Josh Ederington[1] / Jenny Minier[2] / C. Jill Stowe[3]

Risk and Discrimination

[1] Economics, University of Kentucky, Lexington, KY 40506, USA, E-mail: ederington@uky.edu
[2] Economics, University of Kentucky, Lexington KY 40506, USA, E-mail: jminier@uky.edu
[3] Agricultural Economics, Lexington KY 40546, USA, E-mail: jill.stowe@uky.edu

Abstract:
In the traditional Becker model of employer discrimination, discriminatory behavior arises from a utility-maximizing owner who balances firm profits against the disutility of hiring workers from the disadvantaged demographic group. However, in the modern firm, many human resource decisions are made by agents of the owner (managers) whose actions may not reflect the preferences of even profit-maximizing owners. We present a principal-agent model of discrimination with a profit-maximizing owner and a gender-discriminating manager and show that managerial discrimination is increasing with the degree of risk in the firm's revenue stream. Empirical tests using a Colombian plant-level dataset support a prediction of our model that female workers should be under-represented in more revenue volatile firms and industries.

Keywords: labor market discrimination, principal-agent model, gender
DOI: 10.1515/bejeap-2017-0204

1 Introduction

For over 50 years, Gary Becker's path-breaking book on labor market discrimination (Becker 1957) has influenced how economists think about the incidence and persistence of discrimination. This relatively simple model produces some very powerful predictions: for example, in a perfectly competitive market, discriminating employers incur higher costs than non discriminating employers, and so will be driven out of the market due to this cost disadvantage.[1]

One of the more surprising aspects of the Becker model is how little work has subsequently incorporated advances in economic theory into the model. [2] For example, in the version of the model that has received the most attention – the case of employer discrimination – the owner of a firm can indulge his or her discriminator preferences by refusing to hire a member of some demographic group despite the fact that a worker from the group suffering discrimination is just as productive as a worker from the majority group. However, in today's economy, most economic activity takes place in firms where many decisions, including whom to hire, are not made by the owner of the firm, but by *agents* of the owner whose actions the owner cannot perfectly observe. As work on the principal-agent problem has shown, the resulting outcomes do not necessarily reflect the preferences of the principal. The potential for discrimination to arise from managerial discretion has been suggested before in the literature (see the discussion in Ashenfelter and Hannan (1986) and Méon and Szafarz (2011))[3], however the implications of a principal-agent approach on the determinants and prevalence of discrimination has not previously been investigated. Thus, in this paper, we introduce a principal-agent model of discrimination and explore the implications of a principal-agent approach for determining the conditions under which discrimination is more likely.

We extend the Becker model of discrimination by developing a principal-agent model of discrimination with a profit-maximizing principal/owner and a discriminating agent/manager where the owner delegates managerial decisions to the agent. Since it is costly for the owner to observe all the managerial decisions of the agent, it is possible for the manager to indulge his or her discriminatory preferences without being detected by the owner. To align incentives more closely, owners traditionally rely on performance pay schemes.[4] We show that, in the presence of such performance-pay schemes, the manager's discriminatory actions will become a function of the variability of firm profits. Specifically, more noise in the revenue generation process results in more discriminatory behavior by managers. [5] Intuitively, this is due to the fact that increased noise in the production process makes it more difficult to distinguish between positive revenue shocks and non-discriminatory (profit-maximizing) behavior by the manager. Thus, the incentives to engage in non-discriminatory behavior are lessened in an environment of greater risk, and discriminatory behavior is more likely. This mechanism suggests that workers from some demographic groups might be underrepresented in more volatile firms.

There is an obvious parallel between the ideas of this paper and the recent literature on performance pay and wage discrimination. Of course those papers are looking at the effects of performance pay for *employees* on

C. Jill Stowe is the corresponding author.

employee earnings and typically, although not always, find that such performance pay reduces earnings differentials (e.g. see Heywood and Jirjahn (2002), Jirjahn and Stephan (2004), Heywood and O'Halloran (2005), Fang and Heywood (2006), and de la Rica, Dolado, and Vegas (2015)). The typical argument is that performance-pay to employees reduces managerial discretion and thus can reduce wage-differentials generated by managerial discrimination.[6] In contrast, this paper considers the case of performance pay for *managers* as a means of reducing discriminatory behavior by managerial agents (as in Bandiera, Barankay, and Rasul (2009). In this context we derive a novel prediction in which increased risk reduces the effectiveness of managerial performance pay schemes and thus results in greater discriminatory behavior.

Without loss of generality, we present the model in the context of gender discrimination, where the manager has a preference for male employees. The theoretical model predicts an increase in discriminatory behavior on the part of the manager when uncertainty in revenues increases. We then test the implications of our model using Colombian plant-level data and find support for the predictions of our model. Specifically, we find that female workers are, in fact, underrepresented in industries which exhibit greater revenue and profit variability. In addition, we show that, within industries, female workers are underrepresented in plants that exhibit greater revenue and profit variability.[7] These correlations are consistent with our model that discriminatory behavior by managers is more prevalent in situations of greater revenue uncertainty.

Of course, while these correlations are suggestive, we can only claim they are consistent with our underlying story of labor market discrimination in a principal-agent framework. For example, an alternative explanation is that there is some underlying (and unobservable) difference between male and female labor which results in male labor being over-represented in higher variance industries/firms. Thus, our regressions include a host of plant-level characteristics in an attempt to control for differences in technology or the production process across plants. In addition, there is a large empirical literature establishing that women are more risk averse than men, so it is possible female workers are self-selecting into more stable work environments. However, in Section 4, we show our link between profit/revenue variability and female labor share continues to hold, even when we control for wage/employment variability.

2 Model

A well-known result in the principal agent literature is that a more risky environment leads to lower incentives for managerial effort (e.g. see Holmstrom and Milgrom (1987) and Gibbons (1997)). Thus, if one considers (profit-reducing) discriminatory behavior as one aspect of managerial effort, it seems direct to connect greater risk in the firm environment with increased discrimination (although our paper is the first, to our knowledge, to make this link). However, we find it helpful to write down an explicit model to both make clear the assumptions necessary to establish this link as well as provide a template for future work on the determinants of discrimination in a principal-agent context.

Thus, assume the principal (the owner of a firm) gives the task of managing the production process to an agent (upper management of a plant). The production process requires a sequence of tasks to be performed, and the manager hires and assigns a set of applicants (as well as needed inputs) to complete these tasks (i.e. the standard assignment problem in operations research). Assume there are a continuum of different allocations of prospective workers/inputs to tasks available to the manager and index them on an interval, say [0,1] in accordance with decreasing aggregate utility to the prospective male workers. Thus, an allocation f is associated with each point on the interval and, by construction, an allocation with a higher f represents an allocation that is less favorable to male workers and more favorable to female workers.

The owner is a risk-neutral profit maximizer whose revenue, $R(f)$, is determined by the manager's choices (i.e. the allocation chosen) and a random state of nature:[8]

$$R(f) = y(f) + \varepsilon, \text{with} : \begin{cases} y' > 0, \ \ y'' < 0 \text{ for } \ \ f < f^* \\ \ \ y' = 0 \text{ for } \ \ f = f^* \\ y' < 0, y'' < \ 0 \text{ for } \ \ f > f^* \end{cases}. \tag{1}$$

The term $y(f)$ is the (unobservable) level of revenue produced and ε is an independently distributed noise term drawn from normal distribution $\varphi(\varepsilon)$ with zero mean and variance σ^2.[9] While the distribution of ε is common knowledge, the realization of ε occurs after decisions are made. In (1), f^* is the profit-maximizing allocation and, by assumption, we assume revenue is continuously declining as one trends away from the optimal allocation by either favoring male or female workers. Without loss of generality, we assume the marginal cost of production is zero.

