DANILO COELHO
MARCELO FERNANDES
MIGUEL N. FOGUEL

Foreign Capital and Gender Differences in Promotions: Evidence from Large Brazilian Manufacturing Firms

Widespread consensus asserts that occupational segmentation is the chief determinant for gender differences in wages in the Brazilian labor market (see, for example, Oliveira 2001). Although women represent 39 percent of the formal workforce in Brazil, they amount to only 23 percent of the employees in the sectors that pay relatively higher wages in 2004. As well, the predominance of segmentation increases if one restricts attention to managerial positions, of which women hold only 14 percent. This suggests that limited promotion prospects for women in the Brazilian job market may present barriers to their promotion to the upper rungs of the corporate ladder.

This paper contributes to the literature by examining data from the largest firms of the Brazilian manufacturing industry. In particular, our goal is to determine whether obstacles to women's ascension exist by tracking the amount of time it takes for a woman to get a promotion to a managerial position. The aim is to complement the literature on gender differences in promotions, whose papers mainly focus on computing the promotion rates. Time to promotion offers a different angle in that samples with a large time span could well present similar likelihoods of promotion for both genders even if

We are grateful to Alexandre Carvalho, Kiko Corseuil, Gabriela Bertol Domingues, Giulio Fella, Claudio Ferraz, Winfried Koeniger, Marco Manacorda, Pedro Martins, Aureo de Paula, and Fabio Veras for helpful discussions as well as to seminar participants at Fundação Getulio Vargas–São Paulo, Pontifícia Universidade Católica do Rio de Janeiro, the 2007 meeting of the Associação Nacional de Pós-graduação em Economia, the 2008 meeting of the Latin American and Caribbean Economic Association, and the twenty-sixth Economía Panel (Lima, Peru, 2012). We also thank Hugo Ñopo, the editor of this volume, for exceptionally constructive remarks. The usual disclaimer applies.

the promotion durations are different. We construct the duration variable by following from January 1996 to December 2005 every individual who enters the Annual Report of Social Information (RAIS) database between January 1991 and December 1995. We denote as nonpromoted the individuals who do not obtain a promotion to a managerial position within the sample (that is, the right-censored observations), whereas we assign duration zero to the workers who already are in managerial positions in January 1996. Accordingly, we label as promoted the individuals within our sample who rise to a managerial position (that is, those with uncensored positive durations). Investigating differences in promotions is particularly difficult because observationally similar male and female workers may display equal promotion rates and durations even in the presence of gender differences (for example, promotions may differ in quality).

As far as we know, this is the first study to examine gender differences in promotions using microdata from a developing country. Apart from the obvious interest in finding whether the main stylized facts hold in a Latin American country, we were also motivated by the quality and availability of the Brazilian data. In contrast to the many studies that employ data from individual firms (see, for example, Cabral, Ferber, and Green 1981; Gerhart and Milkovich 1989; Baker, Gibbs, and Holmstrom 1994), our study relies on a homogeneous sample of recently hired workers from the Brazilian manufacturing industry in the period from 1996 to 2005. Our data set is particularly convenient. First, the data include a wide array of controls for worker and firm characteristics. Second, unlike Blau and deVaro (2007), we observe multiple workers per establishment at different occupations and hierarchical levels as well as their career paths in terms of occupation and wage. A potential limitation is that our data set does not include any direct measure of on-the-job productivity and hence we must come up with indirect controls.

Interestingly, we find that there are significantly fewer gender differences in the time to promotion within foreign-owned firms than within domestic firms in the Brazilian manufacturing industry. Our findings complement to some extent the literature on the differences between multinationals and domestic firms (Doms and Jensen 2006; Greene and others 2009) as well as the evidence that gender differences may depend on the nature of the firm.[1]

1. In this paper, we employ *multinational* and *foreign-owned firm* interchangeably even if, in recent years, we have been witnessing the rise of many Brazilian multinationals, such as Vale (mining and metals), Petrobras (oil and gas), Gerdau (steel), and Embraer (aviation). See Amann (2009) for a historical perspective and a number of case studies.

We argue that such gender differences are consistent with statistical discrimination and self-selection. We reason as follows. Suppose there are relatively more women than men who prefer to dedicate more time to their family than to their careers. Berk (2001) shows that gender differences in average career concern do not ensure statistical discrimination against the group with lower average career concern if employees optimally choose which jobs to apply for. A career-minded woman who applies for a job in a firm with a preference for dedicated employees does so only because she rationally believes that her odds of being hired compensate the effort to go through the application process. This means that her qualifications for the job have to be sufficiently good to stand a chance despite the discrimination.

To complete the argument, we rely on two anecdotal observations concerning the impact of multinationals in the Brazilian labor market (OECD 2008). The first is that domestic firms in Brazil offer a more flexible package in terms of working hours and business trips than multinationals. This makes them more appealing to less career-minded individuals (regardless of gender). The second is that multinationals compete more fiercely for highly skilled workers. Under these circumstances, career-minded women prefer jobs in multinationals and so statistical discrimination will become more prominent within domestic firms.

We find some indirect evidence supporting this explanation. On one hand, male workers tend to officially work similar hours in foreign-owned and domestic firms whether or not they have been promoted. As figure 1 illustrates, the main difference lies in the concentration of the distribution: at least 75 percent of the male employees of domestic firms work exactly forty-four hours a week, whereas there is a bit more variation in multinationals, where 75 percent of the male employees work from forty to forty-four hours a week. It also shows some minor differences in the lower support of the distribution. In particular, the minimum number of hours worked is slightly higher for promoted male workers than for nonpromoted male workers. The same applies to female workers. This is consistent with the fact that promoted workers are more likely to be career minded. On the other hand, nonpromoted female employees work relatively much less in domestic firms (ten to forty-four hours) but not in multinationals. This is consistent with domestic firms' tendency to offer more flexible packages. In addition, women in multinationals work much longer hours than women in domestic firms and than male workers in general. This is well in line with self-selection. As the latter alleviates the impact of statistical discrimination, we fail to observe as much gender difference in foreign-owned companies as in domestic firms.

FIGURE 1. Promotion, by Gender, Type of Firm, and Working Hours per Week

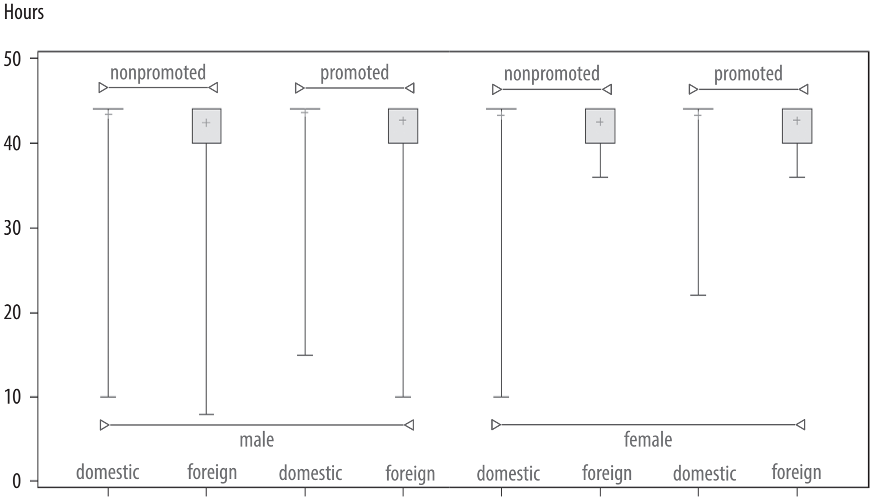

This paper relies primarily on the literature on gender differences in career mobility. The theoretical literature on career mobility mainly focuses on schooling (Sicherman and Galor 1990), abstracting away from gender differences. There are a few exceptions, though. Booth, Francesconi, and Frank (2003) derive a model that distinguishes between the initial pay increase upon promotion and subsequent pay increases. Under the assumption that women have worse market alternatives, the model implications are consistent with their empirical findings that gender does not affect promotion rates in the United Kingdom, though women receive lower wage gains. Baldwin, Butler, and Johnson (2001) identify the effects of occupational segregation in the United States on gender wage gaps using a hierarchical discrimination model in which men dislike supervision by female managers. They predict an exponential decline in the relative proportion of female workers in the top tiers of the job ladder, which is supported by the evidence from a 1988 Current Population Survey sample of workers in the insurance industry.

In contrast, it is common practice in the empirical literature to include gender among the determinants of job mobility and promotion likelihood (see, among others, Groot and van den Brink 1996; Booth and Francesconi 2000; Blau and deVaro 2007) and to note that men and women may differ both in alternative opportunities (Mincer and Ofek 1982; Lazear and Rosen

1990; Royalty 1998) and in job search costs (Meitzen 1986). For instance, women may face more constraints to work longer hours or even to remain in the labor market. If women are more likely to quit, firms will have fewer incentives to train and promote them. Moreover, if women view promotion as unlikely owing to discriminatory promotion practices, they may refrain from putting themselves forward for training programs at the firm (Arrow 1972).[2] The papers closest to ours are McCue (1996) and Pekkarinen and Vartiainen (2006). Using data from the Panel Study of Income Dynamics from 1976 to 1988, McCue (1996) demonstrates that it takes, on average, more time for women and black men to get promotions than for white men. Pekkarinen and Vartiainen (2006) analyze gender differences in time to promotion for workers within the metallurgical industry in Finland. They show that women usually take more time to get a promotion than men with similar jobs, even if they are consistently more productive than men. Our findings are interestingly different in that the Brazilian manufacturing industry appears to feature a greater gender differential for domestic relative to foreign-owned firms.

Data Description

The data set we employ gathers information from several databases. In particular, it combines data from the RAIS, covering the period running from 1991 to 2005, as well as data from the 1996 Foreign Trade Census of the Foreign Trade Secretary and from the 2000 Census of Foreign Capital.

