SWEDISH ECONOMIC POLICY REVIEW 14 (2007) 121-144

The glass ceiling in Europe: Why are women doing badly in the labour market?

Alison L. Booth*

Summary

■ Average gender pay gaps have absorbed the interest of economists for many years. More recently, studies have begun to explore the degree to which observed gender wage gaps might differ across the wages distribution. The stylised facts from these studies, summarised in the first part of the paper, are that the gender pay gap in Europe is typically increasing across the wages distribution. This finding—more pronounced in the private than the public sector—has been interpreted as a glass ceiling effect. The existence of this glass ceiling suggests that the average gender pay gap in Europe is mainly due to the gender gap towards the top of the wages distribution. What explains these stylised facts? We briefly outline some relevant hypotheses in the second part of the paper. A fundamental challenge for labour economists is to identify the extent to which these stylised facts are due to policies and institutions, discrimination, to other unobservable factors, or to fundamental differences between men and women. Finally, we briefly summarise the policy initiatives that might be introduced to deal with gender wage gaps. ■

JEL classification: J31, J70.
Key words: gender, discrimination, glass ceilings, sticky floors, quantile regression decompositions.

* *Alison L. Booth is Professor of Economics at the University of Essex and the Australian National University, and Research Fellow at the CEPR and IZA.*

SWEDISH ECONOMIC POLICY REVIEW 14 (2007) 121-144

The glass ceiling in Europe: Why are women doing badly in the labour market?

Alison L. Booth*

The purpose of this paper is to address a number of questions about the gender pay gap in Europe. The questions are necessarily focused. They concentrate on the gender pay gap rather than on other gender gaps in, for example, unemployment or participation in the informal sector.[1] The first question asks what are the stylised facts. Are women really doing badly in terms of pay in European countries' labour markets? Does the gender gap vary over the wages distribution and which European countries are doing better than others in this regard? The second question asks what are the candidate explanations for these stylised facts. What might cause observed gender wages gaps to vary over the wages distribution? Why are there cross-country differences? Do policies and institutions play a role? And are there any other factors—apart from policies and institutions—that might explain these stylised facts? The final question in this paper relates to policy. Given the state of our knowledge, what further policy initiatives might be introduced to deal with gender wage gaps?

The paper is set out as follows. Section 1 provides some stylised facts obtained from recent studies investigating the extent of glass ceilings and sticky floors across a number of different European countries. In Section 2, we briefly outline how policies and institutions might be correlated with these stylised facts. In Section 3 we summarise other factors that might explain the glass ceiling. Section 4 considers policy prescriptions and Section 6 draws some conclusions.

** I am grateful to Mark Bryan and an anonymous referee for helpful comments and to the ARC for financial support under Discovery Project Grant No. DP0556740.*

[1] See Azmat et al. (2006).

1. Stylised facts about the glass ceiling in Europe

Mean gender pay gaps have absorbed the interest of economists for many years and much important work has documented these. In the European context, many authors have estimated gender wage gaps for particular countries and consistently find an average gender wage gap. Average gender wage gaps have also been extensively charted in cross-country studies, and some examples of these are found in Blau and Kahn (1992, 1996, 2003), Smith and Westergaard-Nielsen (1988), and Datta Gupta, Oaxaca and Smith (2006). Recently labour economists have begun to explore the degree to which gender wage gaps might differ across the wages distribution, using quantile regression (QR) techniques. These facilitate an investigation of the extent to which gender affects the location, scale and shape of the conditional wage distribution. The pioneering work in this regard is that by Albrecht, Bjorklund and Vroman (2003) for Sweden. [2]

1.1. The gender pay gap in Sweden

Albrecht et al. (2003) use 1998 Swedish data to show that the gender pay gap is increasing across the wages distribution and accelerating in the upper tail. They interpret this finding as a "glass ceiling" effect, where they define a glass ceiling to be "the phenomenon whereby women do quite well in the labor market up to a point after which there is an effective limit on their prospects". Clearly to observe this empirically one needs to investigate the gender pay gap across different parts of the distribution. The fact that they observe this glass ceiling effect in Swedish data suggests that the average Swedish gender pay gap in the 1990s was mainly due to the gender gap towards the top of the wages distribution. They also use 1999 data for the US and find a different result for that country. In particular, they find that the

[2] Arulampalam et al. (2007) also investigate these issues for a number of European Union countries, and this current paper draws on their results extensively. See also de la Rica, Dolado and Llorens, (2005) who use 1998 European Community Household Panel data for Spain. They stratify by education group, and find that the gender wage gap is expanding over the wage distribution only for the group with tertiary education. Miller (2005) and Kee (2006) employ a similar QR approach for workers in Australia, and Kee (2006) also stratifies her sample by public and private sector workers. Baron and Cobb-Clark (2006) use alternative estimation methods to estimate gender gaps across the wages distribution for Australian workers. Lucifora and Meurs (2004) use both QR and kernel techniques to estimate the public-private sector pay gap for men and women in France, Britain and Italy.

mean gender wage gap in the US is larger than in Sweden but that the gender wage gap at the top of the Swedish wage distribution is far larger than the corresponding gap in the US. To some, this is a surprising result, since the family-friendly policies encapsulated in the Swedish welfare state have been the envy of women everywhere. However, Datta Gupta, Smith and Verner (2006) suggest that the Nordic model might have "boomerang" effects that can exacerbate gender wage gaps.[3]

The study by Albrecht et al. (2003) is important, since it shows not only that the mean gender pay gap disguises differences across the wage distribution, but also that the gender pay gap is increasing across the distribution. But are these results for Sweden peculiar to that country, or are they also found in other European countries? To answer this question, we now turn to the study by Arulamapalam et al. (2007) using harmonised data from the European Community Household Panel Survey.