Consistent with the standard principal-agent model we assume the manager may have a taste for gender discrimination but that the actions of the manager (i.e. the allocation f) are unobserved by the owner.[10] The owner has full bargaining power and offers the manager a take-it-or-leave-it wage $w = s + bR(f)$.[11] The term s represents the salary component of the wage, and b represents the piece rate, which is based on observed revenue.

The manager's utility from his outside option is assumed to be zero. There is no cost of effort, but the manager does bear disutility when setting allocations which advantage female employees at the expense of male employees. This disutility is represented by c(f), where $c'(f) > 0$ and $c''(f) > 0$. The manager's utility is given by the constant absolute risk aversion function $u(x) = -\exp(-rx)$, where net wage is written as $x = s + b(y(f) + \varepsilon) - c(f)$ and $r > 0$ is the coefficient of absolute risk aversion.

The manager chooses f to maximize expected utility:

$$\max_f \int -\exp(-r(s + b(y(f) + \varepsilon) - c(f))\phi(\varepsilon)d\varepsilon.$$

The expected utility-maximizing allocation chosen by the manager, $\hat{f}$, satisfies the necessary first-order condition:

$$by'(f) = c'(f). \tag{2}$$

Equation (2) is the standard optimality condition equating the marginal benefit and marginal cost of engaging in discrimination. Note that, because the manager has a taste for discrimination (i.e. $c'(f) > 0$), the manager will choose an non-optimal allocation that is biased towards male employees (i.e. $f < f^*$). While it is the case that performance-pay will reduce the degree of discrimination (consistent with Bandiera, Barankay, and Rasul (2009)), our model parallels that of Méon and Szafarz (2011) in that the use of performance-based contracts does not fully eliminate managerial discrimination. An interior solution always exists because the second-order condition is satisfied:

$$by''(f) - c''(f) < 0. \tag{3}$$

Since the firm has full bargaining power, maximizing profit is identical to maximizing total surplus; total surplus is determined by the sum of the firm's expected profit and the agent's certainty equivalent. The firm's expected profit is expected revenue minus expected wage, or $(1-b)y(\hat{f}) - s$. The agent's certainty equivalent is his expected wage less the disutility from hiring females and the cost of bearing risk, or $s + by(\hat{f}) - c(\hat{f}) - \frac{1}{2}rb^2\sigma^2$. Hence, the firm chooses the piece rate to maximize profits:

$$\max_b \; y(\hat{f}) - c(\hat{f}) - \frac{1}{2}rb^2\sigma^2.$$

The profit-maximizing piece rate, b^*, then satisfies the first-order condition:

$$y'(\hat{f})\frac{d\hat{f}}{db} - c'(\hat{f})\frac{d\hat{f}}{db} - rb\sigma^2 = 0. \tag{4}$$

From (2), we know that $c'(\hat{f}) = by'(\hat{f})$ and $\frac{d\hat{f}}{db} = \frac{y'(\hat{f})}{c''(\hat{f}) - by''(\hat{f})} > 0$. Substituting into (4) and simplifying yields

$$[r\sigma^2 y''(\hat{f})]b^2 - [(y'(\hat{f}))^2 + r\sigma^2 c''(\hat{f})]b + (y'(\hat{f}))^2 = 0. \tag{5}$$

The optimal piece rate, b^*, solves (5), and the second-order condition is satisfied, ensuring an interior solution:

$$2br\sigma^2 y''(\hat{f}) - (y'(\hat{f}))^2 - r\sigma^2 c''(\hat{f}) < 0. \tag{6}$$

This setup leads to the following proposition.

Proposition 1.
Discrimination is increasing in risk in the production process (i.e. an increase in σ^2 leads to a decrease in $\hat{f}$).

Proof.
Total differentiation of (5) with respect to σ^2 yields:

$$\frac{db^*}{d\sigma^2} = \frac{rb(c''(\hat{f}) - by''(\hat{f}))}{2br\sigma^2 y''(\hat{f}) - (y'(\hat{f}))^2 - r\sigma^2 c''(\hat{f})}.$$

By (3), the numerator is positive, and by (6), the denominator is negative. Consequently, we have $\frac{db^*}{d\sigma^2} < 0$. This leads immediately to the result

$$\frac{d\hat{f}}{d\sigma^2} = \frac{d\hat{f}}{db^*} \cdot \frac{db^*}{d\sigma^2} < 0. \tag{7}$$

The intuition behind Proposition 1 follows a standard result in the principal-agent literature: that risk adds noise, and greater noise results in (optimally) lower incentives for profit-maximizing managerial effort (see Holmstrom and Milgrom (1987) and Gibbons (1997)). In this case, the piece rate influences the degree of discriminatory behavior by the agent. Specifically, an increase in risk in the revenue generation process, σ^2, results in reduced managerial incentives, b, since increased noise makes it more difficult to distinguish profit-maximizing (non-discriminatory) behavior from random revenue shocks. Taken together, as suggested in Proposition 1, greater revenue uncertainty decreases the profit-maximizing piece rate, which in turn results in more discrimination. In other words, in an environment with a higher degree of risk, it is more difficult to attribute a good or bad realization of revenue to any deviation from the optimal allocation f^*. Consequently, since the manager receives utility from favoring male workers, he chooses an allocation more favorable to male employees when faced with greater revenue uncertainty.

This result has several implications for predicting the female labor share of the workforce as it suggests that female workers will be underrepresented in industries and plants that exhibit greater revenue and profit volatility, as managers have more leeway to indulge their discriminatory preferences.[12] Intuitively, this link between revenue volatility and the female share of the labor force arises from two channels. First, if managers are more likely to choose a more discriminatory allocation in situations of greater revenue volatility, that implies a discriminating manager will hire fewer female workers. Second, it implies that the manager will also be more likely to assign any female workers to the "less desirable"' tasks as well as assigning them the "less desirable" inputs. If female workers are aware that situations of greater uncertainty will result in a worse job environment (due to discriminatory behavior by the manager) they will self-select into firms or industries that exhibit greater stability and predictability (where the chosen allocation will be more favorable to female workers).[13]

It should be stressed that our proposed mechanism only applies to discriminatory behavior that directly reduces firm profits. Thus, our model is more about employment discrimination than wage discrimination. Specifically, simply paying female workers less for the same task does not reduce firm profits and thus would not be deterred by incentive contracts in a principal-agent framework. Indeed, wage discrimination would only be indirectly affected by incentive contracts through the fact that paying female workers lower wages could lead to the loss (or misallocation) of talented female employees and thus lower revenues (i.e. through its effect on employment discrimination). This distinction is discussed in more detail in the appendix.

