Gender Differences in Wage Growth and Promotion

(Abridged and very incomplete draft for submission to the LOWER-2005 Conference)

Philippe Van Kerm
CEPS/INSTEAD
Differdange, G.-D. Luxembourg
philippe.vankerm@ceps.lu

Keywords: wage dynamics, promotion, gender gap, glass-ceiling
JEL-classification: J16, J31, J62

Abstract

This paper provides new evidence about gender differences in wage growth and promotion using nationally representative survey data for Britain. Particular focus is placed on careful examination of a gender gap at the top of the wage scale in order to assess the existence of a glass ceiling above high-skill women. Nonparametric methods are advocated to condition the examination on the level of wage. No strong support for the glass ceiling hypothesis is found. On the contrary, the main result of this paper is that gender differences in wage growth are important at the bottom of the wage scale where women's wage evolution is significantly (and substantially) worse than men's.

1. Introduction

There is no dispute about the fact that women tend to fare worse than men in the labour market. Evidence of a wage gap in pay is abundant. This paper attempts to complement this static picture with an examination of gender differentials in wage 'dynamics' using nationally representative panel survey data for Britain.[1]

There is on-going debate about the existence (or otherwise) of a 'glass ceiling' above women in the labour market, that is an invisible barrier that inhibits promotion opportunities for women (but not men) and prevents women from reaching top positions. Do such 'glass ceilings' exist and exacerbate the male-female wage differential? Or, on the contrary, are women able to take advantage of promotions and other opportunities for wage growth to catch up partially to male wage rates?

Existing evidence based on surveys representative of large populations (rather than based on personnel record data) is relatively scarce, and results are mixed. There is no strong support for the standard 'glass ceiling' hypothesis: promotion probabilities often turn out to be similar between males and females.[2] However, the return to promotion in terms of associated wage growth may differ more substantially between men and women. Booth et al. (2003) coined the term 'sticky floors' (as opposed to 'glass ceiling') to describe such a situation. They developed a general model of promotion and wage growth allowing such (and other) patterns to emerge by combining elements of the classical Lazear and Rosen (1990) model (women having better non-market opportunities), discrimination in the form of different reaction to outside offer threats, and gender differences in the distribution of outside offers.

This paper provides additional empirical evidence about these issues. It first attempts to identify if there are observable differences in wage growth between males and females. Differences in promotion rates (intra-firm job mobility) and in quit rates (inter-firm job mobility) are then considered. Crucially, and this is where the approach is novel, local kernel weighting methods are applied to assess men's and women's wage growth rates, as well as quit and promotion probabilities, conditionally on base period wage level. It makes it possible to examine the issue in greater detail and identify what happens at different points on the wage scale. This is a relevant approach in this context. If 'glass ceilings' are in effect, we may indeed anticipate differences in wage growth and promotion probabilities to appear only for higher wage classes. Similarly, some of the forces at play in the 'sticky floors' model of Booth et al. are very likely to have different effects for men and women at different points on the wage scale (e.g. the distribution of outside offers, or of non-market opportunities). This paper shows that looking at the mean, as typically done in empirical analyses, is too restrictive.

Results based on the British Household Panel Survey (BHPS) 1991-2001 indicate that gender differences in wage growth and promotion are much less substantial than differences in wage levels. The dynamics of pay in continuous employment do not appear to be driving male-female differentials. Rates of wage growth and inter-firm job change fall with wage level, while promotion rates increase. This pattern holds for both men and women. As opposed to what 'glass ceiling'

[1] A companion paper has examined this issue with data for Luxembourg (Van Kerm, 2004).

[2] I look directly at wage growth and promotion to assess the 'glass ceiling' hypothesis. An alternative strand of literature has assessed 'glass ceilings' by comparing the top quantiles of men's wage distribution with the top quantiles of women's wage distribution; see Albrecht et al. (2003). Evidence of a 'glass ceiling' with this alternative approach is strong.

arguments suggest, it is at the bottom of the wage scale that women tend to fare worse than men in terms of wage growth and promotion. Intra-firm and inter-firm mobility patterns explain no difference in wage growth, but human capital and job characteristics mitigate the difference in wage growth: Women would actually do somewhat worse if they had men's jobs and human capital characteristics.

As this research is in progress, this paper is very incomplete. It describes the data used for the analysis, and very briefly sketches the statistical methods applied. The first set of results are described and commented but no detailed discussion of the implications of these results is provided here. A brief (provisional) conclusion ends the paper.

