



TASTES FOR DISCRIMINATION IN MONOPSONISTIC LABOUR MARKETS

A grayscale photograph of a large, multi-story building with a classical facade, featuring arched windows and a prominent entrance. The image is slightly tilted and serves as a background for the author's name and paper information.

BERNARDO FANFANI

Working paper No. 54 - May 2018

Tastes for Discrimination in Monopsonistic Labour Markets*

Bernardo Fanfani[†]

(University of Torino)

Abstract

We study a model where wage differences between men and women arise from taste-based discrimination and monopsonistic mechanisms. We show how preferences against women affect heterogeneity in firms' pay policies in the context of an imperfect labour market, deriving a rigorous test for the presence of taste-based discrimination and of other firm-level mechanisms driving the gender wage gap, in particular compensating wage differentials. These results inform an analysis of sex pay differences in the Italian manufacturing sector showing that taste-based discrimination and preferences for workplaces providing more flexible schedules are two significant determinants of the gender wage gap.

JEL Codes: J00, J16, J31, J71.

Keywords: Gender Wage Gap; Taste-Based Discrimination; Monopsonistic Discrimination; Compensating Wage Differentials; Firm Wage Policy; Matched Employer-Employee Data.

*I would like to thank Fabio Berton, David Card, Francesco Devicienti, Ignacio Monzon, Ronald Oaxaca, Giovanni Sulis and Andrea Weber for their excellent comments and the support they have provided me with. I also would like to thank for their comments participants at Collegio Carlo Alberto and UC Berkeley seminars and at the 2018 IAAEU Workshop on Labour Economics (Trier). Part of this work has been written during a visiting period spent at the Department of Economics of UC Berkeley, which I would like to thank for the hospitality. I would like to acknowledge funding support received by Collegio Carlo Alberto and through the PROWEDEC (Productivity, Welfare and Decentralized Bargaining) project, financed by the Compagnia San Paolo.

[†]University of Torino, Dipartimento di Scienze Scienze Economico-sociali e Matematico-Statistiche, corso Unione Sovietica 218b, 10133 Torino, Italy. Email: bernardo.fanfani@unito.it

1 Introduction

Gender wage gaps are one of the most persistent economic regularities, on which many hypotheses have been formulated and, at least since the seminal work by Oaxaca [1973], many regression approaches have been proposed. In this paper we combine elements of several of the existing theories, building a model where differences between men and women are determined by two main factors:¹ *Becker-type* (or so-called taste-based) and *Robinsonian* (or so-called monopsonistic) discrimination. Based on this theoretical framework, we develop an empirical approach that allows to test for the presence of taste-based discrimination and of other firm-level mechanisms driving the gender pay gap, in particular compensating wage differentials.

According to the theory of Becker [1957], taste-based discrimination arises because some employers have a dis-utility in working with women, so that either they are able pay them less than their productivity, or they avoid hiring them, reducing the aggregate female labour demand. As a consequence, firms where discriminatory preferences are small enough can employ a given quantity of female workers at a lower wage than the one needed to hire the same quantity of men. Instead, *Robinsonian* discrimination is a mechanism arising when firms have monopsonistic power in the labour market. If the assumption of price taking behaviour is relaxed, employers minimize costs not only on the extensive margin, by adjusting quantities, but also on the intensive margin, by adjusting wages. In this context, according to the *Robinsonian* discrimination hypothesis, gender wage differences are driven by employers' greater monopsonistic wage-setting power against women, provided that, on average, the female labour supply to the firm is more rigid than the male one.²

As shown in this paper, a model of taste-based discrimination in which employers can set both, wages and employment levels, provides interesting insights on the nature of gender differences in *firms' wage policies*. Such policies represent pay heterogeneity across em-

¹See Blau and Kahn [2017] for a recent literature review on the main theories and existing evidences on the gender wage gap.

²The original model of monopsony dates back to the 1930s (Robinson [1933]), but interest on the relationship between the labour market structure and gender discrimination has emerged only more recently (see Boal and Ransom [1997] and Manning [2003]).

ployers conditional on workforce composition, and substantial gender differences in this wage component have been documented for several countries by a recent and growing literature (see in particular Card et al. [2016], Sin et al. [2017] and Bruns [2018]). We show that, if tastes against women arise in the context of an imperfect labour market, even highly discriminatory employers hire female workers, but they offer them lower wages to compensate for the dis-utility associated to working with them. This outcome is different from the predictions of the original Becker’s model, according to which all employers below a marginal level of discrimination prefer hiring women, while those above this threshold avoid employing them. Given that in our setting firms pay women below the monopsonistic benchmark more the stronger their prejudices, we are able to show that workplaces’ compensation policies embed taste-based discrimination and we derive empirical tests for the presence of this mechanism.

Using these results, we provide a rigorous assessment on the existence of taste-based discrimination and on its impact on the gender wage gap. We determine the extent to which different measures of employers’ preferences toward women are associated to sex pay differences, providing a characterization of more prejudiced employers and improving our understanding on what drives this behaviour. Furthermore, following a similar approach we also test for the presence of gender differences in compensating wage differentials, focusing on preferences for more flexible working schedules.³ Our empirical modelling strategy adapts a method derived from Card et al. [2016] and applies some of the theoretical insights provided by Card et al. [2018].

Studying the impact of taste-based discrimination in the context of monopsonistic labour markets can be considered a realistic choice for several reasons. First, from a theoretical perspective nothing suggests that the two discriminatory mechanisms (employers’ preferences and wage setting power) should be viewed as mutually exclusive. On this respect, Black [1995] builds a dynamic model where taste-based discrimination itself produces monopsonistic discrimination against minority groups, showing that one of the two

³The availability of part-time work has been found to be an important determinant of female labour force participation (*e.g.* Del Boca [2002]) and it tends to be considered a desirable job characteristic by women (*e.g.* Booth and van Ours [2013]). Thus, it is interesting to study whether there are effects on female wages when more flexibility is available at a given workplace.

factors may even strengthen the other. Secondly, there are important theoretical and empirical considerations suggesting that most labour markets are not perfectly competitive. Several studies, using different approaches in a variety of contexts, have documented the presence of some degree of employers' wage setting power and of substantial differences in the female and male labour supply elasticities to the firm.⁴ Even if exceptions to this trend can also be found in the empirical literature, to the best of our knowledge there are no overwhelming evidences against the hypothesis of imperfect labour markets.⁵ Also from a more theoretical perspective, Boal and Ransom [1997] show that a monopsony is implied by standard dynamic search models in which larger firms face dis-economies of scale in hiring workers. The same authors discuss several other reasonable mechanisms that could induce imperfect competition, among which limited information about vacancies and costs associated to mobility from a job to the other. Finally, in a recent contribution Card et al. [2018] argue that firms' monopsonistic power can also be considered a driver of heterogeneity in wages between observationally similar workplaces, an evidence documented by an extensive literature following the seminal work by Abowd et al. [1999].

The application of this paper is based on data covering the population of private sector workers in the Veneto region of Italy. We measure the relative importance of firm-specific heterogeneity as a determinant of the overall gender pay gap, following the methodology of Card et al. [2016]. However, we focus the analysis on manufacturing *local labour markets* only, since these can be considered groups of firms characterized by relatively homogeneous labour market structures.⁶ For this reason, and given that gender differences in human capital and returns to skills are controlled for by the regression model, we

⁴Among others, see Barth and Dale-Olsen [2009], Hirsch et al. [2010], Ransom and Sims [2010], Ransom and Oaxaca [2010], Depew and Srensen [2013], Muehlemann et al. [2013], Webber [2015], Webber [2016] and, for Italy, Sulis [2011]. All of these studies provide either indirect or direct support to the hypothesis of monopsonistic labour markets.

⁵One notable exception supporting the perfectly competitive hypothesis is Matsudaira [2014]. This study exploits quasi-experimental variation in the size of firms induced by a minimum staffing law introduced in California's caring sector. Results show that there was no growth in wage differences between nursing homes affected and not affected by this reform. However, since around 75% of all nursing homes in the State were under-staffed according to the new law, the resulting strong growth in aggregate demand for caregivers could have produced general equilibrium effects that make it difficult to draw conclusive evidences on the presence of firms' wage setting power.

⁶Local labour markets (or *districts*) are geographical and economic entities that, in Veneto, are characterised by a high density of small-sized manufacturing-oriented firms. We have constructed them using a definition of the Italian national statistical office based on census' commuting data.

provide reasonable conditions for testing whether firm wage policies embed discriminatory preferences and whether they are influenced by other employer-specific mechanisms driving the gender pay gap.

Our results show novel evidences on the importance of taste-based discrimination. In particular, when approximating preferences against women by adopting two commonly used measures, *i.e.* the presence of women at the top of the occupational hierarchy and the female employment share,⁷ we find that, consistently with our theoretical results, both mechanisms contribute to firm-level gender wage differences. A 10 percentage points increase in the female employment share within firms implies up to a 0.5 percentage points reduction in the pay gap conditional on workers' productivity and monopsonistic discrimination effects. Instead, the presence of women at the top of the corporate hierarchy, while determining around one percentage point increase in overall conditional gender wage differences, has no significant effects on the conditional level of female wages at the bottom of the firms' structure. Using the same approach, we also document the presence of gender differences in compensating wage differentials, as women earn relatively less at firms that offer more flexible working schedules.⁸

Being able to distinguish among the sources of wage differences between men and women is not merely a theoretical exercise, but it has important implications on the choice of the most effective policies to implement in order to achieve greater equality. The method proposed here can also be helpful when testing the implications of Becker's theory in the data. For example, according to short-run predictions less discriminatory employers should have lower costs and be more profitable,⁹ while, in the long run, gender differences should reduce, given that incumbent non-discriminatory competitors would find it

⁷Partly due to difficulties in building appropriate tests, evidences on whether these two variables can be considered valid approximations for employers' tastes are not abundant. Among the most convincing evidences on homophily among managers, Giuliano et al. [2009] and Giuliano et al. [2011] find that managers' race characteristics tend to be correlated with the characteristics of new hires and promotions, while Gagliarducci and Paserman [2015] find that, female managers are more likely to work at establishments where women-friendly policies are in place. Instead, for what concerns the employment share of minorities, apart from theoretical considerations, evidences on its link to affirmative action policies are provided by Miller [2017].

⁸On this last respect, we document up to one percentage point growth in the gender wage gap conditional on employers' wage setting power and workers' productivity for each 10 percentage points increase in the share of total days worked part-time within firms.

⁹This outcome is studied, among others, by Hellerstein et al. [2002] and Kawaguchi [2007]

profitable to enter in the market.¹⁰ Most of the existing contributions on similar topics face the challenge of finding a reliable firm-specific parameter for discriminatory preferences. Indeed, such information is sometimes explicitly available from surveys only at an aggregate level or in particular contexts (as for example in Charles and Guryan [2008] and in Glover et al. [2017]). More often, discriminatory preferences are approximated by the female share of workers within firms (*e.g.* Weber and Zulehner [2014]) and by the presence of women in executive boards (*e.g.* Cardoso and Winter-Ebmer [2010], Flabbi et al. [2014], Gagliarducci and Paserman [2015]). Our contribution to this literature is to provide an empirical test on the relevance of these firm-level measures of taste-based discrimination, which is grounded on both, theoretical and empirical considerations.