The Annual Report of Social Information is the administrative registry of the Ministry of Labor that provides socioeconomic information regarding the employees of every firm in the Brazilian formal sector. It reports the employees' identifying security number, age, gender, schooling level, job tenure, average monthly salary (including performance bonus and commissions), occupation (as reported by the employer), number of hours at work, type of labor contract, and month of admission. In addition, it also documents the firm's identifying fiscal number, sector of activity, and location. We make use of the 1996 Foreign Trade Census data to gather information on how

2. See, among others, Cabral, Ferber, and Green (1981); Spurr (1990); Cannings and Montmarquette (1991); McCue (1996); Barnett, Baron, and Stuart (2000); Ransom and Oaxaca (2005); Blau and deVaro (2007); and Acosta (2010) for supporting evidence; and Lewis (1986); Powell and Butterfield (1994); Paulin and Mellor (1996); Petersen and Saporta (2004); and Giuliano, Levine, and Leonard (2011) for evidence against gender differences.

much each firm exports as a proxy for productivity. Unfortunately, it is hard to gather exports data for sectors other than the transformation industry. This is the main reason we focus on the latter.

The Central Bank of Brazil publishes the Census of Foreign Capital every five years, collecting information on the origin of the shareholders' capital for every firm in Brazil. We employ this census to classify firms as either domestic or multinational. We define as multinational a firm in which more than 50 percent of the shareholders' capital is foreign.[3] Matching data from the RAIS and the 2000 Census of Foreign Capital reveal that women account for 21 percent of the employees in multinationals of the Brazilian manufacturing industry and occupy 13 percent of their managerial positions. The figures are similar for domestic firms (that is, firms with less than 50 percent of foreign capital): 25 and 15 percent, respectively.

To form a homogeneous RAIS sample, we focus on individuals who meet the following criteria. First, the individual must work in a profit-seeking private firm with 500 or more employees in the Brazilian manufacturing industry. We focus exclusively on large firms because smaller firms have not enough internal turnover at the managerial level. Although these large firms account for only 5.29 percent of the firms in the Brazilian manufacturing industry, they employ 32.9 percent of the workers in the sector, of which a quarter are female. Second, we consider only individuals with a university degree. The proportion of workers who hold a university degree has increased from 9.58 percent in 1996 to 13.93 percent in 2005 within our sample of large Brazilian manufacturing firms. Third, the individual must have joined the firm between January 1991 and December 1995. Fourth, the individual must work as an accountant, administrator, director, economist, engineer, intermediate manager, lawyer, manager, or purchase or sales supervisor. Fifth, the individual must have a labor contract with no expiration date.

The resulting sample comprises 1,422 firms, of which 297 (20.9 percent) are multinationals, that altogether employ 23,737 male and 3,552 female workers. The average individual in our sample is about thirty-four years old and works around forty-three hours a week. As for occupations, engineers are the mode, representing 32.5 percent of our observations. This is in line

3. The threshold at 50 percent is arbitrary but also pretty harmless. The fraction of foreign capital concentrates either at zero (relative frequency of 65.83 percent) or at one (18.35 percent) and hence varying the cutoff point from 25 percent to 75 percent does not change any of the findings we report in the subsequent sections.

TABLE 1. Sample Size, by Gender and Occupation[a]

	Female		Male		Total	
Occupation	*Absolute*	*Relative (%)*	*Absolute*	*Relative (%)*	*Absolute*	*Relative (%)*
Engineer	689	19.4	8,171	34.4	8,86	32.5
AAEL group[b]	864	24.3	2,637	11.1	3,501	12.8
Intermediate manager	773	21.8	3,334	14.0	4,107	15.1
Supervisor	562	15.8	2,721	11.5	3,283	12.0
Manager and director	664	18.7	6,874	29.0	7,538	27.6
Total	3,552	100	23,737	100	27,289	100

Source: Annual Report of Social Information (RAIS), 1991 to 2005.

a. The label *absolute* refers to the number of observations in that cell, whereas *relative* corresponds to the relative sample size as a percentage of the total number of observations in that column.

b. Accountants, administrators, economists, and lawyers.

with the predominance of male workers, given that only 19.4 percent of the women in our sample are engineers. Whereas 29.0 percent of men in our sample occupy top management positions, only 18.7 percent of women do so. Table 1 stratifies the sample according to occupation and gender, and table 2 reports the sample averages of the individuals' main characteristics according to censoring and gender.

Note that the promotion duration variable in table 2 considers only promotions within the same firm (that is, intrafirm promotions). We thus treat interfirm promotions as right censoring. We prefer not to examine promotions through firm switches because of selection issues. The sample we have allows us to track the career of the worker only within large firms in the transformation industry. This would most likely bias any duration analysis if male workers tended to hold more sector-specific positions than female workers. For instance, engineering-related jobs involve much more sector-specific skills than human resources jobs. In addition, interfirm analyses would entail biased results in the event that male and female workers had different reactions to promotion offers from smaller firms. This would happen because our sample considers only large firms and hence we would incorrectly classify as right censored a worker who gets a promotion to a top management position in a smaller firm. To avoid such biases, we restrict attention to intrafirm promotions.

Table 3 examines gender differences in greater depth by looking at career progression according to whether the firm is domestic or multinational. The first panel shows that nonpromoted female workers receive an average monthly salary of about BRL (Brazilian real) 4,420 in domestic firms and

TABLE 2. Sample Statistic, by Censoring and Gender[a]

	Zero duration				*No censoring*				*Right censoring*			
	Female		*Male*		*Female*		*Male*		*Female*		*Male*	
Variable	*Mean*	*S.D.*	*Mean*	*S.D.*	*Mean*	*S.D.*	*Mean*	*S.D.*	*Mean*	*S.D.*	*Mean*	*S.D.*
Time to promotion	0	0	0	0	41.1	35.1	41.9	34.8	33.7	33.0	35.2	33.9
Engineer	0	0	0	0	0.22	0.41	0.44	0.50	0.24	0.43	0.49	0.50
AAEL group[b]	0	0	0	0	0.27	0.44	0.13	0.34	0.3	0.46	0.16	0.37
Intermediate manager	0	0	0	0	0.35	0.48	0.27	0.44	0.26	0.44	0.19	0.39
Supervisor	0	0	0	0	0.16	0.37	0.16	0.36	0.20	0.40	0.16	0.37
Manager and director	1	0	1	0	0	0	0	0	0	0	0	0
Age	34.4	7.2	39.6	7.5	30.0	5.8	33.2	7.6	30.2	6.1	33.1	7.7
Hours at work	43.0	2.5	43.3	1.9	42.9	2.5	43.1	1.8	43.0	2.5	42.9	2.5
Multinational	0.34	0.47	0.32	0.47	0.48	0.50	0.49	0.50	0.37	0.48	0.41	0.49
Size	2,314	2,983	2,419	3.37	3,033	3,081	3,046	3,604	3,555	6,177	3,568	6,125
Exports	0.038	0.120	0.043	0.131	0.076	0.184	0.073	0.163	0.059	0.151	0.085	0.194
University ratio	0.19	0.14	0.16	0.12	0.20	0.13	0.18	0.13	0.17	0.13	0.17	0.13
Turnover	0.56	0.24	0.56	0.24	0.52	0.21	0.52	0.25	0.53	0.24	0.52	0.24
Mean wage	2,107	1,108	2,017	961	2,238	973	2,280	1,003	2,180	1,055	2,291	1,030
Modern	0.37	0.48	0.37	0.48	0.37	0.49	0.41	0.49	0.38	0.49	0.49	0.50
Number of observations	664		6,874		161		1,465		2,727		15,398	

Source: Annual Report of Social Information (RAIS), 1991 to 2005; 1996 Foreign Trade Census of the Foreign Trade Secretary (SECEX); and 2000 Census of Foreign Capital (CEB).

a. We split the sample into three subsamples, which we label as zero duration, no censoring, and right censoring. The first subsample considers only individuals who start at managerial positions. The second subsample restricts attention to individuals who get a promotion within our sample, whereas the third subsample consists of nonpromoted individuals. Means and S.D. correspond to the sample averages and standard deviations, respectively. We gauge time to promotion in months, age in years, and hours at work in hours per week. All other individual-specific variables are binary, assuming value either one or zero according to the occupational group. As for the variables relating to the firm at which the individual works, size and exports denote the number of employees and how much the firm exports in USD billions, respectively; university ratio is the fraction of employees with a university degree; turnover gauges the job flow intensity of the firm; mean wage of the firm is the average monthly earnings across employees. Finally, modern is a binary variable that takes value one if the firm is within a technology; intensive sector.

b. Accountants, administrators, economists, and lawyers.

TABLE 3. Career Progression by Gender and Firm Ownership[a]

	Domestic		*Multinational*	
	Nonpromoted	*Promoted*	*Nonpromoted*	*Promoted*
Salary (BRL)				
Female	4,420	8,133	6,146	10,444
Male	5,993	9,755	7,494	11,428
Promotion likelihood				
Female	0.0463		0.0718	
Male	0.0768		0.1009	
Time to promotion (weeks)				
Female	38.17		39.57	
Male	44.29		44.38	

Source: Annual Report of Social Information (RAIS), 1991–2005; 1996 Foreign Trade Census of the Foreign Trade Secretary (SECEX); and 2000 Census of Foreign Capital (CEB).

a. The first panel reports average salary (in Brazilian reals) for promoted and nonpromoted workers by gender and by firm ownership—that is, whether the firm is domestic or multinational. The second and third panels reveal similar averages for the probability of getting a promotion and for their durations (in weeks), respectively. The latter is over uncensored observations, corresponding to the time it takes to get a promotion conditional on getting a promotion.

roughly BRL 6,146 in multinationals. These figures increase respectively to BRL 8,133 and BRL 10,444 after promotion. The salary ratios of the promoted to the nonpromoted are then 1.84 in domestic firms and 1.70 in multinationals (not shown). We find a similar pattern for male workers, though with lower promoted-to-nonpromoted salary ratios: namely, 63 percent salary increase after promotion in domestic firms as opposed to 52 percent in foreign firms. Although career progression has a stronger impact in domestic firms, the average salary in multinationals is much higher on average—about 39 percent and 25 percent higher for female and male nonpromoted workers, respectively, and 28 percent and 17 percent, respectively, for those in top management positions. It is also interesting to observe that gender differences in salary are much lower in multinationals. Male managers receive less than 10 percent more than their female counterparts in multinationals. This gap is more than twofold for domestic firms. The same pattern arises, though with much higher gender differences, for workers who are not in top management positions (36 percent in domestic firms against 22 percent in multinationals).