A potential caveat to our proposed positive link between risk and discrimination is the potential that firm owners might systematically adopt "voluntary affirmative action" schemes in more risky environments.[14] Specifically, as discussed earlier, while many managerial actions (e.g. assigning female workers less desirable tasks or inputs) may be unobservable by firm owners, other actions (e.g. hiring few female workers) are more observable. This raises the potential for owners to engage in input-based incentive schemes to monitor and reward certain managerial actions that counter discriminatory behavior (e.g. rewarding the interviewing and hiring of female workers). As discussed in the principal-agent literature in the presence of multitasking (see Holmstrom and Milgrom (1991), Amuedo-Dorantes and de la Rica (2006), and Schnedler (2008)) the weight placed on different performance measures depends on the distortion and risk characteristics of the measure. Intuitively, firms do not want to incentivise performance measures with a lot of distortion (i.e. that distort managerial effort towards less productive actions) or a lot of noise (i.e. that provide little information about managerial productivity). Indeed, this perhaps explains why directly tying managerial pay to diversity targets is so rarely explicitly practiced. However, to the extent this distortion/noise trade-off systematically differs across firms with different risk environments, it could result in more volatile firms endogenously choosing to reward non-discriminatory actions which would mitigate our proposed link. Indeed, Prendergast (2002) argues that the decision to adopt output-based incentive pay (rewarding managers for firm performance) versus input-based pay (monitoring and rewarding managers for undertaking certain actions) depends precisely on the degree of risk and uncertainty in the firm's environment. Of course the main argument of Prendergast (2002) is that firms will rely more on output-based pay schemes and managerial discretion in more risky environments which would simply reinforce our proposed link between risk and discrimination. However, Prendergast (2002) also provides a situation in which input-based schemes would be favored by more risky firms (basically a situation where greater uncertainty results in input-based measures being better measures of managerial performance) and thus the overall link between risk and discrimination is not unambiguous. Thus, in the following empirical section, we test these two predictions: that industries and firms exhibiting more volatility also exhibit higher levels of managerial discrimination and thus lower participation by female workers.

3 Empirics

We introduce our data in Section 3.1. The main empirical prediction of the preceding section is that higher variability results in more discrimination, which we test by measuring variability at both the industry level and at the plant level. The unit of observation in all of our specifications is the plant. At the industry level, presented in Section 3.2, we test whether (plant-level) discrimination is more common in industries with higher variability in performance. Our plant-level results are presented in Section 3.3, in which we examine whether plants with more variable performance over time discriminate more. In Section 4, we check these results for robustness.

3.1 Data

Our data are from the plant-level dataset based on the Colombian Manufacturing Census by DANE (National Statistical Institute), described further in Roberts (1996) and in our Appendix A.2. The dataset includes data on all manufacturing plants in Colombia with at least ten employees.[15] Throughout our analysis, we drop approximately 2 % of the sample classified as "other" types of enterprise (e.g. collectives, cooperatives).

Table 1: Summary statistics: plant-level

	Mean (Std. dev.)	**Minimum**	**Maximum**
Female share	0.363 (0.266)	0	1
Exports	0.042 (0.147)	0	1
Firm age	18.16 (12.79)	0	91
Employment (log)	3.51 (1.09)	0	8.40
Salary (log)	6.86 (0.46)	−0.69	15.76
Skill ratio	0.270 (0.208)	0	1
Energy use (log)	4.66 (1.28)	0	9.84
Productivity (log)	8.18 (0.89)	2.76	12.73
Capital/labor (log)	6.56 (1.55)	−2.54	12.46
Office equipment	0.094 (0.126)	0	0.997
Corporation	0.188 (0.391)	0	1
Proprietorship	0.112 (0.315)	0	1
Partnership	0.700 (0.458)	0	1
Variability: revenue (millions)	0.043 (0.196)	0	8.925
Variability: net income (millions)	0.020 (0.178)	0	9.028
Variability: profit (millions)	0.050 (1.045)	0	67.48
Variability: employment (thousands)	0.001 (0.019)	0	1.466
Industry: 1 – (MS corr)	0.221 (0.224)	0	1.708
Industry: Abs MS chg	0.570 (0.185)	0.002	1.222
Industry: Mean, TR var	0.043 (0.061)	0.001	1.433
Industry: 1 – (wage corr)	0.399 (0.230)	0	1.995
Industry: 1 – (emp corr)	0.377 (0.244)	0.003	1.990
Observations	6,208		

Notes: See Appendix A.2 for variable definitions. All variables are measured in 1991, except the variability measures, which are the variance (of revenue, net income, or profits) over 1986–90 scaled by the mean of the same variable, also over 1986–90.

Table 1 presents summary statistics for our data in 1991. The average plant has 72 employees, 36 % of whom are female. On average, a firm exports about 4 % of its output; however, 5,001 firms (81 % of firms) do not export at all. The average exporting firm exports 22 % of its output. There are 29 3-digit industries in the sample; the most heavily represented are food products (ISIC 311, 14.4 %), clothing excluding shoes (322, 12.9 %) and metal products excluding machinery (381, 8.2 %). Our four measures of plant-level variability are in the middle of the panel of the table: these are constructed as the variance of the variable (revenue, net income or profits) over 1986–90 at the plant level, scaled by the (plant-level) mean of the variable.[16] The bottom panel of the table includes our industry-level measures of variability (the unit of observation remains the plant, although these measures are constructed at the 4-digit ISIC industry level). These measures include the industry-level correlation between market share in 1984 and 1991, the sum over the industry of the absolute value of changes in market share between 1984 and 1991, and the industry mean of the total revenue variability measure. The industry-level correlation is also constructed for average wage and employment. The correlations are all subtracted from one; for all of our variability measures, higher levels indicate more variability.

3.2 Empirics at the Industry Level

In Table 2, we present results from OLS regressions in which the dependent variable is the female share of workers; the unit of observation is the plant.[17] In this section, our focus is on the correlation between industry-level variability and firm-level discrimination. Specifically, we test the hypothesis that plants in more variable (or unpredictable) industries have lower female shares due either to discriminatory hiring decisions by firm managers or to self-selection by female workers (who are aware that more stable industries exhibit less discriminatory behavior).

Table 2: Plant-level female share with industry variability

	(1)		(2)		(3)		(4)	
1 – (corr MS)	−0.133*	(0.067)					−0.100*	(0.055)
abs MS chg			−0.109	(0.068)			−0.022	(0.070)
mean, TR var					−0.594*	(0.297)	−0.562**	(0.273)
exports	0.105***	(0.025)	0.103***	(0.026)	0.100***	(0.027)	0.105***	(0.026)
firm age	−0.000	(0.000)	−0.000	(0.000)	−0.000	(0.000)	−0.000	(0.000)
employment (log)	0.014**	(0.005)	0.014***	(0.005)	0.016***	(0.004)	0.016***	(0.004)
salary (log)	−0.043***	(0.008)	−0.042***	(0.008)	−0.041***	(0.009)	−0.042***	(0.009)
skill ratio	0.076*	(0.044)	0.077*	(0.044)	0.072	(0.042)	0.070	(0.043)
energy use (log)	−0.019***	(0.007)	−0.020***	(0.006)	−0.019***	(0.006)	−0.018**	(0.007)
K/L (log)	−0.008**	(0.004)	−0.008**	(0.004)	−0.007**	(0.003)	−0.007**	(0.003)
productivity (log)	−0.011*	(0.005)	−0.011*	(0.006)	−0.009**	(0.005)	−0.009**	(0.004)
office equip	0.035	(0.043)	0.039	(0.044)	0.035	(0.044)	0.035	(0.044)
corporation	−0.035***	(0.012)	−0.035***	(0.012)	−0.033***	(0.011)	−0.034***	(0.011)
proprietorship	−0.011*	(0.007)	−0.012*	(0.007)	−0.012*	(0.006)	−0.012*	(0.006)
concentration	−0.011	(0.103)	−0.003	(0.119)	0.085	(0.124)	0.041	(0.106)
Observations	(6,890)		(6,893)		(6,893)		(6,890)	

Notes: The dependent variable is the plant's female share of workers. All data are for 1991, except where indicated. 29 (3-digit) industry-level fixed effects are also included, and standard errors (in parentheses) are clustered at the 3-digit level. See Appendix A.2 for data definitions and notes. *** indicates statistical significance at the 99 % level or better, ** at the 95 % level, and * at 90 %.