The paper is organized as follows. Section 2 presents the theoretical model and Section 3 discusses its identification. Section 4 presents the data and Section 5 presents the main empirical results of the paper, while the final section contains the concluding remarks.

2 Theoretical framework

2.1 Profit maximization

We consider a static model where *Robinsonian discrimination* arises as the result of third degree price discrimination. Taste-based discrimination is defined as an employer-specific exogenous cost, which is proportional to the female employment level. In this model, an employer chooses a quantity of labour $L = L^m + L^f$ maximizing the profit function, which reads as

$$\pi(L^m, L^f) = pq(L) - w^m(L^m)L^m - w^f(L^f)L^f - \delta L^f \quad (1)$$

Throughout the paper, the subscripts m and f stand for male and female respectively.

The parameter p is the output price and $q(L)$ is the production function, with $q' > 0$ and

¹⁰Studies on Becker's long-run prediction often exploit shocks in product market competition across time, but the magnitude and the extent of their effects on discrimination are often found to be particularly limited (*e.g.* Black and Brainerd [2004] and Heyman et al. [2013]). Among studies adopting different approaches, Ashenfelter and Hannan [1986] find a negative relationship between women's employment shares and market concentration in the banking industry. More recently, Weber and Zulehner [2014] find that firms with larger female shares have better survival rates, while surviving firms tend to increase women's employment levels, an evidence supporting the hypothesis of employers' learning and of convergence toward less discriminatory outcomes.

$q'' < 0$. Male and female workers are perfect substitutes in the technology. w^m and w^f are gender-specific inverse labour supply functions, which are increasing in L^m and L^f , respectively. Finally, δ is a taste-based discrimination parameter.

Under standard assumptions,¹¹ the first order conditions of profit maximization can be written as

$$mp = w^m \left(1 + \frac{1}{\epsilon^m} \right) \quad mp = w^f \left(1 + \frac{1}{\epsilon^f} \right) + \delta$$

where mp is the marginal revenue product and ϵ^g is the elasticity of the labour supply for $g = m, f$. The solution of the model is graphically represented in Figure 1, where the optimality conditions are characterized under different choices of the parameters. Namely, the left panel of the figure represents the solutions when $\epsilon^m \rightarrow \infty$ and $\epsilon^f < \infty$ for the cases of zero and positive taste-based discrimination, while the right panel describes the solutions when $\epsilon^g < \infty$ ($g = m, f$), again for the cases in which $\delta = 0$ and $\delta > 0$.


In general, in this model wages are marked-down with respect to the marginal revenue product (mp), and this mark-down grows as the labour supply becomes more rigid. If $\delta = 0$ the marginal revenue product is set equal to each gender-specific marginal factor cost (MFC^g). When $\delta > 0$, there is a difference between mp and MFC^f in the case of women. In order to adjust for the cost of δ , the employer reduces female employment (L^f) and wage levels (w^f) below the monopsonistic benchmark and this reduction is compensated by only a less than proportional growth in male employment (L^m), as hiring more men is increasingly costly, unless the male labour supply is perfectly elastic.

To sum up, an employer for which $\delta > 0$ produces less output, hires less women, has a lower female share and pays women less than what would be observed at the monopsonistic benchmark (*i.e.* at $\delta = 0$). However, when comparing any two firms, the negative relationships between any of these variables and employers' discriminatory preferences do not necessarily hold, since each firm may have different labour supply functions and production technologies. For this reason, in the next paragraph we characterize employers'

¹¹Two sufficient conditions for optimality are

$$2w^{g'} + w^{g''}L^g > 0 \quad \text{for } g = m, f$$

Figure 1: Graphical Representation of the Model Under Different Market Structures and Discrimination Levels



(a) The equilibrium in the absence of taste based discrimination ($\delta = 0$) is represented by the black dots, where $mp = MFC^f = MFC^m$. If $\delta > 0$, female employment is reduced to the level $L^f(\delta > 0)$, represented by the grey dot where $mp - MFC^f = \delta$. Instead, mp and MFC^m are kept constant, which implies that the difference $L^f(\delta > 0) - L^f(\delta = 0)$ is compensated by an equivalent growth in male employment (L^m).

(b) The equilibrium in the absence of taste based discrimination ($\delta = 0$) is represented by the black dots, where $mp = MFC^f = MFC^m$. The grey dots represent the optimal points on each of these curves when $\delta > 0$. At this equilibrium, female employment is reduced to the level $L^f(\delta > 0)$, represented by the grey dot on the MFC^f curve. As a consequence, mp grows above the curve MFC^m and more male workers are hired until mp and MFC^m are set equal again (grey dots on the respective lines). Since hiring more men is increasingly costly, the growth in male employment is smaller than $L^f(\delta > 0) - L^f(\delta = 0)$ (i.e. for fixed technology and supply functions, more discriminatory firms are smaller in size).

heterogeneity more explicitly.

2.2 The Role of Workplace Heterogeneity

In this paragraph, by introducing heterogeneity in employers' taste-based discrimination and wage setting power, we characterize differences in wages across firms in the context of the profit maximization model discussed above. We also introduce heterogeneity in individual labour productivity, by allowing workers to provide different contributions to firms' revenues. The next paragraph further characterizes the model by imposing restrictions on the labour market structure, deriving some useful implications.

Consider a population of firms indexed by j . We assume that each firm faces arbitrary gender-specific inverse labour supply functions, where the respective elasticities are denoted by ϵ_j^g . The functional form of these labour supplies is discussed below. For the time being, the first order condition of profit maximization can be written as

$$w_j^g = mp_j \left(\frac{\epsilon_j^g}{1 + \epsilon_j^g} \right) \left(1 - 1 [g = f] \frac{\delta_j}{mp_j} \right)$$

Notice that in the above equation δ_j is modelled as an employer-specific discriminatory parameter. This parameter can also be expressed as a percentage of labour productivity and, in the remainder of the paper, we use this latter definition of discrimination in order to rank employers' prejudices. With this approach, a given amount of δ_j is considered more discriminatory at firms that are relatively less productive - which, consequently, pay women proportionally less than men - with respect to firms having a higher marginal revenue product. For notational convenience, we define the following parameter

$$-\hat{\delta}_j \equiv \ln \left(1 - \frac{\delta_j}{mp_j} \right)$$

where $\hat{\delta}_j$ is monotonic and increasing in δ_j , it is constant at the firm level and it describes the percentage of women's productivity that is marked-down due to employer's prejudices. We introduce individual heterogeneity in productivity by assuming that workers provide

different amounts of equally productive units of labour l^i .¹² If employees are *endowed* with such heterogeneous quantities of labour, we can write worker i wage equation as a function of the firms' unitary pay level, that is

$$\begin{aligned}
 w_j^i = l^i w_j^g &= mp_j^i \left(\frac{\epsilon_j^g}{1 + \epsilon_j^g} \right) \exp(-\hat{\delta}_j 1[g = f]) & mp_j^i &\equiv l^i mp_j \\
 \implies \ln w_j^i &= \ln mp_j^i + \ln \left(\frac{\epsilon_j^g}{1 + \epsilon_j^g} \right) - \hat{\delta}_j 1[g = f] & & (2)
 \end{aligned}$$

According to the above equation, log wages are an additively separable function of workers' productivity, of firms' wage setting power and of employers' discriminatory preferences. In order for this model to be considered realistic, we need to introduce the possibility of misspecifications, which could arise due to several firms', workers' or match wage components that we have not considered explicitly. Moreover, in the absence of information on employers' tastes, workers' productivity and firms' wage setting power, the above three elements can be estimated or controlled for only in a longitudinal setting. Thus, we also need to add dynamic considerations to our static framework. Before turning to these problems, in the next section we discuss more carefully the functional form of the labour supply to the firm.

2.3 *The Labour Supply to the Firm*

According to equation (2), monopsonistic *mark-downs* of wages with respect to productivity have an influence on a worker's pay, at least unless we believe such mark-downs to be fairly close to zero. There are several reasons why employers might have some degree of wage-setting power in most labour markets. For example, the presence of imperfect information about vacancies, of direct and indirect costs associated to job switching and, more generally, of dis-economies of scale in hiring are among the main arguments supporting this hypothesis (see, among others, Boal and Ransom [1997] and Manning [2003]).

¹²A richer modelling choice would be to assume imperfect substitutability of workers in the firm's technology along some dimension. We do not consider this extension of the model explicitly, but, in the application (Section 5.3), we test our main results using also a more nuanced empirical specification, where the unit of analysis are specific jobs within a firm (thus, a specification where imperfect substitutability is allowed for different jobs within a workplace).

On this respect, Sulis [2011] provides a direct assessment of the amount of labour market power held by firms in the Italian private sector, which is the market considered in the application, showing evidences on the presence of this mechanism and of relevant gender differences in this parameter.

In this theoretical framework we consider a special case of firms' wage setting behaviour. In particular, we assume that all firms operating in a given factor market face an inverse labour supply to the firm of the following form

$$w_j^g = (L_j^g)^{\alpha^g} (z_j^g)^{\gamma^g} \quad \implies \quad \ln w_j^g = \alpha^g \ln L_j^g + \gamma^g \ln z_j^g \quad (3)$$

In the above equations, w_j^g is the firms' wage paid to each productive unit expressed in levels, z_j^g is a vector of characteristics and a residual term, γ^g is a vector of parameters and a constant, L_j^g is the total amount of *productive units* supplied by gender g at firm j and α^g is a real-valued parameter. For the time being, we consider α^g to be only gender-specific. However, in Section 3.3, by providing a more precise definition of labour markets, we explicitly model heterogeneities in this parameter across firms.

The characteristics included in the vector z_j^g control for all factors determining heterogeneities in availability and quality of productive units in the labour market, as long as they have an influence on wages. Instead, any characteristic *influencing only the demand of labour* (and not its supply) should be considered as excluded from z_j^g , since such element would bias any estimate of α^g towards zero (*e.g.* Manning [2003]). In this setting, the parameter $1/\alpha^g$ can be interpreted as a measure of the elasticity of the labour supply faced by firms, net of any other composition effect influencing the wage-size relationship. As mentioned, we assume that this elasticity is finite, as firms face increasing costs in hiring, the larger their employment level.

