

Industrial & Labor Relations Review

Volume 59, Issue 2

2006

Article 6

Gender Differences in Promotion on a Job Ladder: Evidence from Finnish Metalworkers

Tuomas Pekkarinen*

Juhana Vartiainen†

*Nuffield College, Oxford,

†National Institute of Economic Research, Stockholm,

Copyright ©2006 Cornell University. All rights reserved.

Gender Differences in Promotion on a Job Ladder: Evidence from Finnish Metalworkers

Tuomas Pekkarinen and Juhana Vartiainen

Abstract

This paper, using panel data on Finnish metalworkers for the years 1990–2000, explores gender differences in the allocation of workers across jobs of different complexity. The data provide measures for the complexity of the workers' tasks and for the individual productivity of each worker. The results indicate that women were less likely to be promoted than men who started their careers in similar tasks. A productivity comparison shows that there was no gender-related productivity differential at the time of the initial assignment, but that women became, on average, more productive than men afterward, in the subsets both of promoted workers and of non-promoted workers. The most plausible interpretation of these results, the authors argue, is that women faced a higher promotion threshold than men. Consistent with this interpretation, they find that the quit rate for young female workers was higher than that for young men.

KEYWORDS: gender differences in promotion on a job ladder

GENDER DIFFERENCES IN PROMOTION ON A JOB LADDER: EVIDENCE FROM FINNISH METALWORKERS

TUOMAS PEKKARINEN and JUHANA VARTIAINEN*

This paper, using panel data on Finnish metalworkers for the years 1990–2000, explores gender differences in the allocation of workers across jobs of different complexity. The data provide measures for the complexity of the workers' tasks and for the individual productivity of each worker. The results indicate that women were less likely to be promoted than men who started their careers in similar tasks. A productivity comparison shows that there was no gender-related productivity differential at the time of the initial assignment, but that women became, on average, more productive than men afterward, in the subsets both of promoted workers and of non-promoted workers. The most plausible interpretation of these results, the authors argue, is that women faced a higher promotion threshold than men. Consistent with this interpretation, they find that the quit rate for young female workers was higher than that for young men.

The gender wage gap is a persistent phenomenon. Even though female participation rates have been high for decades and many countries have long experience with anti-discriminatory legislation, women still tend to earn lower wages than men. An important part of this wage gap is explained by occupational segregation. It is common to find that the gender wage gap is consid-

erably reduced when sufficiently narrow occupational categories are controlled for (Blau and Ferber 1987; Groshen 1991; Bayard et al. 2003). Hence, one of the key elements in understanding the gender wage gap is the asymmetric allocation of men and women across occupations.

This paper uses panel data on Finnish metalworkers to study gender differences in the allocation of workers across jobs of different complexity. In the Finnish metal industry, as in many other industries, women typically work on less complex jobs than men. Our aim is to determine whether this happens because women need to meet higher productivity requirements than men

*Tuomas Pekkarinen is Postdoctoral Research Fellow, Nuffield College, Oxford. Juhana Vartiainen is Head of the Research Division at the National Institute of Economic Research, Stockholm. The authors thank Kate Antonovics, David Autor, Roland Fryer, Andrea Ichino, Markus Jäntti, Frank Vella, and seminar participants at the University of Stockholm and IZA for helpful comments. They also thank the Confederation of Finnish Industry and Employers for permission to use its wage records; and Pekkarinen thanks the Academy of Finland, Finnish Work Environment Fund, and Yrjö Jahnsson Foundation for financial support.

The data are archived at the Labour Institute for Economic Research, Helsinki. Permission to use the data is controlled by the Confederation of Finnish Industry and Employers.

in order to be assigned to more complex jobs. We address this question by examining gender differences in the probability of promotion from the initial job assignments. Furthermore, we use a measure of individual productivity that is based on supervisors' performance evaluations to study how promoted and non-promoted women perform by comparison with their male counterparts. We then use this information to infer whether the productivity threshold for promotion differs for men and women.

1. Motivation

The previous empirical literature on gender differences in job allocation has almost exclusively concentrated on estimating gender differences in the probability of promotion.¹ There are studies that use data from large surveys, such as Winter-Ebmer and Zweimüller (1997) using Austrian census data, McCue (1996) using the Panel Study of Income Dynamics, and Booth et al. (2003) using the British Household Panel Study. Other studies have used data from a single industry or firm: Granqvist and Persson (2002) analyzed gender differences in the career mobility of workers in the Swedish retail trade industry, Hersch and Viscusi (1996) focused on one U.S. public utility, and Jones and Makepeace (1996) looked at workers in a British financial institution. A special branch of the literature has studied the career advancement of academics (for example, Ginther and Hayes 1999, 2003; McDowell et al. 1999). The common finding is that women are less likely than men to get promoted.

But gender differences in assignment thresholds, if they exist, would have implications that go beyond promotion probabilities. If the productivity threshold is higher for women than for men, the pro-

motion process will improve the relative productivity of women within jobs. Since, under such conditions, promoted women will have passed a higher threshold than promoted men, they will, on average, be more productive than men who are promoted to the new job; moreover, men who remain in the less complex jobs will have failed to meet a lower productivity threshold and will therefore be less productive than women who remain in the same job.

These are the basic implications of the famous model by Lazear and Rosen (1990), which provides an alternative to the discrimination-based explanations of gender differences in wages and career attainment. In their model, men and women are assumed to have identical distributions of productive ability but to differ in their outside options. Because of their comparative advantage in non-market activities such as housework, women have better outside options than men and are thus more likely to quit their jobs. In a setting where the promotion of workers entails costly firm-specific training, this gender difference in outside options leads employers to set higher ability requirements for promotion of women than of men.

Here, we will directly address the question of whether women have to be more productive than men to be assigned to complex jobs. We focus on the early careers of workers who are initially assigned to jobs of similar complexity. The idea is that the productivity differences between men and women who are initially assigned to the same jobs should be so small that we can draw conclusions about the assignment thresholds by examining the gender differences in the probability of promotions and in productivity among promoted and non-promoted workers. If women are at the same time less likely to be promoted than men who start in similar tasks and also more productive than men among promoted and non-promoted workers, we can conclude that the promotion threshold is higher for women than for men. Finally, we use the estimated male and female quit rates to assess whether the gender differences in outside options could plausibly

¹Exceptions are Booth et al. (2003) and Hersch and Viscusi (1996), who also estimated gender differences in the wage growth upon promotion. Both studies found that women gain less from promotions than men do.

explain the asymmetry in the job assignment.