1.2. The gender pay gap in other European countries

Are there sticky floors as well as glass ceilings?

Although in Sweden the gender pay gap is biggest at the top of the wage distribution, *a priori* there does not seem to be any compelling theoretical reason why this should be the case. Indeed, there are reasons why it might be wider at the bottom rather than—or as well as—the top. For example, women towards the bottom might have less bargaining power or be more subject to firm's market power than comparable men. This could be due to unobservable family commitments, or to social custom whereby the man's career takes precedence. Alternatively minimum wage compliance at the bottom may be unequal across gender, or trade unions might differentially represent the interests of their female electorate at the bottom. The situation

[3] Datta Gupta et al. (2006) argue that, while the Nordic (i.e. Denmark, Finland, Iceland, Norway and Sweden) family-friendly schemes make women more economically independent of their partners, there can be negative boomerang effects since mothers on average take much longer periods of leave than fathers. This reduces their experience capital relative to men. Moreover selection by females into the public sector, with its better family-friendly policy provision, exacerbates the overall gender pay gap, especially at the top of the wage distribution. See Pylkkänen and Smith (2004) for a comparison of the impact of family-friendly policies in Denmark and Sweden on mothers' careers.

where the gender pay gap widens at the bottom of the wage distribution is defined in Arulampalam et al. (2007) as a *sticky floor*, although elsewhere this term has been used in a slightly different sense.[4]

The Arulamapalam, Booth and Bryan (2007) study

This study exploited data from the European Community Household Panel (ECHP), collected annually since 1994 in a standardized format facilitating cross-country comparisons. The authors estimated each specification separately by country and by gender. They reported estimates for both public and private sectors combined, as well as separate estimates by sector. We discuss only the latter here.[5] The public sector, being isolated from the rigours of a market economy, could in principle more easily follow "tastes for discrimination".[6] But the public sector is also subject to government objectives and policies which are typically held to work against discrimination. To the extent that gender wage gaps reflect discrimination, this could work against gender pay gaps being higher in the public sector.[7]

Column (2) of Table 1 shows the male proportion working in the public sector for each of the eleven countries in the sample, while Column (4) shows the corresponding male proportion in the private sector. The estimating sub-samples comprised full-time and part-time employees between the ages of 22-54 years inclusive, who were working at least 15 hours per week, and who were not employed in agriculture.[8] Notice that the public sector has a majority female workforce in

[4] Booth et al. (2003) first defined a sticky floor as the situation arising where otherwise identical men and women might be appointed to the same pay scale or rank, but the women are appointed at the bottom and men further up the scale. Such a strategy can evade some discrimination laws, since the appointment rank is the same. Arulampalam et al. (2007) use the term more generally to describe the situation where the gender pay gap widens at the bottom of the wages distribution.

[5] Women might self-select into the public sector, especially if family-friendly policies are better provided there. However estimates using the combined samples yielded similar results.

[6] See Becker (1957) for the first formalization of discrimination as a reflection of tastes, and for a recent outline see Filer, Hamermesh and Rees (1996).

[7] See Lucifora and Meurs (2004) and references therein for an extensive discussion of public and private sector pay. They use QR and kernel techniques to estimate the public-private sector pay gap for men and women in France, Britain and Italy.

[8] This age restriction was chosen to minimise selectivity issues associated with cross-country differences in labour force participation of younger and older workers, which might depend on country-specific educational and early retirement sys-

7 of the 11 countries. (Only in Austria, Ireland, Italy and Spain are men in the majority and the majority is slim). In contrast, men predominate in the private sector across all countries. In 7 countries they account for 60 per cent or more of the private sector workforce.

Table 1. Raw gender wage gap and male proportion by sector

	Public		Private	
(1) Country	(2) Males (%)	(3) Mean wage gap	(4) Males (%)	(5) Mean wage gap
Austria	51.9	0.135	60.6	0.292
Belgium	47.8	0.073	57.7	0.137
Britain	34.9	0.212	56.4	0.306
Denmark	33.6	0.114	63.6	0.134
Finland	35.0	0.259	60.0	0.167
France	45.1	0.116	58.8	0.202
Germany	43.2	0.128	62.4	0.262
Ireland	52.1	0.110	56.6	0.273
Italy	51.3	*0.006*	63.8	0.153
Netherlands	48.1	0.200	64.3	0.208
Spain	52.7	0.054	65.4	0.230

Source: Arulampalam et al. (2007). Figures in italics are statistically insignificant from zero. All other reported figures are statistically significant at the 1 percent level.

Now consider the mean gender wage gaps, reported in columns (3) and (5) of Table 1. Inspection of these columns reveals that the raw average wage gap is higher in the private sector than the public sector in all countries except for Finland. In the Netherlands, there is virtually no difference across sectors. But in, for example Austria, the gender gap is 0.135 in the public sector and 0.292 in the private sector. The largest gender pay gap in the private sector is in Britain. In the public sector in Italy there is no gender pay gap.

But of course these figures are based only on the raw data. Next we turn to a summary of the QR-based estimates of the *public sector* wage gap reported in Arulampalam et al. (2007). These were calculated following the Machado and Mata (2005) decomposition method. The calculated wage gap measures the effect of different returns to men and women when women's attributes were used in the counterfactual decomposition.

tems. The restriction of working at least 15 hours per week necessary because of data, as explained at length in Arulamapalam et al. (2007).

