

Matching as a Tool to Decompose Wage Gaps

by

Hugo Nopo

December 2003

MIDDLEBURY COLLEGE ECONOMICS DISCUSSION PAPER NO. 04-06



DEPARTMENT OF ECONOMICS
MIDDLEBURY COLLEGE
MIDDLEBURY, VERMONT 05753

<http://www.middlebury.edu/~econ>

Matching as a Tool to Decompose Wage Gaps

Hugo Ñopo*

December 10, 2003

Abstract

In this paper I present a methodology that uses matching comparisons to explain gender differences in wages. The approach emphasizes gender differences in the supports of the distributions of observable characteristics and provides useful insights about the distribution of the unexplained gender differences in pay.

The proposed methodology, a non-parametric alternative to the Blinder-Oaxaca (BO) wage gap decomposition, does not require the estimation of earnings equations. It breaks down the gap into four additive elements, two of which are analogous to the elements of the BO decomposition (but computed only over the common support of the distributions of characteristics), while the other two account for differences in the supports. Using data for Peru in the period 1986-2000, I found that this problem of non-comparability accounts for 23% and 30% of the male and female working populations respectively.

The matching methodology allows us to quantify the effect of explicitly recognizing these differences in the supports. In this way, the 45% gender wage gap in Peru is decomposed as: 11% explained by differences in the supports, 6% explained by differences in the distributions of individual characteristics and the remaining 28% cannot be explained by differences in observable individuals' characteristics. Approximately half of the latter is due to unexplained differences in the highest quintile of the wage distribution.

Keywords: Matching, Non-parametric, Gender Wage Gap, Latin America.

JEL Classification Codes: C14, D31, J16, O54.

*The advice of Chris Taber, Luoja Hu, and Dale Mortensen is deeply acknowledged. This paper benefited from comments and suggestions by Gadi Barlevy, Fabio Caldieraro, Alberto Chong, Juan Jose Diaz, Libertad Gonzalez, Zsolt Macskasi, Rosa Matzkin, Bruce Meyer, Lyndon Moore, Andrew Morrison, Jim Sullivan, Maximo Torero and participants at The Latin American Meeting of the Econometric Society, GRADE, The Joint Labor/Public Economics Seminar at Northwestern University, The Latin American and Caribbean Economic Association Meeting, The Midwest Econometric Group Meeting, IZA, CIDE, ITAM, Middlebury College, the European Association of Labour Economist Meeting and the North-Eastern Universities Development Consortium Meeting. All errors are my own, comments are welcome, please send them to hnopo@middlebury.edu

1 Introduction

Gender differences in the labor market, particularly the gender wage gap, have been a significant area of concern for theoretical and empirical research in economics. On average, males earn more than females in yearly, monthly and per hour terms. These differences in average earnings —the “gender gaps”— are partially explained by gender differences in observable characteristics of individuals that the labor market rewards. The *wage gap decomposition* developed by Blinder and Oaxaca in 1973 has been a key tool in explaining the wage gap and the role that differences in individual characteristics play. This decomposition requires the linear regression estimation of earnings equations for both females and males. Based on these earnings equations, it generates the counterfactual: “What would a male earn if the compensation scheme for his individual characteristics aligned with the compensation scheme for females?” Based on that counterfactual, the difference in average wages between males and females is broken into two additive components: one attributable to differences in average characteristics of the individuals, and the other to differences in the rewards that these characteristics have. The latter component is considered to contain the effects of both unobservable gender differences in characteristics that the labor market rewards and discrimination in the labor market.

There is a potential problem associated with this approach: mis-specification due to differences in the supports of the empirical distributions of individual characteristics for females and males (hereafter called *gender differences in the supports*). This is an issue that Rubin(1977) originally pointed out in the program evaluation literature. There are combinations of individual characteristics for which it is possible to find males, but not females, in the labor force —as is the case for individuals who are in their early thirties, who are married and hold a college degree or superior— while there are also combinations of characteristics for which it is possible to find females, but not males —as is the case for single individuals who are in their late forties and have less than elementary school education.¹ With such combinations of characteristics one cannot compare wages across genders. This problem of comparability is accentuated when job characteristics are included in the explanation of the wage gap. As females tend to concentrate in certain occupations that demand particular skills (e.g., nurses or maids) males are more likely to be found working in risky or managerial occupations for which long tenure is required.²

The traditional Blinder-Oaxaca (BO) decomposition fails to recognize these *gender differences in the supports* by estimating earnings equations for all working females and all working males without restricting the comparison only to those individuals with comparable characteristics. By not considering this restric-

¹As I will show in section 5, empirical evidence for Peru suggests that the sets of non-comparable individuals involve 30% of working females and 23% of working males.

²For a discussion about typically female-dominated occupations and occupational segregation by gender in Latin America during the 90’s see Deutsch *et al.* (2002).

tion, the BO decomposition is implicitly based on an “out-of-support assumption”: it becomes necessary to assume that the linear estimators of the earnings equations are also valid out of the supports of individual characteristics for which they were estimated. Empirical evidence (which I show in this paper) suggests that this assumption tends to over-estimate the component of the gap attributable to differences in the rewards for individuals’ characteristics.

Besides the mis-specification problem associated with *gender differences in the supports*, it is also important to note an informative limitation of the original approach: the BO decomposition is informative only about the average unexplained difference in wages, not about the distribution of these unexplained differences. Exploring a different approach promises to be more fruitful.

In this paper I adapt a tool of the program evaluation literature, matching, to fix the problem of *gender differences in the supports* and provide information about the distribution of the differences in wages that remain unexplained by the characteristics of the individuals after the decomposition, without requiring any estimation of earnings equations and hence, no validity-out-of-the-support assumptions. The proposed approach is to consider the gender variable as a treatment and use matching to select sub-samples of males and females such that there are no differences in observable characteristics between the matched groups. Motivated by this matching approach, I propose a new decomposition of the wage gap that accounts for differences in the distribution of individual characteristics, paying special attention to *gender differences in the supports*.

The proposed methodology is implemented using data from Peru between the years 1986 and 2000. The convenience of using this data set and time frame is two-fold. First, as it is documented by Blau and Ferber (1992), the Latin American region reports the highest levels of gender occupational segregation in the world and there are substantial gender differences in observable characteristics of the individuals. There is reason to think that the problem of the *gender differences in the supports* matters substantially in this region. Second, Peru implemented many of the economic reforms that took place in Latin America during the early nineties. These economic reforms, besides instituting an accelerated privatizing process, included substantial changes in labor market regulations.³

The remainder of this paper proceeds as follows: Section 2 explores the related literature. Section 3 presents the matching approach and its link to a non-parametric extension of the Blinder-Oaxaca decomposition that emphasizes the *gender differences in the supports*. Section 4 discusses the data and reports the main gender differences in characteristics that are related to wages. Section 5 describes the results of the hourly wage gap decomposition, exploring the distribution of unexplained differences in pay and comparing matching with the traditional BO approach based on linear regressions. Section 6 explores

³See Heckman and Pages (2000) for a detailed description of those reforms and their implied costs in the Latin American region. Saavedra (2000) and Saavedra and Torero (2000) provide an analysis for those changes in Peru.

gender differences in participation and unemployment rates. Section 7 concludes and outlines a short term research agenda in the path of this matching approach.