2. Data and methods

The data analysed are extracted from the first 11 waves of the British Household Panel Survey (BHPS) covering the period 1991-2001.[3] The BHPS is a nationally representative sample survey of the population living in private households in Britain. It is a longitudinal household survey. The initial interviews were made in Autumn 1991 and were repeated annually thereafter. The same respondents are followed over time. The original sample is augmented by new entrants in originally sampled households, and reduced by drop-out due to non participation. Refreshment samples were added to the original sample in the course of the survey to compensate for attrition. See Taylor (2003) for more information on the BHPS.

At each interview, the BHPS collects data on activity status and wages, among many other things. It also tracks job changes made by the respondents since the previous interview, including job changes within the same employer. To identify promotion in the BHPS data, I follow Francesconi (2001) who uses the following approach. "For each job ended during the 12-month period between interview dates, individuals give the reason for stopping the job. One of the reasons on the available list is promotion. We use this information to construct our measure of whether or not an individual is promoted in a given year. Therefore promotions are likely to be substantial, in that they cause individuals to perceive themselves to be doing a different job for the same firm."

The analysis is conducted on all 11 waves of data pooled. Each pair of consecutive wage observations for t and t+1 in the data contributes one observation. I focus only on people with strong labour market attachment. Men and women in the data must meet the following criteria: (i) be aged between 16 and 60, (ii) be a full-time worker at both t and t+1, (iii) not being self-employed at either t or t+1. This ensures that people in the sample have strong labour market attachment and are exposed to promotion opportunities. The sample consists of approximately 12,000 observations (t-t+1 pairs) for female workers and 17,000 observations for male workers.

The distinctive feature of this analysis is to use flexible non-parametric methods in order to look at wage growth and promotion rates for men and women, conditionally on the wage level. This permits a more direct assessment of the existence of a 'glass ceiling' above women at the top of the wage distribution.

[3] The BHPS data were made available through the ESRC Data Archive. The data were originally collected by the ESRC Research Centre on Microsocial Change at the University of Essex (now incorporated within the Institute for Social and Economic Research). Neither the original collectors of the data nor the Archive bear any responsibility for the analyses or interpretations presented here.

Methods are only briefly sketched here. The various statistics are presented when conditioning on a base period wage located at (or near) the q^{th} quantile of the base wage distribution. The first step is therefore to estimate the quantile position of each observation (indexed by *it*) in the distribution of base (time t) wages when the men and women samples are merged. Call it p_{it}. In a second step, kernel weights based on p_{it} are computed for all observations and for a grid of fixed quantile points. For any grid point $0<p<1$, the kernel weight is $w_{it}^{p} = \frac{1}{h} K\left(\frac{p - p_{it}}{h}\right)$ where K is the Epanechnikov kernel function, and h is a bandwidth parameter that controls the degree of smoothness of the resulting estimates over the grid points. A small h gives a large weight to p_{it} very close to p but the weight declines rapidly with $|p-p_{it}|$. On the contrary, a large h gives weights that are positive over a much larger range of p_{it} values around p. In the application, h is set to 0.15. This results in the Epanechnikov kernel weight functions illustrated in Figure 1.

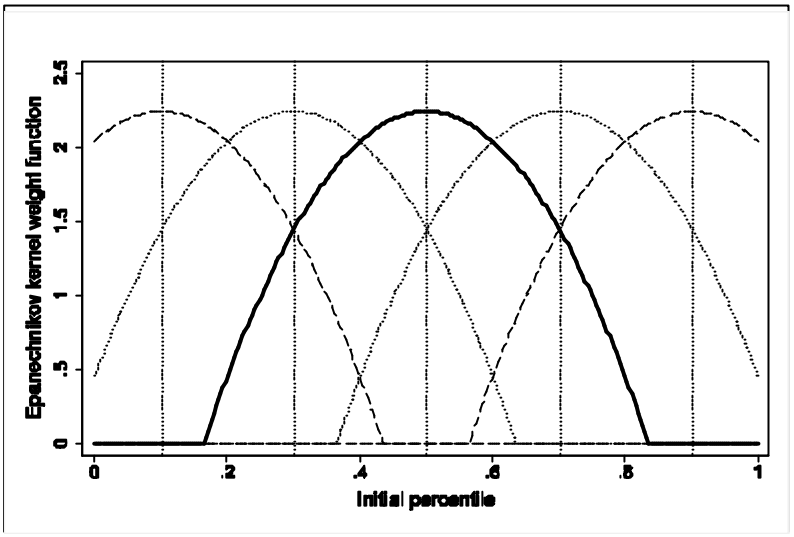

Figure 1: Epanechnikov kernel weight functions with $h=0.15$ for p at 0.10, 0.30, 0.50, 0.70, and 0.90

