Research Paper

I S F O L

# The Gender Wage Gap Among Italian Employees

## Evidences from the ISFOL PLUS Database

di Bernardo Fanfani

ISSN 2281-499X

Collana Isfol Research Paper | numero 26 – maggio 2015 | www.isfol.it

La collana *ISFOL Research Paper* raccoglie i risultati di ricerche o studi dell'ISFOL, anche in corso, e risponde all'esigenza di pubblicare in tempi rapidi dati (o selezioni di dati) ed elaborazioni compiute e rappresentative del patrimonio tecnico-scientifico dell'Istituto, per favorirne l'immediata visibilità e fruibilità da parte dei target di utenza.

*L'ISFOL, Ente nazionale di ricerca, opera nel campo della formazione, del lavoro e delle politiche sociali al fine di contribuire alla crescita dell'occupazione, al miglioramento delle risorse umane, all'inclusione sociale e allo sviluppo locale.*

*Sottoposto alla vigilanza del Ministero del lavoro e delle politiche sociali, l'ISFOL svolge e promuove attività di studio, ricerca, sperimentazione, documentazione e informazione, fornendo supporto tecnico-scientifico ai Ministeri, al Parlamento, alle Regioni, agli Enti locali e alle altre istituzioni, sulle politiche e sui sistemi della formazione e apprendimento lungo tutto l'arco della vita e in materia di mercato del lavoro e inclusione sociale.*

*Fa parte del Sistema statistico nazionale e collabora con le istituzioni europee. Svolge il ruolo di assistenza metodologica e scientifica per le azioni di sistema del Fondo sociale europeo ed è Agenzia nazionale del programma comunitario Erasmus+ per l'ambito istruzione e formazione professionale.*

**Presidente**: *Pier Antonio Varesi*
**Direttore generale**: *Paola Nicastro*

**Riferimenti**:
Corso d'Italia, 33
00198 Roma
Tel. +39.06.85447.1
web: www.isfol.it

La Collana *ISFOL Research Paper* è curata da *Claudio Bensi*, responsabile del Servizio per la comunicazione e la divulgazione scientifica

**Contatti:** editoria@isfol.it

Il lavoro utilizza i dati dell'Indagine ISFOL PLUS, curata dalla Struttura Lavoro e Professioni.

Questo testo è stato sottoposto con esito favorevole al processo di *peer review* interna curato dal Comitato tecnico scientifico dell'ISFOL.

**Autore**: Bernardo Fanfani
L'autore ringrazia Carolina Castagnetti, Emiliano Mandrone, Francesco Devicienti e i referee anonimi per aver fornito utili commenti e suggerimenti. Ringrazia, inoltre, le Università di Pavia e di Torino.

**Testo chiuso a** gennaio 2015

**Coordinamento editoriale**: *Costanza Romano*
**Impaginazione ed editing**: *Mafalda D'Onofrio*



ISBN 978-88-543-0090-3

## ABSTRACT


**THE GENDER WAGE GAP AMONG ITALIAN EMPLOYEES - EVIDENCES FROM THE ISFOL PLUS DATABASE**

*In this paper, we measure the amount of gender discrimination in the Italian labour market, using the ISFOL PLUS database and covering the period between 2005 and 2010. We adopt a quantile regression decomposition methodology, in order to measure the gender wage gap across the entire distribution of earnings. Moreover, we test the robustness of our estimates by comparing the quantile regression decomposition results when more dependent variables are added to the model. Finally, by comparing the results obtained from a random effect regression to those obtained using the Hausman-Taylor regression model, we test whether the estimated level of discrimination changes when correlation of time varying characteristics to the unobservable individual heterogeneity is allowed. The results show that, other things being equal, female employees earn around 10% less than men employees. This percentage is increasing with wages, sometimes reaching levels higher than 20% toward the top of the earnings distribution. The gap in earnings is increasing with age. This trend can be attributed to several factors, such as to a positive relation between seniority and discrimination. However, this relation could also be driven in part by an improvement of market potential among the younger generation of women. Moreover, the pay gap was reducing between 2006 and 2008, but it has increased since then. Finally, the results show that most of the gender wage differences are not attributable to individual characteristics, nor to segregation of women into less remunerated occupations. Therefore, most of the gender differences in income have to be attributed to differences in the wage structures. That is, men and women are rewarded differently for a given level of human capital. In order to better understand what could be driving this result, we will review some of the theories and evidences in the recent literature that have attempted to shed light on this phenomenon.*


**KEYWORDS:** *Gender Wage Gap, Discrimination, Quantile Regression, Oaxaca Decomposition.*


**LE DIFFERENZE SALARIALI DI GENERE TRA I DIPENDENTI ITALIANI - EVIDENZE DAI DATI ISFOL PLUS**

In quest'articolo si stima il livello di discriminazione di genere presente nel mercato del lavoro italiano, usando le annualità dal 2005 al 2010 del database ISFOL PLUS. Per questo proposito si utilizza una decomposizione basata sulla quantile regression, metodo che consente di stimare la diseguaglianza di genere lungo l'intera distribuzione dei salari. Inoltre, si conduce un'analisi sulla robustezza dei risultati, confrontando le stime ottenute aggiungendo un maggior numero di variabili dipendenti. Infine, si confrontano i livelli di discriminazione ottenuti da una regressione col metodo dei random effects, rispetto ai risultati ottenuti da una regressione col metodo Hausman-Taylor, metodo che è valido anche nel caso di correlazione tra abilità individuali non osservabili e caratteristiche osservabili che abbiano una sufficiente variabilità temporale. I risultati ottenuti mostrano che le donne, tenuto costante un determinato livello di capitale umano, guadagnano circa il 10% in meno dei dipendenti maschi. Questa percentuale cresce all'aumentare dei salari, raggiungendo livelli talvolta superiori al 20% per i redditi più alti. Inoltre, le differenze salariali sono maggiori al crescere dell'età. Questo fenomeno può essere attribuito a diversi fattori, tra cui la presenza di una relazione crescente tra discriminazione e anzianità, ma potrebbe anche essere determinato, almeno in parte, da un miglior livello di capitale umano tra le giovani donne, rispetto a quello degli uomini della stessa generazione. Dall'analisi dei dati emerge inoltre che le differenze di reddito si sono ridotte tra il 2006 e il 2008, ma sono diventate crescenti tra il 2008 e il 2010, in corrispondenza dei primi anni della recente fase di recessione economica. Infine, i risultati ottenuti mostrano che la maggior parte delle differenze salariali di genere non si può attribuire alle caratteristiche individuali dei lavoratori, e nemmeno alla segregazione femminile in occupazioni meno remunerative. Ne consegue che la gran parte delle differenze di genere è attribuibile a una diversa struttura dei salari. Ciò significa che uomini e donne sono retribuiti in maniera diversa per uno stesso livello di capitale umano. Per comprendere meglio l'origine di queste divergenze, cercheremo di fornire una panoramica quanto più completa delle teorie e delle evidenze emerse nella letteratura più recente riguardo a questo fenomeno.

**PAROLE CHIAVE:** differenze salariali di genere; discriminazione; quantile regression; decomposizione di Oaxaca.

**PER CITARE IL PAPER**: ISFOL, Fanfani B., *The Gender Wage Gap Among Italian Employees - Evidences from the ISFOL PLUS Database*, ISFOL, Roma, 2015 (Isfol Research Paper, 26)

# INDICE

# 1 INTRODUCTION

Nowadays most of the Western Countries, and Italy in particular, are facing the strong challenges posed by demographic trends. Low fertility rates, as well as increased life expectancy rates, are putting pressure on the sustainability of welfare systems. In order to respond to increasing levels of economic dependency[1], the most desirable paths are productivity improvements and a growth in labour force participation, since both solutions are needed to maintain the current level of welfare benefits.

In this context, it is important to investigate whether discrimination against groups of the population is playing an important role in the labour market. Discrimination represents an inefficiency, and it reduces productivity by determining a sub-optimal allocation of resources. Moreover, discrimination reduces labour force participation, since it lowers the incentives to work for those who suffer from it. In particular, discrimination against women is very harmful, since it involves a large proportion of the labour force.

In this paper, we are going to measure the amount of gender discrimination among Italian employees, using the 2005 to 2010 waves of the ISFOL PLUS database.[2] For this purpose, we are going to employ a quantile regression based decomposition, which is a methodology that allows to measure differences in discrimination among higher and lower remunerated job positions. We will review and derive carefully the chosen estimator. Moreover, we are going to discuss in detail which is the definition of discrimination adopted, what are the assumptions needed, and whether the findings are robust when imposing less demanding assumptions. Finally, we will try to derive some conclusions on what is driving the gender wage gap in the Italian labour market.

The estimated level of the gender pay gap in Italy amounts to around 10%, except at the upper tail of the wage distribution, where it increases substantially. Overall, our results are quite similar to the ones found by Christofides, L., A. Polycarpou, and K. Vrachimis (2010) and by Di Tommaso and Piazzalunga (2013). Both studies make use of the quantile regression based decomposition, analyse years that are covered by our data, and apply the Heckman (1979) procedure to take into account selection problems. Instead of the Heckman correction, here we will carry out several tests on the sensitivity and robustness of our results. Taking advantage of the longitudinal structure of the data, we will apply panel techniques, and in particular the Hausman and Taylor (1981) model, in order to take into account the problem of unobserved individual heterogeneity. However, one should be aware that correlation between *time invariant* characteristics and individual fixed effects is not allowed under the Hausman-Taylor model. In particular, we won't be able to obtain estimates of discrimination that are consistent under correlation of gender to unobservable abilities. However, we can rely on the fact that the rich set of information contained in the PLUS samples allows us to build a good model of wage prediction, and we will interpret our measures of discrimination as a composite residual effect.

---

[1] Economic dependency is defined as the ratio between the working population, which is financing welfare systems, and the inactive population.

[2] The database is available at http://www.isfol.it/open-data-delle-ricerche/isfol-microdati.

We will show that the amount of gender discrimination in the Italian labour market is substantial, and explains almost all of the gender differences in earnings. We will show some interesting patterns that may help understanding better where such gender differences originate. First, we will show that the gender wage gap is influenced by the business cycle, since it has not been constant across years. Second, we will show that it is increasing with age. Our model shows that this pattern is determined partly by a lower level of discrimination among younger workers, and in part by better market potential of younger women. However, the positive relation between seniority and discrimination could be driven by several dynamics, which we can't fully take into account. Finally, we will show that most of the gender differences in pay are determined *within* occupations and *within* sectors. That is, most of the differences in the Italian labour market are not the result of segregation of women into less remunerative sectors or occupations, but rather they are determined by different payment structures for similar jobs.