Table 2. Estimated wage gap, public sector

		Quantile regression estimates				
	OLS	**10%**	**25%**	**50%**	**75%**	**90%**
Austria	0.227	0.153	0.140	0.190	0.239	0.289
Belgium	0.122	0.065	0.072	0.099	0.141	0.209
Britain	0.176	0.109	0.138	0.182	0.192	0.241
Denmark	0.089	0.058	0.063	0.084	0.122	0.181
Finland	0.255	0.158	0.192	0.247	0.298	0.313
France	0.172	0.145	0.130	0.146	0.189	0.273
Germany	0.099	0.058	0.072	0.102	0.137	0.167
Ireland	0.177	0.167	0.153	0.161	0.163	0.186
Italy	0.086	0.031	0.037	0.072	0.125	0.169
Netherlands	0.142	0.049	0.088	0.131	0.183	0.235
Spain	0.077	0.102	0.096	0.082	0.040	0.062

Source: Arulampalam et al. (2007), table.
Notes: (a) Regressions include controls for training, age, education, tenure, marital status, health, unemployment experience, part-time, fixed term & casual contracts, region (where possible), sector, year. (b) All reported coefficients are statistically significant at the 1 per cent level.

The first column of Table 2 gives the ordinary least squares (OLS) estimates of the mean public sector wage gap. Including all the controls reported under the table, the average gender wage gap varies from a low of around 8 per cent in Spain followed by 9 per cent in Denmark, up to a high of 26 per cent in Finland. But do these average estimates disguise differences across the wages distribution and is the gender pay gap increasing across the distribution, as was found for Sweden by Albrecht et al. (2003)? The answer is an unambiguous yes, as inspection of the QR estimates reveals. Even when men and women have the same distributions of characteristics, there is a positive gender gap across the wages distribution due to different returns. And *all* of these are significantly different from zero at the 1 per cent level. As an example, consider Finland, a country that is geographically close to Sweden. This also happens to be the country with the highest gender wage gap at the 90th percentile—a staggering 31 per cent - and with the second highest gap at the bottom. The gap is about 16 per cent at the 10th percentile. Only Ireland is slightly worse at the 10th percentile. Denmark—also geographically close to Sweden—exhibits the same glass ceiling phenomenon. However the Danish gender gap is much smaller than in Finland, ranging from around 6 per cent at the 10th percentile up to around 18 per cent at the 90th.

Notice also that in nine countries (all except Ireland and Spain), the public sector gender gap is highest at the 90th percentile. There is clearly a widespread "glass ceiling" in the public sector in these 11 countries. The wage gap increases monotonically in 7 countries (Belgium, Finland, Britain, Denmark, Germany, Italy, Netherlands). Finally, observe that there is only relatively weak evidence of a public sector sticky floor. In Austria, France, Ireland and Spain the gender pay gap at the 10th percentile is slightly higher than the gender pay gap at the 25th percentile.

Table 3. Estimated wage gap, private sector

		Quantile regression estimates				
	OLS	**10%**	**25%**	**50%**	**75%**	**90%**
Austria	0.251	0.212	0.207	0.215	0.233	0.269
Belgium	0.144	0.090	0.120	0.144	0.174	0.218
Britain	0.247	0.201	0.224	0.246	0.272	0.302
Denmark	0.118	0.045	0.081	0.110	0.163	0.209
Finland	0.211	0.134	0.165	0.207	0.250	0.284
France	0.234	0.197	0.174	0.189	0.236	0.294
Germany	0.162	0.139	0.142	0.146	0.159	0.200
Ireland	0.230	0.185	0.215	0.240	0.256	0.269
Italy	0.172	0.156	0.138	0.146	0.169	0.205
Netherlands	0.127	0.029	0.068	0.107	0.172	0.249
Spain	0.211	0.214	0.211	0.207	0.202	0.205

Source: Arulampalam et al. (2007).
Notes: Regressions include controls for training, age, education, tenure, marital status, health, unemployment experience, part-time, fixed term & casual contracts, region (where possible), sector, year. All reported coefficients are statistically significant at the 1 per cent level.

Table 3 summarises the Arulampalam et al. (2007) estimates of the *private sector* wage gap. Again these measure the effect of different returns to men and women when women's attributes were used in the counterfactual decomposition. The first column of Table 3 gives the ordinary least squares (OLS) estimates of the mean public sector wage gap. The remaining columns give the QR estimates. The other controls included in estimation are reported in the notes under the table. The OLS estimates reveal that the average private sector wage gap gender varies from a low of around 12 per cent in Denmark, up to a high of 25 per cent in Britain and Austria.

A comparison of the OLS and QR estimates reveals that the average estimates disguise differences across the private sector wage distribution. Even when men and women have the same characteristics, there is a positive and increasing gender gap across the wages distribution due to different returns. All of these estimates are significantly different from zero at the 1 per cent level. The other important point to draw from Table 3 is that in the private sector there are very large wage gaps compared to public sector. In most countries - indeed, in all except for Spain—the private sector gender wage gap is highest at the 90th percentile. And the country with the dubious distinction of having the highest wage gap at the 90th percentile is Britain, followed by France and then Finland. The country with the biggest gap between the 10th and the 90th percentile is the Netherlands.

Is there any evidence of sticky floors in the private sector? There is some evidence in Austria, Italy and Spain. In France there is rather stronger evidence, since at the 10th percentile the wage gap is 20 per cent while at the 25th percentile it is 17 per cent. For Italy, the same comparison yields 16 per cent with 14 per cent.