We believe that the Finnish metal industry data are suitable for this kind of analysis. First of all, they provide a measure for the complexity of the jobs. This complexity measure is valid for both within- and between-firm comparisons. In this industry, all jobs are evaluated according to their complexity, and on the basis of this evaluation a minimum wage is attached to each job. We use these minimum wages to construct a complexity ladder of jobs. Second, the panel nature of the data allows us to distinguish between the initial job assignments and subsequent promotions. We can therefore compare the careers of men and women who start in jobs of similar complexity. Finally, the data provide information on bonuses that are based on individual performance evaluations, so that we do not have to rely on final wages when measuring individual productivity.

Even though these unusual features of the data are the main reason to focus on the metal industry, we would argue that this industry is as good as any other for studies of gender career differentials. It includes the entire modern electronics sector and is larger and more successful than any other Finnish manufacturing industry. Its tasks are multifarious but they almost never depend, anymore, on sheer physical force, a requirement that would give male workers a comparative advantage. Finally, the metal industry's overall gender gap is representative of the Finnish situation and of that of many other European economies: the raw gap is 20%, and half of it can be accounted for by Mincer-type variables and detailed job classifications.²

2. Empirical Strategy

Our empirical strategy is based on estimating gender differences in the probability of promotion from the first task assignment and in the productivity of promoted

and non-promoted workers. However, in order to use these estimated gender differences to infer whether the promotion thresholds are different for men and women, we need to spell out the conditions under which the fact that women are at the same time less likely to be promoted and more productive among promoted and non-promoted workers implies that the productivity thresholds of promotion are higher for women than for men.

Denote the productive ability of a worker as η and the thresholds of promotion for men and women as η_m^* and η_f^* .³ Thus, men and women whose productive abilities are higher than η_m^* and η_f^* , respectively, are promoted to more demanding tasks. Assume that the productive abilities of men and women are distributed according to distribution functions $G_m(\eta)$ and $G_f(\eta)$.

Women will be less likely to be promoted to the demanding job if

$$(1) \quad \Pr(\eta \geq \eta_m^* \mid \text{male}) = 1 - G_m(\eta_m^*) \geq 1 - G_f(\eta_f^*) = \Pr(\eta \geq \eta_f^* \mid \text{female}).$$

Furthermore, women will be, on average, more productive within jobs if

$$(2) \quad E[\eta \mid \eta \geq \eta_m^*, \text{male}] = \int_{\eta_m^*}^{\infty} \eta dG_m(\eta \mid \eta \geq \eta_m^*) \leq \int_{\eta_f^*}^{\infty} \eta dG_f(\eta \mid \eta \geq \eta_f^*) = E[\eta \mid \eta \geq \eta_f^*, \text{female}]$$

Under which assumptions concerning the relationship between $G_m(\eta)$ and $G_f(\eta)$ does the fact that (1) and (2) are both true necessarily imply $\eta_m^* \leq \eta_f^*$?

It is easy to see that if $G_m(\eta)$ and $G_f(\eta)$ are identical, both (1) and (2) can hold only if $\eta_m^* \leq \eta_f^*$. When the male ability distribution first-order stochastically dominates the female distribution, $1 - G_m(\eta_m^*) \geq 1 - G_f(\eta_f^*)$, condition (1) can also hold for

²See Vartiainen (2002).

³See the model by Lazear and Rosen (1990) for the derivation of these thresholds.

some values $\eta_m^* > \eta_f^*$, but condition (2) holds only when $\eta_m^* \leq \eta_f^*$. On the other hand, if the female distribution dominates, $1 - G_m(\eta_m) \leq 1 - G_f(\eta_f)$, condition (1) can only hold if $\eta_m^* \leq \eta_f^*$, whereas condition (2) also holds for some values such that $\eta_m^* > \eta_f^*$. Thus, in these cases women can be *simultaneously* less likely to be assigned to the complex job and more productive within jobs only if the female promotion threshold is higher than the male promotion threshold.

Naturally, one cannot draw such conclusions if the male and female ability distributions intersect. Ruling out intersecting ability distributions may seem like a strong assumption. After all, it is often argued that the variance of male ability is higher than that of female ability. However, in our case we are not comparing the “global” male and female ability distributions. Instead, we focus on “ability brackets” of men and women who have been initially assigned to similar tasks by their employers. It seems very unlikely that the male and female ability distributions would systematically intersect within all of these ability subsets. Indeed, in our empirical analysis we will show that the male and female distributions of measured productivity seem to be remarkably similar within the initial jobs. Hence, in these circumstances, observing that women are both less likely to be promoted and, within promoted and non-promoted groups, more productive than men implies that the female promotion threshold is higher than the male one.

3. Data

The data used in this paper come from the wage records of the Confederation of Finnish Industry and Employers (*Teollisuus ja työnantajat*). The wage records contain detailed information on the wages and working hours of all the workers who are affiliated with the confederation. These register data are collected to facilitate wage bargaining, but they also provide the raw material for much of Statistics Finland’s wage statistics on manufacturing. In the Finnish metal industry, the wage records

cover practically 100% of the firms that employ more than 25 workers.

The wage records’ data on wages and working hours can be considered exceptionally reliable since the information comes directly from the firms’ wage accounts. However, the information on the individual characteristics is rather sparse. Basically, only age, gender, and seniority can be identified from the raw data. For the purposes of this paper, however, most serious is the absence of information with which to create variables for the marital status of the worker and the number of dependent children.

In this paper, we use data on all the blue-collar metalworkers who started their first employment relationship in the Finnish metal industry between 1990 and 1995 and whom we can follow for at least five years up to year 2000, when our data end. Thus, we only use information on workers who were newcomers to the industry and whose career lasted for at least five years.⁴ We chose to restrict the sample this way because for our purposes it is essential to observe the workers at their initial job assignments and to follow them for a reasonable number of years. Furthermore, restricting the analysis to workers who stay for more than four years selects the workers with a strong labor market attachment.

This panel of newcomers to the metal industry has 82,012 employee/year observations on 11,472 workers, of whom 2,484 (22%) are women. We have 63,023 observations for which both the current year’s job and next year’s job are known.