Giving to α^g its structural interpretation, as the parameter characterizing the labour supply to the firm, is quite difficult when adopting traditional regression methods. As documented by a vast stream of literature (*e.g.* Oi and Idson [1999]) there could be employer-size wage effects determined not only by monopsonistic mechanisms, but also by the fact that, for example, larger employers could attract workers of better quality,

they could offer inferior working conditions, they could share a larger proportion of rents, or they may pay efficiency wages to deter shirking. However, when the objective is to *control* for monopsonistic mark-downs, rather than estimating their size, the labour supply function just described has two convenient features. First, for any two firms s and j facing the same factor market structure

$$\epsilon_j^g = \epsilon_s^g = \frac{1}{\alpha^g} \quad \forall s \neq j$$

That is, the elasticity of supply is a constant parameter across such firms. Secondly, provided that (3) is an appropriate functional form specification, monopsonistic mark-downs in equation (2) are not only additive, but also independent of employment levels,¹³ so that worker's i wage equation becomes

$$\ln w_j^i \approx \ln mp_j^i - \alpha^g - \hat{\delta}_j 1 [g = f] \quad (4)$$

An useful implication of these two properties is that, whenever firms face the same factor market structure, gender-specific monopsonistic mark-downs can be controlled for in a regression framework by simply adding fixed effects for each of these labour markets. However, this approach is feasible only if the functional form of equation (3) is reasonable and if sets of firms facing an approximately similar gender-specific labour supply function can be identified. Furthermore, when a regression method based on equation (4) is adopted to test for the presence of preferences against women (as in the empirical model of Section 3.3), such prejudices can only be estimated conditioning on the average level of discrimination present in a given labour market.¹⁴

It is worth noticing that Card et al. [2018] present a similar model of monopsonistic wage setting, where a log-log functional form of the labour supply to the firm is derived from

¹³In particular, notice that for all firms j in a given labour market

$$\ln \left(\frac{\epsilon_j^g}{1 + \epsilon_j^g} \right) = \ln \left(\frac{1}{1 + \alpha^g} \right) \approx -\alpha^g \quad \forall j$$

¹⁴This is always true unless for the case in which taste-based discrimination $\hat{\delta}$ can be approximated by a variable that is not correlated with gender differences in firms' wage setting power.

specific assumptions on workers' indirect utility. They show that a simple two-period extension of the model has similar implications in steady state to this static framework and leads to random mobility of workers across firms. Nevertheless, we stress that this conclusion is reached in a simplified framework, where important considerations, such as workers' job switching costs and oligopsonistic strategic interactions between employers, are not taken into account.

3 Empirical Specification of the Model and Identification

3.1 Firms' Wage Policies and Their Estimation

In this section, we show how the wage equation derived in the theoretical framework relates to the two-way fixed effects (or AKM) regression model (Abowd et al. [1999]), discussing the main assumptions required for its consistent identification. For this purpose, we incorporate additional components to equation (4), considering the possibility of measurement error and model misspecifications. Moreover, we also take into account dynamic considerations, allowing for the presence of innovations in workers' productivity as well as in other unobserved wage components, but also introducing some restrictions on these time-varying processes.

We define ρ_j^m and ρ_j^f as *time-constant, gender- and employer-specific* residual terms, representing firms' deviations from the predicted wage schedule defined by equation (4). Such deviations can be attributed to several factors usually linked to heterogeneity in compensation policies across workplaces (see Card et al. [2018] for an overview of the main arguments). In particular, these error terms could be linked to compensating wage differentials, efficiency wages, employers' rent-sharing policies¹⁵ as well as measurement error. Notice that, even if these terms are specific for each gender, part of the above mechanisms could also affect men and women in the same way within firms.

We also define r_{it} as an *individual-specific and time-varying* wage residual (where t denotes discrete periods), which we assume to be normally distributed with mean zero in the

¹⁵With respect to rent-sharing, Card et al. [2016] show that it is a relevant mechanisms influencing firms' compensation policies and they find relevant gender differences in rent-sharing within Portuguese firms.

population and independent from all the other wage components. Adding these elements to the wage equation, we have a model that reads as

$$\ln w_j^{it} = \ln mp_j^{it} - \underbrace{\alpha^g - \hat{\delta}_j 1[g=f]}_{\equiv \omega_j^g} + \rho_j^g + r_{it} \quad (5)$$

As can be noticed by the time index, workers' productivity is allowed to change across periods and, as discussed below, we assume that it can be approximated correctly by unobserved time-constant and observable time-varying individual characteristics. In this setting, the element ω_j^g defined in equation (5) can be interpreted as a time-constant firm wage residual. Throughout the paper, we call this residual *firm wage policy*, or *firm wage premium*.

Under given assumptions on the error term r_{it} , which we discuss below, the identification of compensation policies ω_j^g can be achieved by estimating an AKM regression model separately by gender. In particular, let $j = \iota(i, t)$ index the firm in which worker i is employed at time t . Assume that employees are observed for T time periods and let W_i represent a $T \times 1$ vector of daily wages, while X_i a $T \times P$ matrix of time-varying individual characteristics. Then, the two-way fixed effects model can be specified as follows

$$\ln w_{it} = x_{it}\beta + \eta_i + \omega_j^g + r_{it} \quad (6)$$

where w_{it} and x_{it} are rows of W_i and X_i respectively, β is a $P \times 1$ vector of parameters, while ω_j and η_i are respectively firm-constant and time-constant components of individual wages, which are allowed to be arbitrarily correlated with any of the characteristics in x_i , and which could be not perfectly observable.¹⁶ In the application, we have adopted a specification of the model suggested by Card et al. [2018], including as covariates in x_{it} a cubic polynomial interacted by three occupation dummies, a dummy for fixed-term contracts and a full set of year fixed effects.

¹⁶We implicitly maintain that the regression model is estimated separately by gender, so that each parameter should be considered gender-specific. For notational convenience, we have omitted the subscript g whenever redundant.

The main assumption required for a consistent identification of the parameters in (6) is the absence of correlation between the error term r_{it} and all the other time-varying and (unobserved) time-invariant characteristics included in the model (Abowd et al. [1999] and Card et al. [2013]). This condition must hold also for error terms in periods different from t , so that, for example, mobility towards employers with given firm wage policies can not be correlated with previous idiosyncratic shocks in earnings.¹⁷

Two relevant components entering in r_{it} are innovations in workers' unobserved earning abilities and job match effects associated to given employer-employee pairs. In the context of the model of Section 2.2, match effects could be interpreted as productivity shocks, like innovations in the parameter l^i that are not predicted by the time-varying controls included in x_{it} and that are associated to a match with a given firm j . They could also represent systematic differences in firms' wage policies associated to given worker-employer pairs, which would then enter in the residual term r_{it} . As in Card et al. [2013], we assume that innovations in workers' unobserved earning abilities have mean zero and contain an unit root, while job match effects have mean zero for all i and j in the sample interval. Section 4 presents some tests on the credibility of these restrictions along the lines suggested by Card et al. [2013] and Card et al. [2016]. Moreover, we provide a sensitivity analysis on our main empirical results, by allowing firm fixed effects ω_j to be specific for manual and non-manual workers within workplaces.

3.2 Normalization of Firms' Wage policies

In the previous paragraph, we have shown that, by computing the AKM regression model separately for men and women, we can obtain two estimators ($\hat{\omega}_j^m$ and $\hat{\omega}_j^f$) of the *gender-specific* firms' wage policies. Notice that these parameters are constant at the firm-gender level, and they measure additional wage premiums (discounts) that firms pay to their workforce, conditional on workers' characteristics. Since the main step for isolating factors related to taste-based discrimination, which, according to equation (5), are embedded in $\hat{\omega}_j^f$, involves taking differences between male and female firm wage policies, this section

¹⁷Figure 2 provides descriptive evidences suggesting that this restriction can be considered realistic in our sample.

describes the procedure adopted to make such parameters comparable across gender-specific regression results.

One of the main challenges in comparing $\hat{\omega}_j^m$ and $\hat{\omega}_j^f$ directly is given by the fact that both are computed with respect to an arbitrary reference group.¹⁸ For this purpose, as discussed by Card et al. [2016], we need to define a firm, or a set of firms, as the common reference group across gender, and rescale each firm wage policy accordingly. The normalization choice proposed by these authors involves the use of balance sheet data, in order to set the lowest value-added group of firms as the reference. Here, we adopt a different choice, and select the largest firm (in terms of person-year observations) as the reference one.¹⁹ However, since the largest firm is different in the female and male samples, to identify the common largest employer we have considered the following size function

$$\text{size}_j = \min\left\{\frac{N_j^f}{N_j^f + N_j^m}, \frac{N_j^m}{N_j^f + N_j^m}\right\}$$

where N_j^m and N_j^f are total firms' person-year observations among men and women, respectively. The reference firm that we have chosen (\hat{j}) is the largest according to the above definition of size, which gives more weight to workplaces in which the share of female and male workers is closer to 50%.

Denote by $\hat{\omega}_{\hat{j}}^m$ and $\hat{\omega}_{\hat{j}}^f$ the gender-specific pay policies of the reference firm \hat{j} defined above. Then, for any j we have applied the following differences

$$\omega_j^m = \hat{\omega}_j^m - \hat{\omega}_{\hat{j}}^m \quad \omega_j^f = \hat{\omega}_j^f - \hat{\omega}_{\hat{j}}^f$$

The above normalization allows to express all wage policies with respect to the ones paid by the largest employer. When considering the standardized difference $\omega_j^m - \omega_j^f$, we can measure whether men's pay policy are proportionally higher (positive difference) or

¹⁸Also the number of reference groups depends on the number of *connected sets* among men and women (see Abowd et al. [2002]). For simplicity, we will restrict our analysis on the *largest connected set* only, which is composed of around 98% of all the observations.

¹⁹We have also tested an alternative normalization choice, finding only small differences in the results. In particular, we have tested the results defining firms in the first percentile of the male wage policy distribution to be the reference category, expressing all other policies as (gender-specific) differences between the average firm wage residual of this group.

lower (negative difference) than women’s pay policy at firm j , with respect to the same difference computed at firm \hat{j} .

The standardized difference $\omega_j^m - \omega_j^f$ can be used to rank firms according to their *relative* discriminatory behaviour, identifying groups of employers who provide more favourable work environments for women. Alternatively, it can be tested whether given proxies of preferences against women are actually good predictors of this difference, thus whether they are a relevant mechanism determining the gender wage gap in a given labour market. Sections 3.3, 5.2 and 5.3 provide examples of the above mentioned applications and a further discussion on their interpretation.

3.3 Gender Gap in Firms’ Wage Policies, Taste-Based Discrimination and Compensating Differentials

We now turn to the problem of identifying employers’ discriminatory tastes and other workplace-specific mechanisms driving sex pay differences, starting from the gender gap in firms’ wage policies. Since firms’ premiums can be influenced by the degree of labour market power held by employers, which could differ by gender, we begin by considering more explicitly the role of heterogeneities induced by this mechanism.

Following the discussion of Section 2.3 and adopting a similar notation, we assume that each firm belongs to a given labour market k . Moreover, we assume that employers face market- and gender-specific log-log inverse labour supply functions that read as

$$\ln w_j^g = \alpha_k^g \ln L_j^g + \gamma_k^g z_j^g \quad g = m, f \quad k = 1, \dots, K$$

where w_j^g is the unitary wage and the vector z_j^g contains unobserved characteristics affecting the labour supply and an error term. Given this functional form, it follows that the labour supply elasticity to the firm is determined by the constant parameter α_k^g only. This model also implies that the gender gap in firms’ wage residuals can be written as

$$\omega_j^m - \omega_j^f \approx \hat{\delta}_j + \underbrace{\alpha_k^f - \alpha_k^m}_{\equiv \alpha_k} + \underbrace{\rho_j^m - \rho_j^f}_{\equiv \rho_j} \quad (7)$$

where $\hat{\delta}_j$ represents discriminatory preferences against women. Instead, the composite error term ρ_j reflects *heterogeneities* in residual firm wage components, as long as they affect differently men and women within the same workplace.

