results.

Country	Glass Ceiling Measured by: ^a			Sticky Floor Measured by: ^b		Estimated Profile of Wage Gap along Distribution	Estimated Range of Wage Gap (%)
	90-All Gaps	90-75 Difference	90-50 Difference	10-50 Difference	10-25 Difference		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Industry and Occupation Excluded							
Public Sector							
Austria	✓	✓	✓				14-29
Belgium	✓	✓	✓			Increasing	7-21
Britain	✓	✓	✓			Increasing	11-24
Denmark	✓	✓	✓			Increasing	6-18
Finland			✓			Increasing	16-31
France	✓	✓	✓				13-27
Germany	✓	✓	✓			Increasing	6-17
Ireland		✓	✓				15-19
Italy	✓	✓	✓			Increasing	3-17
Netherlands	✓	✓	✓			Increasing	5-24
Spain		✓					4-10
Private Sector							
Austria	✓	✓	✓			Increasing	21-27
Belgium	✓	✓	✓			Increasing	9-22
Britain	✓	✓	✓			Increasing	20-30
Denmark	✓	✓	✓			Increasing	5-21
Finland	✓	✓	✓			Increasing	13-29
France	✓	✓	✓		✓		17-29
Germany	✓	✓	✓			Increasing	14-20
Ireland			✓			Increasing	19-27
Italy	✓	✓	✓				14-21
Netherlands	✓	✓	✓			Increasing	3-25
Spain						Flat	21

Notes:

^aA glass ceiling is defined to exist if the 90th percentile wage gap exceeds the reference gap by at least 2 points.

and Italy are the only countries in which women's superiority to men in productive characteristics was sufficient to compensate them for different returns across all parts of the distribution.

The estimates just discussed exclude controls for occupation and industry. We have focused on these results because of the possibility that industry and occupation are endogenous. An additional reason for omitting these controls is that employers' or trade unions' discriminatory or exclusionary practices could be highly correlated with occupation and industry. However, the occupational and industry controls might also reflect otherwise unmeasured human capital, and for this reason we may wish to include them. The estimates of the gender pay gap *without*

such controls—and which thereby ignore the potential effect of unobserved human capital—could be viewed as an upper bound for the extent of “discrimination.” Conversely, the estimates of the gender pay gap *with* such controls could be viewed as a lower bound for the extent of “discrimination.”

For comparison with our main results, the right-hand side of Table 4 presents estimates including occupation and industry. The right-hand panel of Table 5 summarizes the pattern of glass ceilings and sticky floors according to the various definitions. Table 5 shows that using occupation and industry to try to account for the gender wage gap does not greatly change the previous conclusions that glass ceilings are widespread, especially in the public sector, and that two countries

Table 5 (continued).

Country	Glass Ceiling Measured by: ^a			Sticky Floor Measured by: ^b		Estimated Profile of Wage Gap along Distribution	Estimated Range of Wage Gap (%)
	90-All Gaps	90-75 Difference	90-50 Difference	10-50 Difference	10-25 Difference		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Industry and Occupation Included							
Public Sector							
Austria	✓	✓	✓		✓		16–27
Belgium	✓	✓	✓			Increasing	5–17
Britain	✓	✓	✓			Increasing	9–21
Denmark	✓	✓	✓				5–14
Finland	✓	✓	✓			Increasing	12–32
France	✓	✓	✓				8–17
Germany			✓			Increasing	1–15
Ireland						Quite Flat	17–19
Italy	✓	✓	✓			Increasing	4–17
Netherlands	✓	✓	✓			Increasing	4–22
Spain							7–10
Private Sector							
Austria	✓	✓	✓				23–30
Belgium	✓	✓	✓			Increasing	15–20
Britain			✓			Increasing	20–22
Denmark	✓	✓	✓			Increasing	10–16
Finland			✓			Increasing	20–22
France	✓	✓	✓		✓		14–17
Germany	✓	✓	✓			Increasing	20–21
Ireland			✓			Increasing	17–26
Italy	✓	✓	✓				16–21
Netherlands	✓	✓	✓			Increasing	14–19
Spain						Flat	21–24