A note on occupation and industry

Albrecht et al. (2003) emphasized the potential endogeniety of occupation and industry, and therefore estimated specifications with and without these controls. Arulampalam et al. (2007) also followed this procedure. On the one hand, one might wish to omit industrial and occupational controls on the grounds of (i) potential endogeneity, and (ii) employers' or unions' "discriminatory" practices are likely to be correlated with occupation and industry (and hence by including them we under-estimate the true effect of discrimination). But on the other hand, one might wish to include industrial and occupational controls, since they may embody otherwise unmeasured industry-specific and occupation-specific human capital. They may also pick up non-wage benefits such as pension rights and days of vacation that might differ across industries and/or occupations due to, for example, collective bargaining agreements. Arulampalam et al. (2007) suggest that estimates *without* such controls (and thereby ignoring the potential effect of otherwise unobserved human capital) might be viewed as upper bound for the extent of "discrimination". Estimates *with* such controls might instead be viewed as a lower bound for the extent of "discrimination".

However, Arulampalam et al. (2007) found that their inclusion does not greatly change result that glass ceilings are widespread. This may suggest that occupational and industry controls not picking up heterogeneity in discriminatory practices towards women. But there were some notable exceptions. For example, with the inclusion of industry and occupation dummies, the private sector pay gap at the 90th percentile in Britain and Finland reduce from 30 per cent and 28 per cent respectively to around 23 per cent and 21 per cent .

The stylised facts about the glass ceiling in Europe

Albrecht et al. (2003) used 1998 Swedish data to show that the gender pay gap is increasing across the wages distribution—the "glass ceiling" effect. Arulampalam et al. (2007) found a similar result in all of the eleven separate European countries they examined. They also showed that, while the magnitude of the glass ceiling differs across sector (public or private), it is everywhere apparent.[9]

But what causes these glass ceilings? Do they arise on the demand side through discriminatory practices? Or are there other forms of unobserved heterogeneity, perhaps arising on the supply side from the traditional female role within the family? Next we turn to an examination of the potential causes of observed glass ceilings and sticky floors.

2. What role do policies and institutions play?

In this section we initially briefly discuss how policies and institutions might contribute to the glass ceiling and sticky floor effects. Then in the subsequent section we shall move on to consider what other factors—apart from policies and institutions—might explain the stylised facts summarised in Section 2 above.

2.1. Policies and institutions

Gender-specific policies—such as equal opportunities and anti-discrimination laws, parental leave provisions and the availability of child care—are likely to affect gender wage gaps, both mean gaps and gaps across the wages distribution. Gender wage gaps are also likely to

[9] In a study also using QR techniques, Kee (2006) finds that the gender pay gap in Australia is constant in the public sector and increasing across the wages distribution in the private sector.

be influenced by wage-setting institutions that do not directly impinge on gender, such as those governing collective bargaining and minimum wages. Differences in such policies and institutions across Europe may well contribute to observed cross-country variations in gender wage gaps across the wages distribution.

2.2. Gender-specific policies affecting women directly

Although equal opportunities and discrimination are proscribed by legislation, they might not be effectively implemented. If only the more articulate and better educated women take legal action to combat discrimination, the impact of these policies might work against glass ceilings. And if less educated women are less likely to take recourse to legal action, the gender pay gap could widen at the bottom of the distribution.

Another important set of gender-specific policies are those that are so-called "family-friendly". There are many types of these, but perhaps the two most important are maternity leave policies and state-provided childcare for pre-school children. We briefly consider each of these in turn.

It is well-recognised that leave 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. We will term this a beneficial effect. But on the other hand, generous leave policies could increase women's time out of workforce for childbearing, resulting in relatively lower experience capital and thus widening the average gender pay gap for that group. We will term this an adverse effect, whilst bearing in mind that it might not be adverse from the child's perspective. In addition, employers incur indirect costs from leaves (hiring a substitute worker, for example), which may impinge on wages. Empirical research tends to find a beneficial effect of short leaves on women's wages but an adverse effect for long leaves (Ruhm, 1998; Waldfogel, 1998).

But why should these leave policies affect gender wage gaps across the wages distribution? We might expect *a priori* that women at the bottom might be less attached to the workforce, and so the beneficial impact of leave policies increasing women's attachment to firms might dominate the adverse effect outlined above. But ultimately it is an empirical question as to what effect dominates—and in which countries.

A second important type of family-friendly policy is formal childcare for pre-school children. This might be expected to have a beneficial effect on women's wages, since it is likely to increase women's attachment to firms. This in turn will increase the incentives to make specific skills investments. Moreover, formal childcare might act to encourage women back to work earlier than otherwise possible. Thus the "experience" capital and human capital of affected individuals is likely to increase. We would therefore expect the provision of formal childcare to reduce the gender pay gap, all else equal. However, subsidised childcare is also likely to attract into the workforce 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 wages distribution—what we have termed the adverse effect.

Arulampalam et al. (2007) performed some cross-country comparisons of correlations between their estimated glass ceilings or sticky floors on the one hand, and several separate proxies for institutional differences on the other hand. In this context, glass ceilings were measured by the difference between the 90th and 50th percentiles while sticky floors were measured by the difference between the 10th and the 50th percentiles. The results are summarised in Table 4 below.

Each cell in Table 4 presents the result of a separate regression. The summary policy variable in Row A is the *OECD Work-family Reconciliation Index* (WRI).[10] Consider the intersection of Row A and Column (1). The entry "positive" indicates that the glass ceiling is increasing in the Work-family Reconciliation Index (the t-statistic is 3.2). Countries with higher family-friendly policies have a bigger pay gap at the top of the wages distribution. This suggests that the adverse effect referred to above dominates the beneficial effect at the top of the distribution in this sample of eleven European countries.