## 2 THEORETICAL FRAMEWORK

In economic literature, discrimination against sub-groups of the population had been an almost neglected topic until Becker seminal work, The Economics of Discrimination, published in 1957. According to his approach, discrimination in the marketplace can be modelled as an implicit transaction cost, by introducing a so-called taste for discrimination:

*Discrimination is commonly associated with disutility caused by contact with some individuals, and this interpretation is followed here. [...] To the employer [discrimination] represents a non-monetary cost of production, to the employee a non-monetary cost of employment, and to the consumer a non-monetary cost of consumption* (Becker, 1973, p. 15).

An empirical methodology to quantify the amount of discrimination in the marketplace has been introduced by Oaxaca (1973) and Blinder (1973). Since then, a growing number of publications have addressed this problem, and so called decomposition methodologies have been applied in many different contexts[3]. The Oaxaca-Blinder approach is based on the following definition of discrimination against sub-groups of the population

$$D = \frac{(\mathrm{w}_m/\mathrm{w}_f) - (\tilde{\mathrm{w}}_m/\tilde{\mathrm{w}}_f)^{ND}}{(\tilde{\mathrm{w}}_m/\tilde{\mathrm{w}}_f)^{ND}} \tag{1}$$

where $\mathrm{w}_i$ are wages observed for the (gender) group $i \in \{m, f\}$, and the superscript ND denotes the hypotetical and unkown wage ratio that would be observed in the absence discrimination. In a competitive equilibrium, due to the well known theory of cost minimization, the non-discriminatory ratio of wages $(\tilde{\mathrm{w}}_m/\tilde{\mathrm{w}}_f)^{ND}$ would equate the ratio of male and female marginal products. One way of approximating such marginal contributions is by estimating a mincerian wage equation, where labour income is considered a true measure of productivity, and it is predicted using a series of controls for human capital and other relevant individual attributes. Such wage equation usually takes the following semi-logarithmic functional form $\ln w_i = x_i \beta_i + \varepsilon_i$, where $x_i$ is a $n_i \times p$-matrix containing a constant, while $\beta_i$ is a $p$-vector of coefficients.

If there was no discrimination in the labour market, one would expect that individual characteristics were rewarded equally across gender groups. Stated differently, in a non-discriminatory market the wage structure faced by males would also apply to females[4]. Then, using the properties of OLS, it is straightforward to decompose differences in income as follows

[3] Hereafter, when using the term wage structure, we will refer to the pay schedule faced by individuals, given their set of skills. As we will see in a moment, we are assuming no general equilibrium effects.

[4] Hereafter, when using the term wage structure, we will refer to the pay schedule faced by individuals, given their set of skills. As we will discuss in more detail in a moment, we are assuming no general equilibrium effects. Broadly speaking, we are assuming that males are rewarded according to the competitive prices for skills, while women are discriminated. Instead, one could also construct the counterfactual distribution of wages where men skills are evaluated according to the female pay schedule, or according to a weighted combination of the two pays schedules.

$$\ln w_m - \ln w_f = \overline{x_m}\widehat{\beta_m} - \overline{x_f}\widehat{\beta_m} + \overline{x_f}\widehat{\beta_m} - \overline{x_f}\widehat{\beta_f} = (\overline{x_m} - \overline{x_f})\widehat{\beta_m} + \overline{x_f}(\widehat{\beta_f} - \widehat{\beta_m}) \qquad (2)$$

where bars represent mean values, while parameters estimated by applying OLS separately in the male and in the female samples are denoted with a hat. $\overline{x_f}\widehat{\beta_m}$ is a counterfactual wage, which measures the average wage that women would earn, had they been paid as men are. Equation (2) is the classical Oaxaca-Blinder decomposition. The term $(\overline{x_m} - \overline{x_f})\widehat{\beta_m}$ is the so called characteristics effect, which is coherent with the wage prediction model, since it is driven by mean differences in individual skills among the two groups. Instead, the second addend of (2) is the coefficient (or wage structure) effect. It measures differences in the way gender groups are rewarded for the same characteristics.
Using the decomposition (2), a logarithmic equivalent of the coefficient D in (1) can be defined as[5]

$$\ln(D+1) = \ln(\mathrm{w}_m/\mathrm{w}_f) - \ln(\tilde{\mathrm{w}}\,\mathrm{m}/\tilde{\mathrm{w}}\,\mathrm{f})^{ND} = (\ln w_m - \ln w_f) - (\overline{x_m} - \overline{x_f})\widehat{\beta_m} = \overline{x_f}(\widehat{\beta_f} - \widehat{\beta_m})$$

It is important to discuss what are the relevant assumptions that have to be satisfied, in order to correctly identify this discrimination coefficient. In particular, there are three conditions that have to be imposed.

- **No general equilibrium effects** We are assuming that the counterfactual (non-discriminatory) wage structure $\beta*$ will not be affected by the removal of discrimination from the labour market. In particular, in the context of the decomposition (2), we have assumed that the male pay schedule would prevail in a fair labour market, so that $\beta* = \beta_m$.[6]
- **Overlapping support** Either we assume that the estimated regression coefficients can be extended to combinations of covariates not observable in the data, or we need to restrict attention to combination of characteristics observable among both men and women.[7]
- **Ignorability** This is the most important assumption and, in its more general formulation, can be stated as follows. Let $F(.)$ represent the conditional distribution of the error term. Then

$$F(\varepsilon_m|x_m) = F(\varepsilon_f|x_f) = F(\varepsilon|\mathrm{x}) \qquad (3)$$

Classical exogeneity is not required, as long as the conditional distribution of the error term is the same among men and women[8]. However, imposing exogeneity of the regressors in the

[5] This definition of discrimination follows from the fact that $\overline{x_f}\widehat{\beta_m}$ is choosen as the counterfactual wage. See footnote 4 and the description of the no general equilibrium effects assumption for a more detailed discussion of this point.
[6] See Oaxaca and Ransom (1994) for a more detailed discussion of this assumption, and for possible alternatives.
[7] This problem has been addressed explicitly by Nopo (2008), who introduces a method based on a matching algorithm. He proposes a four-fold wage decomposition, where the additional components represent wage differences between matched and unmatched observations in each gender group.
[8] Such assumption is satisfied in the context of exogenous policies, when the members of each group we want to compare are randomly selected. Whenever this is not the case, a possible strategy is to adopt some form of correction for self-selection bias, such as the procedure developed by Heckman (1979). See Buchinsky (1998) for an analogous procedure in the context of quantile

wage equation of both groups is a sufficient condition for the ignorability assumption to be satisfied. Chernozhukov, Fernàndez-Val, and Melly (2013) have shown that, under ignorability, the discrimination coefficient can have a causal interpretation, in the sense that it reflects solely gender differences attributable to the wage structures.

It is important to recognize that many factors are likely to induce a violation of condition (3). On one hand, group-differences in unobservable individual heterogeneity could contribute substantially to the determination of the wage gap. To some extent, this bias can be limited through the use of panel techniques, like a within transformation of the data. However, these methods, which will be discussed and implemented in Section 6, have further limitations. Indeed, the more robust estimators often require the use of some instrument, which might be difficult to be found[9].

A second reason why condition (3) could be violated is due to the presence of firm-specific heterogeneity in wage compensation schemes. In their seminal work, Abowd, Kramarz, and Margolis (1999) have found evidence suggesting that the more profitable firms are those who tend to pay higher wages. In our context, a bias would arise if there were gender differences in the way workers sort into high-wage firms.

Card, Cardoso, and Kline (2013), using a Portuguese matched employer-employee database, have found that, even after controlling for firm fixed effects, the relation between firm profits and wages is weaker for women. This suggests that women do not systematically work for less profitable firm, but rather that they gain less than men from firms' profits. However, it is not clear whether this tendency could be found also in the Italian labour market, especially because female labour force participation is higher in Portugal. Using a similar matched employer-employee database for Italy, which covers the period 1996-2003, Matano and Naticchioni (2013) have found no evidence of under-representation of women in more profitable firms. Both of the evidences above suggest that the role played by firm fixed-effects in the determination of the gender earnings gap might be limited.

Finally, a reason that could explain the gender pay gap, but that is not accounted for by our framework, is given by the possible presence of gender-specific differences in labor supply elasticity to the firm. Indeed, one driver of discrimination could arise in monopsonistic labour markets,[10] if women had a rigid labour supply at the firm level. The monopsonistic employer could extract more rents from employees whit a rigid supply, and women could be disadvantaged for this reason. This type of

regression. See Di Tommaso and Piazzalunga (2013) for a recent estimation of the Italian gender wage gap using the the Heckman-correction.

[9] In particular, by using a simple a fixed effect regression, the coefficients of time invariant regressors, such as schooling, can't be recovered. A possible solution is to estimate the regression model proposed by Hausman and Taylor (1981), and to control the endogeneity of time invariant regressors using some instruments. See Polachek and Kim (1994) for a review of panel techniques for the estimation of the gender earning gap.

[10] A monopsony is a market with many sellers and only one buyer, who can extract a rent in a similar way to that of a monopolist. The monopsony could be a realistic framework for modelling labour markets, mainly due to the presence of frictions which reduce the mobility of workers and to the presence of excess labour supply. For a review of the recent literature on monopsony see Ashenfelter, Farber and Ransom (2010).

discrimination is usually referred to as *robinsonian discrimination*[11], and is conceptually different from the kind of discrimination defined by Becker. Indeed, according to the monopsonistic framework, discrimination would be an equilibrium outcome, rather than the product of inefficient and discriminatory markets. Sulis (2011) calculates the wage elasticities to the recruitment and separation rates for Italy, analysing the period between 1985 and 1996. He finds that women labour supply to the firm is significantly more rigid than the one of men. This result seems to provide indirect evidence supporting the hypothesis of *robinsonian discrimination*, even if the link between wage elasticities to the firm and discrimination, to the best of our knowledge, has not yet been established in the literature.

In general, given the nature of our data, which lack information about individual employers, we will not be able to test more elaborate hypothesis that could explain the presence of gender discrimination. Therefore, we should be aware that a decomposition of the wage structure effect into a component due to sorting of workers, a component due to labour market frictions, and a component attributable to pure bargaining effects, is not feasible. In general, the estimated discrimination coefficient should be interpreted more like a composite residual effect, rather than a simple employers' disutility parameter. Therefore, policies designed to reduce the levels of wage discrimination should take into account the complexity of this problem. Gender differences in pay likely originate from a variety of labour market characteristics, all of which should be taken into account in order to design effective policies.