^aA sticky floor is defined to exist if the 10th percentile wage gap exceeds the reference wage gap by at least 2 points.

also have sticky floors. This may suggest that these occupational and industry controls are not picking up heterogeneity in discriminatory practices toward women. However, there are some notable exceptions: with the inclusion of industry and occupation, the private sector glass ceilings in Britain and Finland disappear using the column (1) and (2) definitions. This could indicate that the glass ceilings in these countries primarily reflect occupational or industrial segregation, but we do not push this interpretation too far given the possible endogeneity issues and potential unobserved productivity effects. Finally, notice that a public sector sticky floor using the 25–10 differential is now found for Austria and a private sector sticky floor remains in France.

Discussion and Conclusions

The European Institutional Setting

Our estimated gender pay gaps show that even if the distribution of characteristics were the same across gender, men and women would receive different returns across the wage distribution. We now speculate as to why the estimated gender pay gaps in Europe are (i) larger at the top of the wage distribution (*glass ceilings*) in most countries, and (ii) larger at the bottom of the wages distribution (*sticky floors*) in two countries.

Gender-specific policies—such as equal opportunity and anti-discrimination laws, parental leave provisions, and the availability of childcare—are likely to affect gender

wage gaps, including both mean gaps and gaps across the wage distribution. Gender wage gaps are also likely to be influenced by wage setting institutions that do not directly impinge on gender, such as those governing collective bargaining and minimum wages. Cross-country differences in such policies and institutions across the eleven European Union countries for which we have data may well contribute to observed variations in gender wage gaps. While clearly with just eleven observations we cannot hope to provide a conclusive test of the impact of different institutions on the gender pay gap, we are able to provide some interesting correlations between summary measures of various important institutions and our observed gender pay gaps. In what follows, we use our estimated gaps (taken from Table 2) for the combined samples (in which public and private sector workers were pooled). We do so because our summary measures of various institutions are only available country-wide and are not disaggregated by sector.

Gender-Specific Policies

First we consider how gender-specific policies might affect pay. Member states of the European Union prohibit discrimination based on gender, race (including nationality and citizenship), and—in Northern Ireland only—religion.¹⁷ Nonetheless, although discrimination may be proscribed by legislation, whether or not it is still practiced may depend on the effectiveness of the law's implementation and the willingness of individuals to take violations to the courts. To the extent that only the more articulate and better educated are willing to take legal action against breaches of the law, we might expect the impact of these policies to work against glass ceilings as defined above.

Parental leave provisions and state provision of childcare for under-school-age children vary considerably across countries

(OECD 2001; Jaumotte 2003).¹⁸ These institutions are likely to influence the behavior of men and women differently and hence affect gender wage gaps. Blau and Kahn (2003) noted that the expected impact of these policies is unclear *a priori*. On the one hand, women who are *not* subject to parental leave provisions might give up—or lose—their jobs on having a child, and might re-enter subsequently at lower-level jobs providing shorter hours and lower pay, while women who *do* have access to parental leave might enjoy higher relative earnings through the fact that these policies preserve their ties with the firm and thereby increase incentives to invest in specific human capital. These factors will lead to a correlation between parental leave policies and higher female pay. We will term this the positive effect of parental leave policies. On the other hand, however, generous leave policies could increase women's time out of the work force for childbearing, thus exacerbating the average gender pay gap for that group. This is the potential negative effect of parental leave policies.

Empirical research tends to find a positive effect of short leaves on women's wages but a negative effect for long leaves (Ruhm 1998; Waldfogel 1998). But why should these leave policies affect gender wage gaps across the wage distribution? We might expect, *a priori*, that women at the bottom will be less attached to the work force, and so the positive impact of leave policies increasing women's attachment to firms might dominate the negative effect outlined above. Ultimately, however, which effect dominates—and in which countries—is an empirical question.

What is the extent of statutory parental leave policies in the EU countries for which we have data? According to the OECD (2001, Table 4.7), the country with the highest total duration of maternity/childcare leave (weeks) is Finland, followed by Spain, France, Germany, and Austria; the country with the

¹⁷In light of the European Framework Directive, categories covered by anti-discrimination legislation were extended in 2003 to include religion and sexual orientation, and in 2006 to include the additional categories of age and disability.