Next consider the intersection of Row A and Column (2). The entry "negative" indicates that the sticky floor is declining in the index *(*the *t*-statistic is -2.7). The fact that, across countries, the work-family index is negatively correlated with sticky floors suggests that the beneficial effect dominates at the bottom of the distribution.

[10] The OECD Work-family Reconciliation Index is the sum of indicators for the coverage of the under-3s in formal childcare, maternity leave, flexi-time, voluntary part-time and one half of the extra-statutory leave by firms indicator (see OECD, 2001, p. 152).

Table 4. Correlations—glass ceilings/sticky floors by institutions

Row	Institution	(1) Glass ceiling (90-50 diff)	(2) Sticky floor (10-50 diff)	(3) Ave. gender wage gap
A	**Work-family reconciliation index**	Positive	Negative	Negative
B	**Wage dispersion**	Negative	Positive	Positive
C	**Union coverage**	Positive	Positive	Positive

Note: Summary of results in Arulampalam et al. (2007). There were 11 country-observations for the wage dispersion measure, and 9 for the union measure (no data available for Ireland and Italy).

The dispersion of the wage distribution

Albrecht et al. (2003) suggested an additional reason for the Swedish glass ceiling phenomenon. This 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 child care", especially for the very young children under 12 months who cannot be admitted into daycare. For this reason, women might be found in less-demanding jobs and thus fall substantially behind men towards the top of the distribution.

Arulampalam et al. (2007) investigated this by looking at cross-country correlations between the magnitude of the glass ceiling and the dispersion of the wages distribution. This is shown in Row B of Table 4. Wage dispersion is measured by the 90^{th}-10^{th} 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.)

2.3. Pay-bargaining institutions

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 workforce—than of the male electorate. In addition, collective bargaining and associated institutions affect the wage structure

in general.[11] To the extent that the wages distribution is compressed, they may thus impinge *indirectly* on women's wages and through this mechanism affect the gender pay gap. Minimum wages and high wages floors might increase the likelihood women stay in workforce, because of the higher opportunity cost of time out, and they might therefore have higher levels of work experience and skills acquisition. Whether these effects on the gender pay gap vary across the wages distribution is ultimately an empirical issue

The institution in Row C of Table 4 is union coverage. The results summarised in Row C reveal a positive correlation between the magnitude of the glass ceiling and union coverage, and also between the magnitude of the sticky floor and union coverage. However, none of these estimated relationships is statistically significant.

Next we turn to a brief overview of other potentially important causes of glass ceilings.

3. What other factors might explain glass ceilings?

A fundamental challenge for labour economists is to identify the extent to which observed gender differences in labour market outcomes are due to discrimination, or to other unobservable factors, or to fundamental differences between men and women. We briefly outline some relevant hypotheses in this section.

Many labour markets are hierarchical, and promotions and appointments procedures can exacerbate gender pay gaps. While promotions are typically subject to well-defined procedures, especially in larger organisations, exactly *where* in the rank-specific salary scale a successful candidate is appointed can depend on discrimination and individual negotiation in addition to experience (Booth et al., 2003).[12] If promotions procedures favour men rather than women towards the

[11] Countries with higher levels of unionisation and more centralized or coordinated bargaining also tend to have lowest wage dispersion (Blau and Kahn, 1992, 1996, 2003; Boeri, Brugiavini and Calmfors, 2001). This is likely to lower the gender pay gap—perhaps especially at the bottom of the wages distribution. In all the countries Arulampalam et al. examined, the female wage lies below the male across the entire wages 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.

[12] Booth et al. (2003), using data on promotions from the British Household Panel Survey, found that women gained less from promotions than did men, ceteris paribus.

top of the wages distribution, then the gender pay gap might be bigger towards the top. Moreover, promotion criteria can act to perpetuate gender gaps. Landers, Rebitzer and Taylor (1996) show, in their study of US law firms, how criteria for promotion like excessively long hours of work can exacerbate gender pay gaps towards the top of the lawyers' wage distribution.

Discrimination at the hiring stage can also matter, not least because it can affect women's willingness to bargain over offered wages. Suppose that women have a lower probability of being offered a particular job—for example playing in an orchestra. In an interesting paper, Goldin and Rouse (2000) demonstrated that blind audition procedures can result in a higher proportion of female members of symphony orchestras. Although Goldin and Rouse did not address relative salaries, consider the following scenario. Suppose a woman has overcome the hiring barrier for the organisation where she wishes to work and has actually got offered a job. Given this is hard—hiring procedures are not always impartial, as Goldin and Rouse demonstrated—it is possible that the woman is so grateful for the job offer that she will not bargain as aggressively as comparable men for her starting salary. The book by Babcock and Lashever (2003)—*Women Don't Ask*—has numerous instances of interviewed women who were in exactly this situation.

This example serves to emphasise the potential importance of bargaining. And it is well known from bargaining theory that one's share of the cake in a bargained outcome is increasing in one's fall-back option. If women are less likely to get outside offers, then they are in a weaker position with regard to bargaining. Later we shall return to the question of whether or not they are willing to bargain on their own account anyway. But first we shall consider an empirical study attempting to get a handle on outside offers in one particular labour market—that for academic economists.