[11] See Ransom and Oaxaca (2010) for a discussion and a more precise definition of the concept of *robinsonian discrimination.*

# 3 ECONOMETRIC METHODOLOGY

We are now going to illustrate the econometric methodology used to identify the discrimination coefficient described in the previous section. The chosen estimator, originally developed by Machado and Mata (2005) and Melly (2005), is usually called quantile regression decomposition. The main advantage of this technique, with respect to the traditional OLS decomposition, is that it allows us to study the wage gap along the distribution of income, and not only at the average level.
Quantile regression, as developed by Koenker and Bassett (1978), is based on the notion of conditional quantiles of the dependent variable y (log wages), given the covariates x (individual characteristics)[12]. Such relation is described by the following function

$$Q_y(\theta|x) = x\beta(\theta) \tag{4}$$

where $Q_y$ is used to indicate θth conditional quantile of y. The parameter β is estimated as the solution to the following problem

$$\hat{\beta}(\theta) = \operatorname{argmin}_{\beta \in R^p} \sum_{k=1}^{n} \rho_\theta(y_k - x_k\beta)$$

where n represents the sample size, the parameter space is given by $R^p$, and the $\rho_\theta$ loss function takes the form

$$\rho_\theta(u) = u[\theta - I(u < 0)] \qquad \theta \in [0, 1]$$

where $I(.)$ represents the indicator function, equal to one if the term inside the brackets is true and zero otherwise. In a finite sample, the number of distinct regression coefficients $\hat{\beta}(\theta)$ that can be estimated is finite. Let $\Theta = \{\theta_0 = 0, \theta_1,...,\theta_J = 1\}$ be the set of points where the solution changes, and notice that $\hat{\beta}(\theta_j)$ prevails from $\theta_{j-1}$ to $\theta_j$ for j = 1,..., J. Moreover, let $\hat{\beta} = \left[\hat{\beta}(\theta_1), \dots, \hat{\beta}(\theta_j), \dots, \hat{\beta}(\theta_J)\right]$ be the vector of all different quantile regression coefficients. Using such solutions, we can build a model for the estimated conditional quantiles $\hat{Q}_y$, in order to recover a conditional distribution of income[13]. Then, following Melly (2005), we will be able to construct a non-discriminatory wage structure using a particular unconditional (marginal) distribution. Let $F_Y$ represent the distribution function of the random variable y. The θth quantile of y is defined as $y_\theta = F^{-1}(\theta)$. Therefore

[12] Here, y is a n×1 vector and x is a n×p matrix.

[13] Using the entire set of solutions increases the risk of quantile-crossing. That is, $\hat{Q}y(\theta|x)$ could be non-increasing in θ when evaluated at a given x. In general, the larger the number of solutions used to approximate the conditional distribution of y, and the smaller the number of observations available, the greater becomes the risk of quantile-crossing. Notice however that using the entire set of solutions is not necessary, since a sufficiently large set $\Omega \subseteq \Theta$ of quantiles, which can be drown from a uniform on [0, 1], will produce valid results.

$$F_Y\left(y_\theta\right) = P\left(y \leq y_\theta\right) = \int_{-\infty}^{y_\theta} dF_Y(z) = \theta$$

The above result can also be obtained by integrating θ over the interval $[0, F_Y(y_\theta)]$[14]. Indeed, we have that

$$\int_0^{F_Y(y_\theta)} d\theta = \theta$$

Consider now the following indicator function, defined as

$$I[F_Y^{-1}(\theta) \leq y_\theta] = \begin{cases} 1 & if\ \theta \in \{z \in (0,1): F_Y^{-1}(\theta) \leq y_\theta\} \\ 0 & otherwise \end{cases}$$

Notice that integrating such function, with respect to θ, over the interval [0,1], is equivalent to integrate θ over the interval $[0, F_Y(y_\theta)]$. Given the above definitions, we can conclude that

$$F_Y\left(y_\theta\right) = \theta \Leftrightarrow \int_0^1 I\left[F_Y^{-1}(\theta) \leq y_\theta\right] d\theta = \theta \tag{5}$$

The probability distribution function of conditional quantiles will be denoted with $\hat{F}_{Y|X}$. It is obtained by taking the integral in (5), and substituting $F_Y^{-1}(\theta)$ with the expression for the estimator of conditional quantiles Qy(θ), which were defined by equation (4)

$$\hat{F}_{Y|X}\left(q_y \middle| X = x\right) = \int_0^1 \mathrm{I}\left[x\hat{\beta}(\theta) \leq q_y\right] d\theta = \sum_{j=1}^{J}(\theta_j - \theta_{j-1})\mathrm{I}\left[x\hat{\beta}(\theta) \leq q_y\right] \tag{6}$$

The shift from integration to summation is possible in a finite sample, since, as we have noticed earlier, there is a finite number J of distinct conditional quantiles, which are characterized by the J-vector of distinct solutions $\hat{\beta}$. For a given x, we can consider (6) to be the conditional distribution of income implied by quantile regression. It follows from their definition that conditional quantiles can be estimated from the distribution of y given x as

$$\hat{\mathrm{Q}}_{\mathrm{y}}(\theta|\mathrm{x}) = \inf\left\{q_y: \sum_{j=1}^{J}(\theta_j - \theta_{j-1})\mathrm{I}\left[x\hat{\beta}(\theta) \leq q_y\right] \geq \theta\right\} \tag{7}$$

Equation (7) is a convenient expression for conditional quantiles, since it is derived from their estimated conditional probability distribution. The next step is to obtain the marginal distribution of income, using

[14] We have to change the variable of integration, so that the random element will become the length of interval over which we integrate.

some basic properties of probability. Denote the marginal density of the variable y by $f_Y(z)$. Notice that such density can be written as a function of the conditional density $f_{Y|X}(z|w)$, and of the covariates' density $f_X(w)$ as follows

$$f_Y(z) = \int_{-\infty}^{\infty} f(z,w)dw = \int_{-\infty}^{\infty} f_{Y|X}(z|w) f_X(w)dw$$

where $f(z,w)$ is used to denote the joint probability. In our framework, the conditional density of income can be derived from (6). By integrating such function with respect to $F_X$, the result will be the unconditional distribution of y, denoted by $\hat{F}_Y$[15]

$$\hat{F}_Y = \int \hat{F}_{Y|X}\left(q_y \middle| X\right) dF_X = \int \left(\int_0^1 \mathrm{I}\left[x\hat{\beta}(\theta) \leq q_y\right] d\theta\right) dF_X \tag{8}$$

We now introduce some more notation. Let $i \in \{m, f\}$ represent male and female observations, so that we have two samples $\{(y_{ki}, x_{ki}): k = 1, \dots, n_i\}$ and the vector $\hat{\beta}$ can be estimated separately for the two groups. Moreover, consider $x_{ki}$ as the kth row of the $n_i \times p$ matrix $x_i$. The distribution $F_{X_i}$ of group i covariates can be approximated by the empirical distribution function as follows[16]

$$\hat{F}_{X_i}(z) = n_i^{-1} \sum_{k=1}^{n_i} I(x_{ki} \leq z), \qquad i \in \{m, f\}$$

where z is a given p-vector. Using the above distribution to evaluate $F_{X_i}$, we can estimate the unconditional distribution, expressing (8) as

$$\hat{F}_{Y_i} = \int \hat{F}_{Y_i|X_i}\left(q_{yi} \middle| X_i\right) dF_{X_i} = n_i^{-1} \sum_{k=1}^{n_i} \sum_{j=1}^{J} (\theta_j - \theta_{j-1}) \mathrm{I}\left[x_{ki}\hat{\beta}_i(\theta) \leq q_{yi}\right] \tag{9}$$

[15] We are exploiting the definition of marginal (unconditional) distribution, which is given by

$$F_Y = \int_{-\infty}^{\infty}\left(\int_{-\infty}^{\infty} f_{Y|X}(z|w) f_X(w)dw\right)dz = = \int_{-\infty}^{\infty}\left(\int_{-\infty}^{\infty} f_{Y|X}(z|w)\, dz\right) f_X(w)\, dw = E_X\left[\int_{-\infty}^{\infty} f_{Y|X}(z|w)\, dz\right]$$

[16] Machado and Mata (2005) propose a random sampling method to approximate the covariates distribution. The two approaches are compared by Melly (2006), who shows that they become identical as the number of covariates' drows in the Machado-Mata procedure tends to infinity.

In equation (9), each element of $\hat{\beta}_i$ is weighted by the length of the interval over which it prevails. Moreover, each row of $x_i$ contributes to the cumulative probability only if the resulting conditional quantile is lower than $q_{yi}$.[17] Unconditional (marginal) quantiles can now be estimated as

$$\hat{y}_i(\theta) = F_{Y_i}^{-1}(\theta) = inf\left\{q_{yi}: n_i^{-1} \sum_{k=1}^{n_i} \sum_{j=1}^{J} (\theta_j - \theta_{j-1}) \mathrm{I}\left[x_{ki}\hat{\beta}_i(\theta) \leq q_{yi}\right] \geq \theta\right\} \tag{10}$$

Using these results, we can turn to the problem of decomposing group-wage differences. Assume that in a non-discriminatory labour market, females would have the same income's conditional distribution of males, that is, their characteristics would be rewarded as if they were males. The next step is to build the counterfactual quantile, in order to have a distributional measure of what would be female income, had the wage structure been the same as the male one. Such non-discriminatory quantile of income can be estimated by integrating $\hat{F}_{Y_m|X_m}$, that is, the male conditional income distribution, with respect to the female distribution of characteristics, $\hat{F}_{X_f}$. Using this procedure, we can estimate the following counterfactual income distribution, denoted by $\hat{F}_C$

$$\hat{F}_C(q_c) = \int \hat{F}_{Y_m|X_m}\left(q_{ym}\middle|X_m\right) dF_{X_f} = n_f^{-1} \sum_{k=1}^{n_f} \sum_{j=1}^{J} (\theta_j - \theta_{j-1}) \mathrm{I}\left[x_{kf}\hat{\beta}_m(\theta) \leq q_C\right] \tag{11}$$

Equation (11) represents the non-discriminatory distribution of female income, assuming that the pay schedule faced by males would prevail in a fair labour market. Using the analogous procedure of equation (10), from the counterfactual distribution (11) we can obtain an estimator of the θth marginal quantile as follows

$$\hat{y}_C(\theta) = inf\left\{q_C: n_f^{-1} \sum_{k=1}^{n_f} \sum_{j=1}^{J} (\theta_j - \theta_{j-1}) \mathrm{I}\left[x_{kf}\hat{\beta}_m(\theta) \leq q_C\right] \geq \theta\right\} \tag{12}$$

Using the estimators in equations (10) and (12), we can carry out a wage gap decomposition similar to the traditional Oaxaca-Blinder decomposition, which was defined in equation (2). However, using the quantile regression approach, wage differences can now be evaluated at any θth quantile of the income distribution. More precisely, the pay gap between males and females can be divided in two parts, one representing the effect of different characteristics between the two groups, the other representing differences unexplained by the quantile regression model. For a given θ we can estimate

$$\hat{y}_m(\theta) - \hat{y}_f(\theta) = [\hat{y}_m(\theta) - \hat{y}_C(\theta)] + [\hat{y}_C(\theta) - \hat{y}_f(\theta)] \tag{13}$$

[17] Notice that the probability distribution (9) is not well defined if quantile-crossing occurs.