¹⁸Paternity leave entitlements are still relatively uncommon, and where they are found are typically of short duration (OECD 2001:145). Of our eleven countries, Spain, the Netherlands, Belgium, and France mandate three or fewer days' paternity leave, while Denmark mandates fourteen days.

lowest is the United Kingdom, followed by Ireland. There is also some evidence of extra-statutory provision of similar family-friendly arrangements by firms. According to the OECD (2001:147) and Evans (2001), the countries with the highest coverage of extra-statutory provisions are Austria and West Germany, followed by Italy, Greece, and Spain, and coverage is lowest in the Nordic countries, Ireland, and the United Kingdom.¹⁹

We now consider formal childcare policies. These are likely to have a positive effect on women's wages, not only because they are likely to increase women's attachment to firms, thereby increasing incentives to make investments in firm-specific skills, but also because they may enable women to return to work earlier than would otherwise be possible. We would therefore expect childcare provision to reduce the gender pay gap, *ceteris paribus*. However, an additional possibility is that the presence of subsidized childcare attracts into the work force those women who are the least committed to market production. The associated selectivity effect may then actually increase the gender wage gap at the bottom of the wage distribution.

According to the OECD (2001), the countries with the largest proportions of children under three years old using formal childcare are Denmark at 64%, followed by Ireland at 38%, Britain at 34%, France at 29%, and Finland at 22%. In four other countries, in contrast—Italy, the Netherlands, Spain, and Austria—the figure is below 7%.

Figure 1 presents three scatter plots illustrating the cross-country correlation between the OECD (2001) work-family reconciliation index and (a) sticky floors defined as the 10-50

pay difference, (b) glass ceilings defined as the 90-50 pay difference, and (c) the mean gender pay gap. The OECD work-family reconciliation index is the sum of indicators for the coverage of children under three years of age in formal childcare, maternity leave, flexi-time, voluntary part-time, and one half of the extra-statutory leave by firms indicator (see OECD 2001:152). The plots show that, across countries, the work-family index is negatively correlated with sticky floors and positively correlated with glass ceilings, with *t*-statistics of -2.7 and 3.2, respectively. Thus countries with more "generous" work-family policies have a lower wage gap at the bottom of the wage distribution and a wider gap at the top. For example, Denmark and the Netherlands have the most liberal work-family policies, and they have a big pay gap at the top and a small pay gap at the bottom of the wage distribution.


We noted above that family-friendly policies could be a double-edged sword. On the one hand, they might raise women's relative earnings by preserving their ties with the firm, thereby increasing incentives to invest in specific human capital and leading to higher female pay. But on the other hand, family-friendly policies could increase women's time out of work force for childbearing, thus widening the average gender pay gap for that group. Our scatter-plot in Figure 1 (b)—showing that the gender pay gap is higher at the top of the wage distribution in countries with generous family-friendly policies—suggests that the negative effect dominates at the top of the distribution in our sample of eleven European countries.

What about women at the bottom of the wage distribution? If such women are, on average, less attached than other women to the work force, it is plausible that generous family-friendly policies might increase their attachment to firms. Our data are certainly consistent with that hypothesis. The scatter-plot in Figure 1 (a)—showing that the gender pay gap is lower at the bottom of the wage distribution in countries with generous family-friendly policies—suggests that the positive effect of family-friendly policies dominates at the bottom of the distribution.

Albrecht et al. (2003:172), using a QR

¹⁹Evans (2001) identified four main types of family-friendly arrangements by firms: leave from work for family reasons; changes to work arrangements for family reasons; practical help with childcare and eldercare; and the provision of training and information. OECD (2001) summarizes these extra-statutory provisions along two dimensions: average percentage of female employees reporting extra family leave, and percentage of female employees reporting provision/subsidies for child daycare. The Netherlands has especially high levels of firm-provided daycare relative to the amount of extra-statutory leave.