Because plants may differ in their technology choices, and men and women may not be perfect substitutes for different types of technology, we control for many firm characteristics in addition to the variability measures of interest. Indeed, a previous literature has established that the characteristics of female-oriented plants are different from the typical plant (e.g. see Ozler (2000), Catagay and Berik (1991), and Ederington, Minier, and Troske (2010)). This is a concern if male (or female) labor is over-represented in production processes that are correlated with more stable firms/industries. Thus, we include a host of plant-level production characteristics including plant age, total employment, productivity and average wage of workers. We also include energy share, capital share, skilled labor share, export intensity, and office equipment share to control for technology differences across plants.[18] As can be seen in Table 2, and consistent with earlier results (e.g. see Ederington, Minier, and Troske (2010)) female labor shares are higher in plants that pay lower wages, are larger, less capital- and energy-intensive, and are more likely to export.

Our variables of interest are the industry-level measures of variability (note again that the unit of observation is the plant). Given the lack of a single, widely recognized measure of industry-level variability, we include three different measures of variability. First, we use the industry-level correlation between market share in 1984 and

1991, subtracted from one so that larger numbers reflect increased variability.[19] The correlation coefficient of market share is a common measure of industry mobility in the industrial organization literature, dating back to the work of Gort (1963), and is often seen as a proxy for the degree of predictability of firm performance over time within an industry. Our second measure of industry variability is the sum over the industry of the absolute value of changes in market share. The absolute value of market share changes is also a common measure of industry-level variability, dating back to Hymer and Pashigan (1962). Our final measure is the industry mean of plant-level variability in total revenue.[20] Since these measures of industry variability are at the 4-digit industry level, we include 3-digit industry fixed effects in each regression.

The results in Table 2 provide support to our model's prediction that discrimination is more likely in industries with more variability. In Regression 1, a higher female share in 1991 is negatively correlated with one minus the industry-level correlation between market share in 1984 and market share in 1991, suggesting that in industries with higher values of this measure (lower correlation in market shares), it is easier to discriminate, resulting in plants with lower female shares. In Regression 2, the measure of industry variability is the sum over an industry of the absolute value of changes in market share of the plants in that industry. This coefficient estimate is negative (in industries with larger changes in market share plants hire proportionally fewer female workers) although not statistically significant at conventional levels. Finally, in Regression 3 we use the industry mean of the variability of net income, which is also negatively correlated with plant-level female share. We include all of the measures of industry variability in Regression 4, and each remains negative with both the market share correlation and variability measures being statistically significant.

In addition, the magnitudes of the coefficient estimates are not trivial. Using the more conservative estimates in Regression 4, an increase of one standard deviation in the mean of total revenue variability (0.061) generates a decrease in the female share of 3.4% points; with an average female share of 36.0, this is slightly more than a 9 % decrease. The effects of a one-standard-deviation increase in the other measures is smaller, but both generate an estimated decrease in the female share of more than one percentage point.

3.3 Empirics at the Firm Level

Table 2 provides some evidence that plants in more volatile industries tend to hire fewer women. In this section, we examine how plant-level variability affects the female share of labor. We consider three different measures of plant-level volatility, each measured as the variance scaled by the mean, following our theory: total revenue, net income, and profits.

Table 3: Plant-level female share with plant-level variability

	(1)		**(2)**		**(3)**		**(4)**	
Var: revenue	−0.050***	(0.013)					−0.050**	(0.020)
Var: net income			−0.031**	(0.014)			−0.002	(0.016)
Var: profit					−0.000	(0.002)	0.001	(0.001)
exports	0.102***	(0.022)	0.101***	(0.022)	0.100***	(0.021)	0.102***	(0.022)
firm age	−0.000*	(0.000)	−0.000*	(0.000)	−0.000*	(0.000)	−0.000*	(0.000)
employment (log)	0.014***	(0.005)	0.013***	(0.004)	0.012***	(0.004)	0.014***	(0.005)
salary (log)	−0.041***	(0.008)	−0.040***	(0.008)	−0.041***	(0.008)	−0.041***	(0.008)
skill ratio	0.059	(0.042)	0.059	(0.042)	0.060	(0.042)	0.059	(0.042)
energy use (log)	−0.017***	(0.006)	−0.017***	(0.006)	−0.017***	(0.006)	−0.017***	(0.006)
capital-labor (log)	−0.004	(0.002)	−0.004	(0.003)	−0.004	(0.002)	−0.004	(0.002)
productivity (log)	−0.005	(0.004)	−0.007	(0.004)	−0.008*	(0.004)	−0.005	(0.004)
office equip	0.028	(0.039)	0.029	(0.039)	0.029	(0.039)	0.028	(0.039)
corporation	−0.021**	(0.009)	−0.022**	(0.009)	−0.022**	(0.009)	−0.021**	(0.009)
proprietorship	−0.016**	(0.007)	−0.017**	(0.007)	−0.017**	(0.007)	−0.016**	(0.007)
Observations	6,208		6,208		6,208		6,208	

Notes: The dependent variable is the plant's female share of workers. All data are for 1991, except the scaled variances are the variance over 1986–90 for all plants with at least two years of non-missing data. 95 4-digit industry-level fixed effects are also included, and standard errors are clustered at the industry level. See Appendix A.2 for data definitions and notes. *** indicates statistical significance at the 99 % level or better; ** at 95 %; and * at 90 %.

The coefficient estimates on the control variables remain very similar to those in Table 2. Our main focus, however, is on the correlation between the female share of labor and the measures of firm variability: the variance of revenue, net income, and profits (each scaled by the mean of the same variable). As predicted, the coefficient estimates on these variables are negative, and highly statistically significant in the case of revenue and net income.[21] This supports our theory that plants with high income variability are more likely to discriminate by hiring fewer female workers. When we include all three measures of plant-level variability together, only the revenue variability measure remains statistically significant; the magnitude is only slightly smaller than the estimates in Table 2.

4 Discussion

The results presented in Table 2 and Table 3 support the main prediction of our model that the female labor share of a firm is decreasing in the variability of firm (or industry) performance. We interpret this correlation as reflecting the fact that managers have greater incentives to engage in non-discriminatory behavior when firm performance is more stable and predictable. Thus, managers in more stable firms (or industries) will hire more female workers, and female workers will be more likely to self-select into such firms (or industries).