2 Matching and Wage Gap Decompositions: A Literature Review

Matching comparison techniques aim to find matched samples with “similar”⁴ observable characteristics (or a linear combination of them) except for one particular observable variable, the “treatment”, which is used to group observations into two sets: the treatment and the control group. Having controlled for observed characteristics, the comparison techniques are used to measure the impact of the treatment on these groups under different sets of identifying assumptions. These studies, concerned with the comparison of groups with similar characteristics, has been of especial interest to experimental design and statistics for many years. However, not until the introduction of propensity scores in experimental designs by Rosenbaum and Rubin (1983) did the matching subject enter into the discussion of estimation of causal effects in economics. As a result of their seminal work, a debate started in the economic literature about the widespread use of matching not only in experimental, but also in non-experimental designs (LaLonde (1986), Meyer (1995), Heckman, Ichimura and Todd (1997), Dehejia and Wahba (1998) and Smith and Todd (2000) among others).

Almost thirty years ago, Blinder (1973) and Oaxaca (1973) proposed a methodology to decompose wage gaps in terms of explained and unexplained components. The method is based on the separate estimation of earnings equations for the two groups being compared, namely females and males: $\bar{y}^F = \hat{\beta}^F \bar{x}^F$ and $\bar{y}^M = \hat{\beta}^M \bar{x}^M$. Thus, the wage gap can be expressed as $\bar{y}^M - \bar{y}^F = \hat{\beta}^M \bar{x}^M - \hat{\beta}^F \bar{x}^F$. Then, the method requires the addition and subtraction of the term $\hat{\beta}^F \bar{x}^M$ (or alternatively, $\hat{\beta}^M \bar{x}^F$) which can be interpreted as the counterfactual situation, “What would the earnings for a male (female) with average individual characteristics be, in the case that he (she) is rewarded for his (her) characteristics in the same way as the average female (male) is rewarded?” After some algebraic manipulations, the wage gap takes the form: $\bar{y}^M - \bar{y}^F = \hat{\beta}^F (\bar{x}^M - \bar{x}^F) + (\hat{\beta}^M - \hat{\beta}^F) \bar{x}^M$. Which has a natural interpretation: the first component of the right-hand side, $\hat{\beta}^F (\bar{x}^M - \bar{x}^F)$, is attributed to differences in average characteristics between males and females, while the second component, $(\hat{\beta}^M - \hat{\beta}^F) \bar{x}^M$, is attributed to differences in average rewards to the individual characteristics. Juhn, Murphy and Pierce (1993) extended the decomposition “characteristics-rewards” into one that considers three components: observable characteristics, observable rewards and

⁴The precise way in which these similarities can be computed varies. The literature provides propensity scores, Euclidean distances, and Mahalanobis distances among others. The precise type of matching proposed in this paper will be introduced in Section 3.

unobserved heterogeneity.

Dolton and Makepeace (1987) and Munroe (1988) pointed out an informative limitation of the original BO approach. The wage gap decomposition is only informative about the average unexplained differences in pay but not about the distribution of such unexplained differences. One strategy for overcoming that distribution limitation has been the estimation of quintile earnings equations (Buchinsky (1994)). Another strategy, proposed by Jenkins (1994) and Hansen and Wahlberg (1999), has been the use of Generalized Lorenz Curves (GLC) for both observed earnings and predicted counterfactual earnings. These strategies suffer from the same drawback of ignoring the problem of *gender differences in the supports* that this paper addresses.

The idea of extending the BO decomposition to a semi-parametric setup in order to explore the distribution of unexplained differences can be found in DiNardo, Fortin and Lemieux (1996). In a setup in which they analyze the role of labor market institutions, DiNardo *et al.* estimate earnings equations non-parametrically by means of kernel estimations, facing the “curse of dimensionality” that arises when there are many explanatory variables in non-parametric setups.⁵ Another related semi-parametric approach is proposed by Donald, Green and Paarsch (2000). By adapting techniques from the duration literature to the estimation of density functions, they explore differences in wage distributions between Canada and the United States. In the same line of exploring differences in density functions, Bourguignon, Ferreira and Leite (2002) adapt tools from the micro-simulation literature to generate sequences of counterfactual densities and compare earnings distributions in Mexico, Brazil and the U.S.

A setup closely related to the one I use in this paper is proposed by Barsky, Bound, Charles and Lupton (2001). In their paper, Barsky *et al.* decompose the black-white wealth gap in the U.S. based exclusively on one explanatory variable (income), avoiding in this way the dimensionality problem that the non-parametric literature faces. They recognize the importance of *differences in the supports* and restrict the comparison to the common support. In this paper I take a step further and propose a decomposition that accounts for *differences in the supports* measuring gaps in and out of the common support.

Pratap and Quintin (2002), also using a matching approach, measured wage differences between the formal and informal sectors in Argentina. My approach differs from that matching approach in two ways: one is the explicit assumption that the supports of the distributions of individual characteristics are different—which for the case of gender differences matters substantially—and the other is the use of matching on characteristics instead of propensity scores.⁶

The next section of the paper shows the details of the link between the matching approach and the

⁵Unfortunately for the purposes of this paper, matching also suffers from the same dimensionality problem.

⁶The “ignorability of treatment” assumption required by Rosenbaum and Rubin (1983) in order to allow to match by propensity scores instead of characteristics is not likely to be satisfied in the gender setup of this paper.

wage gap decomposition proposed.

3 A Link Between Matching and Wage Gap Decompositions in a Non-Parametric Setup

Let Y denote the random variable that models individual earnings and X the n -dimensional vector of individual characteristics (such as age, education, occupational experience, occupation, firm size, etc.) presumably related to these earnings. Furthermore, let $F^M(\cdot)$ and $F^F(\cdot)$ denote the conditional cumulative distribution functions of individual characteristics X , conditional on being male and female respectively, and $dF^M(\cdot)$ and $dF^F(\cdot)$ denote the implied probability measures. For a correct definition of the measures and integrals that will be introduced later in this section it is enough to assume that $F^M(\cdot)$ and $F^F(\cdot)$ are measurable functions from \mathbb{R}^n to \mathbb{R} (in the Borel sense). Consequently, $\mu^F(S)$ denotes the probability measure of the set S under the distribution $dF^F(\cdot)$, that is, $\mu^F(S) = \int_S dF^F(x)$ and analogously $\mu^M(S) = \int_S dF^M(x)$.