It is useful to remark that the first addend is the so-called characteristics effect, since it is the consequence of the different distribution of covariates for the two groups. On the other hand, the second addend in (13) represents the so-called coefficient effect, since it is obtained by evaluating female characteristics using two different conditional distributions. The asymptotic distribution of the counterfactual estimator has been studied by Chernozhukov, Fernández-Val, and Melly (2013), who have shown the validity of exchangeable bootstrap inference procedures to estimate the covariance matrix[18].

[18] Drowing n × r observations with replacement from the empirical distribution to compute r estimates, is an example of a valid exchangeable bootstrap inference procedure.

# 4 DATA

To analyse the Italian gender wage gap, we will use the 2005, 2006, 2008 and 2010 waves of the ISFOL Population, Labour, Unemployment Survey (PLUS). These data are collected through telephone interviews, which are conducted during the first quarter of the year. Since the first PLUS survey of 2005, each year includes a proportion of panel observations, some of which are present in all four waves[19]. The target population is composed of individuals between 15 and 64 years old[20], and the total sample size for each year is reported in the top part of Table 1.

The questionnaire is composed of specific sections designed to collect information on the following sub-groups of the population: young individuals between 15 and 29 years old; women between 20 and 49 years old; elderly population between 50 and 64 years old; unemployed individuals; employed population. A rich set of information for each of these categories is included, ranging from family characteristics to individual skills and personal history. It is then possible, for properly specified subsets of the sample, to provide a detailed explanation of some individual decisions (such as the choices of working, having children, studying) and to investigate which social environment factors may have influenced them.

**Table 1:ISFOL PLUS Samples Size and Composition**

| | Year | 2005 | 2006 | 2008 | 2010 |
|---|---|---|---|---|---|
| Total ISFOL PLUS | Males | 16.292 | 16.825 | 15.277 | 17.817 |
| | Females | 24.094 | 20.688 | 18.653 | 20.858 |
| | Total | 40.386 | 37.513 | 33.930 | 38.675 |
| Employee Sample Criteria of selection: full-time (30-72 h/week), wage at least 2 Euro/h, wage not imputed. | Males | 4.101 | 3.765 | 3.399 | 3.948 |
| | | 51,13% | 64,89% | 51,16% | 39,97% |
| | Females | 3.340 | 2.716 | 3.245 | 2.923 |
| | | 48,68% | 66,27% | 41,20% | 39,51% |
| | Total | 7.441 | 6.481 | 6.644 | 6.871 |
| | | 50,03% | 65,47% | 46,30% | 39,78% |

Percentages refer to panel observations (i.e. employees whose wages are reported in more than one ISFOL PLUS wave).
*Fonte: Dati ISFOL PLUS*

The bottom part of Table 1 reports the composition of the sample that we will be studying. Despite the fact that also self-employed and those with project-linked job positions are present in the PLUS data sets, for our analysis we have considered only salaried employees, which form the largest category of workers. This choice is motivated by the fact that data on income is not harmonized among the above categories of job position. Moreover, in the context of gender wage gap estimation, we want to study a sample of workers that should be as homogeneous as possible. Indeed, the lack of overlapping support, an assumption discussed in Section 2, may become more relevant as we increase the

[19] For a detailed illustration of the features of this survey and of the sampling design, see Mandrone (2012), chapter 9.
[20] In 2005, there were 38, 827, 322 individuals aged 15-64 in Italy, while the same population was 39, 655, 921 in 2010.

categories of job position included. In order to gain further accuracy in our results, we have decided to include only full-time workers, and we have excluded all outliers and those who have chosen to not report their wages[21].

We have used log-hourly net income (adjusted to the 2010 level) as the dependent variable. In Figure 1 we have graphed the Kernel density estimates of the wage distribution by year and by gender[22]. We can see that the modal observation is always lower for women, and that their income distribution tends to be shifted toward the left with respect to the male one. Both tendencies are a preliminary evidence of the presence of a gender pay gap. Such impression is confirmed when looking at Table 2, where we can see that the average level of log hourly net wages is always lower for women. Notice also that 2008 seems to be the year with the lowest gender pay gap.

Table 2 reports some descriptive statistics, by year and by gender, for the most important controls of our model. Notice that, in all years, women tend to be younger and better educated than men. Moreover, the proportion of long-term contracts is lower for the female group, while the proportion of public employees is lower in the male group. Notice also that mean log wages are always lower for female workers, while the dispersion of wages, as measured by the standard deviation, is similar for the two groups. Finally, there are no major differences in the sample composition across years, at least with respect to these variables.

As mentioned in Section 2, to carry out the decomposition exercise, we need to build a valid model of wage prediction. For this purpose, we have selected a rich group of independent variables. Specifically, the controls of the model are: years of schooling; a quadratic term for market experience, as approximated by age; a dummy for tenure, denoting employees who have been in their current job position for less than two years[23]; type of contract (long- or short-term); family characteristics (marital status, presence of pre-schooling age children, education of the mother); sector (services/goods production and public/private sectors); four occupation dummies and a dummy for firms with more than 50 employees; geographic variables (denoting people living in urban areas, people living in the North, with an interaction for North-West, and those living in the South, with an interaction for insular regions). In Section 6 we will compare the results obtained with this model specification to the ones obtained by adding nine occupation dummies, nineteen sectoral dummies and the entire set of Italian Regions.

[21] More precisely, we have kept only individuals who were working between 30 and 72 hours a week, and whose net hourly wage was at least 2 Euro.

[22] All figures are placed in the Appendix.

[23] The coefficients associated to dummies controlling for other levels of tenure were always not significant. Moreover, there were no substantial gender differences in the average levels of tenure.

**Table 2: Mean and St. Dev. For Selected Variables**

| Group | Statistic | Log Hourly Wage | Age | Schooling | Long-Term Contract | Public Sector |
|---|---|---|---|---|---|---|
| | | | **2005** | | | |
| **Males** | *Mean* | 2,108 | 41,532 | 12,473 | 0,894 | 0,367 |
| | *St. Dev.* | *0,373* | *12,728* | *3,369* | *0,308* | *0,482* |
| **Females** | *Mean* | 1,968 | 38,628 | 13,336 | 0,854 | 0,43 |
| | *St. Dev.* | *0,336* | *11,623* | *3,239* | *0,354* | *0,495* |
| **Total** | *Mean* | 2,045 | 40,228 | 12,86 | 0,876 | 0,396 |
| | *St. Dev.* | *0,363* | *12,328* | *3,339* | *0,33* | *0,489* |
| | | | **2006** | | | |
| **Males** | *Mean* | 2,096 | 40,991 | 12,578 | 0,881 | 0,355 |
| | *St. Dev.* | *0,371* | *12,893* | *3,259* | *0,324* | *0,478* |
| **Females** | *Mean* | 1,959 | 37,52 | 13,413 | 0,811 | 0,399 |
| | *St. Dev.* | *0,331* | *11,704* | *3,171* | *0,391* | *0,49* |
| **Total** | *Mean* | 2,038 | 39,536 | 12,928 | 0,852 | 0,373 |
| | *St. Dev.* | *0,361* | *12,525* | *3,248* | *0,355* | *0,484* |
| | | | **2008** | | | |
| **Males** | *Mean* | 2,062 | 41,409 | 12,766 | 0,852 | 0,326 |
| | *St. Dev.* | *0,377* | *13,259* | *3,275* | *0,355* | *0,469* |
| **Females** | *Mean* | 2,013 | 39,953 | 13,442 | 0,828 | 0,389 |
| | *St. Dev.* | *0,373* | *12,678* | *3,181* | *0,377* | *0,488* |
| **Total** | *Mean* | 2,038 | 40,698 | 13,096 | 0,84 | 0,356 |
| | *St. Dev.* | *0,376* | *12,998* | *3,247* | *0,366* | *0,479* |
| | | | **2010** | | | |
| **Males** | *Mean* | 2,079 | 41,872 | 12,751 | 0,848 | 0,315 |
| | *St. Dev.* | *0,385* | *13,524* | *3,244* | *0,359* | *0,464* |
| **Females** | *Mean* | 1,964 | 38,688 | 13,638 | 0,805 | 0,384 |
| | *St. Dev.* | *0,348* | *12,057* | *3,185* | *0,396* | *0,486* |
| **Total** | *Mean* | 2,03 | 40,517 | 13,129 | 0,83 | 0,344 |
| | *St. Dev.* | *0,374* | *13,016* | *3,249* | *0,376* | *0,475* |

*Fonte: Dati ISFOL PLUS*

# 5 RESULTS

Our analysis of the gender wage gap is based on the method of equation (13). We have estimated the following decomposition

$$\hat{y}_{m,t}(\theta) - \hat{y}_{f,t}(\theta) = \left[\hat{y}_{m,t}(\theta) - \hat{y}_{C,t}(\theta)\right] + \left[\hat{y}_{C,t}(\theta) - \hat{y}_{f,t}(\theta)\right] \quad (14)$$

for t = 2005, 2006, 2008, 2010. In order to approximate the various conditional distributions of income, we have estimated the quantile regression coefficients at 300 randomly drown percentiles.[24] The wage gap has been computed at 19 distinct quantiles. More precisely, we have estimated the decomposition model every 5 percentiles, starting from the 5th quantile until the 95th quantile. Finally, standard errors have been computed using 200 bootstrap replications.

The results of the decompositions of equation (14) are reported in Figures 2 and 3. For each year, the left graph represents the estimated total earning gap, as measured by the difference between the male and female conditional distributions of income. Instead, the graph on the right, is a decomposition of the total gender wage gap in a part attributable to individual characteristics, and a part attributable to differences in the estimated quantile regression coefficients (the wage structure or discrimination effect).

Notice that the estimated difference between male and female earnings is always positive and significant. With the exception of 2008, its magnitude is around 10 percentage points almost everywhere. Moreover, it tends to increase toward the top of the wage distribution, an evidence that is usually described as a glass ceiling effect[25]. This tendency seems to be stronger in 2005 and 2006, since, in such years, the wage gap is above 20 percentage points for the highest percentiles of the earning distribution. Notice however that the shape of the gender earnings gap distribution is similar across years. This implies that women are discriminated more in jobs where the wages are higher.

It is interesting to compare the magnitude of the estimates across years. Notice that the earning gap was reducing between 2006 and 2008, but it has been increasing between 2008 and 2010, after that an economic downturn phase had begun. For example, the median estimate of the gap was around 10% in 2006, it dropped to less than 5% in 2008, and it increased again to more than 9% in 2010. These strong variations indicate that the wage gap is influenced by the economic cycle, and that while, before 2008, there was some progress during a phase of small growth, discrimination has increased again since the beginning of the economic crisis.

---

[24] Using a subset of quantile regression solutions reduces the risks of quantile crossing. See footnote 10 on this point.