As to what influences the probability of getting a promotion, the second panel of table 3 documents two findings. First, multinationals promote more than domestic firms, probably because they are larger on average and hence have more managerial positions to offer than domestic firms. Second, gender differences are very large, with a much higher fraction of promoted men than women. In particular, less than 5 percent of the female workers in domestic firms receive a promotion as opposed to 7.68 percent of the men. These fractions increase in multinationals to 7.18 percent and 10.09 percent for female and male workers, respectively. Finally, the third panel looks at the average time it takes to get a promotion by gender and firm ownership. Interestingly, the time to promotion is less for women than for men, especially in domestic firms. Note that, as we restrict attention to uncensored observations, this average duration is conditional on getting a promotion within the sample period, and, as mentioned earlier, women are less likely to get promoted in domestic firms than in multinationals. To appreciate the overall effect, we must consider not only the probability of getting a promotion but also the censored observations in the right tail of the duration distribution. Similarly, even if a promotion is more likely in a multinational, it takes on average more time to obtain in a multinational than in a domestic firm. This is consistent with a fiercer competition in multinationals.

Altogether, women seem less likely to receive a promotion, but it takes less time for them to accomplish that. In the absence of affirmative action, the latter provides indirect evidence that promoted women are on average more productive than promoted men (despite their lower average salary). In contrast, although promotions in multinationals are more frequent, they take on average five to six weeks more to take place than in domestic firms.

Duration Models for Time to Promotion

We describe below the duration model that we estimate to address gender differences in time to promotion. Although we also consider a more general semiparametric duration model later, we start with a simple linear regression specification for the log of the time to promotion:

$$\ln T_i^* = \mathbf{X}_i\beta + \varepsilon_i, \tag{1}$$

where T_i^* gauges how much time it takes for the individual i to obtain a promotion, $\mathbf{X}_i$ is a vector of control variates, and ε_i is an error term with scale and shape parameters σ and ς, respectively. In the context of duration models, equation 1 corresponds to an accelerated failure time (AFT) specification.

We construct the duration variable by following from January 1996 to December 2005 every individual who enters the RAIS database between January 1991 and December 1995. As some individuals do not obtain a promotion to a managerial position within the sample, we do not observe the time they take to get promoted, and hence we classify them as nonpromoted. In contrast, we classify as zero duration individuals whose first position within the firm is at the top managerial level (either as a manager or as a director). We initially exclude these individuals from the AFT regression, since equation 1 involves a log transformation of the duration. Altogether, these sample criteria ensure that promotion durations may exhibit only right censoring. Under right censoring, instead of observing the time to promotion T_i^* for each individual in the sample, we have information only on the following promotion duration variable:

$$T_i = \begin{cases} T_i^* & \textit{in the absence of censorship} \\ R_i & \textit{under right censoring} \end{cases}$$

where R_i corresponds to how much time the individual i has had on the job up to December 2005. If the individual i exits the firm before December 2005, then the right-censoring variable R_i denotes tenure on the job, without receiving a promotion, up to the exit date.

The control variables at the individual level come from the RAIS database and refer to the month at which the individual starts at the firm. Using first-month data for individual-specific controls avoids further endogeneity issues, but it has the disadvantage of ruling out hours worked as a control, given that these data are available only from 1995. In contrast, information at the firm level stems from the RAIS database of January 1996, the 1996 Foreign Trade Census data, and the 2000 Census of Foreign Capital.

We construct the binary variable *male* to control for the individual's gender. We additionally include the dummy variables *multinational* and *modern* that take value one for firms with more than 50 percent foreign capitalization (*multinational*) and for firms from a technology-intensive sector (*modern*). Also, we consider the interaction dummy *male* × *multinational* as well as several continuous variables to control for firm-specific factors such as productivity and exposure to international markets. More specifically, *size* and *exports* correspond to the natural logarithm of the number of employees and of the total exports (in USD billions) of the firm that the individual works for, respectively. University ratio is the proportion of employees in the firm with a university degree, and mean wage is average monthly stipend within the firm. We also consider a measure of *turnover* that gauges the firm's job flow intensity by means of the ratio of job flow (hires plus dismissals) to the number of employees in the firm.

Let F denote the cumulative distribution function of the error term in equation 1, with density function f and survival function $S = 1 - F$, and let $\theta = (\beta, \sigma, \varsigma)$ denote the parameter vector. The log-likelihood function then reads

$$L_N(\theta) = \sum_{i=1}^{N}\left(1 - I_i^{(RC)}\right)\ln\left[f(\varepsilon_i)/\sigma\right] + \sum_{i=1}^{N} I_i^{(RC)} \ln S(\varepsilon_i),$$

where $I_i^{(RC)}$ is the indicator function that takes value one if there is right censoring and zero otherwise. In particular, we assume a generalized gamma distribution, with scale and shape parameters σ and ς. The generalized gamma distribution is very flexible, encompassing both the lognormal and Weibull distributions ($\varsigma = 0$ and $\varsigma = 1$, respectively). To account for heterogeneity, we specify duration models with frailty, treating the unobserved individual effects as random draws from a gamma distribution with variance λ.

We initially suppose that censoring is independent of the regressors. This is a very strong assumption in that it rules out the situation in which women are more likely than men to quit their jobs, as in Lazear and Rosen's (1990) model.[4] The data in table 2 indicate that 6.5 percent of the durations relating to male workers exhibit no censoring, whereas 64.6 percent display right censoring and 28.9 percent zero duration. These figures are respectively 4.7 percent, 76.7 percent, and 18.6 percent for female workers. These differences suggest to some extent that censoring depends on gender, in violation of the independence assumption. We thus control for covariate-dependent censoring by using Khan and Tamer's (2007) partial rank estimator. The partial rank estimator entails distribution-free estimates of the regression coefficients without imposing any parametric specification on the link function. In the empirical analysis, we show that accommodating for this more general form of censoring is paramount to examining promotion durations within the Brazilian manufacturing industry.

Khan and Tamer (2007) propose a partial rank estimator for duration models that imposes no parametric specification on the baseline hazard function and allows for general forms of censoring. For instance, in the right-censored version of Ridder's (1990) generalized accelerated failure time model, one observes that

$$Y_i = \left(T_i, I_i^{(RC)}\right)',$$

4. The empirical evidence is conflicting at best. Pekkarinen and Vartiainen's (2006) results confirm that women quit more often than men, whereas Blau and Kahn (1981) and Ransom and Oaxaca (2005) find no evidence supporting such gender differences.

where the duration is $T_i = \min\{\ell^{-1}(\mathbf{X}_\mathbf{i}\beta + \varepsilon_i), \ell^{-1}(R_i)\}$, with R_i denoting the censoring variable and $\ell(\cdot)$ some unknown monotone link function, and

$$I_i^{RC} = I\left\{\ell^{-1}\left(\mathbf{X}_\mathbf{i}\beta + \varepsilon_i\right) \le \ell^{-1}\left(R_i\right)\right\} = I\left\{\mathbf{X}_i\beta + \varepsilon_i \le R_i\right\}$$

indicates whether there is right censoring.[5] As before, $T_i = T_i^*$ for uncensored observations; otherwise $T_i = R_i$.

Similarly to Han's (1987) maximum rank correlation estimator, the idea is to find a transformation $F_{ij} = f(Y_i, Y_j)$ such that

$$\mathrm{E}\left\lfloor I\left\{F_{ij} \ge 0\right\} \middle| X_i, X_j \right\rfloor = \mathrm{E}\left\lfloor I\left\{F_{ji} \ge 0\right\} \middle| X_i, X_j \right\rfloor$$

if and only if $X_i\beta \ge X_j\beta$. Han (1987) considers $F_{ij} = T_i^* - T_j^*$ in the context of uncensored transformation models, which turns out to produce inconsistent estimates if the censoring variable R_i depends somehow on the covariates X_i. Instead, Khan and Tamer use

$$F_{ij} = \overline{T}_i - T_j,$$

where $\overline{T}_i = I_i^{(RC)}\, T_i + (1 - I_i^{(RC)}) \times (+\infty)$ with $0 \times (+\infty) = 0$. As such,

$$I\left\{F_{ij} \ge 0\right\} = I\left\{\overline{T}_i - T_j \ge 0\right\} = 1 - I_i^{(RC)} + I_i^{(RC)} I\left\{T_i \ge T_j\right\}.$$

It then follows that

$$\ell\left(T_i\right) \le \mathbf{X}_\mathbf{i}\beta + \varepsilon_i \le \ell\left(\overline{T}_i\right)$$

and hence that

$$\mathbf{X}_\mathbf{i}\beta \ge X_j\beta \Rightarrow \Pr\left(\overline{T}_i \ge T_j\right) \ge 1/2$$

by monotonicity of $\ell(\cdot)$.

Identification is possible only up to scale, given that the function $\ell(\cdot)$ is unknown and hence it is more convenient to reparameterize the model by setting $\beta = (1, \theta')'$. The partial rank estimator of θ then is

$$(2) \qquad \hat{\theta} = \arg\max_{\theta \in \Theta} \frac{1}{n(n-1)} \sum_{i \ne j} I_i^{(RC)} I\left\{T_i < T_j\right\} I\left\{\mathbf{X}_\mathbf{i}\beta < X_j\beta\right\},$$

5. Note that it is also possible to extend the partial rank framework to consider double censoring.

where Θ is the parameter space. Note that the rank correlation function depends only on uncensored observations, though their ranks consider all observations with longer or equal durations. The partial rank estimator thus combines the information on both censored and uncensored observations, just as in the partial maximum likelihood method put forth by Cox (1972, 1975). In addition, it is straightforward to observe that Khan and Tamer's (2007) partial rank estimator is numerically equivalent to Han's (1987) maximum rank correlation estimator in the absence of censoring as well as in the case of fixed censoring (for example, $R_i = R$).