However, a second possibility is that female workers self-select into more stable firms and industries, not because stability reduces discriminatory actions on the part of managers, but because women prefer more stability in their employment situation. Indeed, there is a large literature in economics establishing that women are more risk-averse than men in their preferences (see, for example, Croson and Gneezy (2009)). Needless to say, it is difficult to distinguish between different motivations for self-selection in the labor market. However, in our context, if self-selection were driven by risk preferences, it seems reasonable to assume that a female employee would care more about the variability of wages or employment – which are more likely to directly affect her – than about variability in output, market share, or revenue of the firm. Likewise, as our simple model establishes, it is precisely variability in firm performance (i.e. output, market share and revenue) that influences the degree of discriminatory behavior by firm managers. Thus, to address this distinction, we constructed firm-level scaled variance measures of employment and average unskilled wages analogous to the firm-level scaled variance measures used in Table 3. [22]

In Regression 1 of Table 4, we present results of the employment measure; results with the wage measure are similar (the coefficient estimate on wage variability is also negative and statistically insignificant). In Regression 2, we use the industry-level correlation between plants' average unskilled wage in 1984 and average unskilled wage in 1991; in Regression 3, we construct this correlation for total employment. As before, we subtract the correlation from one so that a higher (absolute) value reflects greater variability. While the coefficient estimate on the wage correlation measure remains statistically not different from zero, the coefficient estimate on the employment measure is negative and statistically significant. This is also true in regression (4) when all measures are included together. Industries in which the correlation between (plant-level) employment in 1984 and employment in 1991 is higher (i.e. employment is more stable) have higher female shares, controlling for other determinants of female employment shares. Thus, although the evidence is not as strong as it was for our measures of firm performance, we do find evidence consistent with women selecting into more "stable" industries as measured by stability in firm employment.

Table 4: Plant-level female share with wage, employment, and income volatility

	(1)		(2)		(3)		(4)	
Var: emp	−0.036	(0.029)					−0.034	(0.033)
1 – (wage corr)			−0.007	(0.078)			0.016	(0.073)
1 – (emp corr)					−0.057	(0.039)	−0.058*	(0.031)
Var: TR	−0.050***	(0.013)	−0.066***	(0.019)	−0.066***	(0.019)	−0.065***	(0.019)
exports	0.102***	(0.022)	0.111***	(0.025)	0.115***	(0.028)	0.114***	(0.025)
firm age	−0.000*	(0.000)	−0.000	(0.000)	−0.000*	(0.000)	−0.000*	(0.000)
employment (log)	0.014***	(0.005)	0.016***	(0.005)	0.015***	(0.005)	0.016***	(0.005)
salary (log)	−0.041***	(0.008)	−0.042***	(0.010)	−0.042***	(0.010)	−0.043***	(0.010)
skill ratio	0.059	(0.042)	0.070	(0.044)	0.068	(0.043)	0.069	(0.043)
energy use (log)	−0.017***	(0.006)	−0.027***	(0.006)	−0.027***	(0.006)	−0.027***	(0.006)
capital-labor (log)	−0.004	(0.002)	−0.006*	(0.003)	−0.006*	(0.003)	−0.006*	(0.003)

productivity (log)	−0.005	(0.004)	−0.007	(0.007)	−0.007	(0.007)	−0.007	(0.007)
office equip	0.028	(0.039)	0.037	(0.043)	0.037	(0.043)	0.036	(0.043)
corporation	−0.021**	(0.009)	−0.026**	(0.011)	−0.026**	(0.011)	−0.026**	(0.011)
proprietorship	−0.016**	(0.007)	−0.011	(0.006)	−0.013*	(0.006)	−0.012*	(0.006)
Observations	6,208		6,196		6,206		6,196	

Notes: The dependent variable is the plant's female share of workers. All data are for 1991, except where indicated. The industry fixed effects depend on the variability measures included; variability of employment and total revenue are measured at the plant level, so 4-digit industry effects are included in Regression 1. Employment and wage correlations are measured at the 4-digit industry level, so 3-digit industry effects are included in Regressions 2, 3, and 4. Standard errors are clustered at the respective industry level. Variability of revenue is measured at the plant level in all regressions. See Appendix A.2 for data definitions and notes. *** indicates statistical significance at the 99 % level or better; ** at 95 %; and * at 90 %.

While, we do find some evidence that the female labor share is higher in more "employment stable" industries, the question is whether this is driving our results that women are over-represented in more "revenue-stable" industries. Thus, we also include the firm revenue variability measure in all regressions. If selection driven by risk aversion about employment were driving our results, we would expect to see this coefficient decrease in magnitude. However, it does not: even in Regressions 3 and 4, where both employment correlation and revenue variability are included, the coefficient estimate on net income variability remains negative and statistically significant. The magnitudes of the employment (selection) effect and the revenue effect are very comparable: a one standard deviation increase in the employment correlation measure (0.240) implies a 1.4% point decrease in female share, while a one standard deviation increase in plant-level revenue variance (0.196)[23] implies a decrease of 13% points in the female share.

It should be noted that a third possibility is that, due perhaps to female workers being more risk-averse, the female labor share could be directly correlated with greater firm stability directly due to less risky behavior by workers. Such a direct link would not be captured by either our approach of including firm control variables or controlling for measures of employment stability. Thus, while our empirical results are suggestive of our proposed link (that a risky environment provides discriminatory managers more leeway to engage in discriminatory behavior), we can not necessarily rule out all alternative explanations.

5 Conclusion

As a complement to the traditional Becker framework, we derive a model where discriminatory behavior arises, not from the preferences of firm owners, but from the preferences of firm managers whose actions can only be imperfectly observed by firm owners. We show that a natural prediction arising from this principal-agent framework is that discriminatory behavior is more prevalent in situations of greater uncertainty. This prediction is a function of the well-known theory that greater noise in the firm generation process reduces incentives for profit-maximizing behavior by managers. In the context of gender discrimination, we utilize a Colombian plant-level dataset and show that, consistent with this prediction, a negative correlation exists between female labor share and various measures of industry-level and plant-level volatility.

The prediction we derive and investigate about the link between risk and discrimination obviously parallels the long assumed link between risk and managerial effort. In a sense this is not surprising as discriminatory behavior can be considered simply another component of managerial effort as they both involve a direct tradeoff between a manager's preferences and firm performance. There exists a large and growing literature on managerial incentives and managerial effort and it seems a potentially productive line of research is to consider what lessons can be drawn from this literature in understanding the conditions under which discriminatory behavior is more prevalent.[24] We view this paper as a first step in this direction.

Acknowledgements

We are grateful to Ken Troske, two anonymous referees and seminar participants at the Unversité Libre de Bruxelles for many comments and suggestions, as well as to Mark Roberts for making the Colombian manufacturing data available to us.

Appendix

A

A.1 Hiring Model

As an application, consider the case where the principal (the owner) assigns the task of interviewing and hiring workers to the agent (upper management). Initially, the firm has no workers other than the agent. Each period, the agent is faced with an exogenous pool of applicants consisting of i_m male applicants and i_f female applicants. The profit $\pi_i = r_i - w_i$ earned for the principal by any worker, i, (male or female) is the result of an *iid* draw from a uniform distribution with support $[\pi_L, \pi_H]$ where r_i is the amount of revenue generated by the agent while w_i is the agent's wage demand.[25] The agent observes each applicant's expected profitability draw before making the hiring decision.

The agent chooses whether or not to hire each applicant to maximize his expected utility. The agent has exponential utility $u(x) = -\exp(-rx)$, where x is the agent's net wage and $r > 0$ is the agent's coefficient of absolute risk aversion. The net wage is determined by $x = w - \alpha(f)$. The variable w is the wage paid to the agent by the principal. The number of female applicants hired by the agent is denoted by f, and $\alpha(f)$ represents the agent's taste for discrimination, or the disutility from having more female workers in the firm. Purely for tractability, we assume that $\alpha(f) = \alpha f$, where $\alpha > 0$.