In Table 1, we present the descriptive statistics on this sample and compare these men and women with the rest of the population of workers in the metal industry between 1990 and 2000. It is clear that this is

⁴To ascertain that an individual was a newcomer, we checked that he or she did not appear in the metal industry’s wage register in any previous year (we have records from year 1988 onward). We double-checked this condition by using an experience variable that is also available in the data set.

Table 1. Descriptive Statistics.

Variable	Men in Our Sample		Women in Our Sample		Men in the Population		Women in the Population	
	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation	Mean	Standard Deviation
Age	32.25	8.79	35.00	9.52	39.33	10.34	40.66	10.61
Complexity	34.80	3.07	30.84	2.74	35.59	3.11	31.09	2.79
Seniority	4.34	2.47	4.20	2.37	14.91	9.95	12.15	8.63
Bonus	8.12	3.92	8.83	3.75	9.44	4.06	9.47	3.96
Avg. Hourly Earnings	47.58	7.60	40.53	6.31	49.13	8.02	41.02	6.21
Individuals	8,988		2,484 (22%)		85,750		24,099 (22%)	
Ind./Year Observations	81,799							

Our sample consists of all the 1990–2000 observations on workers who entered the metal industry during 1990–1995 and stayed for at least five years. The descriptive statistics are overall means and standard deviations for years 1990–2000. The population consists of all the blue-collar workers in the metal industry during 1990–2000. Complexity refers to the occupation-related wage in FIM 1990. Bonus is reported as a percentage of the occupation-related wage. Average hourly earnings refer to the worker's average final wage in FIM 1990.

a strongly male-dominated industry. This fact should be taken into account when interpreting the results reported below.

3.1. Wage Determination in the Finnish Metal Industry

The sample was restricted to include only metalworkers because the peculiar wage determination process in this industry provides particularly interesting information on the complexity of the jobs. As is typical for the Nordic labor markets, the general wage determination guidelines for the Finnish metal industry are spelled out in the industry's collective agreement, which applies to all the firms in the industry.⁵ According to the metal industry's collective agreement, wages should be determined by the complexity of the job, the individual performance of the worker, and various individual and firm-specific factors. These

ground rules for wage determination apply to the entire industry. Thus, the purpose of the collective agreement is to ensure that each and every worker of the industry is subject to the same wage determination criteria.

A minimum wage for each job is specified based on the job's complexity. This minimum wage is called the *occupation-related wage*. A worker's individual performance on the job affects the wage outcome through a *personal bonus* of 2–17% of the occupation-related wage. The sum of the occupation-related wage (which should only reflect the task) and the personal bonus (which should only reflect the individual) is called the person's "basic wage."⁶ In this paper, we use occupation-related wages as a measure of the complexity of the job and personal bonuses as a proxy for the productivity of the individual.

⁵See Metalliteollisuuden keskusliiton ja Metallityöväen liiton välinen työehtosopimus 1.2.2003–15.2.2005, downloadable at <http://www.metalliliitto.fi/pdf/metallites2003.pdf> (in Finnish and Swedish).

⁶The basic wage, however, is usually not the final wage outcome. The firm is of course free to set an individual's final wage as long as it exceeds the basic wage. For a clear majority of metalworkers, the final wage indeed exceeds the basic wage.

3.2. Job Complexity

The complexity of the job is evaluated with a grading system similar to the ones used in some large establishments in the United States. The evaluation is carried out by a group of experts, who consider various aspects of the jobs and assign them points according to their complexity.

Complexity is evaluated according to three criteria: how long it takes to learn the tasks, the degree of responsibility associated with the tasks, and the working environment and working conditions associated with the tasks. According to the guidelines for the application of the collective agreement, the relative weights of these factors are, respectively, 50%, 17%, and 33%. This recommendatory weighting need not be rigorously enforced at every workplace, and the relative weights can be adjusted according to the specific characteristics of each establishment. In our conversations with the union and the employer officials, we were assured that these weights have remained more or less constant since the adoption of the grading system in 1990.⁷

Based on the evaluation of jobs, an occupation-related wage is determined for each job. The more demanding the job—that is, the more complexity points it gets—the higher is the corresponding occupation-related wage. There is a one-to-one correspondence between the complexity points and the occupation-related wages, as spelled out in the collective agreement. The occupation-related wages can therefore be interpreted as a continuous variable that measures the complexity of the job. Typically, around 50 different occupation-related wage levels are defined each year.

The data also include occupational category codes. Using this information, we excluded from our sample the workers in support jobs, such as clerical workers (who traditionally receive a monthly salary and have not been covered by the metalwork-

ers' collective agreement), janitors, and canteen workers. This exclusion, which yields a final sample that can be regarded as consisting solely of metal industry production tasks, did not change the results reported below in any meaningful way.

3.3. Individual Performance

A worker's individual performance is evaluated by his or her immediate supervisor. The performance evaluation is based on three criteria: how well the tasks are carried out, the worker's output relative to what is considered normal in the job, and how well the worker follows instructions and regulations.

The evaluation is relative in nature. The goal is to evaluate individual performance relative to what is considered normal on the job in question. Hence, if, for example, the output of all the workers in the firm increases by an equal amount, the performance evaluations of the individual workers should not be affected. The guidelines of the performance evaluation are set out in the collective agreement, but the ultimate evaluation is based on the supervisors' subjective assessment of each worker's performance.

Based on the supervisor's evaluation, each worker is paid a personal bonus that should amount to 2–17% of the worker's occupation-related wage. The collective agreement states that bonuses should be symmetrically distributed around the mean of 9.5% within coarser complexity categories that split the firm's complexity axis into three groups. The purpose of these requirements is to force supervisors to use the whole scale of bonuses every year and at all complexity levels, and to discount the effects of any firm- or time-specific shocks on individual productivity.

In order to check whether these principles of the collective agreement are also followed in practice, we examined the variation of bonuses across time, firms, and different levels of complexity. First, the distributions of bonuses in yearly cross-sections were virtually identical. The median was 9.5% in each year and there was only slight

⁷These weights, as well as the rest of the wage determination guidelines, are spelled out in the collective agreement and its appendices.

variation in top and bottom quintiles. Second, across the whole metalworker population, firm dummies explained only 2.5% of the variation in bonuses. Even though this suggests, in our opinion, that the bonuses are indeed used in accordance with the guidelines of the collective agreement, it is still true that exactly what is evaluated by the supervisor may vary across firms and tasks. In order to account for this variation, we measure individual performance using the deviation of the worker's personal bonus from the yearly firm-complexity-level cell mean. We will refer to this productivity measure as the *personal bonus deviation*.