Khan and Tamer (2007) characterize the consistency and asymptotic normality of the partial rank estimator in equation 2 under the usual regularity conditions for rank-based semiparametric estimators.[6]

Promotions in the Brazilian Manufacturing Industry

In this section, we present estimation results for both parametric and semiparametric specifications. We perform a preliminary regression analysis of time to promotion based on parametric duration models. We estimate the models by maximum likelihood assuming a generalized gamma distribution, which nests both lognormal and Weibull distributions. For each of these distributions, we regress the logarithm of the time to promotion on individual and firm characteristics for a subsample that excludes individuals who start at a managerial position (that is, with time to promotion equal to zero). In addition, we report coefficient estimates both with and without controls for the initial occupation at the firm.

We then report partial-rank estimates of the semiparametric time-to-promotion model for different subsamples. To allow for direct comparison with the parametric analysis, we first exclude from our sample individuals who start at a managerial position and then carry out the regression both with and without controls for the initial occupation at the firm. Finally, we also estimate a semiparametric time-to-promotion model for every individual in our sample, regardless of whether already starting at a managerial position.

Preliminary Analysis

Table 4 reports the estimation results of the accelerated failure time models with gamma frailty for the different error distributions as well as log-likelihood

6. See, for instance, Sherman (1993) for similar regularity conditions in the context of maximum rank correlation.

TABLE 4. Maximum Likelihood Estimates of the AFT within Gamma Frailty[a]

	Generalized gamma		*Lognormal*		*Weibull*	
Covariate	*(1)*	*(2)*[b]	*(3)*	*(4)*[b]	*(5)*	*(6)*[b]
Male	−0.3979	−0.4524	−0.4077	−0.4683	−0.3962	−0.4524
	(0.0928)	(0.0930)	(0.0940)	(0.0944)	(0.0925)	(0.0930)
Male × multinational	0.0736	0.0935	0.0872	0.0954	0.0746	0.0932
	(0.1350)	(0.1347)	(0.1402)	(0.1400)	(0.1347)	(0.1348)
Multinational	−0.4142	−0.4195	−0.4547	−0.4436	−0.4164	−0.4195
	(0.1307)	(0.1302)	(0.1351)	(0.1349)	(0.1304)	(0.1303)
Age	−0.0087	−0.0045	−0.0135	−0.0086	−0.0090	−0.0045
	(0.0030)	(0.0031)	(0.0030)	(0.0031)	(0.0030)	(0.0031)
Ln(size)	0.0254	0.0255	0.0348	0.0345	0.0246	0.0256
	(0.0272)	(0.0269)	(0.0292)	(0.0291)	(0.0271)	(0.0270)
Exports	0.3975	0.3649	0.4101	0.4168	0.3974	0.3657
	(0.1440)	(0.1424)	(0.1482)	(0.1485)	(0.1432)	(0.1425)
University ratio	−1.2538	−1.0992	−1.3779	−1.2028	−1.2478	−1.1004
	(0.2776)	(0.2778)	(0.301)	(0.3011)	(0.2773)	(0.2780)
Turnover	−0.2599	−0.2274	−0.2484	−0.2066	−0.2650	−0.2272
	(0.093)	(0.0942)	(0.102)	(0.1022)	(0.0933)	(0.0942)
Mean wage	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Modern	0.465	0.4515	0.5053	0.4830	0.4686	0.4518
	(0.0529)	(0.0518)	(0.0546)	(0.0546)	(0.0521)	(0.0518)
Engineer		0.0339		0.0896		0.0344
		(0.0640)		(0.0663)		(0.0640)
Intermediate manager		−0.3222		−0.3027		−0.3219
		(0.0676)		(0.0709)		(0.0677)
Supervisor		−0.1572		−0.0979		−0.1564
		(0.0746)		(0.0782)		(0.0746)
Constant	6.0976	5.8156	6.2913	6.2496	5.9026	5.8795
	(0.3465)	(0.3323)	(0.2840)	(0.2879)	(0.2694)	(0.2749)
Scale σ	0.9904	0.6521	1.5919	1.5838	0.7385	0.7305
	(0.2983)	(0.2303)	(0.0280)	(0.0278)	(0.0223)	(0.0218)
shape ς	0.6912	1.1263	0	0	1	1
	(0.3146)	(0.4028)				
Frailty variance λ	0.2247	1.7164	0.0000	0.0001	1.2212	1.3379
	−1.061	−1.2459	(0.0013)	(0.0024)	(0.4301)	(0.4225)
Sample size	19,751		19,751		19,751	

Source: Annual Report of Social Information (RAIS) from 1991 to 2005; 1996 Foreign Trade Census of the Foreign Trade Secretary (SECEX); and 2000 Census of Foreign Capital (CEB).

a. The dependent variable is the natural logarithm of the time to promotion. We report maximum likelihood estimates of the regression coefficients as well as of the distributional and frailty parameters. Figures within parentheses correspond to standard errors.

b. Accountants, administrators, engineers, and lawyers are excluded.

TABLE 5. Likelihood Ratio Tests for Linear Restrictions in the AFT Coefficients[a]

Distribution	Log likelihood: Unrestricted	Log likelihood: Restricted	Likelihood–ratio statistic	p value
H_0: male + male × multinational = 0				
Generalized gamma	–5,929.35	–5,936.36	14.01	0.0002
Lognormal	–5,950.79	–5,957.32	12.97	0.0003
Weibull	–5,929.39	–5,936.40	14.02	0.0004
H_0: multinational + male × multinational = 0				
Generalized gamma	–5,929.35	–5,949.19	39.68	0.0000
Lognormal	–5,950.79	–5,970.37	39.15	0.0000
Weibull	–5,929.39	–5,949.23	39.68	0.0000

Source: Annual Report of Social Information (RAIS) from 1991 to 2005; 1996 Foreign Trade Census of the Foreign Trade Secretary (SECEX); and 2000 Census of Foreign Capital (CEB).

a. The restricted log-likelihood value refers to constraining the maximum likelihood estimator so that the sum of the regression coefficients is zero. The first panel tests whether there are gender differences in time to promotion within multinational firms, whereas the second examines whether it takes less time for men to get a promotion in multinationals than in domestic firms. All results are for specifications that include occupational dummies.

ratio tests for the lognormal and Weibull distributions.[7] The estimates of the regression coefficients are very similar regardless of the distribution assumption. They indicate that, within firms of domestic capital, there are very significant gender differences in promotions. In particular, male employees wait for significantly less time than female employees to get a promotion. The time to promotion for men is on average between 32.71 percent and 37.39 percent shorter than that for women, depending on the specification. Gender differences are less pronounced in multinationals, though still significant. Male workers now take on average between 27.42 percent and 31.13 percent less time to get a promotion than women. In addition, time to promotion is relatively shorter, though not significantly so, within multinationals than within domestic firms regardless of gender.[8] Table 5 validates these claims through formal hypothesis testing. The first panel of table 5 indicates that, though multinationals display fewer gender differences than domestic firms, the differences are still significant at the 1 percent level of significance regardless of

7. Although we report only nonrobust standard errors, clustering by firm changes only marginally the confidence intervals of the coefficient estimates. All qualitative results thus remain valid for cluster-robust standard errors.

8. We compute these effects by exponentiating (the sum of) regression coefficient(s) and then subtracting one. For instance, the second specification in table 4, under the heading "No censoring," gives way to a change of $\exp(-0.4524) - 1 = -36.39$ percent in the average time to promotion if the individual is male. This effect reduces to $\exp(-0.4524 + 0.0935) - 1 = -30.16$ percent if the individual is male and works at a multinational.

the distribution one uses. The second panel also shows that males tend to get a promotion faster within multinationals than within domestic firms.

As for the other controls, it is interesting to observe that the size effect is insignificant, though time to promotion increases with productivity and exposure to international markets, given the sign of the coefficient estimates for *exports, mean wage,* and *modern.* This is consistent with a more competitive environment within more productive firms (including Brazilian multinationals; see note 1). In contrast, time to promotion decreases with *university ratio* and *turnover,* reflecting competition effects within the firm and the sector, respectively. The remaining regression coefficients are all as expected. For instance, there is a significant negative relationship between time to promotion and age, which is not surprising given that age acts as a proxy for experience. In addition, intermediate managers and supervisors wait substantially less time to obtain promotion to managerial positions.

The frailty parameter that regulates the variance of the individual random effects does not differ from zero as long as one considers either a generalized gamma or lognormal distribution.[9] In contrast, the frailty variance is quite close to one for the Weibull distribution, indicating to some extent the presence of individual random effects. However, it is not clear whether this is really a material indication of heterogeneity, given that a frailty-implied gamma mixture of Weibull variates results in a Burr distribution, which is very similar to the generalized gamma distribution (Rodriguez 1977). In fact, in that we restrict our attention to individuals who satisfy a number of criteria, our sample is relatively homogeneous, and hence it is not surprising that we find little evidence of individual random effects. Table 6 indicates that, despite the similarity of the regression coefficients, the statistical evidence favors the extra flexibility of the generalized gamma distribution (or Weibull with frailty) over the more parsimonious lognormal distribution.[10]

Finally, it is worth stressing that the robustness of the parametric results is not only to the specification of the error distribution. Interacting the occupational dummy with gender yields insignificant regression coefficients and does not lead to any qualitative difference in the results. Splitting the firms into exporting and nonexporting by means of a dummy variable does not change much, either. It turns out that 97 percent of the foreign-owned firms export.

9. Not surprisingly, the coefficient estimates for the duration models without frailty are almost identical to the ones in table 4 regardless of the distribution. These results are available from the authors upon request.

10. We report the p value coming from the usual χ^2 distribution for the likelihood-ratio test of the generalized gamma versus the lognormal even if the shape parameter is on the boundary.

TABLE 6. Likelihood Ratio for the Shape Parameter[a]

Distribution	*Log likelihood* *Unrestricted*	*Log likelihood* *Restricted*	*Likelihood–ratio statistic*	*p value*
Lognormal ($\varsigma = 0$)	−5,929.35	−5,950.79	42.88	0.0000
Weibull ($\varsigma = 1$)	−5,929.35	−5,929.39	0.08	0.7831

Source: Annual Report of Social Information (RAIS) from 1991 to 2005; 1996 Foreign Trade Census of the Foreign Trade Secretary (SECEX); and 2000 Census of Foreign Capital (CEB).

a. The unrestricted log-likelihood value refers to the maximum likelihood estimation under the generalized gamma distribution, whereas the restricted log-likelihood value corresponds to constraining the shape parameter ς either to zero or to one, so that the generalized gamma distribution reduces to the lognormal or Weibull distributions, respectively.