[25] This terminology has been introduced by Albrecht, Bjorklund, and Vroman (2003). To test the hypothesis of a glass ceiling, we have performed several tests on the equality of the estimated coefficients of the total wage gap at different percentiles of the wage distribution. In doing so, we have used the fact that each estimator is normally distributed around zero. The test on the equality of the estimated total wage gap between the 90th and at the 50th quantiles, as well as the test on the equality between the 75th and the 50th quantiles, is rejected in all four years.

For what concerns the decomposition exercise, it is quite evident that the role played by individual characteristics is seldom significant. Indeed, the *characteristics line* of Figures 2 and 3 is always very close to zero. This evidence shows that, in Italy, there has been a substantial convergence in human capital accumulation between men and women, so that gender differences can't be attributed to the levels of observable market potential. Moreover, the shape of the total wage gap resembles that one of the wage structure effect in all four years. This evidence suggests that the amount of discrimination is significant among Italian employees, and it is the main driver of the gender earnings differential. As we have noticed in Section 2, a significant wage structure effect can't be interpreted simply as a pure employers' disutility parameter, and it should be better described as a composite residual effect played by several unobservable market characteristics (which include discrimination as well). We can't identify directly the sources of this discrimination effect. However, in this and in the next Section, we will rule out some explanations and highlight some interesting tendencies.

To gain some more knowledge on what lies behind the wage gap, we have estimated the same decompositions of equation (14), dividing the sample between older employees, defined as those above age 35, and younger employees. The results of this exercise are plotted in Figure 4. It is quite evident that the gender earnings differential increases with age. Indeed, the estimated total wage gap almost doubles for older workers. Moreover, the *glass ceiling effect* is more pronounced for older workers, an evidence suggesting that women have difficulties in reaching those highly remunerated job positions which require more experience. Notice also that the pay gap was very close to zero for those with less than 35 years old in 2008, but it increased again in 2010. Also the level of discrimination is higher with age, as can be seen from the fact that the dotted line of Figure 4 (denoting discrimination among older workers) lies always above the dash-dotted line (which indicates discrimination among younger workers). Again, such findings suggest that women pay additional penalties as the level of market experience increases[26].

With the possible exception of 2008, the distance between the wage structure effects and the total gap (represented by the solid and by the dashed lines) is higher among younger workers. This result is driven by the fact that, in our sample, younger female employees have better characteristics than older female workers, when compared with male employees of the same age. Therefore, there seems to be a higher level of human capital accumulation among young women, which could partly explain why the wage gap is smaller at earlier stages of the career. However, we should notice also that the *discrimination effect* continues to play an important role for younger women, an evidence that can't be fully explained by our data.

From the results presented here, it is quite evident that the wage structure effect is explaining most of the pay gap. In Section 6 we are going to test some hypothesis and provide further evidences on the existence of a gender pay gap. In particular, we will show that the earnings gap is mostly driven by

[26] In a recent article, Goldin (2014) explains the increasing relation between seniority and the gender pay gap by looking at the cost of flexibility. When women are older, their responsibilities in the informal labour market increase, and they become less willing to supply many hours of work. According to the author, this demand for flexibility in hours worked is costly for the employer. From our data, we could see that women supply less hours then man on average, but we could not find evidence of a positive relation between hours worked and hourly wages, even if measurement error might be negatively biasing our estimates.

*within occupation* and *within sector* differences in compensations. Indeed, we will construct several models with more covariates, and we will show that the estimated discrimination levels do not change much. Finally, we will deal with the problem of correlation between unobservable individual heterogeneity and *time-varying* covariates, showing that this issue is not affecting our results. However, we should be aware that the exogeneity of gender with respect to unobservable abilities remains an untestable assumption.

# 6 ROBUSTNESS CHECKS

In Section 2 we have stated that, in order to have a meaningful decomposition of the gender pay gap, the ignorability assumption should be satisfied. Unfortunately, such assumption can't be tested directly. For this reason, we have stressed the fact that our measure of discrimination is best interpreted as a composite residual effect, which can't be fully explained by our data. Nevertheless, we have performed some robustness checks, in order to gain some knowledge on the quality of the wage prediction model estimated in the previous Section. Moreover, we have tested whether our estimates of discrimination are affected when correlation between *time-varying* dependent variables and individual fixed effects is taken into account, by looking at the direction and size of eventual biases[27].

As a first robustness test, we have repeated the estimations carried out in Section 5, this time adding more explanatory variables to our model. In addition to the usual set of independent variables (which amounts to 26 covariates), we have included 9 dummies for occupational position, around 20 dummies for each Italian administrative Region, 17 dummies for different categories of economic sectors, and some more variables on the family background (education of the father). The results associated to this model, which contains a total of around 65 covariates, are represented by the dashed lines in Figures 5 and 6. The most important feature of this specification is that it allows to measure linearly any wage difference between a very detailed set of occupations and sectors. Therefore, our estimates become more suitable to take into account the effect segregation of women into less remunerative job positions. Finally, we have estimated the quantile regression decomposition over a model where each of the 9 occupational dummies, and each of the 17 dummies denoting sectors were interacted. This model contains a total of around 220 covariates, and is represented by the dotted lines in Figures 5 and 6. This last specification allows us to capture also any non-linearity in the way occupational positions are rewarded across sectors.

In Figure 5, we have plotted the estimated gender wage differences that are attributable to differences in observable characteristics across groups. Each line represents a different model specification. In general, adding more covariates increases the amount of discrimination that is explained by characteristics. However, the difference between the 26-covariates and the other models are quite modest in magnitudes. The greatest difference is observable in 2005, where, for higher quantiles of the wage distribution, it reaches 5 percentage points. We should also stress the fact that this part of the income distribution is associated with the highest levels of estimated discrimination, and that this peak now seems to be in part driven by the composition of the sample in terms of occupational positions. For all other years, the differences in estimations across models seldomly exceed 2 percentage points.

In Figure 6, we have reported the estimated level of discrimination across model specifications. As can be noticed, the differences between the 26-covariates and the other models are smaller than in the previous case. In general, the wage structure effect is reduced when more covariates are added to the model. However, such characteristics effect lies always around or above 10%, with the only exception

[27] This exercise is prone to type two errors. That is, whenever the more robust estimator (the Hausman-Taylor model in our case) is biased, no knowledge can be gained on the error associated to the less consistent estimator (the random effect model).

of 2008. Thus, we can conclude that most of the gender differences in pay originate *within* sectors and *within* occupations. With the partial exception of 2005, only a small percentage of the gender pay gap can be attributed to segregation effects, while discrimination continues to determine most of the differences in wages. In summary, the results presented in the previous section are quite robust when more detailed wage prediction models are considered.

As a final test on the validity of our model, we have exploited the longitudinal structure of the ISFOL PLUS data. We have tried to take into account individual characteristics, which are not observable in our data, but which may be correlated both with our explanatory variables and with wages. Indeed, a reason why the estimates of discrimination could be biased comes from the fact that some variables, such as schooling, experience or tenure, could be correlated with individual abilities not observable to the researcher. For example, using the ISFOL PLUS database, Borgna and Struffolino (2015) have shown that there are persistent gender differences in the dropout rates from secondary schooling. Thus there's the possibility that, due to the presence of unobservable dynamics that we fail to take into account, the same education level determines systematically different wage potential between women and men. Moreover, experience is approximated by age in our model, while actual experience could be influenced by time spent out of the working force. Alternatively, less productive workers could self-select into the category of workers with less years of tenure. All the above situations would lead to biased estimates of the coefficients associated to these variables, and could in turn undermine our decomposition exercise.

In general, correlation between individual fixed effects and *time varying* covariates is a quite common outcome[28], and is confirmed in our data when performing a cluster robust Hausman test, which leads to a rejection of the null hypothesis. Unfortunately, the fixed effect model is not suitable for the estimation of the gender wage gap. Indeed, this model can't estimate the effect of *time invariant* characteristics on wages. In our context, the only feasible alternative is the Hausman-Taylor (HT) regression (Hausman and Taylor, 1981). This model allows for arbitrary correlation between *time varying* regressors and individual unobservable effects. For example, in this model returns to experience are measured as the average marginal effect on income of one more year of seniority, with respect to *individual-specific* wage means. However, the HT model also relies on a more demanding exogeneity assumption of the *time invariant* characteristics, which can't be tested in any way. Therefore, the effect of characteristics with limited, or no time variability at all, such as schooling and gender, are not consistently measured when individual fixed effects are correlated with them.

In order to see how much the estimated amount of discrimination is influenced by correlation of the *time varying* dependent variables with individual unobservable heterogeneity, we have compared the results of the HT model to the ones obtained from the inconsistent GLS estimator. The GLS estimator (also called random effects) is comparable to the OLS regression, but it is more efficient when dealing with a panel sample. This estimator relies on the assumption of exogeneity of all variables with respect to individual unobservable heterogeneity, an assumption that is not met in our data. Then, to obtain a

[28] See Polachek and Kim (1994) for a detailed discussion of the problem of unobservable individual heterogeneity in the context of the gender wage gap estimation.

consistent estimator for the coefficients associated to *time varying* variables, we have estimated the HT model. That is, we have first applied fixed effects regression, which is a consistent estimator in this context. Then, we have regressed the *time invariant* variables using the residuals of the first-stage fixed effects regression, employing the efficient GLS estimator.

In performing the longitudinal estimations described above, we have used our 26-covariates model, and we have pooled all four available years, in order to construct a panel sample. Finally, we have included a dummy equal to one for male observations. The coefficient associated to this variable represents the marginal effect of gender on income, keeping all else constant. Therefore, it can be interpreted as a measure of discrimination similar to the *wage structure effect*. Table 3 compares the estimated gender effect, together with a few selected covariates. As can be noticed by looking at the estimated effect of age on wages, the difference between the coefficients associated to *time-varying regressors* is quite significant across the two models. This result is confirmed by the Hausman test, which we have conducted over all such *time-varying* regressors. Instead, notice that the difference in the estimated gender effect is negligible. This result suggests that correlation between *time varying* dependent variable and individual fixed effect is not biasing our estimations of gender discrimination. However, we must also interpret this result with caution. The possibility that also the HT model is providing biased estimates can't be neither tested nor excluded. Moreover, as we have already stressed at the beginning of this Section, unfortunately the validity of our *ignorability assumption* can't be tested. Nevertheless, gender discrimination seems to be an evidence very robust and persistent across models.