Using a unique UK data source on academic economists' labour market experiences, Blackaby, Booth and Frank (2005) investigated gender differences in pay and promotions. They found a gender promotions gap and a within-rank gender pay gap, controlling for a host of factors including career breaks, best career-publications, and a measure of outside offers. They also found that men receive more outside offers and gain higher pay responses. Why might this be the case? Universities in the UK are non-profit institutions, and hence more able to follow tastes for discrimination on the one hand. But on

the other hand, the UK government and the EU have adopted a strong position in favour of equal opportunities. Blackaby et al. suggested their findings supported the *loyal servant* hypothesis. Women might be less aggressive in asking for pay rises. The Blackaby et al. results were also consistent with a model in which universities are paying women less as a cost-miminising strategy rather than as a taste for discrimination. They suggested that market economy is thus unlikely to eliminate these differentials.

Any rational employer will pay its individual workers as little as it can get away with provided productivity is unaffected. But is the gender pay gap declining with the feminization of occupation? And is it declining if there are more high-level women in the organization? We have already seen, from Albrecht et al. (2003) and the cross-country estimates of Arulampalam et al. (2007), that the inclusion of occupational and industry controls made little difference to the glass ceiling results with only a couple of exceptions. But what about the proportion of women within the organization? Women may do better in organisations in which there is already a high proportion of females. They may prefer to work with similar individuals; they may gain from mentoring opportunities and they may benefit from female networks.

Bell (2005) uses the US ExecuComp dataset that contains information on total compensation for the top five highest-paid executives of a large group of US firms for the period 1992-2003. She uses these data to estimate the impact of women-leaders on the careers of other executive women. She finds that women executives working in women-led firms earn 15-20 per cent more in total compensation than women working in other firms, ceteris paribus. Women-led firms also hire proportionately more top women executives. This is consistent with the notion of mentoring or networking by women.[13] She concludes in favour of "affirmative action at the very top of the corporate hierarchy".[14]

[13] Bell also notes her findings could be consistent with heterogeneity across firms in their "women-friendliness" and that her observed positive correlation may thus not be causal. However, she also finds that the relationship between female headed firms and women exec's outcomes is independent of the share of female directors. She also rejects the notion of sorting of higher-quality women into women-led firms on the grounds that there is no evidence for this in the data. For estimation of gender wage gaps from the ExecuComp dataset for the period 1992-7, see Bertrand and Hallock (2001).

[14] See Holzer (2007) for a survey of studies exploring the effectiveness of affirmative action in the US.

We mentioned above that there might be gender differences in willingness to bargain over wages. These would have to become more important towards the top of the wage distribution for this hypothesis to contribute to explaining the glass ceiling effect. It is possible that women towards the bottom have their wages set by pay-bargaining awards while those towards the top have their wages set by individual negotiation. If this is the case, then the fact that "women don't ask" could contribute to the glass ceiling. In their thought-provoking and important book, Babcock and Lashever (2003) provide a battery of evidence from psychology studies and their own interviews with women to support their thesis that women are unwilling to bargain on their own account, although being very competent at doing so for others. Babcock and Lashever argue that historically women were accustomed to work without pay at a type of work devalued by every objective *financial* measure—home production. Hence women are relatively unaccustomed to evaluating their time and abilities in economic terms. They suggest that society needs to change its attitudes towards women who assert themselves and encourage women to speak up for what they deserve. Perhaps it comes as no surprise (at least one of the authors is at a business school) that they suggest negotiation courses to help women to negotiate pay. They also argue that companies could benefit from adopting affirmative action, not least since it would reduce turnover.[15]

Next we turn to a consideration of possible personality differences between women and men. A relatively recent and rapidly growing literature aims to investigate if women and men differ systematically in some unobserved characteristic that might contribute towards observed gender pay gaps. Examples of such unobservables might include risk aversion, competitiveness, and cooperation. Of course it is extremely hard to disentangle to what degree such attributes are formed by society rather than being innate—the old nurture vs. nature debate. Croson and Gneezy (2004) survey the experimental economics literature that investigates preferences differences between women and men. They focus on risk aversion, cooperative behaviour (which they refer to as "social preferences") and competitiveness. They find that men tend to be more risk-taking in general, although

[15] They cite the experience of the accountancy and consulting firm Deloitte and Touche, who in 1991 decided to embark on cultural change in favour of women. This makes a fascinating case study of how within-company attitudes towards women can be changed.

an exception was found for female financial advisors who were no different from men.

Suppose now that women really are innately more risk averse than men. Then, so the argument goes, women will be less in evidence in areas where risk-loving is desirable. CEOs are the classic example of where taking risks is viewed as being efficient for the firm, and the rationale for stock options for CEOs is precisely to induce this sort of behaviour. So we might expect to see—and indeed we do—that there are relatively few female CEOs. But there is an important corollary which has not been highlighted in this literature. There are also a number of high-level jobs where risk-aversion is valuable. Some obvious examples are flying aircraft, operating space flights, running a country, or being in charge of the nuclear deterrent button. Yet we rarely observe women in these jobs. This could be due to cultural factors. Or it could be that women do not apply for them, perhaps due to social conditioning. Or it could be because men are inherently more competitive and wish to keep these jobs for themselves.[16]

What do experimental studies find about gender differences in competitiveness? Gneezy, Niederle & Rustichini (2003) conducted a laboratory experiment in which university students (half the group were female, half male) were asked to solve mazes on a computer. The rewards could take the form of either "winner takes all" (a tournament), or piece rates. When men and women were paid piece rates, there were no significant gender differences in performance. But in a mixed-sex tournament, they found men performed better relative to the benchmark, but women's performance was unaffected. In contrast, in a single-sex tournament, the mean performance of women increased. Thus women under-perform only when competing against men and not in same-sex scenarios. Gneezy et al. suggested that women might dislike competing with men, or that perhaps they feel less competent than men and this depresses their performance.[17] They argue that single sex tournaments represent a strong form of affirmative action, since the proportion of women among winners reflects the gender composition of the participant pool (50 per cent),

[16] For an interesting perspective on the position of women from hunter gatherer societies through to agriculturally based societies, see Paul Seabright's 2005 Royal Economic Society lecture. See also John Stuart Mill (1869).