The relationship governing these random variables is modeled by the functions $g^M(\cdot)$ and $g^F(\cdot)$, representing the expected value of earnings conditional on characteristics and gender. Being the case that $E[Y|M, X] = g^M(X)$ and $E[Y|F, X] = g^F(X)$.⁷ It follows that

$$\begin{aligned} E[Y|M] &= \int_{S^M} g^M(x) dF^M(x), \\ E[Y|F] &= \int_{S^F} g^F(x) dF^F(x), \end{aligned}$$

where S^M denotes the support of the distribution of characteristics for males and S^F the support of the distribution of characteristics for females. In such a way, the wage gap, defined as

$$\Delta \equiv E[Y|M] - E[Y|F],$$

can be expressed as

$$\Delta = \int_{S^M} g^M(x) dF^M(x) - \int_{S^F} g^F(x) dF^F(x). \tag{1}$$

⁷This is a generalization of the linear model in which $E[Y|X] = \beta X$, where β is a $1 \times n$ parameter vector and X is an $n \times 1$ regressor vector.

Considering the fact that the support of the distribution of characteristics for females, S^F , is different than the support of the distribution of characteristics for males, S^M , each integral is split over its respective domain into two parts: one on the intersection of the supports and one out of the common support, in the following way:

$$\Delta = \left[\int_{\overline{S^F} \cap S^M} g^M(x) dF^M(x) + \int_{S^F \cap S^M} g^M(x) dF^M(x) \right] - \left[\int_{S^F \cap S^M} g^F(x) dF^F(x) + \int_{S^F \cap \overline{S^M}} g^F(x) dF^F(x) \right]$$

Since the measures $dF^F(\cdot)$ and $dF^M(\cdot)$ are identically zero out of their respective supports (by definition), the domains for the first and fourth integrals (the “non-common support” integrals) can be extended to $\overline{S^F}$ and $\overline{S^M}$ respectively without affecting their corresponding values. Also, every integral can be adequately re-scaled in order to obtain expressions involving expected values of $g^F(X)$ and $g^M(X)$ conditional on their respective partitioned domains, as is shown below.

$$\Delta = \left[\int_{\overline{S^F}} g^M(x) \frac{dF^M(x)}{\mu^M(\overline{S^F})} \right] \mu^M(\overline{S^F}) + \left[\int_{S^M \cap S^F} g^M(x) \frac{dF^M(x)}{\mu^M(S^F)} \right] \mu^M(S^F) - \left[\int_{S^M \cap S^F} g^F(x) \frac{dF^F(x)}{\mu^F(S^M)} \right] \mu^F(S^M) - \left[\int_{\overline{S^M}} g^F(x) \frac{dF^F(x)}{\mu^F(\overline{S^M})} \right] \mu^F(\overline{S^M}).$$

Now, replacing $\mu^F(S^M)$ by $1 - \mu^F(\overline{S^M})$ and $\mu^M(S^F)$ by $1 - \mu^M(\overline{S^F})$, the gap decomposition can be expressed (after some rearrangement) in the following way

$$\begin{aligned} \Delta = & \left[\int_{\overline{S^F}} g^M(x) \frac{dF^M(x)}{\mu^M(\overline{S^F})} - \int_{S^F} g^M(x) \frac{dF^M(x)}{\mu^M(S^F)} \right] \mu^M(\overline{S^F}) + \\ & \left[\int_{S^M \cap S^F} g^M(x) \frac{dF^M(x)}{\mu^M(S^F)} - \int_{S^M \cap S^F} g^F(x) \frac{dF^F(x)}{\mu^F(S^M)} \right] + \\ & \left[\int_{S^M} g^F(x) \frac{dF^F(x)}{\mu^F(S^M)} - \int_{\overline{S^M}} g^F(x) \frac{dF^F(x)}{\mu^F(\overline{S^M})} \right] \mu^F(\overline{S^M}). \end{aligned}$$

Finally, the second pair of integrals in this expression (those that are computed over the common support) can be decomposed in an analogous way as is done in the Blinder-Oaxaca setup by adding and subtracting the element that permits them to evaluate the counterfactual mentioned above, $\int_{S^M} g^M(x) \frac{dF^F(x)}{\mu^F(S^M)}$, obtaining⁸

$$\begin{aligned} \Delta = & \left[\int_{\overline{S^F}} g^M(x) \frac{dF^M(x)}{\mu^M(\overline{S^F})} - \int_{S^F} g^M(x) \frac{dF^M(x)}{\mu^M(S^F)} \right] \mu^M(\overline{S^F}) + \\ & \int_{S^M \cap S^F} g^M(x) \left[\frac{dF^M}{\mu^M(S^F)} - \frac{dF^F}{\mu^F(S^M)} \right] (x) + \int_{S^M \cap S^F} [g^M(x) - g^F(x)] \frac{dF^F(x)}{\mu^F(S^M)} + \\ & \left[\int_{S^M} g^F(x) \frac{dF^F(x)}{\mu^F(S^M)} - \int_{\overline{S^M}} g^F(x) \frac{dF^F(x)}{\mu^F(\overline{S^M})} \right] \mu^F(\overline{S^M}). \end{aligned}$$

Which I denote by

$$\Delta = \Delta_M + \Delta_X + \Delta_0 + \Delta_F.$$

The typical interpretation of the wage gap decomposition applies, but in this new construction, only over the common support. In this construction, two new additive components have been included, leaving us with a four-element decomposition.

⁸We are denoting by $\frac{dF^M}{\mu^M(S^F)} - \frac{dF^F}{\mu^F(S^M)}$ the measure (with signal) induced by the original measures dF^M and dF^F and the corresponding arithmetic operations.

The first component,

$$\Delta_M \equiv \left[\int_{S^F} g^M(x) \frac{dF^M(x)}{\mu^M(S^F)} - \int_{S^F} g^M(x) \frac{dF^M(x)}{\mu^M(S^F)} \right] \mu^M(\overline{S^F}), \quad (2)$$

is the part of the gap that can be explained by differences between two groups of males — those whose characteristics can be matched to female characteristics and those who cannot. This component accounts for that part of the gap that would disappear in the case that there were no males with combinations of characteristics X that remain entirely unmatched by females, or alternatively, if those males with individual characteristics that are not matched by females were paid, on average, as the average matched males. It is computed as the difference between the expected wage of males out of the common support minus the expected wage of males in the common support, weighted by the probability measure (under the distribution of characteristics of males) of the set of characteristics that females do not reach.

The second component,

$$\Delta_X \equiv \int_{S^M \cap S^F} g^M(x) \left[\frac{dF^M}{\mu^M(S^F)} - \frac{dF^F}{\mu^F(S^M)} \right] (x), \quad (3)$$

is the part of the wage gap that can be explained by differences in the distribution of characteristics of males and females over the common support. In the linear BO setup this corresponds to the component $\widehat{\beta}^M \cdot (\overline{x}^M - \overline{x}^F)$.