Once kernel weights based on p_{it} are computed for all observations and across a set of grid points spanning the range of quantile positions, standard weighted estimation methods are used to construct the various statistics conditional on base wage level. Conditional promotion rates, job change rates, and expected wage growth rates are estimated as simple weighted means (see Figures 2 and 3 below). Conditional quantiles of the wage growth distribution are obtained by inversing a weighted empirical distribution function (see Figures 3 and 4 below). Conditional wage returns to promotion and job change are estimated by weighted least squares regression (see Figure 5 below). Standard errors are estimated from a grouped jackknife resampling procedure taking the correlation due to the pooling of observations for the same individuals into account.

The next step is to consider wage growth, promotion and quits (and the wage return thereof) jointly. In order to do so, I extend and adapt the methodology developed by DiNardo et al. (1996) to try and explain the observed differences in wage growth between male and female (conditionally on base period wage) by (i) differences in human capital and job characteristics, (ii) differences in promotion and quit rates, and (iii) differences in wage return to promotion and quit (see Figures 6 and 7 below). The basic principle of these methods is to re-weight each observation in one of the two groups (here women) with weights calibrated so that the characteristics of the weighted group are identical to that of the reference group. The methodology applied here is described in greater detail in Van Kerm (2003).

3. Results

Results from the estimation procedures are best presented graphically. In all pictures, men are plotted in green and with solid lines, and women are plotted in orange and with dashed lines. (Orange appears lighter than green on black and white printing.) All point estimates are surrounded by pointwise 1.64*standard-error bands. The thick horizontal bands are the 1.64*standard-error bands for the *unconditional* estimates.

Estimated promotion rates and rates of job change for men and women, conditionally on base period wage are reported in Figure 2. Promotion rates are in general slightly higher for women than for men. These estimates are similar to those of Booth et al. (2003). Promotion rates tend to increase with the level of wage, but the gradient is small. The picture is completely different with rates of job change. For all but the highest wage levels, men are more likely to change employer than women. The rate of inter-firm obility falls drastically with the level of wage. For men, it ranges from about 17 percent near the 10th percentile down to about 7 percent at the 90th percentile. For women the range is about 13 percent at the bottom of the wage scale, down to 7 percent as for men. The steep gradient of the likelihood of changing employer along the wage scale would be completely concealed with the standard approach of looking at the mean. Additionally, 'at the mean', there is not much difference between men's and women's rate of job change, although the differences at the bottom of the wage scale are in fact striking.