4. Complexity Ladder of Jobs

As explained above, occupation-related wages are ordered by job complexity. In this paper, we use this ordering of jobs as a job ladder, such that within-firm upward movement on the ladder is interpreted as a promotion and downward movement as a demotion.

The fact that the occupation-related wages are a component of the final wages makes their use as a complexity measure somewhat problematic. There is some year-to-year variation in the occupation-related wages that is not related to changes in the complexity of the jobs. Collectively agreed-upon wage increases affect the wage determination process by moving the entire scale of occupation-related wages upward by the same percentile increase, typically once a year, to account for the effects of inflation and the general productivity growth. This means that the scale with which the complexity of the jobs is measured is not constant in time.⁸

We corrected for these changes by descaling the occupation-related wages in the following way. We first grouped the workers according to their occupation-re-

lated wages within each year and examined the within-group distributions of changes in the occupation-related wages. This analysis revealed that for most of the workers in these groups the year-to-year changes in occupation-related wages were identical. We interpreted the group mode changes of the occupation-related wages as increases that were not related to changes in the complexity of the jobs but, instead, reflected new collective wage bargains. The occupation-related wages were then corrected by subtracting the annual mode changes from occupation-related wages. After this correction, the occupation-related wages form a continuous measure of task complexity for all the years 1990–2000.

To confirm that using occupation-related wages as a job ladder is reasonable, we examined the distribution of changes in occupation-related wages to see whether they resemble stylized facts about promotions and demotions. In the whole population of metalworkers, on average 84% of the workers remain at the same level of complexity between two typical years, whereas 11% move to more complex jobs and 4% move to less complex jobs. Furthermore, the movement between levels rarely leaps over many complexity levels. In our view, these patterns resemble the stylized facts about promotions and demotions and are in line with the movement of workers on job ladders in studies such as Baker et al. (1994) and Treble et al. (2001). It seems appropriate to interpret the complexity axis as a job ladder.

5. Gender Differences in Job Assignments

Where do men and women wind up on this ladder? Figure 1 plots the kernel estimates of the male and female complexity distributions in the 1990 cross-section.⁹ The

⁸Whether this upward updating reflects national-level centralized wage bargaining or a round conducted solely between the two metal industry parties is immaterial for our analysis, since in both cases the end result is just a general upward shift of the occupational wage scale.

⁹The estimation was done using the Epanechnikov kernel with a bandwidth of .5. Different bandwidths were tried without relevant changes in the qualitative results. The same note applies to Figures 2 and 3 below.




Figure 1. Kernel Estimates of the Male and Female Distributions of Complexity: 1990 Cross-Section.

figure reveals that women were concentrated at the low end of the complexity axis, while most men worked on more complex jobs. This pattern is repeated throughout the years 1990–2000.

Perhaps the most common approach to studying gender differences in job allocation is to examine the movement of workers between jobs. If the assignment thresholds are identical for men and women, we should not see any differences in the movement of men and women once the productive characteristics are controlled for. In this section, we study the gender differences in the allocation of workers across jobs of different complexity at the initial job assignment and in the promotion process.

5.1. Initial Job Assignments

In our sample, we observe the jobs in which workers started their careers and their subsequent promotions and demotions. In Figure 2, we have plotted the kernel estimates of the distributions of the first job assignments of the men and women in our sample.

The asymmetry in Figure 2 is striking.

Women seem to have started their careers in clearly less complex jobs than men. The Duncan and Duncan index of dissimilarity for the initial job assignments gives a value of 49.3. That is, the initial job assignment would have had to be changed for nearly half of the women in our sample in order for women's job allocation to match men's.

However, Figure 2 should not be interpreted as evidence of gender differences in assignment thresholds. After all, it is possible that the asymmetry in Figure 2 only reflects differences in the ability distributions of men and women who chose to work in the metal industry. If the underlying productivity of women in this industry was lower than that of men, it is not surprising that women started their careers in less complex jobs than men. Still, whatever the reason for the asymmetry in Figure 2, it is clear that the initial job assignment has to be taken into account in the analysis of gender differences in promotion and productivity.

5.2. Promotions

Were women, then, less likely than men to move upward on the complexity ladder?



Figure 2. Kernel Estimates of the Male and Female Distributions of Complexity at the Initial Job Assignment.

We follow two approaches to answer this question. First we take advantage of the fact that the occupation-related wage is a continuous variable. Thus, changes in occupation-related wages conveniently summarize both the extent and direction of change in the complexity of the jobs that the worker was performing.

In Table 2, we present the results from the following regression:

$$(3) \quad \Delta c_{it} = \beta F_i + X_{it}'\gamma + \delta \hat{c}_i + \zeta J_i + \varepsilon_{it}$$

where $\Delta c_{it} = c_{i,t+1} - c_{it}$ is the difference in log occupation-related wages, F_i is the female dummy, and X_{it} is a set of productive characteristics, including age and seniority and their squares, firm size, and the personal bonus deviation of the workers at t . The complexity of the worker's initial job is denoted by \hat{c}_i and measured with the log of his or her first occupation-related wage, and J_i is a dummy for the firm in which worker i is employed.

In the first column of Table 2, we report the simple gender difference in the mean changes in complexity. Both men and women moved, on average, .011 log points up the complexity ladder. There is no

statistically significant gender difference in the mean change of complexity. In the second column, we add to the regression productive characteristics such as age and seniority, along with our measure of individual productivity, personal bonus deviation. The gender difference is still statistically insignificant.

Even though the results in the first two columns of Table 2 might suggest that men and women moved up the complexity axis in the same way, recall that Figure 2 indicated very different first task assignment allocations for men and women. Women tend to be crowded in the low complexity tasks, where promotions are likely to be more frequent. When we add a control for the complexity of the initial task assignment to the regression in the third column, the coefficient of the female dummy turns negative, -0.007 , and clearly statistically significant. Finally, in column (4) we add a full set of firm dummies (404) to the regression to account for the possibility that women were crowded into firms that simply did not promote their workers as often. The coefficient on the female dummy is not affected by this modification. Thus, the

Table 2. Changes in Complexity: Regression Results.