[17] Of course it is also possible that women have been conditioned to believe that they should not do better than men in mixed sex environments, in case it affects their marriage/partnering prospects.

whereas in mixed sex tournaments only about 20 per cent of winners are women. They also found that this change in the composition of winners involved no loss of performance. In contrast to the study summarized below, this was not due to gender differences in risk aversion.

Datta Gupta, Poulsen and Villeval (2005) distinguished between competition and risk aversion in their experiment based on a number of French-based undergraduate students. Before performing a task, the subjects chose whether to perform under a competitive payment scheme (tournament) or a non-competitive payment scheme (piece rate). Women were less likely than men to choose the competitive payment scheme. While many men and women were found to be overconfident about relative ability, this did not affect their choices. Instead, risk aversion was found to matter for women in their choice of payment scheme but not for men. A man's choice depended on whether he interacted with a male or female co-participant. When facing a woman, the man competed more if he believed that women compete too. But if the co-participant was male, the man competed regardless of his beliefs about men's entry rate into the competition.[18]

What about cooperative behavioural differences between men and women? The Croson and Gneezy (2004) survey shows that the experimental evidence on social preferences is mixed. Clearly the jury is still out on this. While these studies are interesting and important, we are still a long way off being able to conclude that the glass ceiling in Europe is due to different male and female social preferences.

Do studies based on individual-level survey data with information about preferences and attitudes offer a way forward? Clearly the use of contemporaneous measures of risk-aversion, self-esteem and competitive/collaborative behavioural traits is dogged by potential endogeneity. For interesting studies attempting to address these issues, see Vella (1994), Swaffield (2000) and Manning and Swaffield (2005). Manning and Swaffield (2005) aim to explore the issue of gender differences in psychological factors. For some of their estimation, they focus on otherwise identical, fully "work-committed" women and men (those with no children, no intention of having children, and with continuous full-time work experience). They find that, on labour market entry, the gender wage gap is zero. But after 10 years, there is

[18] The authors speculate that this may be due to social norms and possibly evolutionary factors.

a 12 log points wage gap that is unaffected by occupation. Manning and Swaffield then exploit information on self-esteem measured at age 10, which shows that males at this age have higher self-esteem than females. But self-esteem at age 10, while statistically significant, explains only a small proportion of the subsequent gender gap. While we clearly need more studies investigating these issues, it does seem probable that gender differences in these psychological factors will be unable to explain all of the gender pay gap.

Finally, we turn to discrimination. Expectations of family formation and fertility are private information. Employers base their behaviour on averages. For this reason women may not get pay increases, they may not get promoted when they deserve to, and they may not get offered the jobs they deserve. Moreover their willingness to pursue outside offers and their ability to accept these may also differ from that of men, reinforcing their poorer position in the labour market.

4. What are the policy implications?

There is no unique policy solution to the problems of the gender pay gap and the glass ceiling in Europe. Instead it would seem that policies should be formulated in a number of areas. For instance, affirmative action has the potential to change the proportion of females at higher levels in the public and private sectors. This could have knock-on effects through mentoring of more junior females, networks and the like. Companies' personnel departments clearly have a lot of scope for action in this regard if they wish to change intra-firm cultural attitudes. Pay bargaining and negotiation skills might also be offered by colleges and universities to final year students (it seems unlikely firms would introduce these). Moreover, childcare for pre-school children could be expanded in those countries in which it is weak, which would potentially have the additional effect of improving a country's stock of human capital. Tax incentives might also be offered in this regard, as introduced recently by the UK government.[19]

[19] In the UK recently a scheme has been introduced whereby Childcare Vouchers are Tax and National Insurance exempt for the amount £55 per week. See http://www.hmrc.gov.uk/childcare/. In addition, since April 2004 all three and four year olds have been entitled to a free, part-time early education place. Initially capped at being free for twelve and a half hours per week, this is soon to be extended to fifteen hours. See http://www.surestart.gov.uk/improvingquality/guidance/freenurseryeducation/.

(These policies might also have an additional effect of raising fertility rates, since they lower the cost of children.)

5. Conclusion

The paper summarised some stylised facts about the gender pay gap in a number of European Union countries, and showed that in each country this varies over the wages distribution. While some countries are doing better than others in this regard, almost without exception the gender pay gap was largest towards the top of the wages distribution. This glass ceiling effect is bigger in the private sector than in the public. The paper also outlined a number of possible explanations for these stylised facts. Policies and institutions appear to play a role. There may well be other factors with which economists have only relatively recently begun to grapple—such as culture, social custom, bargaining skills and perhaps preferences—that could play a part. Nonetheless, it also seems highly likely that discrimination contributes towards gender pay gaps and the glass ceiling in Europe.

References

Albrecht, J., Bjorklund, A. and Vroman, S. (2003), Is there a glass ceiling in Sweden?, Journal of Labor Economics 21, 145-177.

Arulampalam,W., Booth A.L. and Bryan, M.L. (2007), Is there a glass ceiling over Europe? Exploring the gender pay gap across the wages distribution, Industrial and Labor Relations Review 60, 163-186.