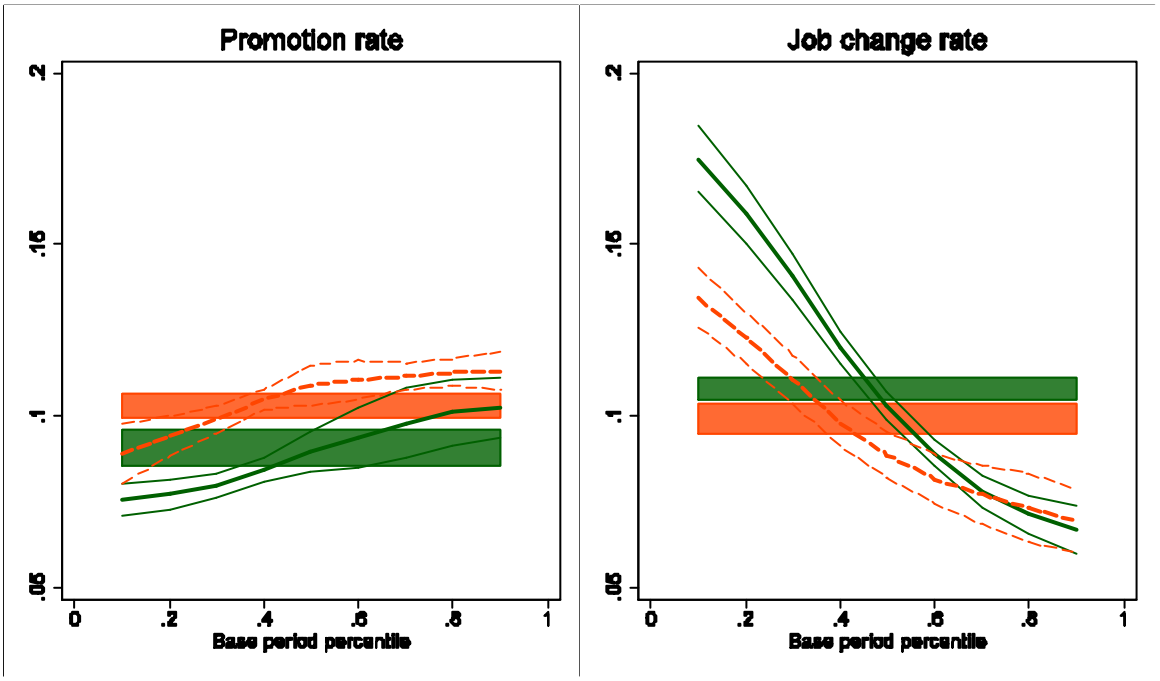

Figure 2: Promotion rate (left) and job change rate (right)

The results for promotion confirm that traditional models of within firm promotion and gender need to be augmented by a more general modelling of between firm mobility too.

How do the results for wage growth parallel those for promotion and job change? Figure 3 clearly indicates that the pattern of wage growth resembles the pattern for the rate of job change. Mean growth rate is higher for men than for women in the bottom half of the wage distribution, then women catch up to men's growth rates in the top half. Women's point estimates are above those of men but standard error bands overlap. But the most striking feature of the picture is the steepness of

the gradient. Mean wage growth rates range from 16 percent down to 2 percent for men (12 percent down to 2 percent for women). Again, this would be completely concealed by the unconditional estimates that are about 6 percent for both men and women.

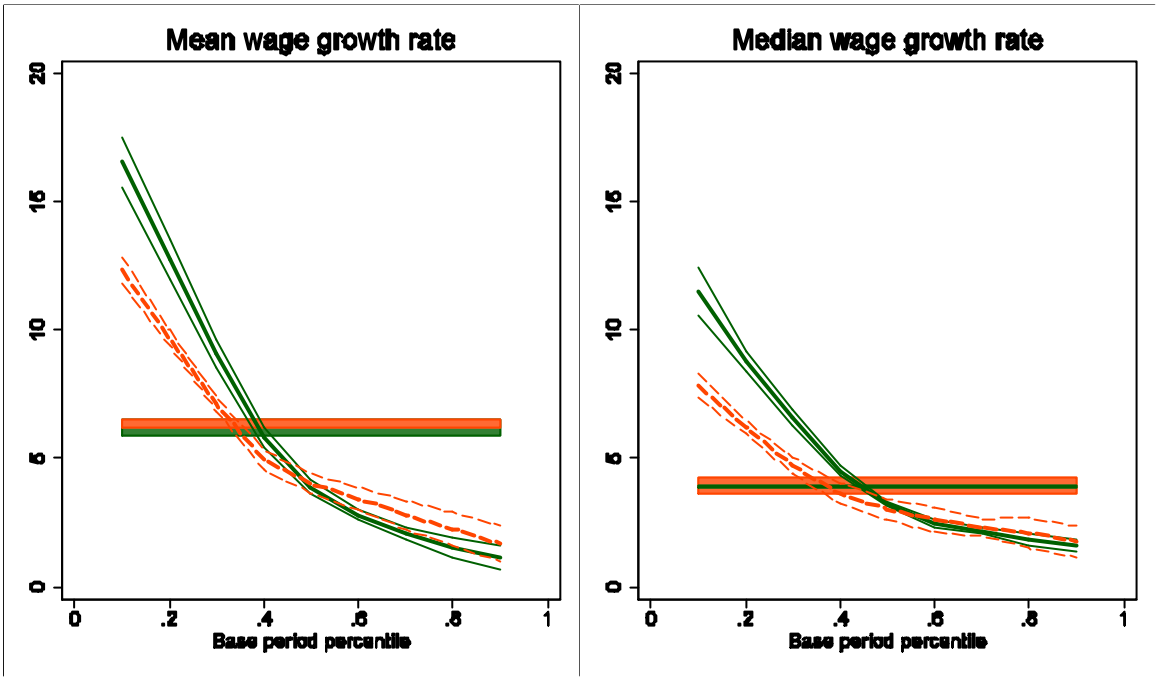

Figure 3: Mean wage growth rate (left) and median wage growth rate (right)

The right panel of Figure 3 and Figure 4 describe the conditional distribution of growth rates in greater detail. Unsurprisingly, the median growth rate is smaller than the mean and the gradient is smaller. But the general pattern remains the same. The top and bottom quartiles indicate that the distribution of growth rate is more concentrated for women than for men. The top quartile is higher for men than for women across the whole range of base period wage. This means that the chances of getting the biggest wage increases are higher for men than for women. Conversely, the bottom quartile is negative but substantially higher for women than for men for all wage levels but the very bottom. This suggests that women are less likely to experience the largest pay losses. The distribution of women's wage growth can be described as less 'risky' than men's.