Variable	(1)	(2)	(3)	(4)
Female	.001 (.000)	.000 (.000)	-.007*** (.000)	-.007*** (.000)
Age/10	—	-.022*** (.001)	-.012*** (.001)	-.011*** (.001)
(Age/10) ²	—	.002*** (.000)	.001*** (.000)	.001*** (.000)
Seniority/10	—	-.092*** (.003)	-.097*** (.003)	-.100*** (.003)
(Seniority/10) ²	—	.070*** (.003)	.074*** (.003)	.076*** (.003)
Personal Bonus Deviation	—	.001*** (.000)	.001*** (.000)	.001*** (.000)
Firm Size/1,000	—	-.002 (.001)	.001 (.001)	-.006 (.008)
Initial Task Complexity	—	—	-.002*** (.000)	-.002*** (.000)
Firm Dummies	—	—	—	404
Constant	.011*** (.000)	.070*** (.002)	.126*** (.003)	.135*** (.003)
R ²	.000	.046	.077	.104
N	63,023	63,023	63,023	63,023

The dependent variable is the difference between the log of the occupation-related wage in the next period and the log of the occupation-related wage in the current period. Seniority is measured as the number of years the individual has been present in the metal industry. Personal bonus deviation refers to the deviation of the individual's personal bonus from the yearly mean in the firm-complexity cell. Initial task complexity is the log of the occupation-related wage at the initial job assignment. Numbers in parentheses are robust standard errors.

***Statistically significant at the 1% level.

results in Table 2 clearly indicate that women took smaller steps on the complexity ladder than men who started in jobs of similar complexity.

Another way of looking at gender differences in the promotion patterns is to examine duration to promotion. Because a worker's productivity is likely to be affected by learning on the job, the higher female productivity threshold of promotion would imply a longer promotion wait for women than for men of similar ability.

We defined a promotion indicator that takes a value of one if the individual experienced a positive change in the occupation-related wage within the same firm and zero otherwise. In Table 3, we report the Kaplan-Meier estimates of the survivor function for men and women. The first column shows the estimates for the whole sample.

Although the female survivor function is consistently above the male one, the differences are relatively small.

However, the gender differences become clear once we focus on workers who started their careers in jobs of more or less similar complexity. This is done in columns (2)–(4) of Table 3, where we report the estimates of the survivor function within subsamples of workers. The groups were formed by dividing the workers into “low-complexity,” “medium-complexity,” and “high-complexity” groups according to the complexity of the initial job assignment. Female survivor functions are clearly above the male ones in the low and medium groups of initial job complexity. A straightforward log-rank test for the equality of the Kaplan-Meier estimates, results of which are reported in the last row of Table 3, rejects the

Table 3. Kaplan-Meier Estimates of the Survivor Function by Gender.

Seniority	(1) The Whole Sample		(2) Low Complexity Tasks		(3) Medium Complexity Tasks		(4) High Complexity Tasks	
	Men	Women	Men	Women	Men	Women	Men	Women
1	0.710	0.738	0.596	0.699	0.680	0.804	0.860	0.943
2	0.566	0.617	0.411	0.570	0.536	0.704	0.771	0.862
3	0.474	0.518	0.328	0.469	0.440	0.612	0.697	0.813
4	0.413	0.454	0.275	0.403	0.384	0.551	0.626	0.766
5	0.372	0.410	0.238	0.361	0.344	0.505	0.584	0.723
Log-Rank Test	X ² (1) = 10.69 (p = 0.001)		X ² (1) = 84.19 (p = 0.000)		X ² (1) = 102.30 (p = 0.000)		X ² (1) = 2.53 (p = 0.110)	

Table 3 reports the Kaplan-Meier estimates of the survivor function, where the hazard is the probability of promotion conditional on not being promoted up to the year of seniority. Low complexity refers to tasks that have occupation-related wages less than or equal to 30. Medium complexity refers to tasks that have occupation-related wages larger than 30 but less than or equal to 35. High complexity refers to tasks that have occupation-related wages higher than 35. The last row reports the log-rank test statistic for the hypothesis that the male and female survivor functions are equal.

null in all cases, with the exception of high-complexity jobs.

In order to control for observable productive characteristics in the duration analysis, we ran a discrete-time proportional hazards model of promotion, using the log-likelihood suggested by Jenkins (1995),

$$(4) \quad \log L = \sum_{i=1}^N \sum_{t=1}^T y_{it} [h_{it} / (1 - h_{it})] + \sum_{i=1}^N \sum_{t=1}^T \log(1 - h_{it}),$$

where N is the number of individuals in our sample, T is the number of years, and y_{it} is a dummy that takes the value 1 for individual i at t if he or she is promoted at that period and zero otherwise. We use a complementary loglog specification for the hazard rate,

$$(5) \quad h_{it} = 1 - \exp\{-\exp[\theta(t) + \beta F_i + X'_{it}\gamma + \delta \hat{c}_i]\},$$

where the baseline hazard $\theta(t)$ is left unspecified and the rest of the variables are as in (3) with the exception that the firm dummies are excluded. The results are presented in Table 4. The estimated coefficient of the female dummy is negative (-0.3) and clearly statistically significant.

We interpret the results in Tables 2–4 as implying that women were initially assigned to jobs that were less complex than those assigned to men but in which promotions were also more frequent. However, compared to the male workers who started their careers on those same jobs, women were much less likely to move to more complex tasks. Thus, women tended to “get stuck” in the initial jobs of low complexity, whereas men were more successful in advancing to more complex jobs.

6. Gender Differences in Productivity

By controlling for productive characteristics, the regressions of the previous section clearly support the hypothesis of a higher promotion threshold for women. However, we can complement our analysis with another approach that is based on comparisons of male and female productivity, before and after eventual promotion. As was explained in Section 2, if the threshold for promotion is higher for women than for men, both promoted and non-promoted women should, on average, be more productive than their male counterparts, after selection for promotion has taken place.

In our view, such an analysis usefully

Table 4. Duration to Promotion:
Discrete-Time Complementary Log-Log
Proportional Hazards Estimates.

<i>Variable</i>	<i>Estimate</i>
Female	-.346*** (.023)
Age/10	-.237** (.086)
(Age/10) ²	.003 (.013)
Personal Bonus Deviation	.052*** (.003)
Firm Size/100	.432*** (.076)
Initial Task Complexity	-.097*** (.003)
Seniority Dummies	10
Log-Likelihood	-27,162
N	63,023

Estimates are from a discrete-time proportional-hazards model of promotion. Complementary log-log specification was chosen for the hazard of promotion. Standard errors are reported in parentheses.