**Table 3: Panel Estimates of Coefficients for Selected Variables**

| **Model** | **(1)** | **(2)** |
|---|---|---|
| *Variable* | *coeff./st. err.* | *coeff./st. err.* |
| Male | 0.1195*** | 0.1200*** |
| | (0.0043) | (0.0042) |
| Schooling | 0.0201*** | 0.0265*** |
| | (0.0007) | (0.0007) |
| Age | 0.0183*** | 0.0267*** |
| | (0.0015) | (0.0042) |
| Age Squared | -0.0001*** | -0.0002*** |
| | (0.0000) | (0.0000) |
| Observations | 27,429 | 27,429 |

*Legend of the models:*
*(1): GLS regression model, evaluated over the entire panel sample. Not consistent.*
*(2): Hausman-Taylor regression, consistent to correlation between time-varying covariates and individual unobservable heterogeneity.*
*Fonte: Dati ISFOL PLUS*

# 7 CONCLUSIONS

In this paper we have measured the gender wage gap using four waves of the ISFOL PLUS sample. This database is interesting for several reasons. First, it contains detailed information on workers' history, and it has a panel structure, which allows us take into account unobserved individual heterogeneity. Moreover, the chosen waves, which range from 2005 to 2010, cover a period that is interesting to study, due to the economic downturn which begun during 2008.

For our analysis, we have considered full-time private and public employees. We have adopted a quantile regression based methodology, which has allowed us to show that the wage gap is stronger in higher remunerated job positions. By dividing the sample between younger and older employees, we have shown that this earning gap is also increasing with age, a fact that can be in part attributed to an improvement of the level of human capital among younger women, while in part is the result of discrimination levels that are increasing with seniority.

In general, we have shown that women earn around 10% less than men, despite having similar market potential. This percentage is increasing with wages, and it reaches the level of 20% at higher quantiles of the income distribution. This *glass ceiling effect* is also stronger among older workers, a fact suggesting that women struggle to reach especially those well-remunerated positions that become available as experience increases. Moreover, we have shown that the level of discrimination is influenced by the economic cycle. Indeed, the gender wage gap was reducing in 2008, but, in 2010, it has reached again levels similar to the 2005 ones. This suggests that the economic recession, which was particularly severe in Italy during 2009, has had a negative impact on discrimination, worsening the position of women in the labour market, in terms of wages earned, more than that of men[29].

To further test our results, we have carried out some robustness checks. We have shown that increasing the information on workers' occupation, sector and location has only a small effect on the wage gap estimates. When including a model with a full set of sector-occupation interactions, the estimated effect of individual characteristics on the gender wage gap changed by no more than 5%, with this percentage being almost everywhere around 1-2%. Moreover, the impact of additional dependent variables on the estimated wage structure effect was even smaller. Therefore, pay differences seems to be more relevant *within* similar job positions and sectors. Moreover, we have used the Hausman-Taylor regression model to control for endogeneity problems that could arise with *time-varying* dependent variables. Our results show that this kind of endogeneity problem does not seem to be affecting the estimated gender wage gap. However, we were not able to test directly the assumptions required by the Hausman-Taylor estimator, and endogeneity problems could still be quite relevant even under this model.

We can conclude that women earn less than men, even when their characteristics are similar. The fact that gender differences in pay could not be attributable to the observable characteristics of our sample

[29] On this issue, Di Tommaso and Piazzalunga (2013), which analyse Italian data until 2011, suggest that the interruption of wage increases in the public sector, which has been implemented since 2010, could have had further negative effects on women's wages relative to the men's ones, given that women are traditionally over-represented in the public sector.

leaves the question on the source of this *wage structure gap* open. Based on previous literature[30], several answers can be suggested, even if they can't be fully tested from our data. In general, there could be a sorting mechanism, which makes men more likely to work for firms who pay higher wages[31]. Moreover, there could be differences in the way men and women are promoted, so that similar workers are paid differently due to internal compensation schemes. Women could be less willing to bargain and less competitive than men. They could prefer more flexible working hours, which could represent a cost for several firms. Finally, employers could be discriminating against women because of some form of disutility in hiring them, or because they could be in a monopsonistic position allowing them to pay lower wages and extract higher rents from female workers.

[30] See for example Goldin (2014), Buser, Niederle, and Oosterbeek (2014), Booth, Francesconi, and Frank (2003), Sulis (2011).

[31] As noted in Section 2, the literature on this topic, even if not vast, has not yet found evidences supporting this hypothesis.