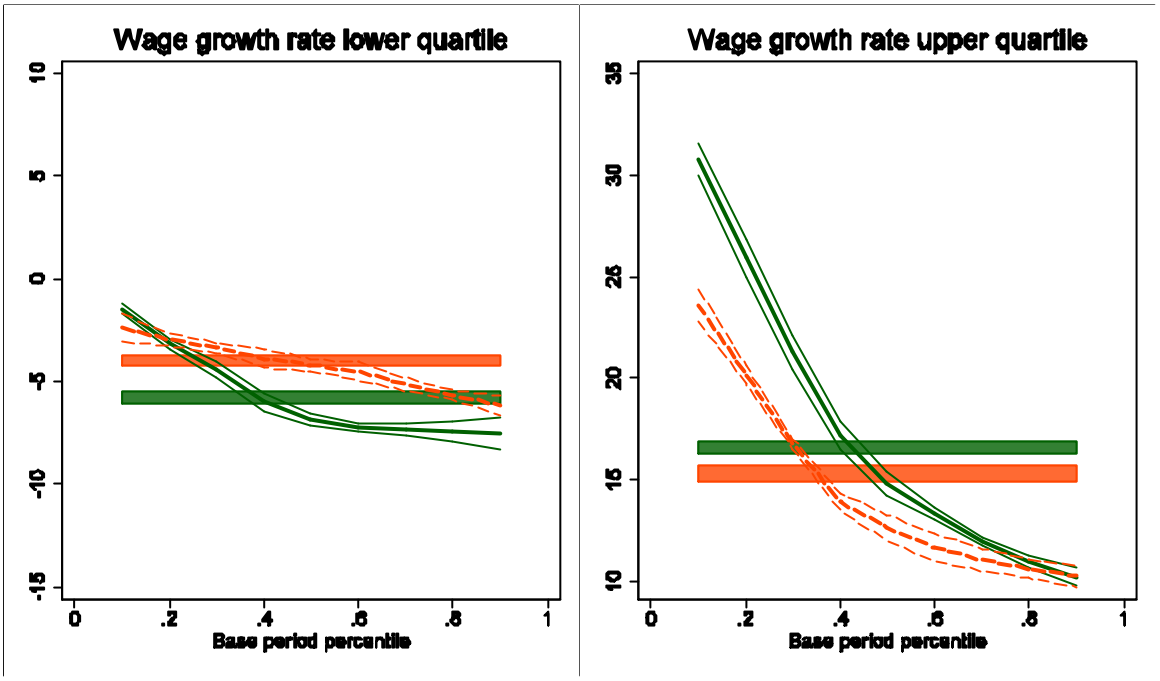

Figure 4: Lower quartile (left) and upper quartile (right) of wage growth rate

Earlier studies on the BHPS reported that there are no substantial differences in wage growth between men and women (see e.g. Booth et al. (2003) or Manning and Robinson (2004)). If this is true 'at the mean', closer scrutiny reveals that there *are* substantial differences between men and women in the dynamics of wage. Glass ceiling? Irrespectively of the definition adopted for the glass ceiling, be it in terms of wage growth or be it in terms of promotion rates, the results reported provide no support for it. On the contrary, much of the gender differences are to be found for low wage levels.

The previous results do not establish clearly the link between mobility (within firm and between firm) and wage growth. Some evidence is given in Figure 5 that presents estimates of the return to promotion and job change in terms of wage growth. These results are derived from a weighted least squares regression of wage growth (measured by the annual change in log wage) with explanatory variables including dummy variables for promotion and between-employer job change (and interaction thereof) as well as a series of human capital variables (age, education) and job-related characteristics. The kernel weighting method is again applied to estimate the return on the promotion and job change dummies conditionally on base period wage level.