***Statistically significant at the 1% level.

complements the regressions. The regressions only use the “base year” information on individual characteristics, plus our knowledge of who will be promoted. The productivity comparisons, by contrast, also use information on productivity after the promotion (and non-promotion). More specifically, suppose that the employers can observe some other characteristics that are related to the person’s potential productivity in the more complex job. We cannot observe these characteristics in the base year, since our productivity measure only applies to the person’s current job assignment. But it could still be the case, for example, that men perform better than women or equally well as women in the new, more complex jobs. If that were the case, we could argue that, regression results notwithstanding, employers can see the “potential productivity” of men and there need be no asymmetry in promotion thresholds. If, on the other hand, we observe that women perform better than men

both among the promoted and among the non-promoted workers, we can be much more confident that an asymmetric promotion threshold is indeed operative.

We now use personal bonus deviations to measure individual productivity and examine the productivity of the workers in our newcomer sample. We first measure the productivity of each worker at the initial job assignment and then partition the workers into two groups: those who were promoted up to some specific year (the “promoted” group), and those who stayed at the initial assignment until that year (the “stagnant” group).

In Table 5, we report the mean personal bonus deviations for workers in each of these groups. In the first row, we display the personal bonus deviations of the new entrants during their first year in the industry. The next row depicts the personal bonus deviations of men and women as measured in the second year of their career, separately for those who were promoted after the initial year and those who were not. The following row depicts the mean personal bonus deviations two years after the initial assignment, similarly differentiated between those who had been promoted up to their third year of their career and those who had not; and analogously for the fourth and the fifth row. All of these comparisons are computed on the same set of individuals.

The means reported in Table 5 indicate that there were hardly any gender differences in productivity at the initial task assignments. Although the female mean is slightly higher than the male mean, the difference is not statistically significant. However, once the workers are split into promoted and non-promoted groups, the gender differences become clear. Women tended to dominate in both groups.

Table 5 only reports the differences in means. For a better view of what happens to the productivity distributions, we have plotted the kernel estimators of the cumulative distributions of personal bonus deviations for men and women separately for stagnant and promoted workers in the first and second years of their careers in Figure 3. The

GENDER DIFFERENCES IN PROMOTION ON A JOB LADDER

297

Table 5. Mean Personal Bonus Deviations in the Groups of Promoted and Not-Promoted Workers by Tenure within Firm.

Tenure	Stagnant Men	Stagnant Women	Difference	Promoted Men	Promoted Women	Difference
First Assignment	-2.823 (0.100) N = 8,988	-2.460 (0.267) N = 2,484	0.364 (0.256)			
After 1 Year	-1.411 (0.135) N = 6,383	-0.353 (0.361) N = 1,833	1.058*** (0.346)	-1.616 (0.144) N = 2,605	-0.861 (0.219) N = 651	0.755*** (0.207)
After 2 Years	-0.896 (0.133) N = 5,088	0.302 (0.245) N = 1,510	1.197*** (0.240)	-1.216 (0.151) N = 3,364	-0.016 (0.217) N = 891	1.200*** (0.197)
After 3 Years	-0.459 (0.116) N = 4,254	0.603 (0.199) N = 1,281	1.062*** (0.198)	-0.882 (0.131) N = 3,672	0.354 (0.173) N = 968	1.236*** (0.165)
After 4 Years	-0.243 (0.108) N = 3,710	0.776 (0.171) N = 1,113	1.019*** (0.174)	-0.611 (0.093) N = 3,757	0.586 (0.142) N = 1,027	1.197*** (0.137)

Numbers in cells are means of bonuses as deviations from firm- and task-specific means. "Stagnant" workers are those remaining in their first assignment job. Promoted workers are those moving to more complex jobs. Mean bonus deviations are reported for each group. "Difference" refers to the estimated coefficient of the female dummy in a regression where the personal bonus deviations of the corresponding group are regressed on the female dummy. Numbers in parentheses are robust standard errors. N refers to the sample size of the cell.

***Statistically significant at the 1% level.

male and female distributions of personal bonus deviations are almost indistinguishable at the initial job assignment. After the first year, the workers are split into stagnant and promoted groups. In both groups, the female distributions clearly dominate the male distributions. These gender differences in productivity become even clearer with tenure.

Whether the results presented above, together with the results on gender differences in promotion probabilities, imply that the female assignment thresholds were higher than the male ones depends on the underlying productivity distributions of men and women who chose to work in the metal industry. This implication only follows if the underlying productivity distributions of men and women did not intersect. As we argued in Section 2.3, it seems plausible to assume that this is the case here. First, we focus on men and women who started their careers within jobs of similar complexity. The employers had already

selected the workers they judged to be suited for given tasks. Any differences in the underlying distributions of productivity of men and women within these groups should be small. Furthermore, it seems very unlikely that distributions would intersect within all these groups. Finally, the results in Table 5 and in Figure 3 indicate that there are no visible gender differences in the personal bonus deviations within initial jobs. Hence, in our view, it is justified to interpret these results as indicating that women had to be more productive than men to be assigned to complex jobs.

Personal bonus deviation comparisons also tell us something about the behavior of the employers. First, it is clear that the workers whom the employers chose to promote tended to be more productive already in their initial jobs. In the framework of Lazear and Rosen's model, this implies that training costs did play a role in these firms; if the assignment to more complex jobs had been costless, the employer would have