## APPENDIX: FIGURES

**Figure 1: Kernel Density Estimates of Hourly Wages by Year and Gender**

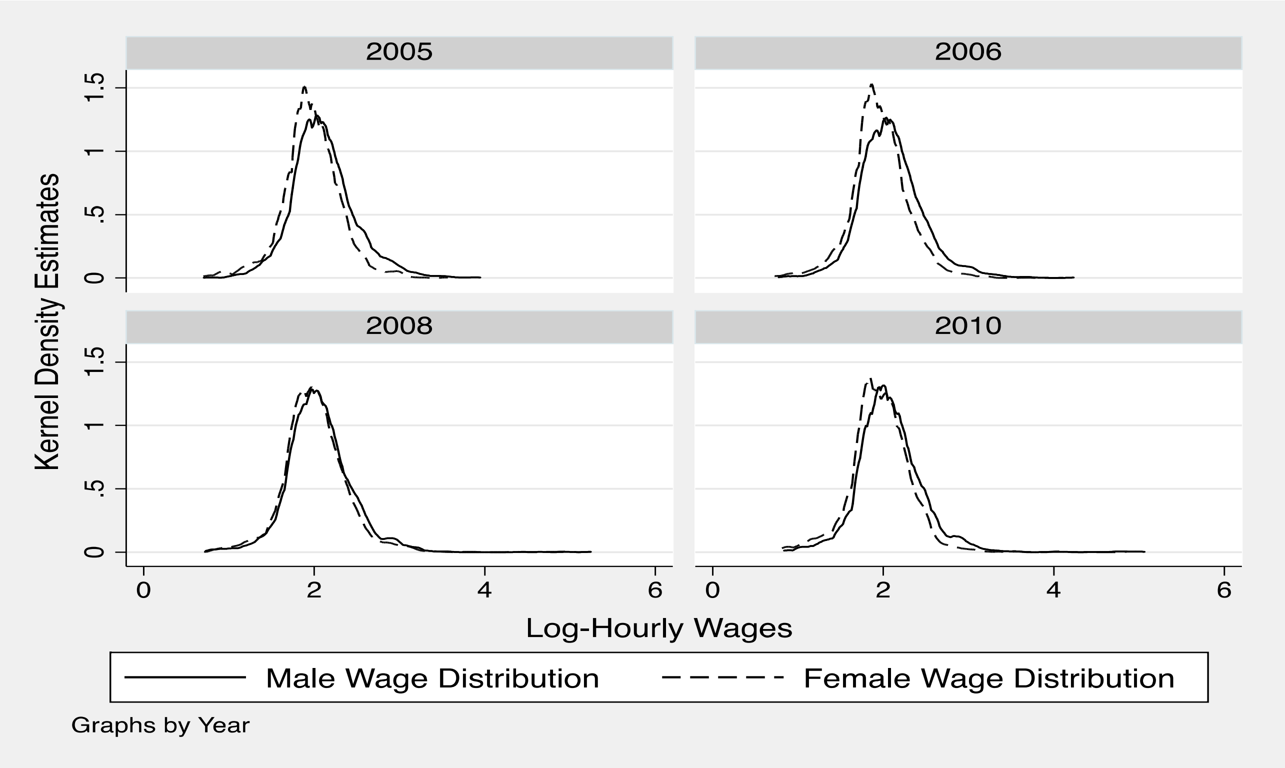


*Fonte: Dati ISFOL PLUS*

**Figure 2: Quantile Regression Decomposition of the Gender Wage Gap by Year (1)**

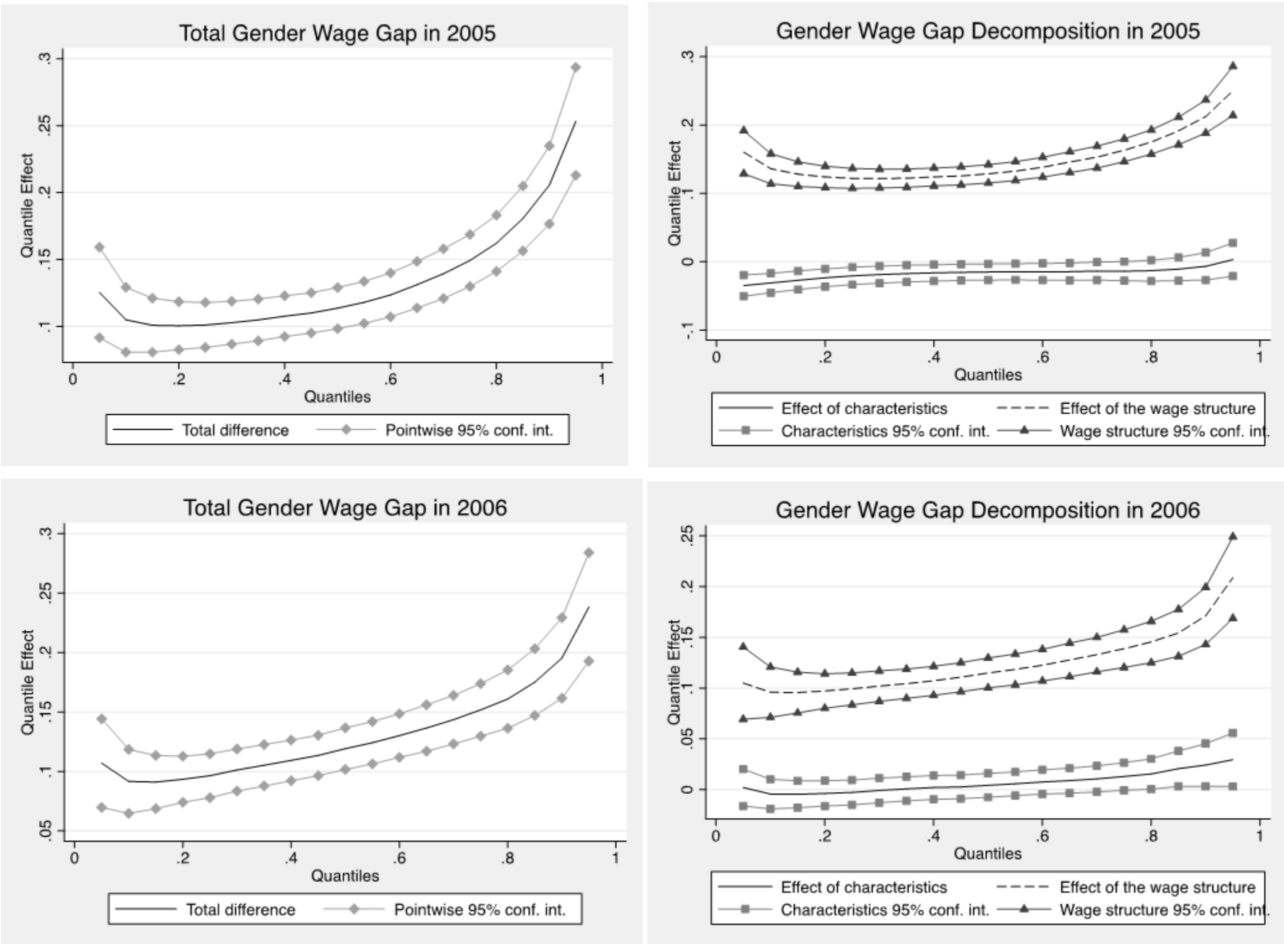


*Fonte: Dati ISFOL PLUS*

**Figure 3: Quantile Regression Decomposition of the Gender Wage Gap by Year (2)**

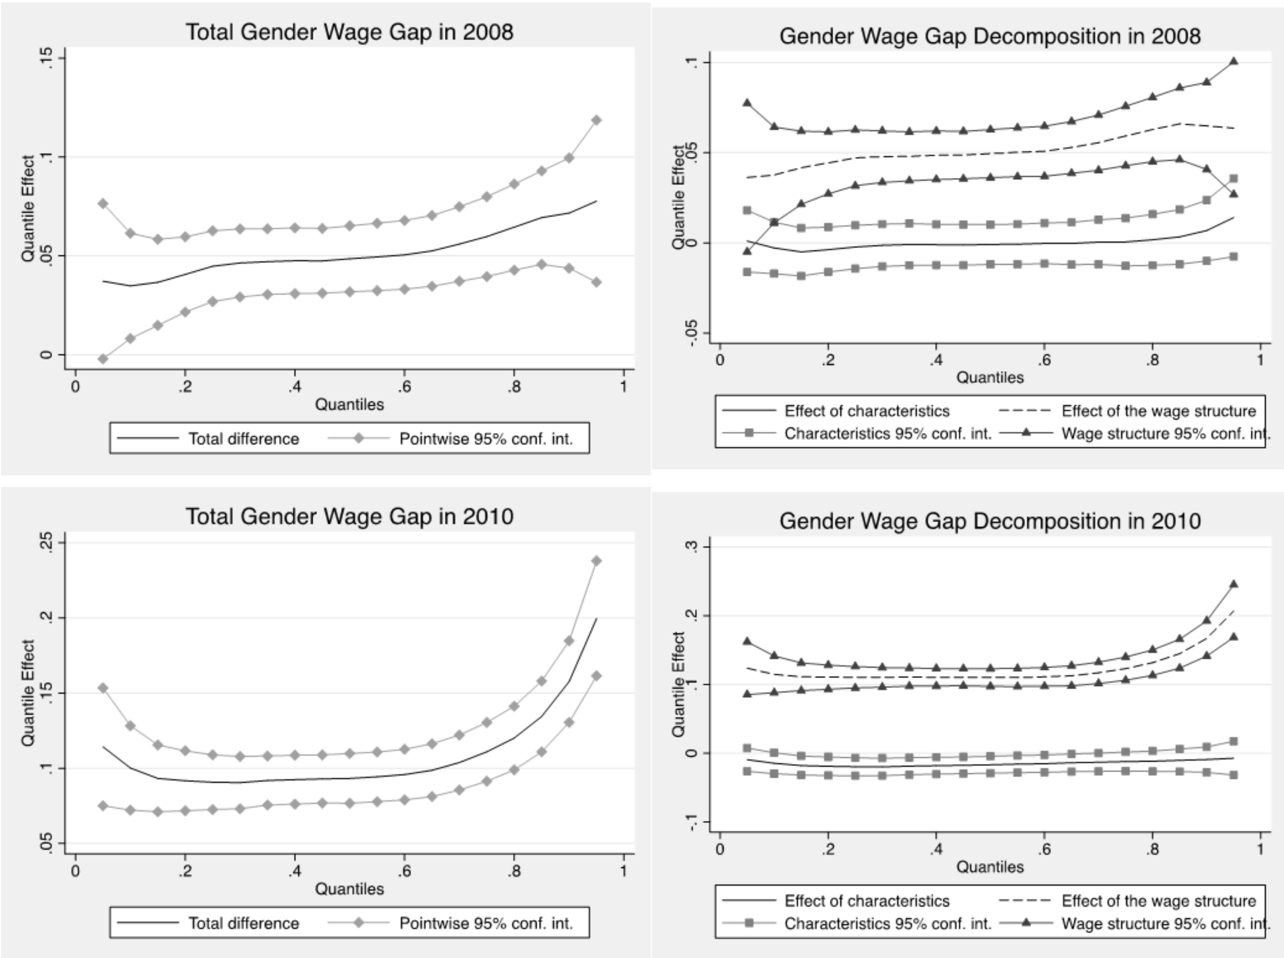


*Fonte: Dati ISFOL PLUS*

**Figure 4: Quantile Regression Decomposition of the Gender Pay Gap by Age Group (Above and Below 35 Years Old)**

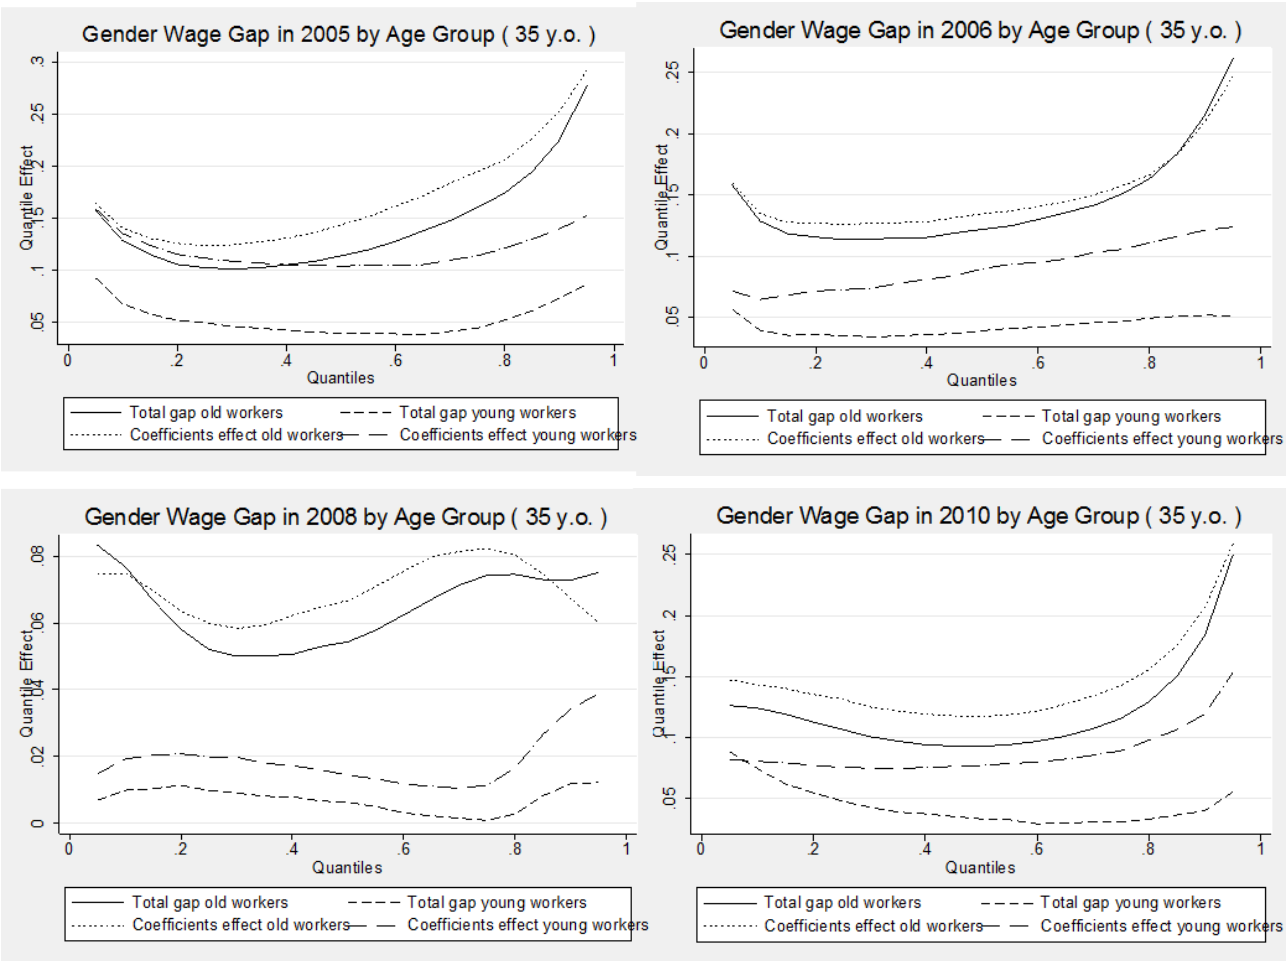


*Fonte: Dati ISFOL PLUS*

**Figure 5: Characteristics effect obtained applying the counterfactual decomposition by year. Each line reports results computed using a different model specification**