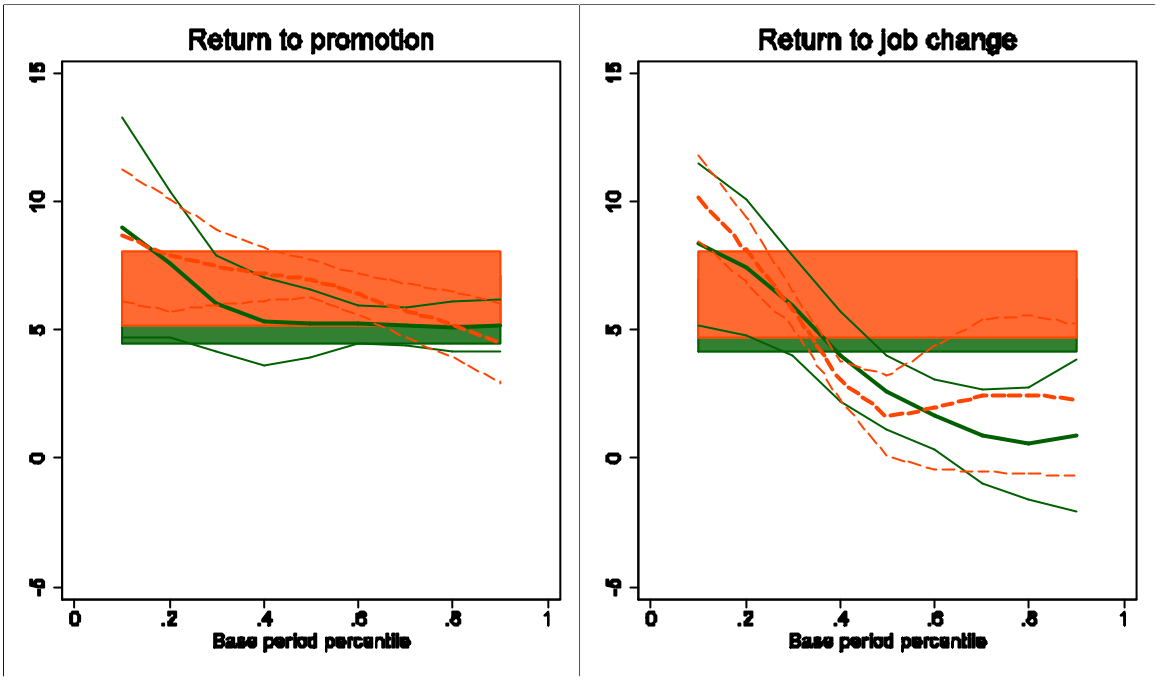

Figure 5: Wage return to promotion (left) and job change (right) (change in log-wage)

No difference emerges in these pictures between men and women. Patterns are similar and the standard error bands are wide and overlap largely. The return to both types of mobility falls with wage level, but the fall seems to be more marked for the job change. Near the bottom decile of the wage distribution, both types of mobility are associated with an additional wage increase of about 9 percent. However, near the upper decile, promotion remains associated with an additional wage increase of about 5 percent while employer change is only associated with a wage increase of about 2 percent. Note that no distinction is made between voluntary change of employer and unvoluntary change of employer. This may explain part of the differences between the effect of promotion and job change.

To complete the exercise, I now attempt to identify how much of the differences in wage growth between men and women can be accounted for (i) by mere compositional differences in terms of human capital and type of jobs held; (ii) by the differences in promotion and job change rates identified above (controlling also for human capital and job characteristics), and (iii) by differences in the return to mobility (although we have seen that no such difference seem to be at play). To this end, the variation of the Di Nardo et al. (1996) approach is applied.

Results are reported in Figures 6 and 7. The left panel of Figure 6 reports the actual mean wage growth for both men and women (as in Figure 3, but with standard error bands removed to improve readability). On the right panel, the added dashed blue line represents the expected growth wage of women that would be observed if the sample of women had men's human capital and job characteristics (but retained women's return to these characteristics, and women's mobility rates). The results are surprising since they do not go in the direction of accounting for the differences. In the bottom third of the wage distribution, where much of the gender differences in wage growth can be observed, adjusting women's attributes to men's does not affect the mean wage growth. Gender differences in attributes do not explain the lower wage growth of women. For higher wages, it is surprising to see that if women had men's characteristics, they would obtain *lower* wage growth

than they actually do. The effect is so large that it can not come as the main explanation for the men/women difference in wage growth.