Figure 1. Work-Family Reconciliation (OECD 2001 index).



framework, charted the extent of the Swedish glass ceiling and speculated as to its causes. They found that gender differences in returns were the primary factor. They rejected the notion of a “taste-based explanation... [whereby] Swedish women prefer to work in family friendly but low-wage jobs,” on the grounds that gender differences arise from differences in rewards even when the analysis controlled for occupation. A more likely explanation, they wrote, is one based on the “work environment faced by Swedish women,” especially “the Swedish parental leave policy and the daycare system,” which provide strong incentives to participate but not to commit strongly to a career.

Albrecht et al.’s conjecture is interesting and is supported by our results for Denmark.²⁰ But the fact that we find the same glass ceiling effect across the majority of other EU countries, with their very different parental leave policies and daycare systems, suggests that this cannot be the primary explanation. For example, Italy has low levels of “work-family reconciliation policies”—and it too exhibits a glass ceiling.

A second reason put forward by Albrecht et al. (2003) for the Swedish glass ceiling phenomenon is the relatively high wages at the bottom of the wage distribution, making it “very difficult for career-oriented women to hire household help or help with childcare,” especially for the very young children under 12 months who cannot be admitted into daycare. For this reason, women might wind up in less demanding jobs and thus fall substantially behind men toward the top of the distribution. Thus cross-country evidence should show a negative correlation between the magnitude of the glass ceiling and the dispersion of the wage distribution. This relationship is illustrated in Figure 2(b) for our sample of eleven countries, where we have measured wage dispersion by the 90th-10th percentile differential of log wages in the full sample of workers in each country. There is indeed a statistically significant

negative correlation that is consistent with this hypothesis: the t-statistic on the log wage dispersion measure in the glass ceiling regression is -4.1 .

Pay-Bargaining Institutions


Next we consider pay-bargaining institutions that may not directly impinge on gender. While collective bargaining institutions, and the degree of coordination and centralization of wage bargaining, might not have direct gender effects, they could well have important indirect effects. For example, trade unions may be less likely to represent the interests of their female electorate—who may be perceived as having a marginal attachment to the work force—than of the male electorate (Booth and Francesconi 2003).

In addition, collective bargaining and associated institutions affect the wage structure in general. To the extent that the wage distribution is compressed, these institutions may thus impinge *indirectly* on women’s wages and through this mechanism affect the gender pay gap. Figure 2 suggests that this may indeed be the case. It shows that sticky floors are increasing, while glass ceilings are declining, in the 90:10 log wage differential. For those countries with greater wage dispersion, in particular Spain, the gender pay gap is relatively high at the bottom of the distribution and relatively low at the top of the distribution compared to the median. Only the latter relationship is statistically significant, however (the former has a t-statistic of only 1.3, while the latter has a t-statistic of -4.1). Quite why the glass ceiling should be declining in wage dispersion is unclear. One hypothesis already mentioned is that high wage dispersion may mean that high-skilled career-focused women can obtain cheaper domestic services.

Countries with higher levels of unionization and more centralized or coordinated bargaining also tend to have the lowest wage dispersion (Blau and Kahn 1992, 1996, 2003; Boeri, Brugiavini, and Calmfors 2001). Centralized bargaining—and more coordinated bargaining—results in lower wage dispersion, and is thus likely to lower the gender pay gap, *ceteris paribus*—perhaps especially at the

²⁰See Pylkkänen and Smith (2004) for a detailed analysis of the impact of family-friendly policies on women’s parental leave behavior in Sweden and Denmark.

Figure 2. Wage Dispersion.



bottom of the wage distribution. Moreover, in all the countries for which we have data, the female wage lies below the male wage across the entire distribution. Hence centralized pay bargaining systems that raise the minimum level of pay regardless of gender are also likely to lower the gender pay gap, *ceteris paribus*.