The third component,

$$\Delta_0 \equiv \int_{S^M \cap S^F} [g^M(x) - g^F(x)] \frac{dF^F(x)}{\mu^F(S^M)}, \quad (4)$$

corresponds to the “unexplained part.” That share of the wage gap that cannot be attributed to differences in characteristics of the individuals and is typically attributed to a combination of both the existence of unobservable characteristics that explain earnings and the existence of discrimination. In the linear Blinder-Oaxaca setup this corresponds to the component $(\widehat{\beta}^M - \widehat{\beta}^F) \cdot \overline{x}^F$.

The fourth component,

$$\Delta_F \equiv \left[\int_{S^M} g^F(x) \frac{dF^F(x)}{\mu^F(S^M)} - \int_{S^M} g^F(x) \frac{dF^F(x)}{\mu^F(S^M)} \right] \mu^F(\overline{S^M}), \quad (5)$$

is the part of the gap that can be explained by the differences in characteristics between two groups of females, those who have characteristics that can be matched to male characteristics and those who cannot. It accounts for that part of the gap which would disappear should it ever be the case that all females reach

at least one possible combination of the set of characteristics X that the population of males reach, or alternatively, if these females were paid, on average, as the matched females are paid. It is computed as the difference between the expected wage of females in and out of the common support, weighted by the probability measure (under the distribution of characteristics of females) of the set of characteristics that males do not reach.

In this way the wage gap has been broken into four additive components:

$$\Delta = \Delta_M + \Delta_X + \Delta_0 + \Delta_F. \quad (6)$$

Being the case that three of them can be attributed to the existence of differences in individuals characteristics that the labor market rewards (Δ_X, Δ_M and Δ_F) and the other (Δ_0) to the existence of a combination of both unobservable (by the econometrician) differences in characteristics that the labor market rewards and discrimination. In that sense, it is convenient to express the wage gap as:

$$\Delta = (\Delta_M + \Delta_X + \Delta_F) + \Delta_0. \quad (7)$$

and interpret it as is traditionally done in the linear BO setup, with two components: one attributable to differences in observable characteristics of the individuals and the other considered as an unexplained component of the gap.

Under this framework, I will introduce the matching procedure in order to estimate these four components.

I will re-sample all females without replacement and match each observation to one synthetic male, obtained averaging the characteristics of all males with exactly the same characteristics x . The matching algorithm in its basic form can be summarized as follows:

- Step 1: Select one female from the sample (without replacement).
- Step 2: Select all the males that have the same characteristics x as the female previously selected.
- Step 3: With all the individuals selected in Step 2, construct a synthetic individual whose characteristics are equal to the average of all of them and “match” him to the original female.
- Step 4: Put the observations of both individuals (the synthetic male and the female) in their respective new samples of matched individuals.
- Repeat the steps 1 through 4 until it exhausts the original female sample.

As a result of the application of this one-to-many-with-zero-discrepancies matching I generate a partition of the dataset. The new dataset contains observations of “matched females”, “matched males”, “unmatched

females” and “unmatched males,” being the case that the sets of matched males and females have the same empirical distributions of probabilities for characteristics X .

The purpose of re-sampling without replacement from the sample of females and with replacement from the sample of males is to preserve the empirical distribution of characteristics for females (being the case that the support for that distribution is finite). This allows us to generate the appropriate counterfactual and interpret the four components as I have done in this section. This generation of a counterfactual — continuing the analogy with the original Blinder-Oaxaca setup — can also be done the opposite way (that is, re-sampling without replacement for males and with replacement for females) with the appropriate changes in the interpretation of the four components derived.

In such a way, the estimation of the four components previously presented is reduced to simple computations of conditional expectations and empirical probabilities without it being necessary to estimate the non-parametric earnings equations $g^M(\cdot)$ and $g^F(\cdot)$.⁹

$$\begin{aligned}
 \Delta_M &= \mu^M(\text{Unmatched}) (E_{M,\text{unmatched}}[Y|M] - E_{M,\text{matched}}[Y|M]) \\
 \Delta_X &= E_{M,\text{matched}}[Y|M] - E_{F,\text{matched}}[Y|M] \\
 \Delta_0 &= E_{F,\text{matched}}[Y|M] - E_{F,\text{matched}}[Y|F] \\
 \Delta_F &= \mu^F(\text{Unmatched}) (E_{F,\text{matched}}[Y|F] - E_{F,\text{unmatched}}[Y|F])
 \end{aligned} \tag{8}$$

The use of this matching criterion allows us to keep away from any type of parametric assumptions that may impose restrictions on the behavior of the random variables involved in the analysis. It is solely based on the modeling assumption that individuals with the same observable characteristics should be paid the same regardless of their sex.

The analysis presented here raises a point to be taken into account in the traditional setup of the BO decomposition, one that has not received considerable attention but plays an important role: the supports of the distributions of characteristics for females and males may not overlap completely, thus it is necessary to restrict the decomposition in terms of “differences in characteristics and differences in coefficients” only to the common support, where the comparison of wages makes sense.

Using the BO decomposition, it is necessary to implicitly make “out-of-the-support” assumptions on the linear estimators obtained by the regressions¹⁰, assumptions that may seem plausible, but for which it is impossible to find evidence in favor or against. By the decomposition proposed here we are not required

⁹The notation used in the following formulae is self-explanatory: the sub-indexes on the expectations denote the distribution according to which the expected value is taken.

¹⁰Namely, the assumption that the fitted regression hyperplane (or surface) can be extended for individual characteristics that have not been found empirically in the data sets, using the same estimators computed with the observed data.

to make these kinds of assumptions, and additionally, I propose a way to compute those components of the gap that correspond to the non-overlapping supports (Δ_M and Δ_F).

As will be shown later in this paper, it is an empirical regularity that the unmatched males have average wages above the average wages of their matched peers. Hence, running regressions to estimate earnings equations for all males without recognizing that empirical regularity tends to over-estimate the unexplained component (Δ_0) in the BO decomposition.

It is important to emphasize the nature of gender discrimination in *pay* that Δ_0 captures (that is, the possibility of having equally productive males and females that are paid differently simply because of gender) and distinguish it from other sorts of discrimination that may play role in the *access* to particular characteristics. The extent to which these differences in access are endogenous or exogenous to the labor market may vary as we may think about discrimination that prevents promotion to high paying occupations as an example of the former and differences in education as (arguably) an example of the latter. While discrimination in access is embodied in the three components attributed to differences in characteristics, I believe that the Δ_M component accounts for the penalization on average wages that females experience by encountering “barriers to the entry” that block their way to certain individual characteristics that males achieve. Unfortunately however, due to unobserved heterogeneity, it is not possible to distinguish whether that Δ_M component is a result of “discrimination” or “choice.”

The next section will explore the data set for which the decomposition just introduced is implemented, analyzing gender differences in some of the observable characteristics that the labor market rewards.