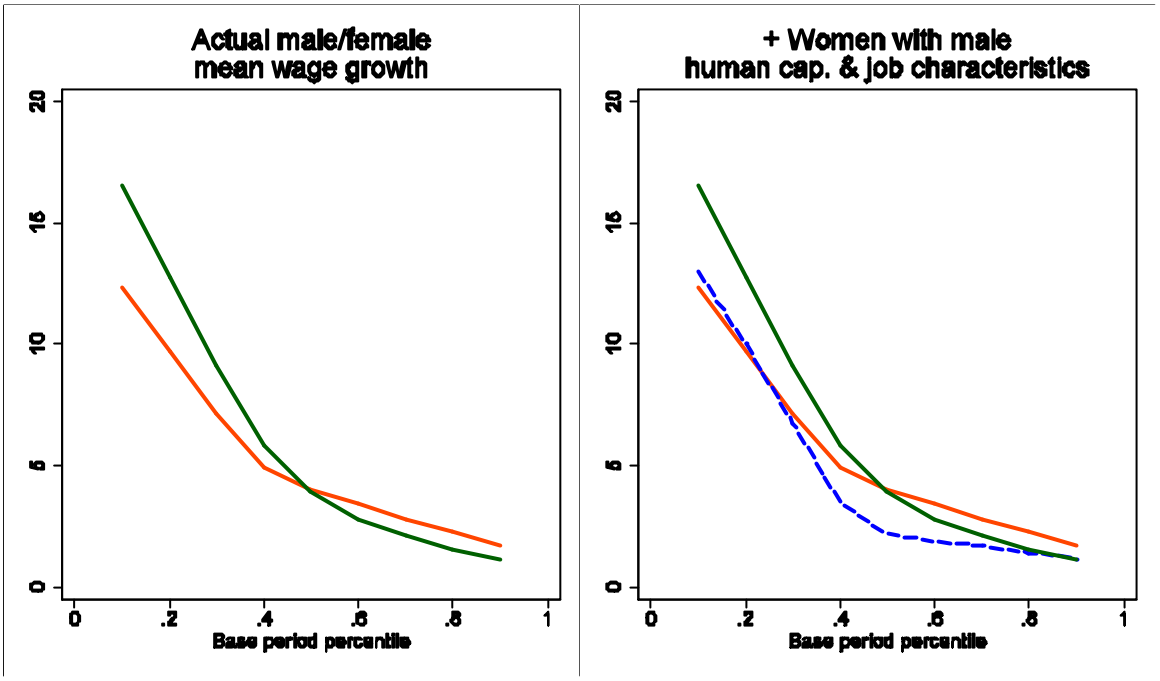

Figure 6: Counterfactual construction of women's mean wage growth as if having men's characteristics

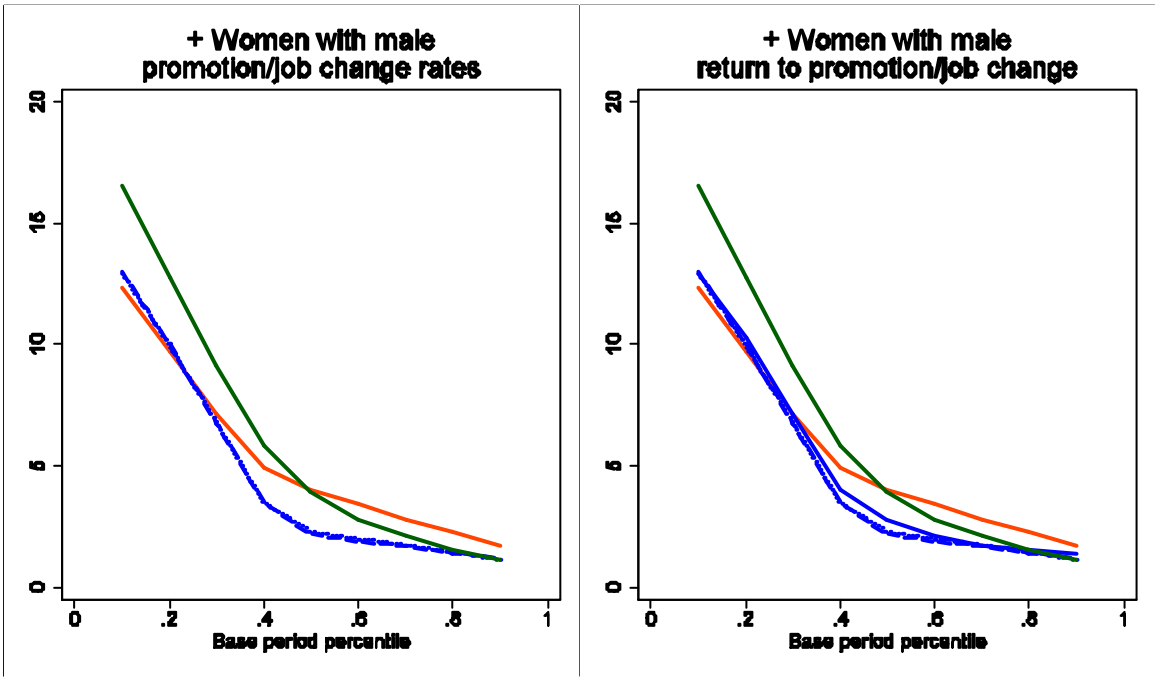

Figure 7: Counterfactual construction of women's mean wage growth as if also having men's promotion rates and job change rates (left), and as if also having men's return to promotion and job change.