The principal wants to incentivise the agent to hire profitable workers, and so the wage paid to the agent, $w = s + bR$ consists of both a salary, s, and a piece rate component, b. The principal will choose s and b to maximize expected profit.

Since the agent is discriminatory, he faces a tradeoff. Rather than hiring to maximize his expected wage, he prefers to hire fewer female workers, even if they exhibit higher expected profitability than male applicants, in order to decrease the number of female workers. However, he may be unwilling to pass up too many high-productivity females, since this would reduce his wage more significantly. Thus, we model the agents decision as choosing minimum productivity cut-offs for both male (π_m) and female applicants (π_f) such that he hires male applicants as long as $\pi \geq \pi_m$ (which happens with probability $\frac{\pi_H - \pi_m}{\pi_H - \pi_L}$), and female applicants as long as $\pi \geq \pi_f$ (which happens with probability $\frac{\pi_H - \pi_f}{\pi_H - \pi_L}$).

In this case, the expected (realized) profitability to the firm of a female-applicant is:

$$E[\pi_f] = \frac{\pi_H - \pi_f}{\pi_H - \pi_L} \frac{\pi_H + \pi_f}{2}$$

with the expected profitability of a male applicant ($E[\pi_m]$) being defined symmetrically (i.e. the product of the probability the applicant is at least the minimum type and the expected profitability of the applicant conditional on at least being the minimum type). While the variance of these profits is given by:[26]

$$\mathrm{Var}[\pi_f] = \frac{4(\pi_H^3 - \pi_f^3)}{12(\pi_H - \pi_L)} - \frac{3(\pi_H + \pi_f)^2(\pi_H - \pi_f)^2}{12(\pi_H - \pi_L)^2}$$

Thus, the central limit theorem suggests that firm revenue can be approximated by $R = y(\pi_m, \pi_f)) + \varepsilon$, where $y(\pi_m, \pi_f) = i_m E[\pi_m] + i_f E[\pi_f]$ is the total expected profit from all the workers and ε is an idiosyncratic shock drawn from a normal distribution $\phi(\varepsilon)$with zero mean and variance $\sigma^2 = i_m Var[\pi_m] + +i_f Var[\pi_f]$. The principal's profit, then, can be represented by $R - w$, where w is the wage paid to the agent. A complication arises from the fact that variance of firm profits is also affected by the agent's decisions (i.e. the choice of the profitability cut-offs for applicants). To solve the model, we impose the assumption that the affect of the agent's decisions on the variance of firm profits is *de minimus* and thus is ignored by the agent in making his decisions.[27]

Since the expected number of females hired is $f = i_f(\frac{\pi_H - \pi_f}{\pi_H - \pi_L})$, the agent's expected utility is given by

$$\int_\varepsilon \exp\left[-r\left(s + b(y + \varepsilon) - \alpha i_f\left(\frac{\pi_H - \pi_f}{\pi_H - \pi_L}\right)\right]\phi(\varepsilon)d\varepsilon. \tag{8}$$

We first maximize (8) wrt π_m and π_f.

$$\frac{\partial EU}{\partial \pi_m} = b\frac{\partial y}{\partial \pi_m} = 0 \tag{9}$$

$$\frac{\partial EU}{\partial \pi_f} = b\frac{\partial y}{\partial \pi_f} - \alpha\frac{\partial f}{\partial \pi_f} = 0. \tag{10}$$

From (9), we find

$$\pi_m = 0. \tag{11}$$

From (10), we have

$$\pi_f = \frac{\alpha}{b}. \tag{12}$$

Since $b > 0$, $\pi_f > \pi_m$. Furthermore, the second-order conditions are satisfied so that π_f and π_m are interior solutions. Thus, the agent hires a male as long as he is a non-negative profit type, whereas a hired female must earn a strictly positive profit for the principal. In addition, in (12), it is clear that the agent acts less discriminatory (π_f decreases) as the principal offers a higher piece rate (b) and acts more discriminatory as the marginal disutility of hiring a female (α) increases.

The above model makes clear that we are only modeling a link between incentive contracts and *employment discrimination* in which discrimination results in potentially profitable hiring of some female workers (for both the worker and the firm) not occurring. This is due to the structure of the model in which the worker makes a take-it-or-leave-it wage offer and the only decision by the manager is the hiring decision. Alternatively, one could consider a model of *wage discrimination* in which the worker and manager engage in wage bargaining and the manager systemically makes lower offers (or bargains harder) with female potential workers. What is interesting is that, to the extent that the manager can successfully hire female workers for lower wages, such wage discrimination would not necessarily be deterred by incentive contracts as it would result in higher profits for the firm. Indeed, this behavior (lower offers to female workers) would only be discouraged to the extent that it resulted in potentially profitable female workers not accepting the position (i.e. to the extent that it resulted in lower than optimal female-shares of the labor force). Thus, we see our model being more about employment discrimination than wage discrimination.

The agent's certainty equivalent (CE) is the wage received less the disutility from hiring females, and the cost of bearing risk:

$$CE = s + by - \alpha i_f\left(\frac{\pi_H - \frac{\alpha}{b}}{\pi_H - \pi_L}\right) - \frac{1}{2}rb^2\sigma^2. \tag{13}$$

The principal's expected profit is $(1-b)y - s$. Since the principal has full bargaining power, maximizing profit is equivalent to maximizing efficiency, or total surplus, which is the firm's expected profit less the agent's CE:

$$TS = y - \alpha i_f\left(\frac{\pi_H - \frac{\alpha}{b}}{\pi_H - \pi_L}\right) - \frac{1}{2}rb^2\sigma^2.$$

Substituting for y, total surplus is:

$$TS = \frac{i_m\pi_H^2 + i_f(\pi_H^2 - (\frac{\alpha}{b})^2) - 2\alpha i_f(\pi_H - \frac{\alpha}{b})}{2(\pi_H - \pi_L)} - \frac{1}{2}rb^2\sigma^2. \tag{14}$$

The principal chooses b to maximize (14). The optimal piece rate is the value of b for which

$$\frac{i_f\alpha^2}{(\pi_H - \pi_L)}[1-b] - r\sigma^2 b^4 = 0. \tag{15}$$

An interior solution is ensured because the necessary second order condition is satisfied:

$$\frac{i_f\alpha^2}{(\pi_H - \pi_L)} - 4b^3 r\sigma^2 < 0. \tag{16}$$

The optimal salary component of the wage, s^*, is chosen to satisfy the agent's participation constraint determined by setting his CE as defined in (13) equal to zero.

Total differentiation of (15) with respect to b and σ^2 gives

$$\frac{db}{d\sigma^2} = \frac{rb^4}{\frac{i_f\alpha^2}{(\pi_H - \pi_L)} - 4b^3 r\sigma^2} < 0. \tag{17}$$

The numerator of (17) is positive, and the denominator is negative by (16). Consequently, we find that an increase in risk (σ^2) results in a smaller piece rate. From (12) this implies that:

$$\frac{d\pi_f}{d\sigma^2} > 0. \tag{18}$$

Hence, an increase in risk (σ^2) ultimately results in greater discrimination (an increase in π_f. Intuitively, an environment with greater profit uncertainty provides discriminating managers less incentive to moderate their tendency to discriminate. In contrast, when the signal about a worker's productivity has less noise, then a discriminating manager has more incentive to reduce discrimination (to increase his own wage).