4 Gender Differences in Characteristics and the Gender Wage Gap in Peru 1986-2000

The data for this study come from the National Household Surveys (Encuestas Nacionales de Hogares) and the Specialized Employment Survey (Encuesta Especializada de Empleo) undertaken by the Peruvian Ministry of Labor and Social Promotion (MTPS) during the period 1986-1995 (except 1988) and by the National Institute of Statistics and Informatics (INEI) for the period 1996-2000. For homogenizing purposes—and taking into account that Lima concentrates almost one half of the Peruvian labor force—only workers fourteen years or older in the metropolitan Lima area have been considered for this study.

Peru, during this time frame is an interesting country to analyze. First, labor markets in Peru are segmented. As mentioned earlier in the introduction, Blau and Ferber (1992) draw attention to the fact that Latin America is the region that reports the highest levels of occupational segregation¹¹ by gender in

¹¹Measured by the Duncan Index of Occupational Segregation.

the world. These high levels of occupational segregation are associated with gender differences in age and schooling of the working population which in turn would presumably imply a severe problem of *gender differences in the supports*.

In addition to the problem of occupational segregation, informality also plays a role in the Peruvian labor markets, as an important fraction of the jobs tend to fail at least one of the formality conditions (formal contract or access to insurance). The formality situation of the working class affects males and females differently: while 55% of males work on informal jobs, the analogous figure for females is 65%.

On the other hand, Peru is one of the Latin American countries that have experienced labor market reforms during the early 1990's.¹² These reforms included dramatic reductions in firing costs that were linked to reductions in formality and a subsequent increase in turnover rates with simultaneous shorter durations of both employment and unemployment spells (Saavedra and Torero(2000)). The theoretical literature has no clear predictions about how these changes in employment dynamics will impact wage differentials. Therefore, it will be interesting to analyze how the gender wage gap evolved during this period.


When explaining gender differences in earnings, it can be argued that “the gender wage gap simply reflects gender differences in some observable characteristics of the individuals that are determinants of wages”. To some extent that argument is valid as there are differences in age, education, occupational experience, and occupations, among others. These differences will partially explain the wage gap. The purpose of this paper is to measure precisely up to what extent these differences in characteristics explain the differences in pay. Exploring some descriptive statistics showing these gender differences will shed some lights.

In terms of the age of the working population, males are, on average, three years older than females. This is in contrast to the whole Peruvian population, where the average age for females is slightly higher than for males (due to females' higher life expectancy). This difference in average ages among workers may reflect earlier entrance or earlier retirement into/from the labor market for females. Both circumstances are expected to have a negative impact on wages. The first is due to the fact that an early entrance into the labor market may imply fewer years of schooling and the second because early retirement implies shorter tenure.

There are also significant differences in gender statistics with regard to educational attainment, as demonstrated in Figure 1. While 16% of working males have an elementary education level or less, 24% of working females fall in this category. In terms of years of schooling there is a related pattern. While working males have an average of 10.75 years of schooling, working females have 9.86 years on average

¹²The two waves of reform occurred in 1991 and 1995.

Figure 1: Educational Attainment by Gender



during the period of analysis.


These average figures for the period 1986-2000 show an evolution that is important to note. The percentage of working females with a college or high school degree increased from 68% to 81%, while for their male peers these percentages moved from 78% to 84%.

The observable characteristic for which the greatest gender difference is found is occupational experience of the working people, measured as years working in the same occupation, illustrated in graph 2.

For the period in consideration, on average, males register between 1.4 and 2.7 more years of occupational experience than females, which represents between 30% and 50% difference. It should be noted, however, that these gender differences in average years of occupational experience have decreased substantially over the period 1986-2000.

Regarding the *differences in the supports* that this paper points out and addresses, I have found that 30%

Figure 2: Evolution of Occupational Experience



of working females exhibit combinations of age, education, migratory condition¹³ and marital status that cannot be matched by any male in the sample. Analogously, 23% of working males report combinations of the same individual characteristics (age, education, migratory condition and marital status) that no female shows. Interestingly, this 23% of working males report wages that are considerably higher than those reported by the rest of working males.


As noted, there are gender differences in some observable characteristics that the labor market rewards. But, it is also noted that these gender differences have been narrowing over the period in consideration. The next section will explore the relationship between the characteristics previously shown and the hourly wage, explaining (partially) the gender wage gap and its evolution during the fifteen years of our analysis in Peru.

Wages evolved considerably during the time frame of analysis. After the rise in real wages that started in 1985 and continued until 1987, there followed a significant fall in real wages in a context of hyperinflation. Real wages reached their minimum level in 1990, after which they improved. During the nineties, real wages increased steadily until the late years of the decade, when they began to decline again. Graph 3 shows the evolution of the hourly wage for males and females during the period 1986-2000. The hourly wages are measured in constant 1994 Peruvian Soles (S/.).

Implicitly, the previous graph is also showing the absolute values (in constant 1994 Soles) of the gender

¹³In this paper, I am distinguishing only those who born in Lima from those who born out of Lima.

Figure 3: Hourly Wages by Gender (in 1994 soles)



wage gap (represented by the difference between any pair of adjacent bars). The next figure explicitly shows the gap in relative terms (average hourly wage gap as multiples of average hourly female earnings)¹⁴. It can be found that the gender wage gap in hourly wages varied around an average value of 0.45 (that is, on average, males earned 45% more per hour than females in Peru during the period 1986-2000) but there are significant fluctuations around that average measure.

The measure of the gap that is reported in this section (multiples of average hourly wages for females, or Δ , as it is called in Section 2) should be taken as “raw” in the sense that it considers all males and females regardless of their differences in observable characteristics, and regardless of whether it is possible to compare them or not. It is necessary to make the appropriate adjustments to that gap in order to obtain a measure of unexplained differences in average earnings for comparable samples of males and females, Δ_0 . That will be the purpose of the next sub-section, but before starting that exercise let us explore how these gender differences in average hourly wages vary according to individual characteristics.

Starting with age, note once the population have reached 30, as they get older, the gender wage gap tends to increase; for people close to retirement age the gap reaches 128%.¹⁵

According to educational attainment, the gender wage gap exhibits a non-monotonic behavior. The gap is bigger both for people with only an elementary education and for people with college degrees. It gets

¹⁴Note that the variable in which the gender gap is measured in this paper is the hourly wage instead the logarithm of the hourly wage as is common place in the literature. In sub-section 4.2 there is a discussion on the convenience of the latter over the former.

¹⁵It is important to note that this basic computation of average wage gaps for different age groups mixes “age effects” and “cohort effects.”

Figure 4: Hourly Wage Gap by Gender (in 1994 soles)




Figure 5: Hourly Wages and Gender Wage Gap for Different Age Groups

PERU 1986-2000
 HOURLY WAGES ACCORDING TO
 GENDER AND AGE
 (In 1994 Soles)

	Less Than 19 Years	20 to 29	30 to 44	45 to 60	60 or more
FEMALES	5.58	10.03	12.97	12.45	10.17
MALES	7.62	11.99	17.11	20.32	23.16
GAP	37%	20%	32%	63%	128%

Figure 6: Hourly Wages and Gender Wage Gap by Educational Attainment

PERU 1986-2000
 HOURLY WAGES ACCORDING TO
 GENDER AND EDUCATION
 (In 1994 Soles)

	NO EDUCATION	ELEMENTARY SCHOOL	HGH SCHOOL	COLLEGE
FEMALES	6.52	6.83	9.56	16.32
MALES	7.79	10.22	11.86	23.81
GAP	19%	50%	24%	46%

smaller for the no-education and high school populations. This fact is in line with the gender differences in the return to schooling for Peru found in Saavedra and Maruyama (1999).