If the AKM assumptions hold, so that ω_j^g is identified consistently, we can consider equation (7) as a valid regression model. In this setting, an estimate of taste-based discrimination could be recovered by controlling for market-constant effects α_k and firm level residuals ρ_j . However, since both of these confounding factors are mostly unobservable, a feasible alternative to this approach, which is followed in Section 5.3, is to test whether reasonable *proxy variables* for preferences against women (or for other potential mechanisms contributing to ρ_j) are significant predictors of the LHS of equation (7).

In our empirical specification, we have considered two measures of $\hat{\delta}$, namely the presence of women at the top of the firm hierarchy and the share of female labour within firms. Both variables are usually associated to taste-based discrimination in the empirical literature. For example, Gagliarducci and Paserman [2015] find that in Germany female managers are more likely to work in female-friendly firms. Similarly, Cardoso and Winter-Ebmer [2010] and Flabbi et al. [2014] show that the gender wage gap tends to be lower at women-led companies.²⁰ Also the female share of employment is typically considered a proxy of preferences against women (see for example Weber and Zulehner [2014]) and its growth has been linked to the presence of affirmative action policies (*e.g.* Miller [2017]). Moreover, this variable is associated to discrimination in the traditional model of Becker [1957]. In the context of our theoretical framework, a lower female share is linked to greater $\hat{\delta}_j$ conditional on firms' size and on the labour market structure. Since, as we discuss below, to some extent we control for both factors in the regression model, women's employment share within firms can be considered a valid proxy for taste-based discrimination also in our context.

The regression model provided by equation (7) can be used to test also other hypotheses on the gender wage gap. A particular interesting one concerns the role of compensating wage differentials. Women could indeed be attracted by given amenities provided

²⁰In the context of racial discrimination, two papers documenting the presence of homophily among managers are Giuliano et al. [2009] and Giuliano et al. [2011].

by employers and this tendency would then be reflected in ρ_j , as some workplaces could be able to lower female wages accordingly. In the application, we have tested whether employers more willing to provide flexibility in working schedules are paying women less due to hedonic considerations.²¹ This is a particularly interesting mechanism, given that part-time work is a relatively scarce resource in the labour market considered in the application, while its availability improves female labour force participation (see in particular Del Boca [2002]). For this purpose, we have included in equation (7) a control for the share of part-time employment within firms.

In order to estimate equation (7) consistently, we need to impose some restrictions on the unobservable components in α_k and ρ_j . As mentioned, given our assumptions on the labour market structure, the inclusion of fixed effects for each market k allows to control for monopsonistic mechanisms embedded in α_k .²² Whether clusters of firms facing the same labour market structure can be identified correctly remains a matter of judgement. In the application (see in particular Sections 4 and 5.3) we have relied on a classification based on firms' product market, geographical location as well as on their employment size.

A second set of assumptions involves unobserved components embedded in the error term ρ_j , which cannot be correlated with our proxies of $\hat{\delta}_j$ and of non-wage amenities. This restriction can be problematic, since, in the case of female managers, there could be simultaneity. Firms' wage policies may be higher for women when one of them is earning the most.²³ Similarly, top earners could have better information on such policies and sort into better workplaces accordingly. Furthermore, the female share of employment could be correlated with gender differences in measurement error of firms' wage policies.²⁴ Given

²¹On this aspect Booth and van Ours [2013] show that married women tend to prefer part-time contracts in the Dutch labour market.

²²However, estimates of market fixed effects α_k can not be given a structural interpretation, as gender differences in monopsonistic mark-downs, as assuming that other unobserved confounding factors were uncorrelated across labour markets would be a too demanding restriction. This aspect of the model is further discussed in Section 5.2.

²³Notice however that firms' pay premiums are estimated conditioning on human capital characteristics of individual workers. Thus, in our context the problem of simultaneity is attenuated by this feature of the AKM regression model.

²⁴On this respect, Andrews et al. [2008] shows that measurement error increases the lowest the number of observed job mobility episodes.

these considerations, taking advantage of the long panel structure of the data available in the application, we have estimated the model by measuring these two proxy variables in lags. The underlying hypothesis made when using lagged proxy variables of preferences toward women is that a firm’s management culture and structure change only slowly over time, while past levels of $\hat{\delta}_j$, being computed in a period different from the one in which ω_j^g is estimated, are less vulnerable to the problems of simultaneity and correlation with measurement error.

Also in the case of non-wage amenities there could be identification problems. In particular, the presence of part-time wage penalties could bias our estimates of $\omega_j^m - \omega_j^f$, due to an over-representation of women in such positions, and the part-time share could be correlated with such a measurement error. We have dealt with this issue by always excluding part-time workers from the AKM estimation sample. Thus, in the application we study only potential “spill-over effects” on full-time workers arising when employers offer more flexible working schedules, testing whether even only the likelihood of having a part-time is a more desirable job characteristic for women.²⁵

Finally, in one specification of the model we have also estimated firm wage policies interacting these fixed effects by occupation (manual and non-manual). Then, we have used gender differences in these parameters as the dependent variable in equation (7). Since the overall female employment share, part-time share and the presence of women at the top of the corporate structure are less correlated with measurement error or third factors embedded in compensation policies toward blue collars, we can consider this approach as a robustness test for our main results, as well as a heterogeneity analysis of our findings across the corporate hierarchy.

4 Data and Sample Selection

We have estimated the models discussed in the previous sections using Italian linked employer-employee data from administrative sources (Veneto Working History database,

²⁵Elsayed et al. [2017] documents the presence of part-time penalties for women in the UK that are declining over time. The presence of negative wage “spillover-effects” on full-time female workers employed by firms offering more part-time contracts may indicate that this penalty could be often under-estimated.

hereafter VWH).²⁶ In particular, we study the population of private sector workers in Veneto during the period between 1996 and 2001. Veneto is an important region of Italy, which represents around 10% of the national GDP. It is a manufacturing-oriented economy and it can be considered as a self-contained labour market, given its relatively limited out-migration.

The data is derived from INPS (National Social Security Institute) social security archives, which cover the population of private sector dependent workers, excluding self employed and public-sector employees. All firms registered at one of Veneto's INPS offices are included in the data,²⁷ which provide demographic and occupational information on their entire workforce and the location and the sector of activity of each company.²⁸ Workers who transit from these firms are observed also if they are employed by a private-sector employer outside of Veneto.

Our analysis is based on gross daily wages, which are inclusive of all pecuniary benefits paid by employers. As mentioned, we have excluded part-time workers from the analysis, as this choice limits the measurement error in actual time worked and it eases the comparison of firms' wage policies by gender. Similarly, given that entry-level contracts providing on-the-job training are often characterized by a very low pay, we have also excluded apprentices from our sample. We have selected one job spell per individual in each year, choosing the longest work episode whenever a person was simultaneously employed at more than one firm. Finally, we have restricted the analysis on firms belonging to the gender-specific largest connected sets, *i.e.* the set of all establishments connected by the mobility of workers, which is a relatively standard procedure in the literature (see for example Card et al. [2013]).²⁹

We have estimated the AKM model by gender on the entire sample of workers defined above. Then, we have analysed the employer-specific gender wage gap, given by equation

²⁶The VWH dataset has been developed by the Department of Economics of the University of Venice Ca' Foscari under the supervision of Giuseppe Tattara.

²⁷Such registration is compulsory for firms hiring dependent workers.

²⁸According to the ISTAT census, in 2001 there were only 1.08 establishments for each firm in Veneto. Thus, the presence of multi-plant activities is quite limited in our data.

²⁹The largest connected set corresponds to around 98% of the observations. See [Abowd et al., 2002] for a discussion of this procedure and a more detailed definition of connected sets.

Table 1: **Mobility of Workers Across Local Labour Markets (1996-2001)**

Industry	Number of Worker-SLL- Industry Pairs	% Observed Out of SLL-Industry	% Changing Firm	% Observed Out of SLL-Industry Among Workers Changing Firm
Manufacturing	828,969	23.3	50.6	47.3
Other sectors	811,969	23.9	51.6	45.3

(7), considering only a more homogeneous group of firms in terms of labour market structure.³⁰ In particular, we have adopted a further sample selection criteria based on firms' geographical location, product market structure and size.

Taking advantage of the comprehensive level of detail in the available data and exploiting also the peculiarities of Veneto, we have considered only firms belonging to one of the region's *local labour markets*. Such geographical entities (also called *districts*) were identified using an official classification, based on census data. In particular, we have used the definition of the Italian statistical office (ISTAT), which calls local labour markets *Sistemi Locali del Lavoro*, or SLL.³¹ Using data on individual commuting habits derived from the census, such SLLs are constructed as a group of municipalities highly connected in terms of employment,³² where the connectivity of each group is maximized considering two main measures: *i*) the proportion of jobs within the districts held by its residents and *ii*) the proportion of residents that work in the local labour market. A map of SLLs within Veneto and its neighbouring regions is provided by Figure 4.

Table 1 shows that SLLs provide a relatively good approximation of a firm's labour market structure, as employers belonging to the same district tend to hire from the same pool of workers. As can be noticed, considering the period 1996-2001, only around 23% of workers employed in a given SLL are observed working also outside of this geographical area or in a different industry, where sectors are broadly defined considering manufacturing

³⁰This choice of starting from the full sample is adopted in order to measure firm wage policies with higher precision. Indeed, as shown by Andrews et al. [2008], measurement error in firms' wage policies reduces as the number of observable job mobility episodes increases.

³¹See Lorenzini [2005] for the details of this procedure. A summary of the methodology employed in the identification of local labour markets from census data is available also online, see (http://www.istat.it/it/files/2014/12/nota-metodologica_SLL2011_rev20150205.pdf).

³²Municipalities are the smallest administrative entities in the Italian territory, as there are around 8000 of them at the national level, of which almost 600 in Veneto.

Table 2: **Workforce Composition and Gender Wage Gap by Sector (1996-2001)**

Number of Observations			Conditional Gender Wage Gap		
Gender	Manufacturing	Other	Independent Var.	Coefficient	S.e.
Male	2,011,273	1,437,580	Male worker	0.166	0.003
<i>Row %</i>	58.3%	41.7%			
Female	941,872	881,630	Male worker * Manufacturing	0.031	0.004
<i>Row %</i>	51.7%	48.3%			
Total	2,953,145	2,319,210	N. Observations	5,272,355	
<i>Row %</i>	56%	44%			

The conditional gender wage gap is estimated by a regression model that includes a quadratic age polynomial together with occupation, year and firm fixed effects. Standard errors are clustered at the firm level.

and non-manufacturing activities. When this proportion is computed considering only the population of workers who switched job during the period of observation, less than 50% of these employees are observed changing their SLL or their industry.