This rules out any interaction between the origin of capital and the dummy relating to whether the firm exports or not, for otherwise multicollinearity kicks in, rendering coefficient estimates insignificant and completely unreliable. As a matter of fact, the same pattern arises if we classify firms by size and interact the resulting large-firm dummy with origin of capital. Our findings seem also specific to top management positions. Reestimating the duration models for the time it takes to get a promotion to intermediate management positions (that is, intermediate managers and supervisors) results in insignificant effects for gender, origin of capital, and their interaction. Although we do not report these robustness checks to conserve on space, they are available from the authors upon request.

Semiparametric Analysis

We now investigate the extent to which the log-linear specification and the assumption of covariate-independent censoring affect the results. The latter is particularly a concern since attrition rates are expected to differ by gender (Lazear and Rosen 1990), and we treat job exits as right censoring. In addition, the semiparametric duration model also accommodates left censoring, allowing us to exploit the information content of workers who start at managerial positions (that is, with zero durations). To better understand the effects of dropping the covariate-independent censoring assumption and of allowing for left censoring, we document in what follows both results.

As frailty does not seem to matter much, we estimate the duration model using Khan and Tamer's (2007) semiparametric estimator. As this estimator relies on rank-based methods, it consistently estimates the relative magnitude of the regression coefficients for any strictly monotonic link function (in particular, we fix the coefficient of *exports* to unit). This results in a semiparametric variant of the AFT model in equation 1 under which the link function is strictly monotonic but otherwise unknown. This is in stark contrast to the simple log transformation that equation 1 imposes. Moreover, the partial rank estimator does not require specification of a parametric family for the error distribution.

We employ a SAS/IML genetic algorithm to compute the partial rank estimator, whereas we obtain standard errors by means of bootstrap methods as in Subbotin (2008). In particular, we compute bootstrap-based standard errors based on 100 artificial samples with the same number of observations as the original sample. We consider two specifications. The first accounts for the position at which the individual starts at the firm, and hence we must exclude every individual who already begins at a managerial position, that is to say, any individual with time to promotion equal to zero. The second specification does not control for the starting occupation, and so we consider samples both with and without individuals with zero duration.[11] Adding individuals who start at a managerial position to the sample improves the precision of the estimates owing to the increase in the sample size but at the expense of inducing a sample selection bias.

Table 7 reveals that the qualitative findings are quite robust to the specification (that is, with and without starting occupation) as well as to the sample (with and without zero durations). We indeed observe no variation in the sign of the coefficient estimates across the different specifications and samples. Accounting for previous occupation does not affect much the coefficient estimates of interest, but it improves considerably their precision. The partial rank estimates reveal that females are at a disadvantage in domestic firms relative to male workers but significantly less so within foreign-owned companies. Including individuals with zero duration in the sample brings about additional information that reinforces the moderating effect of multinationals to the extent of seemingly eliminating gender differences in promotion.[12]

Table 8 confirms these results by means of formal Wald tests for linear restrictions. In fact, we cannot even reject the null of no gender differences within foreign-owned companies for the specifications that do not control for previous occupation. The partial rank estimates also indicate that it takes less time for a male worker than for a female worker to obtain a promotion within multinationals. Furthermore, if we extract the information in the cross-section

11. Note that direct comparison between the parametric and semiparametric duration analyses is possible only for the subsample that excludes workers starting at managerial positions.

12. Adding more than 7,500 observations to the regression has a substantial impact on the precision of the coefficient estimates, with up to 95 percent drops in their standard errors (the only exception is the coefficient for AGE, whose standard error slightly increases).

TABLE 7. Partial Rank Estimates of the Semiparametric Time-to-Promotion Model[a]

	No zero durations		
Covariate	*With occupation*	*Without occupation*	*With zero durations*
Male	−1.6872	−1.5774	−0.4835
	(0.2367)	(0.4310)	(0.0523)
Male × multinational	0.4035	0.4576	0.4838
	(0.1850)	(0.7971)	(0.0640)
Multinational	−1.4063	−1.653	−0.5463
	(0.2079)	(0.8481)	(0.0411)
Age	−0.0598	−0.0937	−0.1486
	(0.0045)	(0.0212)	(0.0297)
Ln(size)	0.1305	0.1241	0.0911
	(0.0686)	(0.1102)	(0.0316)
Exports	1.0000	1.0000	1.0000
University ratio	−5.2572	−6.698	−3.5195
	(0.6845)	(1.0881)	(0.4713)
Turnover	−0.3169	−0.595	−0.1406
	(0.2834)	(0.3894)	(0.0488)
Mean wage	0.0005	0.0007	0.0005
	(0.0000)	(0.0001)	(0.0000)
Modern	2.051	2.2914	0.4490
	(0.1506)	(0.3475)	(0.0866)
Engineer	0.6561		
	(0.9823)		
Intermediate manager	−1.2572		
	(1.4152)		
Supervisor	−0.1089		
	(5.1821)		
Sample size	19,751	19,751	27,289

Source: Annual Report of Social Information (RAIS) from 1991 to 2005; 1996 Foreign Trade Census of the Foreign Trade Secretary (SECEX); and 2000 Census of Foreign Capital (CEB).

a. We estimate a model for time to promotion in which the link function is strictly monotonic but otherwise unknown. We consider two samples. The first excludes individuals who begin at the firm either as managers or directors (that is, with zero duration); the second includes these individuals and hence does not allow one to control for the starting occupation. The partial rank estimator identifies the regression coefficients only up to scale, and so we fix the coefficient for exports to one. We obtain the pointwise estimates by means of a genetic algorithm, whereas we use subsampling to compute the standard errors that we report within parentheses.

of individuals who start in the top management of their firms, we cannot anymore reject at the 5 percent level of significance that time to promotion for men does not depend on firm ownership.

It is worth stressing that the descriptive statistics we report in table 3 are not inconsistent with these regression results. Bear in mind that women take less time to be promoted than men in domestic firms only conditional on getting a promotion. However, the relative frequency of promotions of women

TABLE 8. Wald Tests for Linear Restrictions in the Semiparametric Model Coefficients[a]

	No zero durations		
Distribution	*With occupation*	*Without occupation*	*Zero durations*
H_0: Male + male × multinational = 0			
Partial rank estimate	−1.2837	−1.1198	0.0003
Standard error	(0.2858)	(1.1530)	(0.0831)
p value	[0.0000]	[0.3315]	[0.9966]
H_0: Multinational + male × multinational = 0			
Partial rank estimate	−1.0028	−1.1954	−0.0624
Standard error	(0.1737)	(0.2682)	(0.0359)
p value	[0.0000]	[0.0000]	[0.0817]

Source: Annual Report of Social Information (RAIS) from 1991 to 2005; 1996 Foreign Trade Census of the Foreign Trade Secretary (SECEX); and 2000 Census of Foreign Capital (CEB).

a. We report partial rank estimates for the sum of coefficients with their subsampling-based standard errors within parentheses and with their *p* values within brackets. The first panel tests whether there is gender difference in time to promotion within multinational firms; the second examines whether it takes less time for men to get a promotion in multinationals than in domestic firms.

in domestic firms is very low and, in addition, much lower than the relative frequency in multinationals. That right censoring depends on gender and firm ownership not only confirms the importance of the semiparametric duration analysis but also shows that the reason the results in table 3 seem to contradict our regression results is that they completely ignore the censored observations in the right tail of the duration distribution.

All in all, our findings contribute to the literature that compares multinationals and domestic firms (Doms and Jensen 2006; Greene and others 2009) as well as to the literature showing that gender differences may depend on the nature of the firm (see, for example, deVaro and Brookshire 2007). The impact of individuals who start at managerial positions is also interesting inasmuch as it reflects self-selection effects. Before discussing self-selection issues and the reasons why gender differences are less pronounced in multinationals, we first briefly look into two other dimensions in which gender differences in promotions might arise: pay rise owing to promotion and the likelihood of promotion.

Wage Growth and Promotion Likelihood

In this section, we investigate gender differences in promotion likelihood and wage growth. To examine wage growth, we employ monthly wage information of workers over time to estimate longitudinal models of the effect of promotions on wages. The dummy variable PROMOTION takes the value of one for every month after a promotion and zero otherwise. We estimate both random- and fixed-effects panel regressions for males and females and then compare the coefficient estimates of interest to assess whether the impact of promotions on wages differs according to the gender. Apart from the previous firm-specific covariates, we control for job tenure and starting occupation as well as for interaction effects owing to promotions within multinationals. The latter is to capture differential wage effects, if any, of promotions in foreign-owned firms.

Table 9 reports the coefficient estimates with their robust standard errors as well as Wald tests for the equality of the coefficients of interest as estimated from the samples of males and females. As expected, the results show that promotion increases wages of both males and females. The specification with random effects suggests that women receive on average higher increases in wages after a promotion within domestic firms. However, we cannot reject similar increases for men and women in domestic firms if we restrict attention to fixed individual effects. The same applies to foreign-owned companies regardless of whether the individual effects are random or fixed. Finally, the coefficient estimate of the interaction dummy indicates that pay increases owing to promotions are somewhat smaller in multinationals.

To verify whether the probability of promotion depends on gender, we run logit regressions in which the dependent variable takes a value of one if the individual obtains a promotion to a managerial position, or else it equals zero. We control once more for the gender, occupation, and job tenure of the individual as well as for the same firm-specific covariates as before. The results in table 10 suggest that domestic firms display gender differences in promotions not only in terms of how much time it takes to get a promotion but also in terms of the likelihood of promotion. They also suggest that, apart from featuring fewer gender differences, promotion is more likely in a multinational for both genders. This is similar to the findings of the preliminary parametric analysis concerning time to promotion, reflecting perhaps the weaknesses the two methods share, namely, the restrictiveness of the parametric specification and of the regressors' exogeneity assumption.