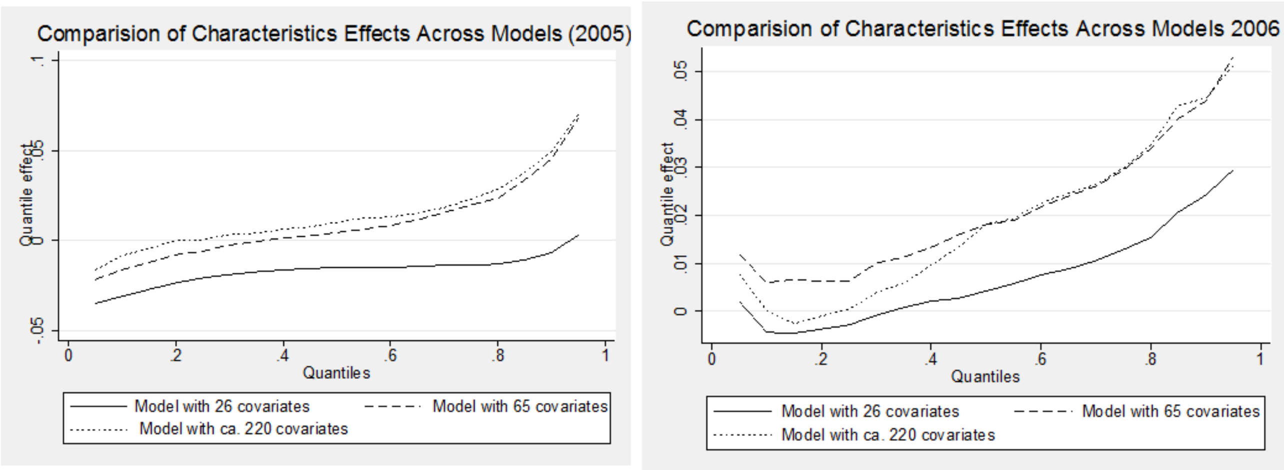


*Fonte: Dati ISFOL PLUS.*

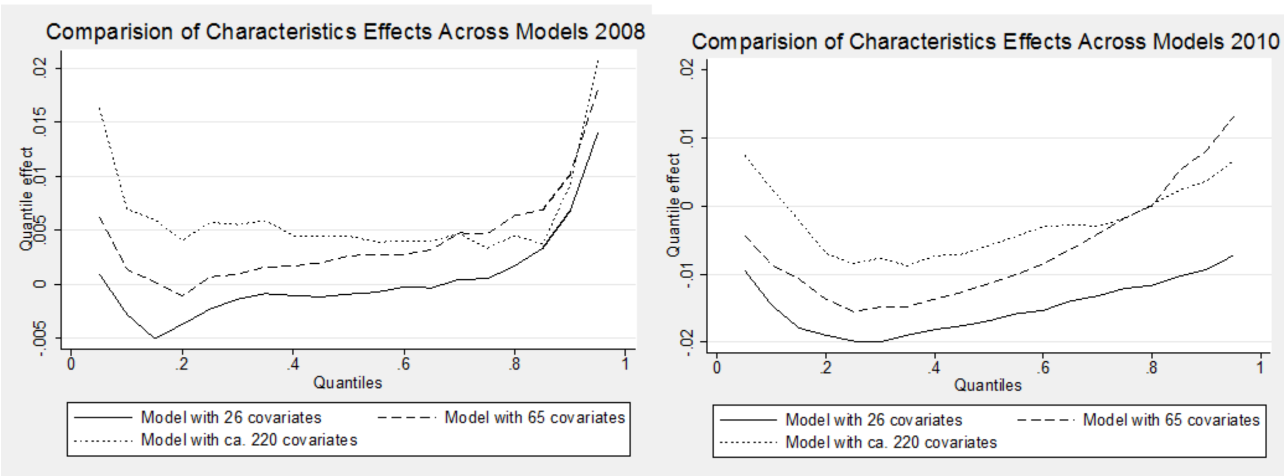


*Fonte: Dati ISFOL PLUS*

**Figure 6: Wage structure effects obtained applying the counterfactual decomposition by year. Each line reports results computed using a different model specification**

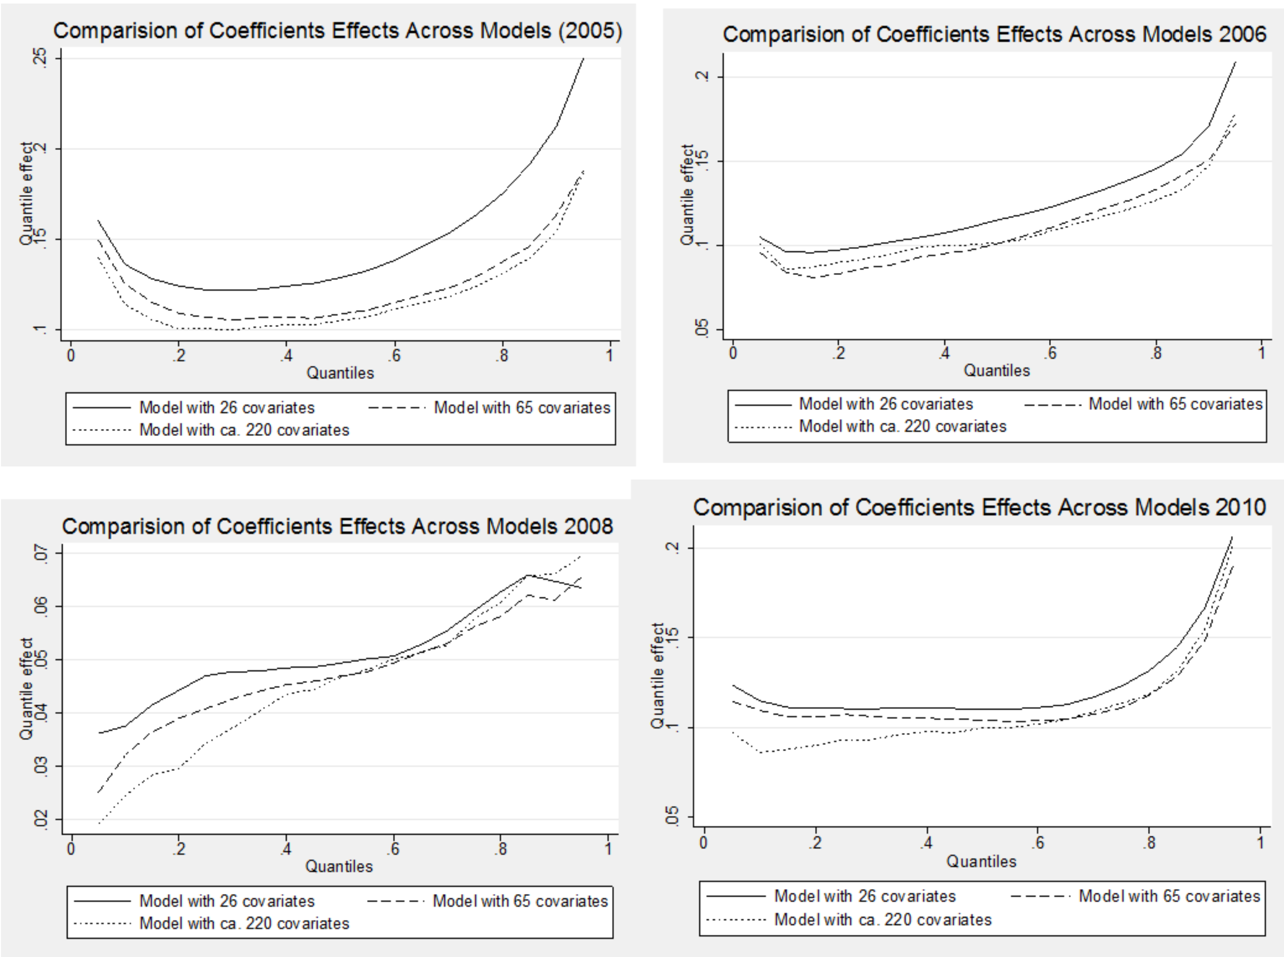


*Fonte: Dati ISFOL PLUS*

## REFERENCES

ABOWD, J., KRAMARZ F., MARGOLIS D.: "*High Wage Workers and High Wage Firms,*" Econometrica, 67(2) (1999), 251–333.

ALBRECHT, J., BJORKLUND A., VROMAN S.: "*Is There a Glass Ceiling in Sweden?*" Journal of Labor Economics, 21, (2003), 145–177.

ASHENFELTER, O.C., FARBER H., RANSOM M.: "*Labor Market Monopsony*" Journal of Labor Economics, 28(2), (2010), 203-210.

BECKER, G.: *The Economics of Discrimination.* University of Chicago Press, Chicago, (1973) second edn.

BLINDER, A.: "*Wage Discrimination: Reduced Form and Structural Estimates*" Journal of Human Resources, 8(4), (1973), 436–455.

BOOTH, A., FRANCESCONI M., FRANK J.: "*A Sticky Floors Model of Promotion, Pay, and Gender*" European Economic Review, 47(2), (2003), 295–322.

BORGNA, C., STRUFFOLINO E.: "*Early School Leaving Dynamics in Italy: The Heterogeneity of Gender Effects*" ISFOL Research Paper, (2015).

BUCHINSKY, M.: "*The Dynamics of Changes in the Female Wage Distribution in the USA: A Quantile Regression Approach*" Journal of Applied Econometrics, 13, (1998), 1-30.

BUSER, T., NIEDERLE M., OOSTERBEEK H.: "*Gender, Competitiveness, and Career Choices*" The Quarterly Journal of Economics, 129(3), (2014), 1409–1447.

CARD, D., CARDOSO A. R., KLINE P.: "*Bargaining and the Gender Wage Gap: A Direct Assessment*" Discussion paper, IZA Discussion Paper, (2013).

CHERNOZHUKOV, V., FERNÁNDEZ-VAL I., MELLY B.: "*Inference on counterfactual distributions*" Econometrica, 81(6), (2013), 2205–2268.

CHRISTOFIDES, L., POLYCARPOU A., VRACHIMIS K. "*The gender wage gaps,'sticky floors' and'glass ceilings' of the European Union*" IZA Discussion Paper No. 5044, (2010).

DI TOMMASO, M. L., PIAZZALUNGA D.: "*It is not a bed of roses. Gender and ethnic pay gaps in Italy*" Discussion paper, University of Turin, (2013).

FORTIN, N., LEMIEUX T., FIRPO S.: "*Decomposition Methods in Economics*" in Labor Economics, ed. by O. Ashenfelter, and D. Card, vol. 4A of Handbooks in Economics, chap. 1, pp. 1–97. North-Holland, Amsterdam, (2011).

GOLDIN, C.: "*A Grand Gender Convergence: Its Last Chapter*" The American Economic Review, 104(4), (2014), 1091–1119.

HAUSMAN, J. A., TAYLOR W. E.: "*Panel Data and Unobservable Individual Effects*" Econometrica: Journal of the Econometric Society, (1981), 1377–1398.

HECKMAN, J.: "*Sample Selection Bias as a Specification Error*" Econometrica, 47(1), (1979), 153–161.

KOENKER, R., BASSETT G.: "*Regression Quantiles*" Econometrica, 46(1), (1978), 33–50.

MACHADO, J., MATA J.: "*Counterfactual Decomposition of Changes in Wage Distributions Using Quantile Regression*" Journal of Applied Econometrics, 20(4), (2005), 445–465.

MANDRONE, E.: *Labour Economics: PLUS Empirical Studies*. Ediguida Editore, collana Temi e Strumenti, (2012).

MATANO, A., NATICCHIONI P.: "*Rent sharing as a driver of the glass ceiling effect*" Economics Letters, 118(1), (2013), 55–59.

MELLY, B.: "*Decomposition of Differences in Distribution Using Quantile Regression*" Labour Economics, 12(4), (2005), 577–590.

MELLY, B.: "*Estimation of Counterfactual Distributions Using Quantile Regression*" Review of Labor Economics, 68, (2006) 543–572.

NOPO, H.: "*Matching as a Tool to Decompose Wage Gaps*" The Review of Economics and Statistics, 90(2), (2008), 290–299.

OAXACA, R.: "*Male-Female Wage Differentials in Urban Labor Markets*" International Economic Review, 14(3), (1973), 693–709.

OAXACA, R., RANSOM M.: "*Discrimination and Wage Decomposition*" Journal of Econometrics, 61, (1994), 5–21.

POLACHEK, S. W., KIM M.-K.: "*Panel Estimates of the Gender Earnings Gap: Individual-Specific Intercept and Individual-Specific Slope Models*" Journal of Econometrics, 61(1), (1994), 23–42.

RANSOM, M. OAXACA R.L: "*New Market Power Models of Sex Differences in Pay*" Journal of Labor Economics, 28(2), (2010), 267-289.

SULIS, G.: "*What Can Monopsony Explain of the Gender Wage Differential in Italy?*", International Journal of Manpower, 32(4), (2011), 446-470.