As mentioned, we have studied the gender wage gap further restricting the sample to manufacturing firms only.³³ This choice is motivated on three grounds. First, skills demanded (thus workers hired) by manufacturing firms tend to be more homogeneous. Such more limited heterogeneity can be noted from an analysis of workers' pay. The standard deviation of log daily wages is 0.34 in the manufacturing sector, while it is 0.43 in other industries. The same pattern holds for both, men (0.34 and 0.44) and women (0.28 and 0.35). A second reason for this choice is that the Italian region under analysis, Veneto, is characterized by a large number of small and manufacturing-oriented firms, which tend to be located in the same areas of the region, forming high-density conglomerates that specialize in narrowly-defined activities. Thus, manufacturing firms within local labour markets tend to be quite similar in terms of product and labour market structures, and they employ a large proportion of Veneto's workforce. The left panel of Table 2 shows indeed that around 56% of our sample is employed at manufacturing firms and that this proportion is relatively high for both, men and women. The third reason motivating our focus on secondary-sector firms is given by the fact that the gender wage gap, even when conditioned on standard controls for human capital, is higher in this industry. The right

³³We have excluded from the analysis also one very marginal sector, tobacco. For this industry, only one firm was observed in the final sample.

Table 3: **Descriptive Statistics by Gender in the Selected Sample**

	Women		Men	
	<i>Mean</i>	<i>St. Dev.</i>	<i>Mean</i>	<i>St. Dev.</i>
Log wage	4.665	0.281	4.900	0.376
Age	33.7	9.2	35.6	9.5
Tenure	6.6	6.9	6.7	6.8
Firm size*	380.7	882.3	343.1	806.7
Fixed-term	6.7%		4.7%	
Blue collar	70.8%		74.1%	
White collar	29.0%		24.2%	
Manager	0.1%		1.7%	
N. firms	11,799		11,799	
N. workers	167,630		237,549	
N. observations	607,759		853,125	

*: Firm size is computed as number of full-year equivalent workers (total days worked in a year by a gender group within the firm, divided by 320).

panel of Table 2 shows indeed that the conditional pay gap between men and women is 3% higher at manufacturing firms.

As a final step, in order to further limit the bias in the measurement of firm wage policies, we have considered only companies where at least 15% of full-time workers employed during the period 1996-2001 were either men or women. Table 3 summarizes the main characteristics of the workforce by gender, considering only employees in the secondary sector, working in Veneto's local labour markets and at a gender-balanced firm. As can be noticed, the raw gender wage gap is of about 23%. Moreover, women are slightly more likely to work in clerical occupations and larger firms, younger, over-represented among fixed-term contracts and less likely to be managers. However, most of these differences are relatively small in magnitude.

5 Empirical Results

5.1 AKM Regression Results

We have estimated the AKM regression model presented in Section 3.1 separately by gender on the entire population of Veneto's private sector workers, considering the six-years

period between 1996 and 2001. Then, in order to analyse the gender wage gap in firms' pay policies, we have considered only gender-balanced manufacturing firms, selected along the lines discussed in Section 4.

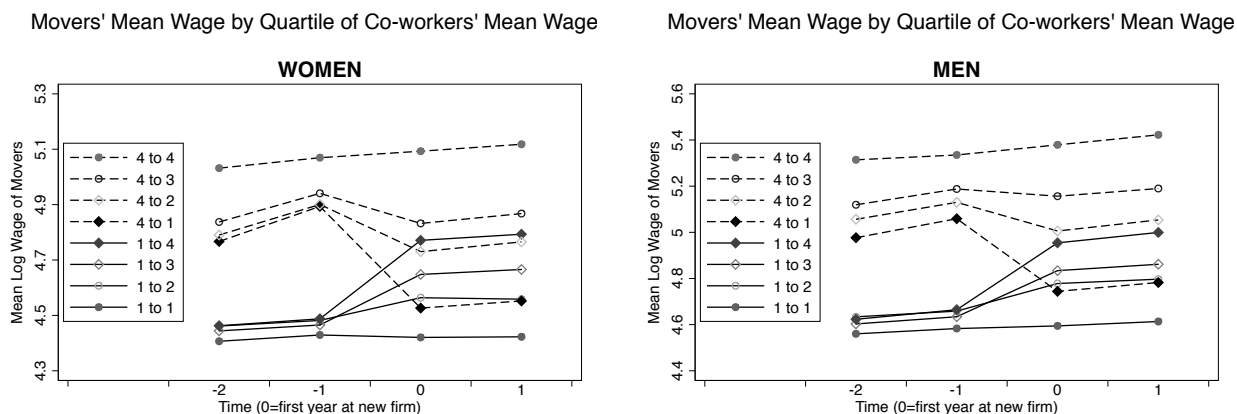
Before showing the AKM regression results, we provide some evidence on whether the assumptions of this model can be considered reasonable in our sample. Following Card et al. [2016], we test whether (gender-specific) co-workers' wages have a good predicting power for the pay of workers who change their job. Figure 2 computes the average wage of employees who work at two different firms for two consecutive years, where such average is computed by quartiles of co-worker wages at origin and destination firms.³⁴ As can be noticed, employees who move upward, from low-paying firms to high paying ones, face wage gains that are relatively symmetric to wage losses faced by workers who move in the opposite direction. Similarly, job movers who stay in the same quartile have a relatively flat wage dynamic.

The evidence provided by Figure 2 suggests that match effects, such as job-specific productivity shocks, have a limited impact on wages. Indeed, firm-specific factors influence individual earnings of both genders in a fairly similar way between job stayers and job movers. A second evidence supporting the AKM assumptions is the relatively parallel trend followed by job movers at origin and destination firms. The only discrepancy is the relatively flatter trend of workers in the fourth quartile of origin who move to a firm in the same quartile, with respect to the more pronounced seniority profile observed for other workers in the fourth quartile of origin who move downward. However, this tendency is present among both men and women and it is quite small in magnitude.

Table 4 summarizes the results of the AKM model, separately for men and women, on the sample of gender-balanced manufacturing firms belonging to one of Veneto's local labour markets. It can be noticed that overall wage dispersion is higher among men. In both cases, the largest contribution to total wage variance is given by the joint effect of individual time-varying and time invariant characteristics and by the returns to such endowments. Moreover, the regression residual is larger among women, implying that the

³⁴The figure shows transitions from the first and fourth quartiles only. Similar patterns are observed also for quartiles at the middle of the co-workers' wage distribution.

Figure 2: Mean Wages of Job Changers at Manufacturing Firms, Classified by Quartile of Co-Worker Average Wages at Origin and Destination Firm



The figure shows the mean wage of workers who change job and work at least two years with two different employers during the period 1996-2001. Jobs are classified by quartile of gender-specific co-worker wages in the last year at the old job (origin firm) and in the first year at the new job (destination firm).

Table 4: AKM Results Among Gender-Balanced Manufacturing Firms

	Women		Men	
$\text{var}(\omega_j)$	0.015	19%	0.014	10%
$\text{var}(x_{it}\beta + \eta_i)$	0.049	63%	0.111	79%
$2 * \text{cov}(\omega_j, x_{it}\beta + \eta_i)$	-0.005	-7%	0.007	5%
$\text{var}(e_{it})$	0.020	25%	0.008	6%
$\text{var}(w_{it})$	0.079	100%	0.141	100%
RMSE of AKM model	0.168		0.107	
RMSE of match model	0.161		0.089	
Variance of match effects	0.002		0.003	
N. Observations	607,759		853,125	

The table presents the wage variance decomposition based on the AKM regression model. The parameters of the regression are estimated separately by gender on the entire database of Veneto's private sector. Results of the table are instead computed considering only the sample of gender-balanced manufacturing firms selected along the lines discussed in Section 4. The variance of match effects is estimated as the difference in mean squared errors of the AKM model and of a regression with separate fixed effects for each worker-firm pair, adjusting for differences in degrees of freedom between the two models.

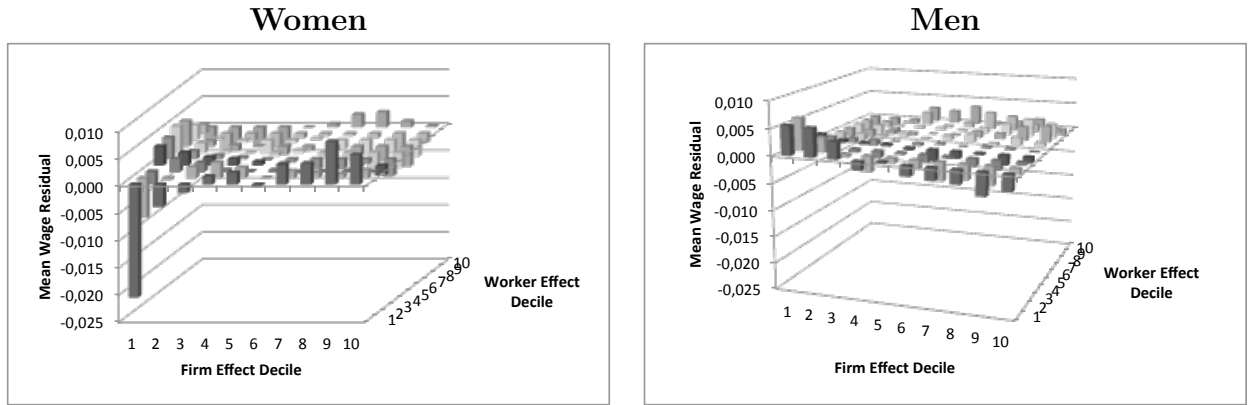
model fits better the data in the case of men. The relatively worst fit of the model in the female sample is also reflected by the negative sorting term observed among women, as this component tends to be biased downward the higher the measurement error (see, among others, by Andrews et al. [2008]).

In the context of the present analysis, the most interesting element of earnings variability is represented by the variance of firms' pay policies. Firm wage premiums provide a larger contribution to wage dispersion for women (19%) than for men (10%). In general, this result is consistent with our theoretical model, given that taste-based discrimination represents an element of variability in firm wage policies that is absent in the case of men. Also a more negative correlation between firm wage policies and human capital, at least in principle, could be considered consistent with our theoretical model, given that firm's wage policies paid to women should be considered distorted by taste-based discrimination. However, it can't be neglected the possibility that this result could be mostly driven by a larger measurement error in women's firm-specific wage residuals.

To investigate this issue, we have compared the fit of the model of the AKM regression with an alternative specification, in which each worker-firm pair effect is estimated by a separate dummy variable. This alternative specification captures any role played by job-match effects, which could be relevant if, due to factors associated to matching quality or to differences in wage posting behaviour along unobserved dimensions, firms were providing workers with highly heterogeneous compensation policies. Results presented in the lower part of Table 4 suggest that the performance of the linearly additive AKM model is quite similar to the one of the job-match model. Indeed, the implied variance of match effects explains at most 2.5% of the total wage variance for both, men and women. Thus, we can conclude that the relatively larger residual observed among women is not linked to an incorrect specification of the linearly additive model, but rather to idiosyncratic shocks that are not related to job match effects.