In sum, we document that wage gains after promotion do not contribute to creating gender differential, at least within foreign-owned firms. The evidence is weaker for domestic firms. The logit regressions for the probability of a promotion indicate that women are at a disadvantage within domestic firms but less so in multinationals, thereby confirming the time-to-promotion results. Finally, multinationals are also characterized by higher likelihood of promotion regardless of gender.

TABLE 9. Wage Regressions with Individual Effects for Both Female and Male Workers[a]

	Random effects		Fixed effects	
Covariate	*Female*	*Male*	*Female*	*Male*
Promotion	0.2569	0.1874	0.1937	0.1304
	(0.0312)	(0.0105)	(0.0387)	(0.0124)
Multinational × promotion	−0.0996	−0.0358	−0.0984	−0.0282
	(0.0418)	(0.0141)	(0.0520)	(0.0168)
Multinational	0.0700	0.0770		
	(0.0240)	(0.0096)		
lnsize	0.0412	−0.0039		
	(0.0123)	(0.0052)		
Exports	0.3079	0.3746		
	(0.0803)	(0.0243)		
University ratio	−0.6454	−0.6342		
	(0.1408)	(0.0559)		
Turnover	0.0606	0.1168		
	(0.0487)	(0.0201)		
Mean wage	0.0003	0.0002		
	(0.0000)	(0.0000)		
Modern	0.0742	0.0613		
	(0.0247)	(0.0097)		
Engineer	0.0550	−0.0101		
	(0.0257)	(0.0120)		
Intermediate manager	0.1084	0.1477		
	(0.0280)	(0.0140)		
Supervisor	−0.1100	−0.0640		
	(0.0304)	(0.0152)		
Tenure	0.1324	0.1217	0.1309	0.1210
	(0.0041)	(0.0016)	(0.0047)	(0.0017)
Tenure^2	−0.0059	−0.0051	−0.0057	−0.0050
	(0.0003)	(0.0001)	(0.0004)	(0.0001)
Constant	7.0700	7.7659	8.0128	8.3169
	(0.1039)	(0.0454)	(0.0111)	(0.0043)
Wald test for equality of coefficients				
Promotion		0.0348		0.119
Promotion + multinational × promotion		0.8487		0.8585
Number of observations	10,343	67,125	10,343	67,125

Source: Annual Report of Social Information (RAIS), 1991 to 2005; 1996 Foreign Trade Census of the Foreign Trade Secretary (SECEX); and 2000 Census of Foreign Capital (CEB).

a. The dependent variable is the logarithm of real monthly wages. The dummy variable promotion takes a value of one for every period after a promotion and zero otherwise, whereas tenure corresponds to the number of years the individual has worked for the firm. We report both random- and fixed-effects estimates with their robust standard errors within parentheses for both female and male wage regressions, as well as the *p* values of the Wald test for the equality of some (linear combinations of) coefficients in both regressions.

TABLE 10. Logit Regressions for the Probability of Promotion[a]

Covariate	*(1)*	*(2)*	*(3)*
Male	0.6029	0.6238	0.6297
	(0.1184)	(0.1191)	(0.1195)
Male × multinational	−0.1409	−0.1596	−0.1628
	(0.1727)	(0.1733)	(0.1737)
Multinational	0.7145	0.7119	0.6912
	(0.1657)	(0.1662)	(0.1669)
lnsize	0.0179	0.0251	0.0086
	(0.0322)	(0.0320)	(0.0323)
Exports	−0.2507	−0.2634	−0.3758
	(0.1801)	(0.1765)	(0.1773)
University ratio	1.6929	1.6144	1.4631
	(0.3106)	(0.3098)	(0.3173)
Turnover	−0.0918	−0.1408	−0.0732
	(0.1249)	(0.1264)	(0.1251)
Mean wage	−0.0002	−0.0002	−0.0002
	(0.0000)	(0.0000)	(0.0000)
Modern	−0.5037	−0.4875	−0.5049
	(0.0649)	(0.0655)	(0.0656)
Engineer		0.1017	0.0609
		(0.0798)	(0.0801)
Intermediate manager		0.4718	0.4593
		(0.0852)	(0.0857)
Supervisor		0.1030	0.1143
		(0.0945)	(0.0948)
Tenure			0.0203
			(0.0022)
Tenure^2			−0.0001
			(0.0000)
Constant	−2.9127	−3.1655	−3.5113
	(0.2943)	(0.2988)	(0.3045)
Wald test for whether the sum of coefficients is equal to zero			
Male + male × multinational	0.0002	0.0003	0.0002
Multinational + male × multinational	0.0000	0.0000	0.0000
Log-likelihood	−5,523.80	−5,502.38	−5,450.79
Number of observations	19,751	19,751	19,751

Source: Annual Report of Social Information (RAIS), 1991 to 2005; 1996 Foreign Trade Census of the Foreign Trade Secretary (SECEX); and 2000 Census of Foreign Capital (CEB).

a. The dependent variable is a dummy that takes a value of one if the individual obtains a promotion to a managerial position, or else it equals zero. We report both random- and fixed-effects estimates with their robust standard errors within parentheses, as well as the *p* values of the Wald test for whether the sum of coefficients is equal to zero.

Foreign Ownership and Gender Differences

It remains to be explained how gender differences arise and why they vary according to ownership. The first point to notice is that a taste for discrimination as in Becker (1957) does not survive long if there is enough competition among firms or employers (Arrow 1972). Of course, one could always claim that, in addition to strong cultural differences, labor-market institutions in Brazil are such that they actually curb the sort of competition that would drive away a taste for gender discrimination. As we have no means to falsify such a conjecture, we turn our attention to the next suspect, namely, statistical discrimination (Arrow 1972; Phelps 1972; Spence 1973).

Assessing the quality of a worker involves costs. Hence some employers might consider costless inference methods, such as considering the worker's gender. These employers would then apply their prior beliefs about the expected qualification of the worker conditional on gender as a hiring or promotion criterion. This would lead to statistical discrimination against a particular gender. Berk (2001) extends the statistical discrimination model to a world in which workers compete for jobs and promotions. It turns out that the self-selection that results from individuals' rationally selecting which jobs or promotions to apply for helps mitigate (and sometimes even overcompensate for) the effects of statistical discrimination. The self-selection mechanism is pretty simple. An individual from the gender with lower average qualifications would apply for a particular job or promotion only if the probability of getting the job or promotion compensated for the application costs. This is more likely to occur if the individual has above-average qualifications.

We next argue that our results are consistent with Berk's (2001) model implications under the assumption that, on average, women face more constraints than men to work long hours or to do business trips. This is enough to generate statistical discrimination against women. A career-minded woman who applies for a job or promotion in a firm with a preference for dedicated employees does so only because she rationally believes that her odds of being hired or promoted compensate her effort in the application process. This means that her qualifications have to be sufficiently high to stand a chance despite discrimination. The latter entails self-selection. To establish the differences between multinationals and domestic firms, we rely on two stylized facts of the Brazilian labor market (OECD 2008). The first is that domestic firms offer a more flexible package in terms of working hours and business trips than do multinationals. This makes them more appealing to less ambitious and less career-minded individuals (regardless of gender). The second is that multinationals compete more fiercely for ambitious, career-minded, highly skilled workers, and thereby promote better pay, on average, than domestic firms (see, for example, Martins and Esteves 2008).

We find some indirect evidence supporting the above claims.[13] On one hand, male workers tend to work similar weekly hours in foreign-owned and domestic firms, whether or not they have been promoted. The only difference is in the lower support of the distribution, where promoted men work much harder than nonpromoted ones (twenty-two to forty-four hours versus eight to forty-four, respectively). This is consistent with the fact that promoted workers are more likely to be career minded. On the other hand, nonpromoted female employees work relatively much less only in domestic firms (ten to forty-four hours), perhaps because domestic firms offer more flexible packages. In addition, that women work longer hours than their male counterparts in multinationals (thirty-six to forty-four hours whether promoted or not) is consistent with self-selection. As the latter alleviates the impact of statistical discrimination, we fail to observe as much gender difference in foreign-owned companies as in domestic firms.

The impact of individuals who start at managerial positions in the coefficient estimates reinforces the self-selection story. These are precisely the workers who are most likely to have the highest qualifications, and hence their inclusion in the sample potentiates self-selection effects. As the latter reduces the imprint of statistical discrimination, it is not surprising that we cannot reject gender differences in multinationals for the sample that includes individuals with zero durations. This is because the repercussion to self-selection is higher in multinationals, given that they compete more intensively for career-minded workers.

Finally, it is important to stress that our sample is homogeneous in observed characteristics, in that we consider only highly skilled individuals, with university degree, and pursuing specific types of career. The impact of career-mindedness in the time to promotion may differ across different levels of observed skills, and hence our findings are completely silent about gender differences for low-skilled workers. However, we have no good reason to believe that the skill distribution is gender specific, and hence differentiating

13. Unfortunately, we observe only the effects of self-selection. We cannot properly model the selection mechanism, given that we do not observe many characteristics that determine individual career-mindedness, such as marital status and number of dependent children. Similarly, we have information neither about household characteristics (for example, the presence of other income earners or dependents at home, marital status) nor leave and vacation periods.