The previous tables revealed substantial differences in the distribution of wages and the gender wage gap according to some individual characteristics, each analyzed independently. Next I will analyze the joint effect of these differences in characteristics on wages by means of matching and the decomposition presented in Section 2.

5 Explained and Unexplained Components of the Gender Wage Gap

5.1 Wage Gap Decomposition. The Matching Approach

Recalling from equation [6], the wage gap, Δ , can be expressed as

$$\Delta = E[Y|M] - E[Y|F] = \Delta_M + \Delta_X + \Delta_0 + \Delta_F.$$

That is, the average wage difference between males and females can be broken into four components. Three of them can be attributed to gender differences in observable individual characteristics (Δ_M , Δ_X and Δ_F) and the fourth component to the existence of both non-observable gender differences in characteristics that determine wages and gender discrimination in pay in the labor market (Δ_0):

Δ_X represents the part of the gap explained by the fact that males and females tend to have individual characteristics that are distributed differently over their common support (for instance, in the Peruvian data sets it is possible to find both males and females with “Masters or Ph.D. degree”, but the proportion of females under that category is substantially smaller than the proportion of males).

Δ_X accounts for the expected decrease in males wages in a hypothetical situation in which their individual characteristics follow the distribution of female characteristics.

Δ_F represents that part of the gap explained by the fact that there are some combinations of female characteristics for which there are no comparable males (for instance, in the Peruvian data sets there are some married females, migrants, with zero or only a few years of schooling and some years of occupational experience, but it is not possible to find comparable males with those combinations of characteristics). Δ_F measures the expected increase in wages that the average female wages will experience supposing all females achieve characteristics that are comparable to those of the males.

Δ_M accounts for that part of the gap that exists because some combinations of characteristics that males have, are not reached by females (for instance, in the Peruvian data sets there are males with high levels of education that have been working for more than ten years at managerial occupations, but it is not possible to find observations for females with such characteristics). Δ_M measures the expected increase in wages that the average female wages would have if females achieve those individual characteristics of males that remain “unreached” by females.

And last, Δ_0 is that part of the wage gap that can not be explained by these differences in observable characteristics. As was mentioned previously, this can be explained as a combination of gender differences in characteristics that are related to productivity but unobservables, and discrimination in pay.


The next chart reports the evolution of the raw¹⁶ gender wage gap accompanied by the wage gap that persists after controlling for age, education, marital status and migratory condition with matching. According to the notation introduced in this paper, the chart is reporting the evolution of the raw, Δ , and controlled, Δ_0 , gender wage gap for the period of analysis.¹⁷

The next additive-components bar chart will be used to represent the wage gaps measured in relative terms (as multiples of female wages) and the decompositions in terms of the four components introduced in this paper. The total height of each bar is proportional to the wage gap in the respective year and the height of each component is proportional to the value of the respective component, such that whenever a component has a negative value, it is represented below the zero line. The first set of decompositions reported below has been calculated using different combinations of explanatory variables such as age (measured in years), education (measured as years of schooling), marital status (a dichotomous variable

¹⁶The measure of wage gap that I am using is $\frac{\bar{y}^M}{\bar{y}^F} - 1$.

¹⁷For this and the next decompositions, I omit the decomposition that corresponds to the year 2000 due to a problem on the coding of one of the explanatory variables.

Figure 7: Gender Wage Gap After Controlling for Observable Characteristics



that takes the value 0 for singles and 1 for married individuals) and migratory condition (a dichotomous variable that distinguishes individuals who were born in Lima from those who were not).

While the gender wage gap without controlling for characteristics, Δ , has an average value of 45% during the period of analysis, the controlled gap, Δ_0 , varies around 28%¹⁸. That is, the mixture between gender differences not considered in the analysis (which may comprise observable and unobservable differences) and discrimination accounts for a differential of 28% in hourly wages for males relative to females. These figures correspond to the use of the particular set of variables specified above. That set does not include variables that are typically considered as being determined endogenously in the labor market. Combinations of these variables are considered for the following decompositions. For these I consider different combinations of age, education, occupational experience (measured in years), informality (a dichotomous variable that distinguishes individuals with formal jobs from individuals with informal jobs¹⁹), occupation (that comprises seven occupational categories) and firm size (with five categories).

The average unexplained gender wage gap (Δ_0) that results after controlling for these endogenous characteristics is slightly below the average that does not consider them. It is around 25%, three percentage points below the gap estimated after matching only on age, schooling, marital status and migratory condition.²⁰ Interestingly, for almost every combination of characteristics I considered in the previous exercises,

¹⁸As will be shown in sub-section 5.3, a 99% confidence interval for this average unexplained gender differences in pay ranges from 24.92% to 31.13%.

¹⁹A job is considered formal if satisfies at least one of the following requirements: being in the Public Sector or being registered on the Social Security System or being affiliated to any private retirement plan or being unionized. Family workers are considered informal workers.

²⁰A detailed spreadsheet with the results for all the decompositions showed here, as well as some other combinations of individual characteristics not reported in this section, is available from the author.

Figure 8: Wage Gap Decompositions for Different Sets of Controls (1)




Figure 9: Wage Gap Decompositions for Different Sets of Controls (2)




Figure 10: Wage Gap Decompositions for Different Sets of Controls (3)







Figure 11: Wage Gap Decompositions for Different Sets of Controls (4)



the controlled gender wage gap shows two peaks, one at the end of the 1980's, during the hyperinflation period, and another in the middle of the 1990's during the recession that followed the stabilization of 1990-1994. Also, the lower values for the gap are found around 1986 and 1993 — years that register significant growth in Peruvian GDP.

Analyzing the role that these four components play in the decomposition, there is an interesting aspect to note: the components Δ_0 and Δ_M explain more than 80% of the wage gap during all years for almost all possible combinations of characteristics. As was mentioned earlier in this section, both components of the gap may be regarded as “noisy” measures of discrimination (or unexplained differences) in the labor market. While the former is linked to differences in pay, the latter is presumably linked to differences in access to particular combinations of characteristics that are well rewarded in the labor market.

In the next two sub-sections I analyze the distribution of the unexplained gender differences in pay that can be obtained thanks to the matching approach. I will analyze first the distribution of wages for females, comparing it to the counterfactual distribution of wages for males when they are re-sampled in order to mimic the distribution of individuals' characteristics of the female population.