To further test the role of the regression error, and its potential impact on the estimated parameters of the AKM model, following Card et al. [2013] we have computed the average AKM error term for 100 bins of firms and workers, classified according to deciles of firm

Figure 3: Mean AKM Residual by Decile of Worker and Firm Effects



The figure shows the mean wage residual derived from the AKM regression model. The residual is computed separately by gender on 100 bins, classified by decile of estimated gender-specific person and firm effects. The sample is composed of 8,859 gender-balanced manufacturing firms considered for the analysis of the gender wage gap.

effects and worker effects. Figure 3 shows the result of this test, separately by gender.³⁵ Wage residuals within each bin do not tend to zero by construction (Card et al. [2013]), but should be reasonably small in absolute value by assumption. Indeed, systematically large positive (negative) errors for given groups of low- or high-wage workers and firms can be interpreted as an indirect evidence of omitted factors that could bias the estimates of the AKM parameters. Figure 3 shows that, even if there are some larger deviations among low-paid male and female employees in the first decile of workers' effects, all of the averages are quite small in absolute value, as they are uniformly below 2%.

5.2 Overall Impact of the Firm-Specific Gender Wage Gap

In this section, we provide descriptive evidences on standardized firms' wage policies estimated through the AKM model, considering gender differences in these parameters. We make employers' compensation policies comparable between men and women by expressing them as deviations from the (gender-specific) firm's wage residual of a common reference group.³⁶ As a result, we are able to analyse differences between how much male

³⁵In order to provide the most tailored picture of the distribution of the AKM residual, Figure 3 is computed considering Veneto's gender-balanced manufacturing firms that were not only active in 1996-2001, but also observed at least once between 1992 and 1995. Indeed, when testing for the presence of taste-based discrimination, we have used lagged proxy variables measured during the period 1992-1995, further restricting the sample.

³⁶As discussed in Section 3.2, we choose the largest employer to be the reference group. Similar results are obtained when considering the lowest paying percentile of firms as the reference.

workers are rewarded at a given workplace (with respect to men the reference group) and how much instead female workers within the same firm are rewarded (with respect to women at the reference group), all conditional on workforce composition.³⁷ Notice however that this is a *relative* measure of sex pay differences. Indeed, even if, in principle, all firms could be on average highly discriminatory toward women or not, this information can not be recovered using this method.

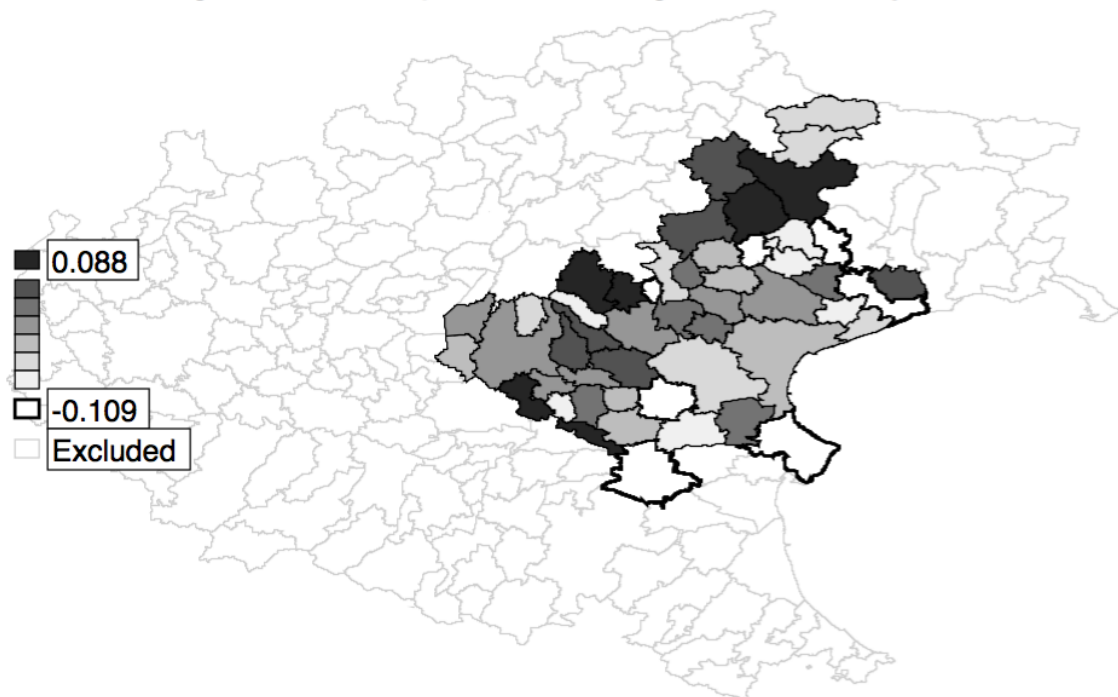
Adopting the notation introduced in presenting the empirical specification of our model, let ω_j^g represent the (standardized) firm wage policy of gender g at firm j . According to the implications provided by equation (7), gender differences in this parameter are driven by a composite effect represented by three main factors: a market-constant effect, α_k , denoting gender differences in employers' wage setting power, a taste-based discrimination parameter $\hat{\delta}_j$ and a residual term ρ_j .

Figure 4 provides a map of Veneto's local labour markets, as well as an overview on the size of $\omega_j^m - \omega_j^f$. For each district, we have computed the average gender gap in firms' wage policies, without weighting for employers' size. As can be noticed, the average level of $\omega_j^m - \omega_j^f$ across districts ranges between -0.1, that is, an average reduction of 10 percentage points in the gender gap in firms' policies with respect to the reference group, and 0.09. Moreover, even if with some exceptions, districts located toward the northern and more mountainous parts of Veneto tend to be darker in colour, *i.e.* they provide less favourable working conditions to women.

A different method to derive descriptive evidences on the size of the firm-specific conditional gender wage gap is to define a set of employers that, according to the metric given by $\omega_j^m - \omega_j^f$, provide less favourable working environments for women. For this purpose, we consider the cumulative distribution function (over firms) of $\omega_j^m - \omega_j^f$, which we denote by $F(\cdot)$. Then, using a standard human capital wage equation, we evaluate the marginal effect on wages of being a female worker employed in one of the firms in the right tail of the distribution $F(\cdot)$. We also evaluate this impact using the metric $\omega_j^m - \omega_j^f - \alpha_k$, where

³⁷In this sense, gender differences in firms' policies are informative about whether women at a given employer are paid differently compared to what they would get at a reference firm, while male firms' wage policies can be interpreted as a counterfactual, as they allow to make comparisons between female wage policies controlling for employer-specific factors equally affecting men and women.

Figure 4: Gender Gap in Firms' Policies by Local Labour Market
Average Gender Gap in Firms' Wage Residuals by District



The gender wage gap in firms' policies is computed as the difference between standardized male and female employer effect estimated through an AKM regression model. The average is computed over firms, without weighting for their size, and represent the percentage wage gain (or loss) experienced by women at a given firm, with respect to the gender gap in employers' policies observed at the reference firm.

market-constant effects α_k are estimated by a regression of $\omega_j^m - \omega_j^f$ on local labour market fixed effects interacted by two-digits sector fixed effects and three firm size dummies.

More precisely, we consider the following regression model

$$\ln w_i = b_1 1[g = m] + b_2 T_\theta + \beta x_i + \gamma 1[g = m] x_i + \eta_j + e_i$$

where η_j is a firm fixed effect (common for both gender groups), x_i is a vector of controls for observable individual characteristics (age and tenure quadratic polynomials, three occupation fixed effects and a fixed effect for open-ended contracts) and e_i is an error term. We interact all the variables in x_i with the gender dummy $1[g = m]$, in order to control not only for human capital characteristics, but also for sex differences in the returns to such characteristics.³⁸ The coefficient of interest in the above model is b_2 , which is associated to T_θ , an indicator variable that we define as

$$T_\theta = 1[g = m] 1[F(\mu_j) > \theta] \quad \mu_j = \begin{cases} \omega_j^m - \omega_j^f \\ \omega_j^m - \omega_j^f - \alpha_k \end{cases}$$

where θ is a given quantile of the distribution $F(\cdot)$. Thus, b_2 can be interpreted as the marginal effect on women of working in a less favourable environment. When the metric $\omega_j^m - \omega_j^f$ is adopted in defining more discriminatory firms, this coefficient provides the effect on the gender wage gap of being employed in workplaces whose pay policies are relatively lower among women than among men, conditional observable characteristics, gender differences in returns to such characteristics, and a gender-constant firm fixed effects. When the metric $\omega_j^m - \omega_j^f - \alpha_k$ is instead adopted, the definition of less favourable working environments becomes also conditional on all factors equally affecting women within a given market, *i.e.* on local labour market effects on women specific for each product-market structure and class of firm size.

As discussed in Section 2.3, we assume that monopsonistic discrimination is embedded in market-constant effects α_k . It would be tempting to attribute to monopsonistic discrim-

³⁸With this specification, not only characteristics effects, but also unexplained coefficient components of a traditional Oaxaca-Blinder gender wage gap decomposition (Oaxaca [1973]) are controlled for.

ination the difference between b_2 estimated using the two alternative definitions of μ_j , but we refrain from this structural interpretation of the parameter, given the well-known identification problems concerning the elasticity of the labour supply to the firm (*e.g.* Manning [2003]). Moreover, it should be noticed that in this model workers' fixed effects are not controlled for, thus also unobserved individual heterogeneity may influence our results. Nevertheless, comparing estimates of b_2 derived from different definitions of μ_j is interesting, as it allows to derive an indirect measure of the impact that factors, which are not specific of a given employer, but that are rather related to the labour and product market structure in which firms operate, have on the gender wage gap.

Figure 5 shows the results of the model discussed above estimated on the cross-section of workers observed in the year 1998, which is the largest in our sample.³⁹ The graph in the top panel shows how the coefficient b_2 varies when T_θ is defined using different percentiles of the distribution $F(\cdot)$ and different metrics μ_j of employers' attitudes toward women. In general, the treatment effect is always strong and significant. Women employed at more discriminatory firms suffer an additional wage loss with respect to men of between 3% and almost 10%, depending on how the treatment variable is defined.


These results imply that the gender wage gap conditional on observable individual characteristics grows substantially, with respect to its baseline level, in work environments less favourable to women. For example, if compared to a conditional gender wage gap of about 20% in the manufacturing sector (Table 2), results presented in Figure 5 suggest that, when θ is set equal to 0.5, this amount grows by around 30% at firms where gender differences in compensation policies are larger and by around 20% due to workplace-specific factors that are independent of the geographical location, product market structure and size of the firm.

Finally, notice that differences in marginal effects across models (given by the vertical distance between each coefficient in the top panel of Figure 5), even if significant for some choices of θ , are quite small in magnitude. This implies that market-constant effects α_k

³⁹We have chosen to apply this model on a cross-section in order to mitigate the problem of serial correlation in workers' unobserved abilities. Results are qualitatively similar when estimating a pooled OLS regression on the full sample.