FIGURE 2. Participation in Top Management Positions, by Gender and Type of Firm[a]

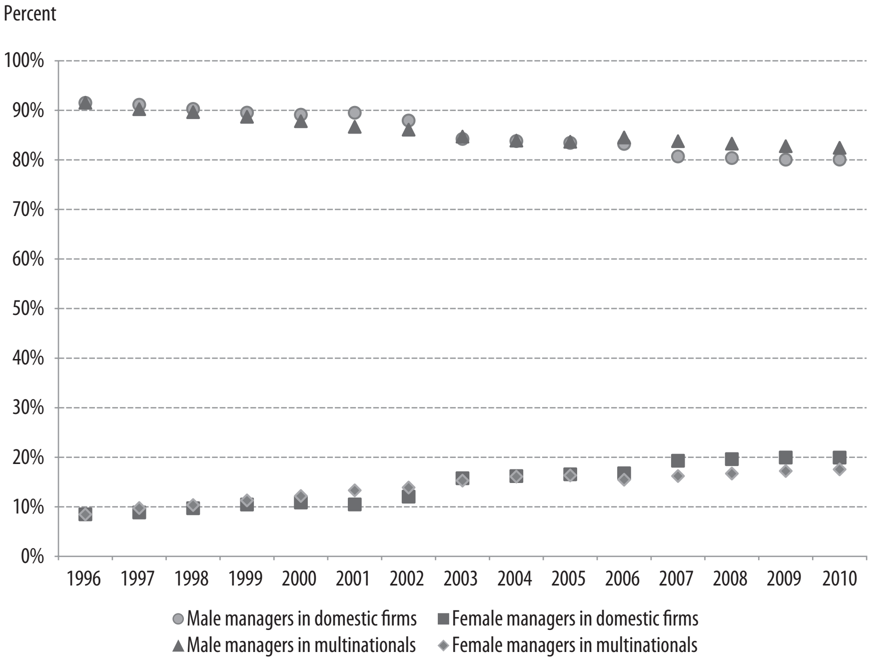

a. Sample is individuals with university degrees.

between low- and high-skilled workers would likely not bring many insights to this study.

Managerial Positions and Gender Concerns

We are so far treating gender concerns in a static manner in that we implicitly assume that labor market conditions did not change much over the sample period. In what follows, we first check whether recent trends in the labor market are responsible for our findings and then examine whether gender concerns drive the creation of new managerial positions.

Figure 2 depicts the female and male participation in top management from 1996 to 2010 for a sample of individuals with university degrees. The female participation increases more than twofold from 1996 to 2010, from 8.5 percent to 19.9 percent in domestic firms and from 8.5 percent to 17.6 percent in multinationals. Interestingly, the participation rate is higher for domestic firms than for multinationals; hence one would expect gender differences to decline, especially for domestic firms. This does not, however, contradict our finding that there are fewer gender differences in multinationals. To appreciate that, one must also look at changes in the number of managerial positions available at domestic and multinational firms from 1996 to 2010. In particular, the number of male workers with a university degree at managerial positions decreases from 11,774 to 9,730 in domestic firms (a 17-percent drop), whereas it increases from 10,226 to 17,684 in multinationals (a 73-percent rise). The corresponding figures for female managers and directors imply increases of 122 percent for domestic firms (from 1,093 to 2,424) and 298 percent for multinationals (from 947 to 3,768). It seems that most promoted women in domestic firms took the place of a male manager or director, thereby increasing their participation in top management. The same does not apply to multinationals, for which promotions result mainly from the increase in managerial positions (most of which are taken by women).[14] Accordingly, the only implication of figure 2 concerning the results presented in the previous section on gender differential in multinational firms is that they are not an artifact owing to the recent trends in female participation in the labor market.

As for the creation of new top management positions, for each uncensored observation, we construct two dummy variables, $\Delta_F^{(ij)}$ and $\Delta_M^{(ij)}$. The former takes a value of one if the number of women in top management positions in firm j has increased after the promotion of individual i and zero otherwise. The latter is the analogous dummy variable for male workers, assuming the value of one if the number of men in top management positions in firm j has increased after the promotion of individual i and zero otherwise. Table 11 reports the relative frequency of these dummy variables by gender and firm ownership.

Consider first the case in which $\Delta_F^{(ij)} = \Delta_M^{(ij)} = 0$. This corresponds to promotions to existing positions so as to replace someone who has quit the firm. This situation accounts for 52 percent of the promotions of female workers in domestic firms and 40 percent in multinationals. The corresponding figures for male workers are 52 and 46 percent, respectively. This means that women are relatively more likely than men to get promotions to newly created top management positions in multinationals. Note also that a female promotion with $\Delta_F^{(ij)} = 0$ means that the woman has replaced another woman; similarly, a male promotion with $\Delta_M^{(ij)} = 0$ implies that the man has replaced another man in the top management position. The frequency of promotions of women in

14. See also Martins and Esteves (2008) for patterns of external and internal hire in Brazil.

TABLE 11. Top Management and Gender Concerns[a]

		Domestic		Multinational	
Gender		$\Delta_F^{(ij)}=0$	$\Delta_F^{(ij)}=1$	$\Delta_F^{(ij)}=0$	$\Delta_F^{(ij)}=1$
Female	$\Delta_M^{(ij)}=0$	0.5238	0.1667	0.4026	0.2078
	$\Delta_M^{(ij)}=1$	0.0595	0.2500	0.0519	0.3377
Male	$\Delta_M^{(ij)}=0$	0.5219	0.0558	0.4663	0.0660
	$\Delta_M^{(ij)}=1$	0.1753	0.2470	0.1348	0.3329

Source: Annual Report of Social Information (RAIS), 1991 to 2005; 1996 Foreign Trade Census of the Foreign Trade Secretary (SECEX); and 2000 Census of Foreign Capital (CEB).

a. The dummy variable $\Delta_F^{(ij)}$ takes a value of one if the number of women in top management positions in firm *j* has increased from the year preceding the promotion of individual *i* and zero otherwise. $\Delta_M^{(ij)}$ is the analogous dummy variable for male workers, assuming a value of one if the number of men in top management positions in firm *j* has increased from the year preceding the promotion of individual *i* and zero otherwise. We report the average values of $\Delta_F^{(ij)}$ and $\Delta_M^{(ij)}$ by gender and firm ownership.

domestic firms with $\Delta_F^{(ij)} = 0$ is 58.33 percent but only 45.45 percent in multinationals. The corresponding figures for promotions of men with $\Delta_M^{(ij)} = 0$ are 57.77 percent and 53.0 percent, respectively. As before, the gap is much wider for women than for men. In addition, the relative frequency is (significantly) below 50 percent only for multinationals, corroborating the evidence that female promotion in multinationals is accomplished mainly through an increase in the number of top management positions.

Conclusion

This paper assesses whether gender matters in the time it takes to get a promotion to a managerial position in the largest firms of the Brazilian manufacturing industry. Our choice of Brazil was based not only on a natural interest in assessing gender differences in promotion within a developing country but also on the general features of the Brazilian data set. This data set includes observations of multiple workers per establishment at different occupations and hierarchical levels as well as their career paths in terms of occupation and wages.

We find that there are significative gender differences in promotions within domestic firms. The evidence for foreign-owned firms is much weaker, though. These findings complement well the recent evidence that the nature of the firm may imply substantial differences in managerial practices and in the role of promotions (Doms and Jensen 2006; deVaro and Brookshire 2007; Greene and others 2009). We argue that the stronger gender differences in domestic firms are a result of the combination of statistical discrimination and self-selection (Berk 2001). If career-minded women prefer to work in multinational firms, statistical discrimination may become more apparent within domestic firms.

The presence of statistical discrimination even for highly skilled workers is obviously a concern. There are a number of policies that one may employ to mitigate gender gaps in the labor market, such as comparable worth initiatives and the imposition of quotas for women. However, these examples seem a bit out of the Brazilian reality at the moment. Ñopo (2012, chap. 17) describes four main policy prescriptions in the case of Latin America: investing in education early in life, boosting productivity and reducing labor market segregation, fostering a more equitable division of household responsibilities, and diminishing stereotypes. Investing more in girls' education will likely have much more impact on low-skilled workers than on the highly skilled. Similarly, reducing occupational segregation seems much more relevant for workers at the bottom of the earnings distribution (Hsieh and others 2013). Promoting a more equitable division of household responsibilities and reducing stereotyping seem much better avenues to tackle statistical discrimination against highly skilled female workers. Policymakers could, for example, establish more equitable maternity and paternity leave, expand early childhood infrastructure, and increase school hours to foster a more equitable division of duties and opportunities within households. As for stereotypes, they survive only in the absence of better information (Altonji and Pierret 2001). Policymakers should seek to improve information in the job market, even if by means of unorthodox methods, such as promoting female role models in Brazilian soap operas (Chong and La Ferrara 2009).

Comment

Renos Vakis: Coelho, Fernandes, and Foguel set about to test whether women face career barriers in the largest manufacturing firms in Brazil. Indeed, gender segmentation and segregation in the labor market is a widespread phenomenon in Latin America, and Brazil is not an exception. Understanding the magnitude of the problem is thus a first step in finding ways to reduce its prevalence.

Typically, studies that look at these issues focus mainly on the cross-sectoral segregation and constraints that women face in moving across low- to high-productivity sectors. The authors take an interesting approach: they look at whether gender matters in the time it takes to get a promotion to a managerial position in the largest firms of the Brazilian manufacturing industry. By focusing on a specific area of the economy, they are able to take advantage of a unique set of information on manufacturing workers at different occupations and hierarchical levels and to track career paths in terms of occupation and wages. In addition, the clever approach to separate the analysis between domestic and international firms also allows them to shed light on the process and potential mechanisms of the career barriers.

Indeed, the approach allows them to confirm large and significant gender differences in career barriers. Specifically, they find that women have a lower probability of being promoted than men, and even when they do get promoted it takes them longer to do so. The authors also find that these results are driven by large differences in promotions within domestic firms as opposed to foreign-owned firms, where the results are weak. This allows them to conclude that preferences, self-selection, and statistical discrimination are at play. On one hand, women with higher career aspirations may optimally choose multinationals, where there are more options for the most career-minded workers, implying that statistical discrimination is less an issue. By contrast, jobs in domestic firms offer more flexibility in terms of hours per week, but they do so at the cost of higher statistical discrimination. In this sense, the self-selection of career-minded women makes statistical discrimination more apparent within domestic firms.

One of the strengths of the paper is that it employs various methods to show the robustness of the results. Controlling for censoring the nonparametric analysis as well as the employment of various indicators for ascension to test the main results are great ways to rigorously validate the results. They all point to the same consistent results mentioned above, further strengthening the findings.

Because it focuses on high-skilled workers in large firms, the paper cannot provide insights about similar gaps in other parts of the economy, something that would be quite interesting to understand in the future. Even within the manufacturing sector, the large firms used in the analysis represent only 5 percent of all the firms in the sector (though these firms employ about one in every three sector employees).