Although we can group our countries by both the extent of union density and the extent of union coverage, we choose to focus only on coverage.²¹ The countries in which at least 75% of the work force is covered are Austria, Germany, Belgium, the Netherlands, and France (Boeri et al. 2001:92). The countries with the highest levels of coordination of bargaining—both of unions and employers—are Austria and Norway, followed by Germany and Finland. The lowest levels of coordination are found in the United Kingdom and Ireland. Of course it is impossible with so few observations to tease out the direct and indirect effects of unions on our estimated gender gaps. We have only eleven observations for the wage dispersion measure, and for the union measure we have only nine observations, as there were no data available for Ireland and Italy. However, Figure 3(a) does reveal a positive correlation between the magnitude of the sticky floor and union coverage. Figure 3(b) reveals a positive correlation between the magnitude of the glass ceiling and union coverage, and Figure 3(c) also shows that the mean wage gap increases with union coverage. However, none of these estimated relationships is statistically significant (the t-statistics for Figures 3a–c are 0.5, 0.3, and 0.2). Perhaps it is unsurprising that the data here are not very informative, given that we have only nine observations.

Statutory minimum wages compress the bottom of the pay distribution, and so are likely to reduce the gap between men and women

at the bottom (Dolado et al. 1996). Institutional pay compression—whether through unions or through minimum wages—may distort skills investment incentives. However, recent studies suggest that, in the presence of labor market imperfections inducing wage compression, firms may be willing to finance work-related training (see, *inter alia*, Stevens 1994; Acemoglu and Pischke 1999, 2003; Arulampalam et al. 2004). If pay returns are reduced by pay compression, women have less incentive to stay in the work force when engaged in childrearing. Conversely, high wage floors might increase the likelihood that women stay in the work force, because of the higher opportunity cost of time out, and women might therefore have higher levels of work experience and skills acquisition. Whether these effects on the gender pay gap vary across the wage distribution is ultimately an empirical issue. Certainly Figures 2(a) and (b) show that countries with low wage dispersion have lower gender pay gaps at the bottom of the wage distribution and higher pay gaps at the top. This correlation is consistent with the hypothesis that women in countries with low pay dispersion are more likely to stay attached to firms at the bottom of the distribution.


Other Factors

Of course, other factors are also likely to be at work. Many labor markets are hierarchical, and promotion and appointment procedures can exacerbate gender pay gaps across the pay distribution. While promotions are typically subject to well-defined procedures, especially in larger organizations, exactly *where* in the rank-specific salary scale a successful candidate is appointed can depend on individual negotiation skills or bargaining power, or employer discretion, in addition to experience. Booth, Francesconi, and Frank (2005), using data on promotions from the British Household Panel Survey, showed that women gained less from promotion than men did, *ceteris paribus*.²² If promotion procedures

²¹The extent of union recognition or coverage is a better measure of union power than is union density, as argued in Booth (1995). In addition, the degree of coordination between unions, and between unions and employers, can also be crucial, as well as the degree of centralization of bargaining. See Boeri et al. (2001) for a discussion and for detailed tables indicating the variation of these measures across the EU.

²²Blackaby, Booth, and Frank (2005), using data on the labor market for academic economists in the United Kingdom, produced further evidence that promotions might exacerbate gender pay inequality.

Figure 3. Union Coverage.



favor men over women toward the top of the wage distribution, then the gender pay gap might be bigger toward the top. Landers, Rebitzer, and Taylor (1996) showed, in their study of U.S. law firms, how such criteria for promotion as excessively long hours of work can exacerbate gender pay gaps toward the top of the lawyers' wage distribution.

Individuals are frequently appointed at a particular rank level of the relevant scale for their occupation or industry and then aim to work their way up the hierarchy. While both promotions and pay are covered by anti-discrimination legislation and equal opportunities policies, there is scope for discretion—or discrimination—in the choice of the particular level within a rank to which an individual is appointed. Thus if men are initially appointed at a higher starting salary (a higher rung) within a particular scale than are comparable women, then the gender pay gap might be bigger toward the bottom of the wage distribution—the *sticky floor*. Another hypothesis is that women toward the bottom might have less bargain-

ing power or be more likely to be subject to firms' market power than comparable men, due perhaps to unobservable family commitments or social custom whereby the man's career takes precedence. Although we are unable to explore these hypotheses with our data, we mention them in passing because it is possible that such mechanisms partly explain our findings.

Conclusions

The analysis in this paper shows that, across our sample of eleven European Union countries in 1995–2001, women were paid less than men, even when the analysis holds the distribution of characteristics constant. The magnitude of the gaps—which can be attributed to differing returns—varied substantially across the different countries and across the wage distributions. We suggest that the considerable heterogeneity in EU countries' institutions is likely to contribute to these differences, as illustrated by the simple correlations in Figures 1–3.