A.2 Data

All data are taken from a plant-level dataset produced from the Colombian Manufacturing census by DANE (National Statistical Institute) for the years 1977 through 1991. Starting in 1983, the census covers industrial production for plants with more than 10 employees. Our empirics are a cross-section of plants operating in 1991, although earlier years are used to construct some of the variables, as noted below. For a thorough description of this dataset see Roberts (1996).

All variables are measured at the plant level unless otherwise noted.

Female share: female share of workers.

Productivity: value added for the plant divided by total employment.

Employment: total employment.

Salary: total payroll divided by total employment.

Wage: Average unskilled wage is calculated by total salary and benefits of unskilled workers divided by total number of unskilled workers.

Firm Age: years since the plant's establishment until 1984.

Exports: plant exports scaled by total sales.

Skill Ratio: share of skilled employment in skilled and unskilled employment.

Capital/Labor Ratio: ratio of fixed capital to total employment. A small number of plants with fixed capital reported as zero are dropped.

Energy Use: one plus the ratio of energy consumed to total employment.

Office Equipment: office equipment's share of total capital equipment.

Variability: Revenue: the variance of total revenue scaled by its mean for the plant over 1986–90, for all plants with at least two years of non-missing data.

Variability: Net income: the variance of net income scaled by its mean for the plant over 1986–90, for all plants with at least two years of non-missing data.

Variability: Profit: the variance of profit scaled by its mean for the plant over 1986–90, for all plants with at least two years of non-missing data.

Market Share Correlation: the correlation between plants' market share in 1984 and market share in 1991, at the 4-digit industry level. Any plant present in only one year is treated as having a market share of zero in the missing year.

Absolute Value of Market Share Changes: the sum of the absolute value of changes in market share from 1984 to 1991, at the 4-digit industry level. Any plant present in one year is treated as having a market share of zero in the missing year.

Mean, Revenue Variability: the mean of the revenue variability measure for the 4-digit industry.

Type of Enterprise: The data set classifies plants by 10 different enterprise types. We omit firms classified as collectives, cooperatives, official entities, and religious communities (overall, these comprise less than 2 % of the sample). We construct dummy variables for *Corporations* (this includes plants classified as corporations, de facto corporations, and joint stock companies), *Proprietorships*, and *Partnerships* (including limited partnerships and joint partnerships).

Notes

1 This simple prediction has generated a large empirical literature on whether discrimination is less prevalent under conditions of greater competition. For examples of studies in this literature see Oster (1975), Ashenfelter and Hannan (1986), Black (1995), Hellerstein, Neumark, and Troske (2002), and Black and Brainerd (2004).

2 There has been some work incorporating the Becker model into an equilibrium search model framework; see, e.g. Black (1995), Bowlus (1997), and Bowlus and Eckstein (2002).

3 Indeed, Méon and Szafarz (2011) provide an explicit principal-agent model in which managers choose between two competing candidates for a position. They demonstrate that performance-based contracts will be unable to fully eliminate managerial discrimination. Our contribution is to show how a similar principal-agent approach can lead to a link between risk (or randomness in plant performance) and the degree of discrimination.
4 For a demonstration of this see the field experiment conducted in Bandiera, Barankay, and Rasul (2009). For an agricultural firm they exogenously switched managerial pay from a fixed wage to a performance pay scheme, and found that managers switched from favoring socially connected workers to high-productivity workers.
5 This result is of course related to the well-known tradeoff between risk and incentives in the principal agent-literature (i.e. that the pay-performance sensitivity is decreasing in the variance of firm returns). For empirical evidence on that relation see, for example, Garen (1994) and Aggarwal and Samwick (1999).
6 The counter argument is that even merit-based pay to employees can be subject to the subjective opinions (see Elvira and Towns (2001) and Castilla (2008)) as well as the differential assignment of inputs (see Madden (2012) by a discriminatory supervisor.
7 A long-standing empirical result in the gender wage gap literature is that female workers are segregated into lower-paying firms (e.g. see Carrington (1998), Bayard et al. (2003), Amuedo-Dorantes and de la Rica (2006) and Javdani (2015)). Indeed, our regressions include average firm wages as a control variable and we find a similar correlation in that female workers are underrepresented in higher-paying firms. However, our results suggest an additional correlation in that female workers appear to be segregated into more stable firms as well.
8 In the appendix, we provide an example of a more concrete model that generates a similar revenue function.
9 Note that our model assumes that the allocation of workers is correlated with average profits but uncorrelated with the variance of profits. Thus, we abstract away from the possibility that, to the extent that female workers engage in less risky behavior, the variance of firm profits is endogenously linked to the female share/allocation of the work force.
10 Note that while some discriminatory actions by the manager may be more observable to owners (e.g. hiring fewer female employees) other actions will be less observable (e.g. assigning female employees to the "less desirable" tasks). This raises the potential for firms to engage in voluntary affirmative action by directly subsidizing the manager for engaging in observable actions that benefit female employees (i.e. subsidizing actions that affect f directly). For now we will maintain the standard assumption that the allocation, f, is unobservable and will discuss later the case where some managerial actions may be observable.
11 We focus on linear wage contracts given their prevalence in the principal-agent literature (see Gibbons (1997)). In addition, Holmstrom and Milgrom (1987) show that such linear contracts are optimal in a standard class of dynamic moral hazard problems.
12 Our model best fits the situation of a single manager of a plant. However, given the diversity of different managerial organizations within plants, our general empirical approach contains the implicit assumption that greater volatility at the industry/plant level results in greater risk at the managerial level.
13 Of course, another explanation is that females self-select strictly due to exogenous risk preferences. This possibility is examined in Section 4; we find that our results are generally robust when controlling for these factors.
14 Of course, another possibility is the the owner has discriminatory preferences in which case our results, which rely on the the principal agent framework, would not necessarily apply. An interesting line of future research would consider the case where the owner has discriminatory preferences while the agent is a profit-maximizer.
15 Prior to 1982, all plants were included.
16 Only plants with at least two years of data over the 1986–90 period are included. There are 448 plants with negative means for profit over this period; 59 of these also have negative means for net income. The absolute value of the mean is used to ensure that the measure captures variability. Results are robust to dropping these observations.
17 These results are with the overall share of female workers, including unskilled, skilled, and managerial employees, as the dependent variable. Results are very similar when the dependent variable is the female share of only unskilled labor; the coefficient estimates on the variability measures remain negative and highly statistically significant, and increase in absolute value.
18 Note that, at least according to our model, average productivity and wages of the firm are likely to be endogenously correlated with the female share of the labor force which could result in our estimates being inconsistent. Thus, we also ran specifications with these control variables omitted (as well as adding them in sequentially). Results were unchanged.
19 In constructing this correlation, we retain plants that enter or exit, giving them a market share of zero for the missing period.
20 This measure is defined as the variance of total revenue for the plant over the period 1986–91, scaled by the mean. Results are comparable when we use variability of profit or net income.
21 Note that a secondary prediction of our model is that firm profits should be higher in the more stable industries/firms given the greater incentive for managers to engage in profit-maximizing behavior. While we do find support for this prediction in some specifications, the results are not at all robust and vary widely depending on the measures of both variability and profitability.
22 While we have data on the numbers of male and female employees, the wage data are only given as a total by type of employment, and not broken down into male/female wages.
23 Not surprisingly, the standard deviation is higher when this variance is at the plant level, rather than the industry level.
24 For example, since Becker (1957) there has been an implicit assumption that discriminatory behavior by firms is less likely in more competitive environments. However, in the principal-agent literature, the theoretical link between product market competition and managerial effort is often seen as ambiguous (e.g. see the reviews in Nickell (1996) and Schmidt (1997)). Thus, if one views discrimination as a principal-agent problem, the link between competition and discriminatory behavior may be more ambiguous than is commonly believed.
25 The common support assumption means that either males and females are equally able and equally costly, or it could mean that members of one group have the ability to to earn higher revenues but are also proportionally more costly to acquire.
26 The variance is calculated by noting that the $Var(\pi_f) = E(\pi_f^2) - (E(\pi_f))^2$.
27 One can show that as the spread between π_L and π_H widens, that $\frac{\partial \sigma^2}{\partial \pi_f}$ tends to zero.