In the left panel of Figure 7, an additional blue line shows the counterfactual mean wage growth profile that women would have if they had men's human capital and job characteristics, and if they also had the same rates of promotion and job change as men. The resulting line is indistinguishable from the previous blue dashed line. The marginal effect of differences in promotion rates and job

change is negligible. Finally the solid blue line on the right panel of Figure 7 represents the counterfactual wage growth pattern of women if they had the same characteristics as men, the same mobility rates, and the same wage return to mobility. The remaining differences between the solid blue line and the green line of men are gender differences that remain unexplained by the distribution of attributes, and by mobility effects. Again the effect of differences in return to mobility is small. It is only noticeable for wages in the middle of the distribution.

The three sets of explanatory factors for the wage growth differences between men and women do not account for much of the actual difference. Of the three, it is the difference in human capital and job attributes that has the largest effect. Rates of within firm mobility and between form mobility and the wage return thereof hardly account for anything of the differences in wage growth pattern.

4. Conclusion

The empirical evidence provided in this paper shows that, contrary to what has been suggested recently, there *are* gender differences in wage growth over time in Britain. But these differences only appear when we condition the investigation on the level of base period wage. Differences are particularly striking for low wage levels, while high wage men and women tend to obtain similar pay increases over time. Looking at the mean largely mitigates the picture. It is interesting to note that there is no sign of a 'glass ceiling' above highly skilled women. Rates of wage growth, as well as promotion probabilities, are similar to men's when we look at the top of the wage distribution. Similarly, no striking evidence of a 'sticky floor' shows up. Promotion chances are indeed similar for men and women, but the wage return to promotion do not vary significantly either. The results also dismiss the contention that women may catch up to men's wage rates because of higher wage growth.

Differences in inter-firm mobility appear to be larger than differences in intra-firm mobility. This may be consistent with the argument made by Booth et al. (2003) that men have better (or more) outside offers than women, or that they have greater flexibility to take advantage of the outside offers. However, much of the differences in wage growth are not explained by mobility differences. Note that this may also be consistent with the outside offer threat argument: men having better outside opportunities have more credible threats of quitting and therefore stronger bargaining power to obtain higher pay increases within the firm.

The empirical evidence reported in this paper raises more questions than it answers. However the main contribution of this paper is to suggest that further research should be better invested in explaining differences between men and women at the bottom of the wage distribution, that is among less skilled workers, rather than looking for potential causes of a 'glass ceiling' among high wage earners. It also indicates that models of intra- and inter-firm mobility need to be augmented by additional hypotheses since much of the growth do not appear to be directly related to the latter.

References

Albrecht J, Björklund A, Vroman S (2003). "Is there a glass ceiling in Sweden?" *Journal of Labor Economics*, 21(1), 145-177.

Booth AL, Francesconi M, Frank J (2003). "A sticky floors model of promotion, pay and gender." *European Economic Review*, 47(2), 99-126.

DiNardo J, Fortin NM, Lemieux T (1996) “Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach.” *Econometrica*, 64(5), 1001-1044.

Francesconi M (2001). “Determinants and consequences of promotions in Britain.” *Oxford Bulletin of Economics and Statistics*, 63(3), 279-310.

Lazear EP, Rosen S (1990). “Male female wage differentials in job ladders.” *Journal of Labor Economics*, 8(2), S106-S123.

Manning A, Robinson H (2004). “Something in the way she moves: A fresh look at an old gap.” *Oxford Economic Papers*, 56(2), 169-188.

Taylor MF, Ed. (2003). British Household Panel Survey User Manual. University of Essex, Colchester, UK. `http://www.iser.essex.ac.uk/bhps`

Van Kerm P (2003). “An anatomy of household income volatility in European countries.” CHER Working paper 16, CEPS/INSTEAD, Differdange, G.-.D Luxembourg.

Van Kerm P (2004). “Gender differences in wage growth and promotion: The case of Luxembourg.” Unpublished paper, CEPS/INSTEAD, Differdange, G.-.D Luxembourg.