DATA APPENDIX

Selection of Estimating Samples and Sample Sizes

The selection criteria are outlined in the text, and result in sample sizes as given in column (5) of Appendix Table A.1. The pooled, economy-wide models use the combined public and private sector samples for each gender. In the separate private and public sector analyses, which include occupation and industry, the various industry and occupation dummies were combined in the following cases: (i) where there were small cell sizes (less than 1% of observations for both sexes); and (ii) where there was strong gender segregation (less than 1% of one sex in a cell).

For the industry dummies, the base case in the public sector is administration, and the base case in the private sector is manufacturing. The table below shows how dummies were combined for each sector and country. For the public sector (where industry structure varies substantially across countries), the table lists the dummies combined and those included separately. Note that when energy is combined with manufacturing, it is simply left in the base case. For the private sector, the table lists only combined dummies. The others are all included separately. The full industry list is as follows: agriculture (all observations dropped), energy, manufacturing, construction, retail, hotels, communications, finance, property, administration, education, social services, other, and missing industry.

For the occupational dummies, the base case is unskilled. The most common form of occupational segregation is in the public sector, where there are very few female craft workers or operatives. These categories were combined with service and shop workers (this was preferred to combining them with the unskilled category). The full occupational list is as follows: manager, professional, associate professional, clerical worker, service worker, agricultural worker (all observations dropped), craft worker, operative, unskilled/other, and missing occupation.

Germany: to account for differences in wage determination between east and west in post-unification Germany, following a pooling test we interacted the following variables with a dummy variable for Eastern Germany: age, education, health status, any experience of unemployment since 1989, fixed term contract, occupation, firm size (private sector), and year.

Appendix Table A1
Sample Selection and Industry/Occupation Definitions

Country (1)	Sector (2)	Industry Dummies (3)	Occupation Dummies (4)	Final No. of Pooled Observations (Men, Women) (5)	Mean Waves per Individual (Men, Women) (6)
Austria	Public	Combined: energy, mfg, constr, retail, hotel, finance. Separate: comms, property, educ, social srvcs, other.	Combined: service worker, craft, operative.	2,389, 2,214	4.3, 4.2
	Private	Combined: (1) energy, mfg; (2) admin, educ, social srvcs.		6,469, 4,205	4.1, 3.6
Belgium	Public	Combined: energy, mfg, constr, retail, hotel. Separate: comms, finance, property, education, social srvcs, other.	Combined: service worker, craft, operative.	2,257, 2,466	4.2, 4.0
	Private	Combined: (1) energy, mfg; (2) admin, educ.		4,271, 3,137	3.8, 3.6
Britain	Public	Combined: energy, mfg, constr, retail, hotel, finance. Separate: comms, property, education, social srvcs, other.	Combined: service worker, craft, operative.	2,099, 3,918	4.4, 4.1
	Private	Combined: admin, educ.		8,980, 6,934	3.8, 3.6

Continued

Appendix Table A1 (continued)