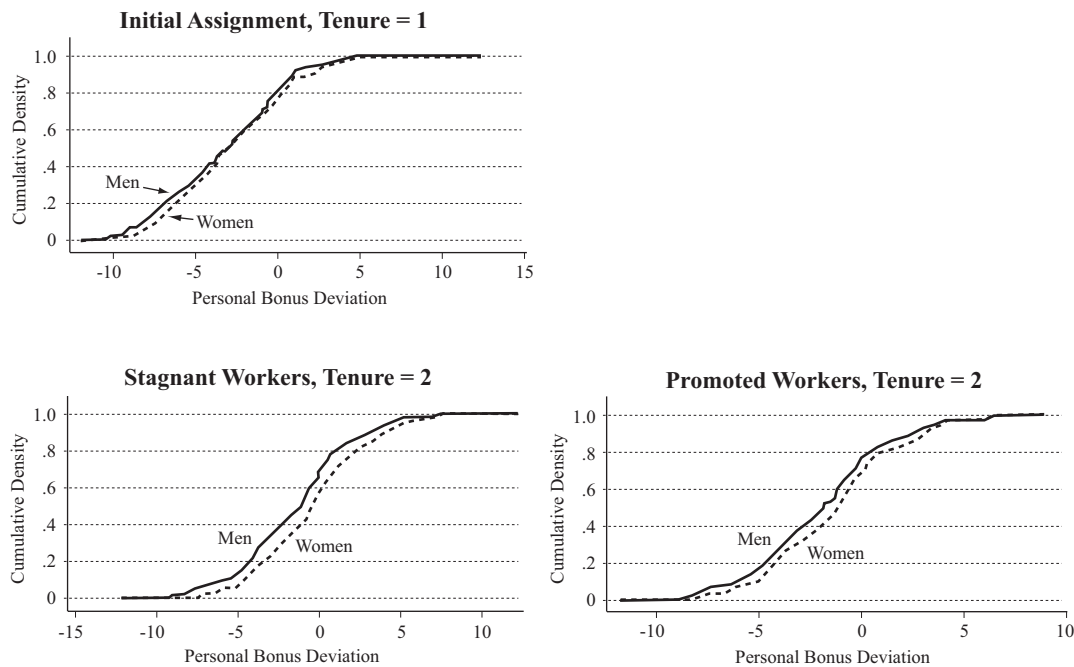


Figure 3. Kernel Density Estimates of the Cumulative Distributions of Personal Bonus Deviations by Tenure and Promotion Status.

assigned all the workers to more complex tasks. However, it seems that women, to get promoted, needed to be more productive at their initial job than men did. Among the women who ended up being promoted during the first five years of their careers, the average personal bonus deviation in the initial job was -2.18 , while for promoted men it was only -2.48 . Thus, promoted women were approximately 12.5% more productive at the initial job than promoted men. Since the Lazear and Rosen model does not assume any differences in the training costs of men and women, the only way to explain these differences in terms of their model is by positing gender differences in quit behavior.

7. Gender Differences in Quit Behavior

According to the Lazear and Rosen model, gender differences in reservation

wages give rise to differences in assignment thresholds. We now examine whether the observed quit behavior of men and women in the Finnish metal industry is consistent with the Lazear and Rosen argument. However, we emphasize that, in our view, these data do not really allow us to test the Lazear and Rosen argument against all alternative models, such as taste-based discrimination in promotions. If there are other forces that hamper women's promotion prospects, women's higher exit rates will also reflect these other mechanisms. That is, any gender differences in quit behavior that we observe may have been caused by gender differences in the assignment thresholds, which themselves could have been a result of taste-based discrimination. Hence, this exercise should be seen as a simple check of whether gender differences in quit behavior could be the underlying cause for the differences in the assignment thresholds in

the most favorable case in which the quit rates are exogenous.

We do not observe the actual reasons for job separations in our data. Therefore it is not possible to distinguish layoffs from voluntary quits directly. However, since voluntary quits are what we are interested in, we tried to tackle this problem by defining as “quits” only those separations that occurred when more than 75% of employees remained in the firm. This definition rules out situations in which a large proportion of the firm’s work force was laid off at the same time, and it is likely to catch more truly voluntary separations. We do not distinguish between job-to-job turnover and job-to-nonemployment turnover, since from the point of the view of the employer these are equivalent.

In the whole metalworker population, the average quit rate was 7.8% for men and 9.4% for women. Gender differences were largest among the younger workers, with young women up to eight percentage points more likely to quit than young men. However, the quit rates of men and women converged among older workers, so that men and women over 45 years old were more or less equally likely to quit.

Table 6 reports the marginal effects from a probit model of quits that uses the whole population of metalworkers. In this model, we allow for different intercepts and age profiles for men and women. The effects are calculated at the mean values of the continuous independent variables. According to these results, women were eight percentage points more likely to leave the firm than were men.

However, from the point of view of the employer who is considering whether to promote one of the workers in our newcomer sample, the most relevant differences are the differences in the quit rates of the younger workers. The predicted quit rates, using the coefficients reported in Table 6, of men and women who were younger than 40 were 8.8% and 12.5%, respectively. The predicted female quit rate was thus approximately 42% higher than the male quit rate among the younger workers. For older workers differences were

Table 6. Probit Model of Voluntary Quits, Marginal Effects.

<i>Variable</i>	<i>Marginal Effect</i>	<i>Standard Error</i>
Female Dummy	.084***	(.013)
Age/10	-.175***	(.002)
Age Squared /100	.024***	(.000)
(Female × Age)/10	-.025***	(.005)
(Female × Age Squared)/100	.002***	(.001)
Seniority/10	-.084***	(.000)
Seniority Squared /100	.018***	(.000)
Firm Size/1,000	.004***	(.000)
14 Complexity Dummies	Yes	
N	617,404	
Log-Likelihood	-160,527	
Pseudo R ²	.076	

Marginal effects of the probit model of quit behavior using the whole population of metalworkers 1990–2000. For the female dummy, the marginal effect is the effect of a 0 to 1 change. Quit is defined as a separation occurring when at least 75% of the employees remained in the firm.

***Statistically significant at the 1% level.

much smaller: the predicted quit rate of women over 40 was 6.7%, compared to 6.5% for men. In our newcomer sample, 75% of the observations come from workers who were younger than 40 years old.

Thus, we find a pattern of exit rates that is certainly consistent with the Lazear and Rosen story but, of course, is consistent with other models as well. Women were more likely than men to quit the firm precisely in the group where one would expect the Lazear and Rosen argument about gender differences in the reservation wages to apply, that is, young workers.

We also examined whether the results on gender differences in the probability of promotion or in the productivity of promoted and non-promoted workers were any different for older workers (over 40 years old) and younger workers (under 40) in our sample of newcomers to the metal industry. Bearing in mind the caveat that there are, of course, very few newcomers who are over 40—only 1,755 individuals in

our case—the results nevertheless seem to go in the right direction. Controlling for productivity and the initial task assignment, the estimated gender difference in the change in complexity for younger workers was $-.008$, whereas for older workers it was $-.004$. Both of these coefficients are very precisely estimated and their 95% confidence intervals do not cross. Furthermore, when the stochastic dominance analysis of Figure 3 is replicated for the population of workers over 40 years old, one does not find a clear-cut pattern of gender productivity differences among promoted and non-promoted workers similar to that in Figure 3. Thus, we conclude that these subsample results are in line with the Lazear and Rosen argument, although the fact that the gender differences do not completely disappear with age indicates that some other factors play a role as well.