Figure 5: Impact of the Gender Gap in Firms' Pay Premiums on the Cross-Sectional GWG (Year 1998)

Marginal Effect of Being Man in the Right Tail of Most Discriminatory Firms



Summary of Regression Results

Effect of Being Man Above the 50th Percentile of Most Discriminatory Firms

Dependent variable: log daily wage

	Model (1)	Model (2)
Firm environment metric (μ_j)	$\omega_j^m - \omega_j^f$	$\omega_j^m - \omega_j^f - \alpha_k$
Coefficients		
$1[g = m] * 1[F(\mu_j) > 0.5]$	0.057***	0.038***
<i>P-value</i>	(0.000)	(0.000)
$1[g = m]$	-0.037	-0.023
<i>P-value</i>	(0.220)	(0.440)
F tests		
Age and tenure polyn.	174.5***	161.7***
Interactions with $1[g = f]$	113.0***	106.1***
Main occupation dummies	1562***	1550***
Interactions with $1[g = f]$	248.0***	242.6***
All covariates	1759***	1924***
Adjusted R^2	0.636	0.635
RMSE	0.212	0.213
N. firm effects	9433	9433
N. of observations	239,295	239,295

*S.e. clustered at firm level. Significance levels: ***: 1%; **: 5%; *: 10%*

Results of regressions of log wages in year 1998 on a dummy for male workers at firms in the right tail of the distribution $F()$ of μ_j . The graph plots treatment effects and 95% CI for different percentiles of $F()$ and different definitions of μ_j . The table summarizes results when the treatment is being male and working in firms above the median of $F()$. Model (1) uses $\omega_j^m - \omega_j^f$ to characterize attitudes of employers toward women, Model (2) uses instead $\omega_j^m - \omega_j^f - \alpha_k$. All regressions include controls for human capital interacted by gender and a full set of firm fixed effects.

explain a relatively small proportion of the within-firms conditional gender gap.⁴⁰ Most of the variability of this gap seems to be employer-specific rather than market-specific, thus more likely to be linked to other factors such as taste-based discrimination or gender differences in compensating wage differentials. In the next section, we provide a more direct and robust assessment of the relevance and size of these latter mechanisms.

5.3 Testing for the Presence of Taste-Based Discrimination and Compensating Wage Differentials

We now discuss the results obtained by estimating the model that was presented in Section 3.3. In particular, we use the gender gap in firms' policies as the dependent variable of an OLS regression, in order to test whether this difference can be predicted by proxy variables for taste-based discrimination and gender differences in compensating wage differentials. The regression equation reads as follows

$$\omega_j^m - \omega_j^f = \alpha_k + b_1 \hat{\delta}_j^1 + b_2 \hat{\delta}_j^2 + b_3 \rho_j^1 + \hat{\delta}_j^r + \rho_j^r$$

where $\hat{\delta}_j^r + \rho_j^r$ is a composite residual, while α_k represent market-constant effects. We approximate α_k by three firm size dummies,⁴¹ two-digits sector fixed effects and thirty dummies for each of the local labour markets of the Veneto region, interacting all of these variables in some model specifications.

We have approximated taste-based discrimination using two variables, $\hat{\delta}_j^1$ and $\hat{\delta}_j^2$, representing the presence of women at the top of the corporate hierarchy and the female share of workers. In order to address the problem of correlation with the residual term, as discussed in Section 3, we have used lagged values of these two proxies, by computing them over the period 1992-1995.⁴² Given that explicit information on firms' ownership and management structures was not available, we have used as independent variable a

⁴⁰This finding seems coherent with the rather limited evidence available on this topic. In particular, Webber [2016] shows that monopsonistic discrimination is more driven by sorting of men and women across labour market structures, rather than by within-firms gender differences in the supply elasticity.

⁴¹We choose the three classes based on symmetric quantiles of the firms' size distribution. In practice, the three groups represent workplaces with less than approximately 7 full-time equivalent workers per year, between 7 and approximately 17, and above the latter threshold.

⁴²As a consequence, only firms observed at least once during the period 1992-1995 enter in our sample.

Table 5: **Summary Statistics on Proxy Variables for Taste Based Discrimination and Compensating Wage Differentials**

Variable	Mean	St. Dev.
Female manager (1992-1995)	15.02%	
Female share (1992-1995)	0.433	0.245
Part-time share	0.072	0.096
Observations	8,859	
All statistics are computed over firms. The number of observations refers to all firms observed during the period 1992-1995 that could be merged with the 1996-2001 sample.		

dummy for the presence of women in non-manual occupations that were receiving the highest observed yearly earning within the firm, where the highest pay is defined over all person-year observations in the period 1992-1995. For firms with more than 60 person-year observations, we have relaxed this definition and considered as female managers also those women in non-manual occupations that were among the top 3% yearly income earners *and* one of the top 10 earners among all person-years observations of a given workplace. Instead, we have computed the female share as the ratio of the weighted monthly average number of workers by gender within firms.

The other variable of interest in the regression model is ρ_j^1 , which we define as the part-time share within firms, measured as the ratio of full-time equivalent days worked part-time in a year over total days worked. Given that we estimate the AKM regression model excluding part-time workers, this control is less subject to endogeneity problems, such as correlation with measurement error, and consequently it is not measured in lags. Table 5 provides descriptive statistics for the independent variables of interest, computed on the sample of analysis. As can be noticed, around 7% of days worked within firms are part-time, the ratio of women is on average slightly above 40% across workplaces, while only 15% of companies are led by women.

In order to test more nuanced hypotheses on the mechanisms driving the gender gap in firms' pay policies, we have also performed an heterogeneity analysis by changing the dependent variable of the model. In particular, we have considered the gender gap in firm

fixed effects interacted by occupation (manual or non-manual), as estimated through an AKM regression by gender. This variable represents sex differences in firm wage policies specific of blue-collar workers, which are standardized and expressed in deviation with respect to those of the larger employer in order to make them comparable between men and women.⁴³ Thus, using this dependent variable we can test whether taste-based discrimination and gender differences in compensating wage differentials have heterogeneous effects on blue collars, or whether they evenly affect the entire workforce.

Table 6 summarizes the results of these regression models. When looking at the left panel of the table, where firm fixed effects are not interacted by occupation, it can be noticed that the presence of women at the top of the hierarchy is a significant predictor of the gender gap in employers' wage policies. In particular, this gap reduces by around 1 percentage point at workplaces lead by females, but the confidence interval is relatively large. When compared to a raw gap of around 20% among manufacturing-sector workers, this point estimate implies a reduction of about 5% in the gender wage gap. Instead, when compared to a within-firms conditional gender gap determined by employers' compensation policies of about 4% (Figure 5), this implies that about one fourth of this residual difference can be linked to the presence of women at the top of the corporate hierarchy. Similarly, a 10 percentage points increase in the female share of workers is associated to a reduction in differences between ω_j^m and ω_j^f of around 0.2 percentage points, which translates into a 1% reduction of the raw wage gap and into a 5% reduction in the conditional gender gap within firms.

Notice however that when the dependent variable of the model is changed and only manual workers are considered (right panel of Table 6), results differ. Indeed, while the point estimates of the impact of being in workplaces with more women becomes more negative and more significant (up to about 0.5 percentage points reduction for each 10 percentage points growth in the female share), the presence of female managers is instead not associated with a significant reduction in the gender wage gap in this case. The reason for

⁴³When the dependent variable is the gender gap firms' wage policies paid to manual workers, the sample size reduces. This is because only firms hiring both, male and female blue collars belonging to the largest connected set can be included in the analysis.

Table 6: Regression Results on Taste-Based Discrimination and Compensating Wage Differentials

Dependent variable Model	Gender gap in firms' premiums			Gender gap in blue collar firms' premiums		
	(1)	(2)	(3)	(1)	(2)	(3)
Coefficients						
Female manager (1992-1995)	-0.010**	-0.012**	-0.010*	-0.002	-0.002	-0.004
<i>P-value</i>	0.050	0.032	0.085	0.700	0.695	0.589
Female share (1992-1995)	-0.023**	-0.019**	-0.020*	-0.038***	-0.042***	-0.039***
<i>P-value</i>	0.011	0.046	0.060	0.000	0.000	0.003
Part-time share	0.054*	0.056*	0.081*	0.074**	0.083**	0.115**
<i>P-value</i>	0.066	0.072	0.090	0.022	0.018	0.035
F tests						
Firm size f.e.	9.97***	7.85***		5.83***	3.84**	
Sector f.e.	2.52***			0.90		
District f.e.	2.47***			2.27***		
District*sector f.e.		1.31***			1.13**	
District*sector*firm size f.e.			1.19***			1.08**
All covariates	3.71***	5.19***	3.63***	2.33***	4.99***	4.43***
Adjusted R^2	0.016	0.033	0.025	0.012	0.016	0.017
RMSE	0.180	0.179	0.179	0.190	0.190	0.189
N. of observations	8,859	8,859	8,859	6,896	6,896	6,896

*S.e. clustered at 2-digit sector and municipality level. Significance levels: * ** : 1%; ** : 5%; * : 10%*

The right panel of the table summarizes the results of regressions of proxies for taste-based discrimination and compensating wage differentials on the gender gap in firms' pay policies. The units of observation are individual firms. Basic controls include fixed effects for three classes of firm size, two-digits sectors and local labour markets, which are then interacted in Model (2) and Model (3). The left panel shows results obtained by changing the dependent variable, which becomes the gender gap in firm wage policies paid to blue collar workers, as estimated by interacting firm f.e. with occupation in an AKM regression model.

this difference in results between models can be mainly attributed to two factors. First, women at the top of the hierarchy may provide more favourable working conditions for female workers only in the case of clerical occupations. However, a second mechanism is related to our sample selection procedure. When we consider AKM firm fixed estimated for blue collars only, we tend to exclude firms where women are mostly employed in clerical occupations and also more likely to be at the top of the corporate hierarchy. For this reason, the lack of robustness in the result related to the presence of female managers may also be attributed to a problem of simultaneity, which could be absent in the specification where only blue collars' wage policies are considered.

Finally, notice that, when studying the role of compensating wage differentials, Table 6 shows that, irrespective of the choice of dependent variable, firms with a relatively higher propensity of providing part-time contracts are able to pay women relatively less. This mechanism can induce a growth in the conditional firm-specific gender gap as high as 1 percentage point for each 10 percentage points increase in the share of days worked part-time at a given workplace. Moreover, this effect seems to be stronger and more significant among blue-collar workers.

Overall, our results can be interpreted as evidence that taste-based discrimination and compensating wage differentials both play a significant role in driving the gender gap in firms' compensation policies within workplaces. This evidence seems coherent with the most recent studies adopting a similar approach. In particular, Bruns [2018] shows that firms opting out from centralized collective agreements, which arguably exert more wage setting behaviour and which are able to better exploit incentives provided by their preferences toward women or by female workers' hedonic considerations, tend to show larger gender gaps in compensation policies in Germany. The presence of wage penalties for full-time women at firms where more flexibility is available is a relatively novel evidence, given that most studies focus instead on penalties among part-timers (*e.g.* Elsayed et al. [2017]), but this finding is coherent with studies on women's preferences for shorter schedules (*e.g.* Del Boca [2002] and Booth and van Ours [2013]).

Partly due to the limited precision in measuring employers' preferences toward women

and employee’s non-wage amenities, both mechanisms that we have documented seem to have a quantitatively small impact on the overall gender wage gap. With respect to discriminatory tastes, our evidence on the importance of having women in managerial positions is mixed. This factor seems relevant for female wages only when women in non-manual occupations are included in the regression. Instead, the female share of workers within firms seems to be a more robust proxy for preferences toward women. However, adopting the theoretical perspective of our model, this parameter can be considered as a coherent predictor of taste-based discrimination only when conditioned on a firm’s labour market structure and size.