<i>Country</i> (1)	<i>Sector</i> (2)	<i>Industry Dummies</i> (3)	<i>Occupation Dummies</i> (4)	<i>Final No. of Pooled Observations (Men, Women)</i> (5)	<i>Mean Waves per Individual (Men, Women)</i> (6)
<i>Denmark</i>	Public	Combined: energy, mfg, constr, retail, hotel, finance. Separate: comms, property, education, social srvc, other.	Combined: service worker, craft, operative.	1,984, 3,922	4.1, 4.0
	Private	Combined: (1) energy, mfg; (2) admin, educ, social srvc.		5,169, 2,955	4.0, 3.8
<i>Finland</i>	Public	Combined: energy, constr, retail, hotel, finance. Separate: mfg, comms, property, education, social srvc, other.	Combined: service worker, craft, operative.	2,240, 4,153	3.7, 3.5
	Private	Combined: (1) energy, mfg; (2) admin, educ, social srvc.		5,413, 3,616	3.3, 3.2
<i>France</i>	Public	Combined: constr, retail, hotel. Separate: energy, mfg, comms, finance, property, education, social srvc, other.	Combined: service worker, craft, operative.	4114, 5017	4.2, 4.1
	Private	Combined: (1) energy, mfg; (2) admin, educ.		10,309, 7,227	3.8, 3.6
<i>Germany</i>	Public	Combined: energy, mfg, constr, retail, hotel, property. Separate: comms, finance, education, social srvc, other.	Combined: service worker, craft, operative.	3,572, 4,698	4.6, 4.2
	Private	Combined: (1) energy, mfg; (2) admin, educ, social srvc.		13,335, 8,031	4.3, 4.8
<i>Ireland</i>	Public	Combined: energy, mfg, constr, retail, hotel. Separate: comms, finance, property, education, social srvc, other.	Combined: service worker, craft, operative.	2,113, 1,945	3.9, 3.6
	Private	Combined: (1) energy, mfg; (2) admin, educ, social srvc.		4,684, 3,597	3.2, 3.0
<i>Italy</i>	Public	Combined: constr, retail, hotel, finance, property. Separate: energy, mfg, comms, education, social srvc, other.	Combined: (1) mgr, professional; (2) service worker, craft, operative.	4,638, 4,404	4.2, 4.5
	Private	Combined: admin, educ, social srvc.	Combined: mgr, professional.	10,255, 5,812	3.7, 3.4
<i>Netherlands</i>	Public	Combined: energy, mfg, constr, retail, hotel, comms, finance. Separate: property, education, social srvc, other.	Combined: service worker, craft, operative	3,125, 3,378	4.8, 4.3
	Private	Combined: admin, educ.		10,491, 5,821	4.6, 3.9
<i>Spain</i>	Public	Combined: energy, mfg, constr, retail, hotel, finance. Separate: comms, property, education, social srvc, other.	Combined: (1) mgr, professional; (2) service worker, craft, operative.	3,155, 2,837	4.1, 3.9
	Private	Combined: admin, educ, social srvc.		11,790, 6,241	3.4, 2.9