References

Aggarwal, R., and A. Samwick. 1999. "The Other Side of the Trade-Off: The Impact of Risk on Executive Compensation." *The Journal of Political Economy* 107 (1): 65–105.

Amuedo-Dorantes, C., and S. de la Rica. 2006. "The Role of Segregation and Pay Structure on the Gender Wage Gap: Evidence from Matched Employer-Employee Data for Spain." *The B.E. Journal of Economic Analysis and Policy* 5: 1–34.

Ashenfelter, O., and T. Hannan. 1986. "Sex Discrimination and Product Market Competition: The Case of the Banking Industry." *The Quarterly Journal of Economics* 101 (1): 149–74.

Baker, G. 2002. "Distortion and Risk in Optimal Incentive Contracts." *The Journal of Human Resources* 37: 728–51.

Bandiera, O., I. Barankay, and I. Rasul. 2009. "Social Connections and Incentives in the Workplace: Evidence from Personnel Data." *Econometrica* 77: 1047–94.

Bayard, K., J. Hellerstein, D. Neumark, and K. Troske. 2003. "New Evidence on Sex Segregation and Sex Differences in Wages from Matched Employee-Employer Data." *The Journal of Labor Economics* 21: 887–922.

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

Black, D. 1995. "Discrimination in an Equilibrium Search Model." *Journal of Labor Economics* 13: 309–34.

Black, S. E., and E. Brainerd. 2004. "Importing Equality? The Impact of Globalization on Gender Discrimination." *Industrial and Labor Relations Review* 57 (4): 540–59.

Black, S., and P. Strahan. 2001. "The Division of Spoils: Rent-Sharing and Discrimination in a Regulated Industry." *American Economic Review* 91: 814–31.

Bowlus, A. J. 1997. "A Search Interpretation of Male-Female Wage Differentials." *Journal of Labor Economics* 15: 627–57.

Bowlus, A. J., and Z. Eckstein. 2002. "Discrimination and Skill Differences in an Equilibrium Search Model." *International Economic Review* 43: 1309–45.

Carrington, W., and K. R.Troske. 1998. "Sex Segregation in U.S. Manufacturing." *Industrial and Labor Relations Review* 51: 445–64.

Castilla, E. 2008. "Gender, Race and Meritocracy in Organizational Careers. *American Journal of Sociology* 113: 1479–526.

Catagay, N., and G. Berik. 1991. "Transition to Export-Led Growth in Turkey: Is There a Feminization of Employment?" *Review of Radical Political Economics* 22: 115–34.

Croson, R., and U. Gneezy. 2009. "Gender Differences in Preferences." *Journal of Economics Literature* 47: 1–27.

de la Rica, S., J. Dolado, and R. Vegas. 2015. "Performance Pay and the Gender Wage Gap: Evidence from Spain." *The Annals of Economics and Statistics* 41–59.

Ederington, J., J. Minier, and K. R. Troske. 2010. "Where the Girls are: Trade and Labor Market Segregation." Manuscript.

Elvira, M., and R. Towns. 2001. "The Effects of Race and Worker Productivity on Performance Evaluations." *Industrial Relations* 40: 571–90.

Fang, T., and J. Heywood. 2006. "Output Pay and Ethnic Wage Differentials: Canadian Evidence." *Industrial Relations* 45: 173–93.

Garen, J. 1994. "Executive Compensation and Principal-Agent Theory." *Journal of Political Economy* 102: 1175–99.

Gibbons, R. 1997. "Incentives and Careers in Organizations." In *Advances in Economics and Econometrics: Theory and Applications*, edited by D. Kreps and K. Wallis. Cambridge, United Kingdom: Cambridge University Press.

Gort, M. 1963. "Analysis of Stability and Change in Market Shares." *Journal of Political Economy* 71 (1): 51–63.

Hellerstein, J., D. Neumark, and K. Troske. 2002. "Market Forces and Sex Discrimination." *Journal of Human Resources* 37: 353–80.

Heywood, J., and U. Jirjahn. 2002. "Payment Schemes and Gender in Germany." *Industrial and Labor Relations Review* 56: 44–64.

Heywood, J., and P. O'Halloran. 2005. "Racial Earnings Differentials and Performance Pay." *Journal of Human Resources* 40: 435–52.

Holmstrom, B., and P. Milgrom. 1987. "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica* 55: 303–28.

Holmstrom, B., and P. Milgrom. 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership and Job Design." *Journal of Law, Economics and Organization* 7: 24–52.

Hymer, S., and P. Pashigan. 1962. "Turnover of Firms as a Measure of Market Behavior." *The Review of Economics and Statistics* 44 (1): 82–87.

Javdani, M. 2015. "Glass Ceilings or Glass Doors: The Role of Firms in Male-Female Wage Disparities." *Canadian Journal of Economics* 48: 529–60.

Jirjahn, U., and G. Stephan. 2004. "Gender, Piece Rates and Wages: Evidence from Matched Employer-Employee Data." *Cambridge Journal of Economics* 28: 683–704.

Madden, J. 2012. "Performance-Support Bias and the Gender Pay Gap Among Stockbrokers." *Gender and Society* 26: 488–518.

Méon, P.-G., and A. Szafarz. 2011. "The Modern Corporation as a Safe Haven for Taste-Based Discrimination: An Agency Model of Hiring Decisions." *Labour Economics* 18: 487–97.

Nickell, S. 1996. "Competition and Corporate Performance." *Journal of Political Economy* 104: 724–46.

Oster, S. 1975. "Industry Differences in the Level of Discrimination Against Women." *Quarterly Journal of Economics* 89: 215–29.

Ozler, S. 2000. "Export Orientation and Female Share of Employment: Evidence from Turkey." *World Development* 28 (7): 1239–48.

Prendergast, C. 2002. "The Tenuous Trade-Off Between Risk and Incentives." *The Journal of Political Economy* 110: 1071–102.

Roberts, M. J. 1996. "*Colombia, 1977-85: Producer Turnover, Margins, and Trade Exposure." In Industrial Evolution in Developing Countries: Micro Patterns of Turnover, Productivity, and Market Structure*, edited by M. J. Roberts, and J. R. Tybout,, 227– 259. Oxford, United Kingdom: Oxford University Press.

Schmidt, K. 1997. "Managerial Incentives and Product Market Competition." *Review of Economic Studies* 64: 191–213.

Schnedler, W. 2008. "When is it Foolish to Reward for a While Benefiting from B." *Journal of Labor Economics* 26: 595–619.