5.2 Differences in Hourly Wages Between Matched Samples


As was mentioned earlier in the paper, another critique of the Blinder-Oaxaca decomposition is that it is informative only about average effects in its basic setup, not about the distribution of such effects. An alternative to address that point has been the use of quantile regressions instead of O.L.S., decomposing gender wage gaps at different quantiles of the distribution of the error term of the earnings equations. This approach suffers from the same problem of the *gender differences in the supports* that this paper is addressing.

This sub-section is devoted to the analysis of the distribution of wages for males and females. For that purpose, the object of analysis will be the cumulative distribution functions of hourly wages for the original and matched samples of females and males.

By plotting the cumulative functions it is easy to verify that not only is it the case that average wages for males are greater than average wages for females as was pointed out before, but also the random variable “wages for females” is stochastically dominated by the random variable “wages for males” in the Peruvian labor market during 1986-2000. That result still holds if the comparison is made between the resampled (by matching) versions of the same random variables. Even after controlling for observable characteristics that the labor market rewards (age, schooling, marital status and migratory condition in this case) there are gender differences in pay that favor males. Figure 12 shows that result.

With the purpose of visualizing better these differences in the cumulative functions, Figure 13 shows

Figure 12: Cumulative Functions of Relative Wages by Gender



an extract of Figure 12.

The differences between the matched versions of the cumulative functions of wages for females and males are smaller than the differences originally found in the cumulative functions of wages for females and males. The gender differences in wages are reduced after matching. The distribution of hourly wages for matched females does not differ too much from the distribution of hourly wages for all females. This is because, by construction of the counterfactual, the re-sampling has been done in order to ensure that the distribution remains unchanged on the common support. The only changes are due to the non-overlapping parts of the support of characteristics for females (and, as it has been shown previously, the Δ_F component of the gap is relatively small compared to the other components). For males, the situation is different. The cumulative distribution of hourly wages for all males differs from the distribution that considers only matched males (with the appropriate re-weighting that is required to mimic the empirical distribution of individual characteristics of females), especially at the upper extreme of the distribution

The previous plot inspires a quantile analysis in the following way: at any height (percentile), the horizontal distance between the two cumulative functions obtained after matching is a measure of the unexplained gender wage gap at the respective percentile. Graph 14 shows, by percentiles, these measures of gender wage gap that remain after matching.

The plot shows that for the first 90 percentiles of the distribution of hourly wages for males and females there are no major differences in hourly wages. The gap is roughly below .2 times the average wage for females. It is in the top 10% of the distributions of hourly wages for males and females that the highest differences are found. At the 99th percentile the gap attains a maximum of 2.2 times the average wage

Figure 13: Cumulative Functions of Relative Wages by Gender (Extract)




Figure 14: Absolute Gender Wage Gap by Percentiles





Figure 15: Relative Gender Wage Gap by Percentiles



of females. The plot shows evidence that the gender differences in pay in the bottom percentiles of the distribution do not contribute considerably to the aggregate measure of gender differences in pay in Peru for the period of analysis. The average gender wage gap in Peru is driven by gender differences in pay at the top percentiles of the wage distributions.

The assertions of the previous paragraph are hiding an important result, namely, the differences in hourly wages that are found in the bottom percentiles of the distributions of wages are small in absolute terms but not in relative terms. The typical male who is in the bottom 10th percentile of the distribution of hourly wages earns a premium of 12% of average female wages over the 10th percentile female (approximately 1.40 Peruvian Soles of 1994). However, this represents a difference of 60% of that female’s earnings. When the same comparison is made at the bottom first percentile the differences are even bigger. The hourly wage gap in absolute terms is approximately 0.70 Peruvian Soles of 1994, but that figure represents a difference of 94% of the corresponding female earnings. The poorest male earns almost twice as much as the poorest female. These percentage differences in hourly wages by percentiles of the wage distributions are shown next in graph 15.

The relative gender wage gap by wage percentiles shows a slight U-shape in which the minimum gap, 18%, is found among those individuals whose wages are between the 8th and 9th deciles. The maximum is found among the poor, 95%.

5.3 Construction of Confidence Intervals for Average Unexplained Gender Differences in Pay

One of the advantages of using matching is that it is straightforward to generate an empirical distribution of unexplained gender differences in pay. It is only necessary to compute the differences in hourly wages between every female observation and her respective matched male observation in the sample. In this section I will take advantage of that simplicity and compute average unexplained gender differences in pay and their respective standard deviations, for the whole sample and for different conditions on characteristics x .

The empirical distribution of unexplained gender differences in pay that is generated by the algorithm described in Section 2 has embedded the random selection of males in the sense that for every female with characteristics x , the matching algorithm picks randomly one individual among all those $n_M(x)$ males that exhibit the same characteristics x . For this sub-section I will modify Step 3 of the matching algorithm to contemplate not only the random match between the female who has characteristics x and the randomly selected male among the $n_M(x)$ who also exhibits the set of characteristics x .²¹ But also contemplate all other $(n_M(x) - 1)$ possible matches that can involve that female with characteristics x .

Formally, for every female with characteristics x , the refined algorithm of this sub-section will compute the gender differences in pay for all the $n_M(x)$ possible matches that can involve this female, each happening with probability $1/n_M(x)$. This approach naturally generates an empirical distribution of unexplained gender differences in pay for all possible matches for each female observation. Aggregating the distributions computed at the individual level for all female observations, I generate an empirical distribution of possible matches at the whole sample level.²²

Such distribution of unexplained gender differences in pay for the whole sample is the object of study of this sub-section. Here, rather than reporting wage gap decompositions, the emphasis will be on reporting results only for the unexplained differences in pay. Specifically, I will report means and standard deviations for the average unexplained gender differences in pay, conditional on some selected observable characteristics. These conditional means and standard errors will be estimated from the empirical distribution of unexplained differences in pay generated by the procedure described in the previous paragraph.

Specifically, the unexplained differences in pay component

$$\Delta_0 = \int_{S^M \cap S^F} [g^M(x) - g^F(x)] \frac{dF^F(x)}{\mu^F(S^M)}$$

²¹I will denote by $n_M(x)$ the number of males in the sample who report the set of characteristics x . Analogously $n_F(x)$ is the number of females reporting such combination of individual characteristics.