A number of additional exercises could have helped shed light on some of the findings and the underlying mechanisms. For example, it is hard to separate self-selection from statistical discrimination in the current approach. It would be important to explore ways to do that so that the magnitude of the difference between these aspects might eventually be understood. In addition, much of the literature on self-selection in the labor market discusses how life-cycle aspects affect employment decisions. In this sense, exploring these results by separating women by marital status or by whether or not they have children may be important. Finally, extending this work to other sectors of the economy, or better, exploring ways to create a labor market–wide analysis, will be extremely useful.

Pushing further into the policy implications, the paper briefly discusses ways to deal with the presence of statistical discrimination. That this exists even for highly skilled workers is a deep concern, as observed by the authors, who also note that traditional policy measures like imposing job quotas may not be realistic for the case of Brazil. They suggest that policymakers explore ways to change social norms, a policy area that could induce deep structural societal change with direct benefits for the high-skilled female population. In the interest of accomplishing this goal, the paper addresses reform areas such as maternity and paternity leaves and increases in the supply of early childhood infrastructure. Promoting female role models may be another channel that provides information while at the same time helps widen the aspirational window for women and break social norms and stigmas. These approaches would undoubtedly be beneficial to all women. Additional work in these areas, including experimental work and impact evaluations to test some of these policy options and their impact on breaking some of these gender barriers and gaps, will be key in the future.

References

Acosta, P. A. 2010. "Promotions Dynamics and the Peter Principle: Incumbents vs. External Hires." *Labour Economics* 17 (6): 975–86.

Altonji, J., and C. Pierret. 2001. "Employer Learning and Statistical Discrimination." *Quarterly Journal of Economics* 116 (1): 313–50.

Amann, E. 2009. "Technology, Public Policy, and the Emergence of Brazilian Multinationals." In *Brazil as an Economic Superpower? Understanding Brazil's Changing Role in the Global Economy,* edited by L. Brainard and L. Martinez-Diaz. Brookings Press.

Arrow, K. J. 1972. "Models of Job Discrimination." In *Racial Discrimination in Economic Life,* edited by A. H. Pascal. Lexington, Mass.: D. C. Heath.

Baker, G., M. Gibbs, and B. Holmstrom. 1994. "The Internal Economics of the Firm: Evidence from Personnel Data." *Quarterly Journal of Economics* 109 (4): 881–919.

Baldwin, M. L., R. J. Butler, and W. G. Johnson. 2001. "A Hierarchical Theory of Occupational Segregation and Wage Discrimination." *Economic Inquiry* 39 (1): 94–110.

Barnett, W. P., J. N. Baron, and T. E. Stuart. 2000. "Avenues of Attainment: Occupational Demography and Organizational Careers in the California Civil Service." *American Journal of Sociology* 106 (1): 88–144.

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

Berk, J. B. 2001. "Statistical Discrimination in a Labor Market with Job Selection." Working Paper. Berkeley: University of California, Haas School of Business.

Blau, F. D., and J. deVaro. 2007. "New Evidence on Gender Differences in Promotion Rates: An Empirical Analysis of a Sample of New Hires." *Industrial Relations* 46 (3): 511–50.

Blau, F. D., and L. Kahn. 1981. "Race and Sex Differences in Quits by Young Workers." *Industrial and Labor Relations Review* 34 (4): 563–77.

Booth, A. L., and M. Francesconi. 2000. "Job Mobility in 1990s Britain: Does Gender Matter?" *Research in Labor Economics* 19:173–89.

Booth, A. L., M. Francesconi, and J. Frank, J. 2003. "A Sticky Floors Model of Promotion, Pay, and Gender." *European Economic Review* 47 (2): 295–322.

Cabral, R., M. A. Ferber, and C. Green. 1981. "Men and Women in Fiduciary Institutions: A Study of Sex Differences in Career Development." *Review of Economics and Statistics* 63 (4): 573–80.

Cannings, K., and C. Montmarquette. 1991. "Managerial Momentum: A Simultaneous Model of the Career Progress of Male and Female Managers." *Industrial and Labor Relations Review* 44 (2): 212–28.

Chong, A., and E. La Ferrara. 2009. "Television and Divorce: Evidence from Brazilian Novelas." *Journal of the European Economic Association: Papers and Proceedings* 7 (2–3): 458–68.

Cox, D. R. 1972. “Regression Models and Life Tables.” *Journal of the Royal Statistical Society, Series B* 34 (2):187–220.

———. 1975. “Partial Likelihood.” *Biometrika* 62 (2): 269–76.

deVaro, J., and D. Brookshire. 2007. “Promotions and Incentives in Nonprofit and For-Profit Organizations.” *Industrial and Labor Relations Review* 60 (3): 311–39.

Doms, M. E., and J. B. Jensen. 2006. “Comparing Wages, Skills, and Productivity between Domestic and Foreign Owned Manufacturing Establishments in the United States.” In *Geography and Ownership as Bases for Economic Accounting: Studies in Income and Wealth,* edited by R. Baldwin, R. E. Lipsey, and J. D. Richardson, pp. 235–55. University of Chicago Press,

Gerhart, B. A., and G. T. Milkovich. 1989. “Salaries, Salary Growth, and Promotions of Men and Women in a Large, Private Firm.” In *Pay Equity: Empirical Inquiries,* edited by R. T. Michael, H. I. Hartmann and B. O’Farrell, pp. 23–43. Washington: National Academy Press.

Giuliano, L., D. I. Levine, and J. Leonard. 2011. “Race, Gender, and Hiring Patterns: Evidence from a Large Service-Sector Employer.” *Journal of Human Resources* 46 (1): 26–52.

Greene, W. H., and others. 2009. “Multinationals Do It Better: Evidence on the Efficiency of Corporations’ Capital Budgeting.” *Journal of Empirical Finance* 16 (5): 703–20.

Groot, W., and H. M. van den Brink. 1996. “Glass Ceilings or Dead Ends: Job Promotion of Men and Women Compared.” *Economics Letters* 53 (2): 221–26.

Han, A. 1987. “Nonparametric Analysis of a Generalized Regression Model: The Maximum Rank Correlation Estimator.” *Journal of Econometrics* 35 (2–3):303–16.

Hsieh, C.-T., and others. 2013. “The Allocation of Talent and U.S. Economics Growth.” Working Paper 18693. Cambridge, Mass.: National Bureau of Economic Research (January).

Khan, S., and E. Tamer. 2007. “Partial Rank Estimation of Transformation Models with General Forms of Censoring.” *Journal of Econometrics* 136 (1): 251–80.

Lazear, E. P., and S. Rosen. 1990. “Male-Female Wage Differentials in Job Ladders.” *Journal of Labour Economics* 8 (1): 106–23.

Lewis, G. B. 1986. “Gender and Promotions: Promotion Chances of White Men and Women in Federal White-Collar Employment.” *Journal of Human Resources* 21 (3): 406–19.

Martins, P. S., and L. A. Esteves. 2008. “Foreign Ownership, Employment, and Wages in Brazil: Evidence from Acquisitions, Divestments, and Job Movers.” Discussion Paper 3542. Bonn, Ger.: Institute for the Study of Labor.

McCue, K. 1996. “Promotions and Wage Growth.” *Journal of Labour Economics* 14 (2): 175–209.

Meitzen, M. E. 1986. “Differences in Male and Female Job-Quitting Behavior.” *Journal of Labor Economics* 4 (2): 151–67.

Mincer, J., and H. Ofek. 1982. “Interrupted Work Careers: Depreciation and Restoration of Human Capital.” *Journal of Human Resources* 17 (1): 3–24.

Ñopo, H. 2012. "New Century, Old Disparities: Gender and Ethnic Earnings Gaps in Latin America and the Caribbean." Washington: Inter-American Development Bank and World Bank.

OECD (Organization for Economic Cooperation and Development). 2008. "Do Multinationals Promote Better Pay and Working Conditions?" Paris: OECD Employment Outlook.

Oliveira, A. M. H. C. 2001. "Occupational Gender Segregation and Effects on Wages in Brazil." In *Proceedings of the XXIVth General Population Conference of International Union of Scientific Studies of Population.* Salvador, Brazil: International Union of Scientific Studies of Population.

Paulin, E. A., and J. M. Mellor. 1996. "Gender, Race, and Promotions within a Private-Sector Firm." *Industrial Relations* 35 (2): 276–95.

Pekkarinen, T., and J. Vartiainen. 2006. "Gender Differences in Promotion on a Complexity Ladder of Jobs." *Industrial and Labor Relations Review* 59 (2): 285–301.

Petersen, T., and I. Saporta. 2004. "The Opportunity Structure for Discrimination." *American Journal of Sociology* 109 (4): 852–901.

Phelps, E. S. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review* 62 (4): 659–61.

Powell, G. N., and D. A. Butterfield. 1994. "Investigating the 'Glass Ceiling' Phenomenon: An Empirical Study of Actual Promotions to Top Management." *Academy of Management Journal* 37 (1): 68–86.

Ransom, M., and R. L. Oaxaca. 2005. "Intrafirm Mobility and Sex Differences in Pay." *Industrial and Labor Relations Review* 58 (2): 219–37.

Ridder, G. 1990. "The Non-Parametric Identification of Generalized Accelerated Failure-Time Models." *Review of Economic Studies* 57 (2):167–82.

Rodriguez, R. N. 1977. "A Guide to the Burr Type XII Distribution." *Biometrika* 64 (1):129–34.

Royalty, A. B. 1998. "Job-to-Job and Job-to-Nonemployment Turnover by Gender and Education Level." *Journal of Labour Economics* 16 (2): 392–443.

Sherman, R. P. 1993. "The Limiting Distribution of the Maximum Rank Correlation Estimator." *Econometrica* 61 (1): 123–37.

Sicherman, N., and O. Galor. 1990. "A Theory of Career Mobility." *Journal of Political Economy* 98 (1): 169–92.

Spence, A. M. 1973. "Job Market Signalling." *Quarterly Journal of Economics* 87 (3): 355–74.

Spurr, S. J. 1990. "Sex Discrimination in the Legal Profession: A Study of Promotion." *Industrial and Labor Relations Review* 43 (4): 406–17.

Subbotin, V. 2008. "Asymptotic and Bootstrap Properties of Rank Regressions." Working Paper. Evanston, Ill.: Northwestern University.