6 Conclusions

In this paper, we have shown that a simple static model of taste-based discrimination in monopsonistic labour markets provides a coherent framework to interpret the gender gap in firms’ wage policies. This component of the earning differential between men and women is estimated through an AKM regression model (Abowd et al. [1999]), and its importance has been documented, in different contexts, by several recent empirical studies (*e.g.* Card et al. [2016], Sin et al. [2017] and Bruns [2018]).

We have provided a theoretical discussion of the conditions under which this residual component of the gender wage gap can be attributed to elements such as taste-based discrimination or compensating wage differentials. Moreover, we have presented an empirical application, introducing methods to test for the presence of such mechanisms while controlling for important confounding factors, most notably gender differences in the labour supply to the firm and workers’ time-constant unobserved heterogeneity.

Using matched employer-employee data on Italy, we have shown that women working in the manufacturing sector suffer wage losses of up to 10%, with respect to men, due to factors that are independent of their characteristics and abilities, as they are instead related to firm-specific mechanisms. By documenting a positive relationship between this residual component of the gender wage gap and traditional proxies associated to discriminatory preferences, namely, the presence of women at the top of the firms’ hierarchy and the

female share of workers within firms, we have provided strong evidence on the presence of taste-based discrimination. However, partly due to the quality of these proxy variables, both of these effects are small in magnitude. Moreover, for what concerns the presence of female managers within firms, its effect on the gender gap was found not significant among women in manual occupations.

We have used the same approach to test for the presence of compensating wage differentials, using the share of the workforce under a part-time contract to approximate for non-wage amenities. We have shown that women tend to prefer work environments where more part-time contracts are available, as they seem willing to accept a lower pay in such places. Moreover, this negative effect on female wages is significant despite the fact that we consider only full-time employees, and seems to be stronger among workers at the bottom of the firm's hierarchy.

Our empirical findings are coherent with the implications of the static model presented in this paper, suggesting that this simple interpretative framework represents a useful tool for future research. Further promising applications of our method concern the design of tests on the implications of Becker's theory, among which the relationship between taste-based discrimination and the product market structure (*e.g.* Heyman et al. [2013]) or firms' survival rates (*e.g.* Weber and Zulehner [2014]), as testing such predictions often requires the definition of employer-specific discrimination parameters. Similarly, impact evaluations on policies such as the introduction of women quota in managerial boards (*e.g.* Matsa and Miller [2013]) can also derive useful insights by the regression approaches discussed in this paper.

More generally, given that the interpretative framework provided by this paper can be used to construct consistent tests on several firm-level mechanisms driving the gender pay gap, we believe that future research on affirmative action policies could derive useful results from this approach, improving our understanding of the most important drivers of discriminatory differences in wages in several contexts.

References

- ABOWD, J., F. KRAMARZ, AND D. MARGOLIS (1999): “High Wage Workers and High Wage Firms,” *Econometrica*, 67, 251–333.
- ABOWD, J. M., R. H. CREECY, AND F. KRAMARZ (2002): “Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data,” Longitudinal Employer-Household Dynamics Technical Papers 2002-06, Center for Economic Studies, U.S. Census Bureau.
- ANDREWS, M. J., L. GILL, T. SCHANK, AND R. UPWARD (2008): “High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias?” *Journal of the Royal Statistical Society Series*, 171, 673–697.
- ASHENFELTER, O. AND T. HANNAN (1986): “Sex Discrimination and Product Market Competition: The Case of the Banking Industry,” *The Quarterly Journal of Economics*, 149–173.
- BARTH, E. AND H. DALE-OLSEN (2009): “Monopsonistic Discrimination, Worker Turnover, and the Gender Wage Gap,” *Labour Economics*, 16, 589–597.
- BECKER, G. (1957): *The Economics of Discrimination*, University of Chicago Press, Chicago.
- BLACK, D. A. (1995): “Discrimination in an Equilibrium Search Model,” *Journal of Labor Economics*, 309–334.
- BLACK, S. E. AND E. BRAINERD (2004): “Importing Equality? The Impact of Globalization on Gender Discrimination,” *Industrial and Labor Relations Review*, 57, 540–559.
- BLAU, F. D. AND L. M. KAHN (2017): “The Gender-Wage Gap: Extent, Trends, and Explanations,” *Journal of Economic Literature*, forthcoming.
- BOAL, W. M. AND M. R. RANSOM (1997): “Monopsony in the Labor Market,” *Journal of Economic Literature*, 35, 86–112.
- BOOTH, A. L. AND J. C. VAN OURS (2013): “Part-Time Jobs: What Women Want?” *Journal of Population Economics*, 26, 263–283.
- BRUNS, B. (2018): “Changes in Workplace Heterogeneity and How They Widen the Gender Wage Gap,” *American Economic Journal: Applied Economics*, forthcoming.
- CARD, D., A. R. CARDOSO, J. HEINING, AND P. KLINE (2018): “Firms and Labor Market Inequality: Evidence and Some Theory,” *Journal of Labor Economics*, 36, S13–S70.
- CARD, D., A. R. CARDOSO, AND P. KLINE (2016): “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women,” *The Quarterly Journal of Economics*, 131, 633–686.
- CARD, D., J. HEINING, AND P. KLINE (2013): “Workplace Heterogeneity and the Rise of West German Wage Inequality,” *The Quarterly Journal of Economics*, 128, 967–1015.

- CARDOSO, A. R. AND R. WINTER-EBMER (2010): “Female-Led Firms and Gender Wage Policies,” *Industrial and Labor Relations Review*, 64, 143–163.
- CHARLES, K. K. AND J. GURRYAN (2008): “Prejudice and Wages: An Empirical Assessment of Becker’s The Economics of Discrimination,” *Journal of Political Economy*, 116, 773–809.
- DEL BOCA, D. (2002): “The Effect of Child Care and Part Time Opportunities on Participation and Fertility Decisions in Italy,” *Journal of Population Economics*, 15, 549–573.
- DEPEW, B. AND T. A. SRENSEN (2013): “The Elasticity of Labor Supply to the Firm over the Business Cycle,” *Labour Economics*, 24, 196–204.
- ELSAIED, A., A. GRIP, AND D. FOUARGE (2017): “Computer Use, Job Tasks and the Part-Time Pay Penalty,” *British Journal of Industrial Relations*, 55, 58–82.
- FLABBI, L., M. MACIS, A. MORO, AND F. SCHIVARDI (2014): “Do Female Executives Make a Difference? The Impact of Female Leadership on Gender Gaps and Firm Performance,” IZA Discussion Papers 8602, Institute for the Study of Labor (IZA).
- GAGLIARDUCCI, S. AND M. PASERMAN (2015): “The Effect of Female Leadership on Establishment and Employee Outcomes: Evidence from Linked Employer-Employee Data,” *Research in Labor Economics*, 41, 341–372.
- GIULIANO, L., D. I. LEVINE, AND J. LEONARD (2009): “Manager Race and the Race of New Hires,” *Journal of Labor Economics*, 27, 589–631.
- (2011): “Racial Bias in the Manager-Employee Relationship An Analysis of Quits, Dismissals, and Promotions at a Large Retail Firm,” *Journal of Human Resources*, 46, 26–52.
- GLOVER, D., A. PALLAIS, AND W. PARIENTE (2017): “Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Stores,” *The Quarterly Journal of Economics*, 132, 1219–1260.
- HELLERSTEIN, J. K., D. NEUMARK, AND K. R. TROSKE (2002): “Market Forces and Sex Discrimination,” *Journal of Human Resources*, 37, 353–380.
- HEYMAN, F., H. SVALERYD, AND J. VLACHOS (2013): “Competition, Takeovers, and Gender Discrimination,” *Industrial and Labor Relations Review*, 66, 409–432.
- HIRSCH, B., T. SCHANK, AND C. SCHNABEL (2010): “Differences in Labor Supply to Monopsonistic Firms and the Gender Pay Gap: An Empirical Analysis Using Linked Employer-Employee Data from Germany,” *Journal of Labor Economics*, 28, 291–330.
- KAWAGUCHI, D. (2007): “A market Test for Sex Discrimination: Evidence from Japanese Firm-Level Panel Data,” *International Journal of Industrial Organization*, 25, 441–460.
- LORENZINI, F. (2005): “Distretti industriali e sistemi locali del lavoro 2001,” in *8 Censimento generale dell’industria e dei servizi*, ISTAT.
- MANNING, A. (2003): *Monopsony in Motion: Imperfect Competition in Labor Markets*, Princeton University Press.

- MATSA, D. A. AND A. R. MILLER (2013): “A Female Style in Corporate Leadership? Evidence from Quotas,” *American Economic Journal: Applied Economics*, 5, 136–69.
- MATSUDAIRA, J. D. (2014): “Monopsony in the Low-Wage Labor Market? Evidence from Minimum Nurse Staffing Regulations,” *The Review of Economics and Statistics*, 96, 92–102.
- MILLER, C. (2017): “The Persistent Effect of Temporary Affirmative Action,” *American Economic Journal: Applied Economics*, 9, 152–90.
- MUEHLEMANN, S., P. RYAN, AND S. C. WOLTER (2013): “Monopsony Power, Pay Structure, and Training,” *Industrial and Labor Relations Review*, 66, 1097–1114.
- OAXACA, R. (1973): “Male-Female Wage Differentials in Urban Labor Markets,” *International Economic Review*, 14(3), 693–709.
- OI, W. Y. AND T. L. IDSON (1999): “Firm Size and Wages,” in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 3 of *Handbook of Labor Economics*, chap. 33, 2165–2214.
- RANSOM, M. R. AND R. L. OAXACA (2010): “New Market Power Models and Sex Differences in Pay,” *Journal of Labor Economics*, 28, 267–289.
- RANSOM, M. R. AND D. P. SIMS (2010): “Estimating the Firm’s Labor Supply Curve in a “New Monopsony” Framework: Schoolteachers in Missouri,” *Journal of Labor Economics*, 28, 331–355.
- ROBINSON, J. (1933): *The Economics of Imperfect Competition*, Macmillan, London.
- SIN, I., S. STILLMAN, AND R. FABLING (2017): “What Drives the Gender Wage Gap? Examining the Roles of Sorting, Productivity Differences, and Discrimination,” IZA Discussion Papers 10975, Institute for the Study of Labor (IZA).
- SULIS, G. (2011): “What Can Monopsony Explain of the Gender Wage Differential in Italy?” *International Journal of Manpower*, 32, 446–470.
- WEBBER, D. A. (2015): “Firm Market Power and the Earnings Distribution,” *Labour Economics*, 35, 123–134.
- (2016): “Firm-Level Monopsony and the Gender Pay Gap,” *Industrial Relations: A Journal of Economy and Society*, 55, 323–345.
- WEBER, A. AND C. ZULEHNER (2014): “Competition And Gender Prejudice: Are Discriminatory Employers Doomed To Fail?” *Journal of the European Economic Association*, 12, 492–521.

DEPARTMENT OF ECONOMICS AND STATISTICS
UNIVERSITY OF TORINO
Corso Unione Sovietica 218 bis - 10134 Torino (ITALY)
Web page: <http://esomas.econ.unito.it/>