REFERENCES

- Acemoglu, Daron, and Jörn-Steffen Pischke. 1999. "The Structure of Wages and Investment in General Training." *Journal of Political Economy*, Vol. 107, No. 3 (June), pp. 539–72.
- _____. 2003. "Minimum Wages and On-the-Job Training." *Research in Labor Economics*, Vol. 22, pp. 159–202.
- Albrecht, James, Anders Bjorklund, and Susan Vroman. 2003. "Is There a Glass Ceiling in Sweden?" *Journal of Labor Economics*, Vol. 21, No. 1 (January), pp. 145–77.
- Arulampalam, Wiji, Alison L. Booth, and Mark L. Bryan. 2004. "Training and the New Minimum Wage." *Economic Journal*, Vol. 114, No. 494 (March), pp. C87–C94.
- Blackaby, David, Alison L. Booth, and Jeff Frank. 2005. "Outside Offers and the Gender Pay Gap." *Economic Journal*, Vol. 115 (February), pp. F81–F107.
- Blau, Francine D., and Lawrence M. Kahn. 1992. "The Gender Earnings Gap: Learning from International Comparisons." *American Economic Review*, Vol. 82, No. 2 (May), pp. 533–38.
- _____. 1996. "Wage Structure and Gender Earnings Differentials: An International Comparison." *Economica*, Vol. 63, No. 25—Supplement: Economic Policy and Income Distribution, pp. S29–S62.
- _____. 2003. "Understanding International Differences in the Gender Pay Gap." *Journal of Labor Economics*, Vol. 21, No. 1 (January), pp. 106–44.
- Blinder, Alan. 1973. "Wage Discrimination: Reduced Form and Structural Estimates." *Journal of Human Resources*, Vol. 8, No. 4 (Fall), pp. 436–55.
- Boeri, Tito, Agar Brugiavini, and Lars Calmfors, eds. 2001. *The Role of Unions in the Twenty-First Century*. Oxford: Oxford University Press.
- Booth, Alison L. 1995. *The Economics of the Trade Union*. Cambridge: Cambridge University Press.
- Booth, Alison L., and Marco Francesconi. 2003. "Union Coverage and Non-Standard Work in Britain." *Oxford Economic Papers*, Vol. 55, No. 3 (July), pp. 383–416.
- Booth, Alison L., Marco Francesconi, and Jeff Frank. 2003. "A Sticky Floors Model of Promotion, Pay, and Gender." *European Economic Review*, Vol. 47, No. 2, pp. 295–322.
- Buchinsky, Moshe. 1998. "Recent Advances in Quantile Regression Models." *Journal of Human Resources*, Vol. 33, No. 1 (Winter), pp. 88–126.
- De la Rica, Sara, Juan J. Dolado, and Vanesa Llorens. 2005. "Ceiling and Floors: Gender Wage Gaps by Education in Spain." Bonn, Germany: IZA Discussion Paper No. 1483.
- Del Rio, Coral, Carlos Gradín, and Olga Cantó. 2005. "The Measurement of Gender Wage Discrimination: The Distributional Approach Revisited." Working Paper No. 25, ECINEQ (Society for the Study of Economic Inequality).
- Dolado, Juan J., Francis Kramarz, Stephen Machin, Alan Manning, David Margolis, and Coen Teulings. 1996. "The Economic Impact of Minimum Wages in Europe." *Economic Policy*, Vol. 23 (October), pp. 319–72.
- Eurostat. 2003. *Eurostat Yearbook 2003: The Statistical Guide to Europe: Data 1991–2001*. Luxembourg: Office for Official Publications of the European Communities.
- Evans, John M. 2001. "Firms' Contribution to the Reconciliation between Work and Family Life." Paris: Labour Market and Social Policy Occasional Papers No. 45, OECD.
- Fitzenberger, Bernd. 1998. "The Moving Blocks Bootstrap and Robust Inference for Linear Least Squares and Quantile Regressions." *Journal of Econometrics*, Vol. 82, pp. 235–87.
- Jaumotte, Florence. 2003. "Female Labour Force Participation: Past Trends and Main Determinants in OECD Countries." Paris: Economics Department Working Paper No. 376, OECD.
- Koenker, Roger, and Gilbert Bassett. 1978. "Regression Quantiles." *Econometrica*, Vol. 46, No. 1 (January), pp. 33–50.
- 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*, Vol. 86, No. 3 (June), pp. 329–48.
- Lucifora, Claudio, and Dominique Meurs. 2004. "The Public Sector Pay Gap in France, Great Britain, and Italy." Bonn, Germany: IZA Discussion Paper No. 1041.
- Machado, Jose, and Jose Mata. 2005. "Counterfactual Decomposition of Changes in Wage Distributions Using Quantile Regression." *Journal of Applied Econometrics*, Vol. 20, No. 4 (May/June), pp. 445–65.
- Oaxaca, Ronald L. 1973. "Male Female Wage Differentials in Urban Labor Markets." *International Economic Review*, Vol. 14, No. 3 (October), pp. 693–709.
- OECD. 1996. "Earnings Inequality, Low-Paid Employment and Earnings Mobility." Chapter 3 of *Employment Outlook*. Paris: OECD.
- _____. 2001. "Balancing Work and Family Life: Helping Parents into Paid Employment." Chapter 4 of *Employment Outlook*. Paris: OECD, June.
- Pylkkänen, Elina, and Nina Smith. 2004. "The Impact of Family Friendly Policies in Denmark and Sweden on Mothers' Career Interruptions Due to Childbirth." Bonn, Germany: IZA Discussion Paper No. 1050, March.
- Ruhm, Christopher J. 1998. "The Economic Consequences of Parental Leave Mandates: Lessons from Europe." *Quarterly Journal of Economics*, Vol. 113, No. 1 (February), pp. 285–317.
- Stevens, Margaret. 1994. "A Theoretical Model of On-the-Job Training with Imperfect Competition." *Oxford Economic Papers*, Vol. 46, pp. 537–62.
- Waldfoegel, Jane. 1998. "The Family Gap for Young Women in the United States and Britain: Can Maternity Leave Make a Difference?" *Journal of Labor Economics*, Vol. 16, No. 3 (July), pp. 505–45.