Azmat, G., Güell, M. and Manning, A. (2006), Gender gaps in unemployment rates in OECD countries, Journal of Labor Economics 24, 1-37.

Babcock, L. and Laschever, S. (2003), Women Don't Ask: Negotiation and the Gender Divide, Princeton University Press, Princeton, NJ.

Baron, J. and Cobb-Clark. D. (2006), A distributional analysis of gender wage differences in Australia, Mimeo, Australian National University.

Becker, G. (1957), The Economics of Discrimination, University of Chicago Press, Chicago.

Bell, L.A. (2005), Women-led firms and the gender gap in top executive jobs, IZA Discussion Paper 1689, Institute for the Study of Labor, Bonn, Germany.

Bertrand, M. and Hallock, K.F. (2001), The gender gap in top corporate jobs, International Labor Relations Review 55, 3-21.

Blackaby, D., Booth, A.L. and Frank, J. (2005), Outside offers and the gender pay gap, The Economic Journal 115, F81-F107.

Blau, F.D. and Kahn, L.M. (1992), The gender earnings gap: Learning from international comparisons, American Economic Review 82, 533-538.

Blau, F.D. and Kahn, L.M. (1996), Wage structure and gender earnings differentials: An international comparison, Economica 63, Supplement: Economic policy and income distribution, S29-S62.

Blau, F.D. and Kahn, L.M. (2003), Understanding international differences in the gender pay gap, Journal of Labor Economics 21, 106-144.

Boeri,T., Brugiavini, A. and, Calmfors L. (eds.) (2001), The Role of the Unions in the Twenty-first Century, Oxford University Press, UK.

Booth, A.L., Francesconi, M. and Frank, J. (2003), A sticky floors model of promotion, pay and gender, European Economic Review 47, 295-322.

Croson, R. and Gneezy, U. (2004), Gender differences in preferences, Mimeo, Wharton University of Pennsylvania.

Datta Gupta, N., Oaxaca, R. and Smith, N. (2006), Swimming upstream, floating downstream: Comparing women's relative wage progress in the United States and Denmark, Industrial and Labor Relations Review 59, 243-66.

Datta Gupta, N., Smith N. and Verner, M. (2006),Child care and parental leave in the Nordic countries: A model to aspire to?, IZA Discussion Paper 2014, Institute for the Study of Labor, Bonn, Germany.

Datta Gupta, N., Poulsen, A. and Villeval, M.-C. (2005), Male and female competitive behaviour: Experimental evidence, IZA Discussion Paper 1833, Institute for the Study of Labor, Bonn, Germany.

de la Rica, S., Dolado, J.J. and Llorens, V. (2005), Ceiling and floors: Gender wage gaps by education in Spain, IZA Discussion Paper 1483, Institute for the Study of Labor, Bonn, Germany.

Filer, R.K., Hamermesh, D.S. and Rees, A. (1996) The Economics of Work and Pay, Harper Collins College Publishers, New York, NY.

Goldin, C. and Rouse, C. (2000), Orchestrating impartiality: The impact of blind auditions on female musicians, The American Economic Review 90, 715-41.

Gneezy, U., Niederle, M. and Rustichini, A. (2003), Performance in competitive environments: Gender differences, Quarterly Journal of Economics 118, 1049-74.

Holzer, H.J. (2007), The economic impacts of affirmative actions in the US, Swedish Economic Policy Review, this issue.

Kee, H.J. (2006), Glass ceiling or sticky floor? Exploring the Australian gender pay gap, The Economic Record 82, 408-427.

Landers, R.M., Rebitzer, J.B. and Taylor L.J. (1996), Rat race redux: Adverse selection in the determination of work hours in law firms, American Economic Review 86, 329-348.

Lucifora, C. and Meurs, D. (2004), The public sector pay gap in France, Great Britain and Italy, IZA Discussion Paper 1041, Institute for the Study of Labor, Bonn, Germany.

Machado, J. and Mata, J. (2005), Counterfactual decomposition of changes in wage distributions using quantile regression, Journal of Applied Econometrics 20, 445-465.

Manning, A. and Swaffield, J. (2005), The gender gap in early-career wage growth, CEP Discussion Paper 700, Centre for Economic Performance, London School of Economics.

Mill, J.S. (1869), The Subjection of Women, Longmans, Green and Company.

Miller, P. (2005), The role of gender among low paid and high paid workers, The Australian Economic Review 38, 405-17.

OECD (2001), Employment Outlook, June 2001, Ch. 4, Balancing work and family life: helping parents into paid employment, OECD, Paris.

Pylkkänen, E. and Smith, N. (2004), The impact of family-friendly policies in Denmark and Sweden on mothers' career interruptions due to childbirth, IZA Discussion Paper 1050, Institute for the Study of Labor, Bonn, Germany.

Ruhm, C.J. (1998), The economic consequences of parental leave mandates: Lessons from Europe, Quarterly Journal of Economics 113, 285-317.

Seabright, P. (2005), Biology and economics of the sex war, The Royal Economic Society 2005 Lecture.

Smith, N. and Westergaard-Nielsen, N. (1988), Wage differentials due to gender, Journal of Population Economics 1, 115-30.

Swaffield, J. (2000), Gender, motivation, experience and wages, CEP Discussion Paper 0457, Centre for Economic Performance, London School of Economics.

Vella, F. (1994), Gender roles and human capital investment: The relationship between traditional attitudes and female labour market performance, Economica 61, 191-211.

Waldfogel, J. (1998), The Family gap for young women in the United States and Britain: Can maternity leave make a difference?, Journal of Labor Economics 16, 505-545.