8. Conclusions

The asymmetric allocation of men and women across jobs is one of the most common explanations for the persistence of gender wage gaps. It is often argued that women have to meet higher productivity requirements than men to be assigned to demanding jobs. Because of these differences in assignment thresholds, women tend to start their careers in less demanding jobs than men and find it more difficult to get promoted than men do.

In this paper, we have examined gender differences in employers' assignment of workers to jobs of different complexity using panel data on Finnish metalworkers. Among these workers, women started their careers in less complex jobs than men, on average. Furthermore, they found it more difficult to move on to complex tasks than men did. Our regressions show that women were clearly less likely to get promoted than men who started their careers in the same jobs, even after the analysis controls for observable productive characteristics.

Using personal bonuses as a measure of individual productivity, we also find that

there were no apparent productivity differences between men and women within the jobs to which the workers were initially assigned. However, when the workers are split into groups of promoted and non-promoted workers, gender differences in productivity become manifest. Promoted and non-promoted women were consistently more productive than their male counterparts.

Since we do not observe clear productivity differences within the initial jobs, we can interpret the combination of the results that women were less likely to be promoted and that women were more productive among the promoted and non-promoted workers as evidence of higher promotion thresholds for women than for men. Although the finding of a lower probability of promotion for women than for men has often been reported, until now it has not been combined with results on gender differences in productivity that fit the asymmetric assignment threshold hypothesis.

Finally, the observed quit behavior of men and women among these workers is broadly consistent with Lazear and Rosen's explanation of gender differences in promotion thresholds: namely, the fact that women are more likely than men to quit the firm may imply that women face higher thresholds because their promotion is more costly to the employer. However, since the quit behavior could conceivably result from the differences in promotion thresholds—a possibility we cannot rule out, given the limitations of the data—we cannot definitively identify the causes of these threshold differences.

To sum up, although the results of our empirical investigations may be consistent with more than one model, they unquestionably accord with the predictions of the celebrated Lazear and Rosen model. Empirical tests of that model have so far been scarce, since it is generally not possible to measure productivity independently of the wage outcome. We hope our results have, to some extent at least, helped fill that gap.

REFERENCES

- Baker, George P., Michael Gibbs, and Bengt Holmström. 1994. "The Internal Economics of the Firm: Evidence from Personnel Data." *Quarterly Journal of Economics*, Vol. 109, No. 4 (November), pp. 881-919.
- Bayard, Kimberly, Judith Hellerstein, David Neumark, and Kenneth Troske. 2003. "New Evidence on Sex Segregation and Sex Difference in Wages from Matched Employer-Employee Data." *Journal of Labor Economics*, Vol. 21, No. 4 (October), pp. 887-923.
- Blau, Francine D., and Marianne A. Ferber. 1987. "Discrimination: Empirical Evidence from the United States." *American Economic Review*, Vol. 77, No. 2 (May), pp. 316-20.
- 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.
- Gibbons, Robert, and Michael Waldman. 1999. "A Theory of Wage and Promotion Dynamics inside Firms." *Quarterly Journal of Economics*, Vol. 114, No. 4 (December), pp. 1321-58.
- Ginther, Donna K., and Kathy J. Hayes. 1999. "Gender Differences in Salary and Promotion in the Humanities." *American Economic Review*, Papers and Proceedings, Vol. 89, No. 2 (May), pp. 397-402.
- _____. 2003. "Gender Differences in Salary and Promotion for Faculty in the Humanities, 1977-1995." *Journal of Human Resources*, Vol. 38, No. 1 (Winter), pp. 34-73.
- Granqvist, Lena, and Helena Persson. 2002. "Gender Differences in Career Mobility—An Example from the Wholesale and Retail Trade." Mimeo, Swedish Institute for Social Research, Stockholm University.
- Groschen, Erica L. 1991. "The Structure of Female/Male Wage Differential: Is It Who You Are, What You Do, or Where You Work?" *Journal of Human Resources*, Vol. 20, No. 3 (Summer), pp. 457-73.
- Hersch, Joni, and W. Kip Viscusi. 1996. "Gender Differences in Promotions and Wages." *Industrial Relations*, Vol. 35, No. 4 (October), pp. 461-72.
- Jenkins, Stephen P. 1995. "Easy Estimation Methods for Discrete-Time Duration Models." *Oxford Bulletin of Economics and Statistics*, Vol. 57, No. 1, pp. 129-38.
- Jones, David R., and Gerald H. Makepeace. 1996. "Equal Worth, Equal Opportunities: Pay and Promotion in an Internal Labour Market." *Economic Journal*, Vol. 106, No. 2 (March), pp. 401-9.
- Lazear, Edward P. 1986. "Salaries and Piece Rates." *Journal of Business*, Vol. 59, No. 3 (July), pp. 405-31.
- Lazear, Edward P., and Sherwin Rosen. 1990. "Male-Female Wage Differentials in Job Ladders." *Journal of Labor Economics*, Vol. 8, No. 1, Part 2 (January), pp. 106-23.
- McCue, Kristin. 1996. "Promotions and Wage Growth." *Journal of Labor Economics*, Vol. 14, No. 2, pp. 175-209.
- McDowell, John M., Larry D. Singell, and James P. Ziliak. 1999. "Cracks in the Glass Ceiling: Gender and Promotion in the Economics Profession." *American Economic Review*, Papers and Proceedings, Vol. 89, No. 2 (May), pp. 392-96.
- Metalliteollisuuden Keskusliiton Ja Metallityöväen Liiton Välinen Työehtosopimus 1.2.2003-15.2.2005. Downloadable at <http://www.metalliliitto.fi/pdf/metallites2003.pdf> (in Finnish and Swedish).
- Treble, John, Edwin Van Gameren, Sarah Bridges, and Tim Barmby. 2001. "The Internal Economics of the Firm: Further Evidence from Personnel Data." *Labour Economics*, Vol. 8, No. 5, pp. 531-52.
- Vartiainen, Juhana. 2002. "Gender Wage Differentials in the Finnish Labour Market." Publications of the Finnish Gender Equality Unit, Ministry of Social Affairs and Health. Downloadable at <http://www.stm.fi/resource.phx/tasa-arvo/julkaisut/julkaisut.htx>.
- Winter-Ebmer, Rudolf, and Josef Zweimüller. 1997. "Unequal Assignment and Unequal Promotion in Job Ladders." *Journal of Labor Economics*, Vol. 15, No. 1 (January), pp. 43-71.