²²In such distribution of matches, the match that involves one female and one male with characteristics x will receive a weight $\frac{1}{n_F} \frac{1}{n(x)}$. (provided that there are n_F observations for females in the sample).

can be expressed as

$$\Delta_0 = \int_{S^M \cap S^F} g^M(x) \frac{dF^F(x)}{\mu^F(S^M)} - \int_{S^M \cap S^F} g^F(x) \frac{dF^F(x)}{\mu^F(S^M)},$$

using $g^M(x) = \frac{1}{n_M(x)} \sum_{i/x_i=x} y_i^M$ and denoting by $n_F(x)$ the number of observations for females with the set of characteristics x in the sample, I construct the sample analog

$$\delta_0 = \sum_x \left(\sum_{i/x_i=x} \frac{1}{n_M(x)} y_i^M \right) \frac{n_F(x)}{n_F} - \sum_{j=1}^{n_F} y_j^F \frac{1}{n_F}$$

which, after denoting $\sum_{i/x_i=x} \frac{1}{n_M(x)} y_i^M$ by $\bar{y}^M(x)$ (the sample average of earnings for males that exhibit the set of characteristics x) and $\frac{n_F(x)}{n_F}$ by $\hat{\omega}^F(x)$ (the sample proportion of females that exhibit the set of characteristics x), can be in turn expressed as

$$\delta_0 = \sum_x \bar{y}^M(x) \hat{\omega}^F(x) - \sum_{j=1}^{n_F} y_j^F \frac{1}{n_F}. \quad (9)$$

From this expression, the asymptotic distribution of the second component of the right-hand side is straightforward to obtain, $\sqrt{\frac{1}{n_F}} (\bar{y}^F - y^F) \xrightarrow[n_F \rightarrow \infty]{} N(0, \sigma_F^2)$. What is not trivial to obtain is the asymptotic distribution of the first component of the right-hand side, but it can be computed applying the δ -method. The details of that computation are shown in the appendix. Applying the same δ -method on restricted samples, according to different sets of characteristics x , I also obtained estimators for the mean and the standard deviation of the unexplained differences in pay for those different sets of characteristics. Next, in Figure 16, I will show the results for the whole population and those that result after conditioning on marital status and migratory condition.

Starting with the gender wage gap for the whole population on the common support of individual characteristics, after controlling for age, schooling, marital status and migratory condition, the average wage gap of 28.03% has a standard error of 1.89%. This can be translated into a 99% confidence interval for the average unexplained differences in pay that ranges from 24.92% to 31.13% of average female wages, while a 90% confidence interval for the same measure would range from 23.17% to 32.89%.

Concerning migratory condition, there is a slight evidence that the unexplained differences in pay are smaller among migrant individuals than among those who born in Lima. With regard to marital status, although there is no clear evidence that the average unexplained gender differences in pay between married and single individuals are different, there is more evidence of dispersion of such unexplained differences among the married than among the singles. That higher dispersion of unexplained wages could be explained in terms of other variables that are considered as endogenous to a model of wage determination in the

Figure 16: Standard Errors for the Average Unexplained Gender Wage Gap

Peru 1986-1999		
Unexplained Gender Wage Gap By Selected Characteristics (After Controlling for Age, Education, Marital Status and Migratory Condition)		
	Mean	Std. Error
All	0.2803	0.0189
Marital Status		
Single	0.2751	0.0242
Married	0.2862	0.0289
Migratory Condition		
Born in Lima	0.3067	0.0300
Born out of Lima	0.2840	0.0269

labor market as occupational experience, tenure, hours worked per week and occupation, as it is more likely to observe higher dispersion in these variables among the married than among the single.

Next I will report the average and standard deviations for unexplained differences in pay conditional on age and marital status. For that purpose I will use box-and-whisker plots to report the confidence intervals. The upper extreme of the whisker corresponds to the maximum of a 99% confidence interval for the average unexplained differences in pay, the upper extreme of the box to the maximum of a 90% confidence interval, the lower extreme of the box to the minimum of a 90% confidence interval and the lower extreme of the whisker to the minimum of a 99% confidence interval. The age groups reported here roughly correspond to the deciles of the age distribution of the employed labor force in Peru.

There is no clear pattern for the evolution of the average unexplained differences by age for single individuals. For the married labor force there is some weak evidence of an increasing evolution of these unexplained differences over the life cycle, but that increase is not significant from decile to decile. However, I note that while the unexplained differences in pay are positive for married individuals above the median age, these differences are not different than zero for those who are younger than the median. The dispersion of such unexplained differences is increasing over the life cycle for singles.

In analyzing unexplained gender differences in pay according to years of schooling there are further interesting issues. These unexplained differences are positive and roughly constant for individuals with less than a high school diploma (less than 11 years of schooling), especially for single individuals. Also, for singles that attended between 4 and 11 years of schooling (which corresponds to 30% of the total employed labor force), there is smaller evidence of dispersion in the unexplained gender wage gap. Among the high

Figure 17: Confidence Intervals for the Unexplained Gender Wage Gap (1)




Figure 18: Confidence Intervals for the Unexplained Gender Wage Gap (2)






Figure 19: Confidence Intervals for the Unexplained Gender Wage Gap (3)



school graduates (those who passed exactly 11 years of schooling and represent 35% of the total employed labor force), there is less evidence for unexplained gender differences in pay, especially among the married.

The biggest unexplained differences are found among those individuals who passed more than 11 years of schooling and represent the remaining 30% of the employed labor force. First, considering only those individuals who attended one to four more years of schooling after having graduated from high school but who did not graduate from college (that is, those who passed between 12 and 15 years of schooling), I find some evidence for positive unexplained differences in pay among the single but not among the married, noting that the dispersion is notoriously higher among the latter group. Next, considering the individuals who have graduated from college (16 years of schooling), the evidence for a positive average measure of unexplained gender differences in pay actually increases, particularly among the married. Finally, the educational group for which I find clear evidence of a positive (and substantial) unexplained gender wage gap is formed by those individuals who graduated from college and continued studying. For this group the unexplained differences seem to represent more than 50% of the average female wage for singles and more than 110% for married females. The dispersion of such unexplained differences among that group is also substantially higher than the dispersion found in any other group.

Figure 20: Confidence Intervals for the Unexplained Gender Wage Gap (4)



5.4 Matching versus Linear Regressions. An Empirical Comparison

After the introduction of the matching approach that decomposes the gender wage gap and avoids linear regressions, there is a comparative question that requires an answer: To what extent do the results obtained by matching differ from those obtained by linear regressions? In some sense, matching is equivalent to Blinder-Oaxaca when the estimations of the earnings equations for males and females are restricted to the common support and performed with the same matching variables and all their possible powers and interactions. We should therefore expect similar results from both. In this paper that question will be answered empirically, comparing the results obtained through matching and linear regressions for two particular years of the sample: 1999 and 2000. I will empirically show that while we obtain similar results (over the common support), the failure to recognize the *gender differences in the supports* accounts for an over-estimation of the unexplained gender differences in pay.

As was pointed out previously, the BO decomposition based on linear regressions depends on the linear specification of the earnings equations. For this purpose it is informative to compare the matching results with those obtained from four different linear specifications. The first specification includes the following variables: age (as a continuous variable), age squared, 3 dummies measuring educational attainment (“elementary school”, “high school” and “college degree”, with “no education” as the base category), 11 dummies measuring years of occupational experience and all the interactions between educational attainment and

tenure. The second uses the same dummies and interactions for educational attainment and occupational experience as in specification 1, but replaces the age and age squared variables with 62 dummies. The third specification changes specification 2 by replacing the 3 dummies for educational attainment with a set of 21 dummies measuring years of schooling (and their interactions with the dummies for years of experience). The fourth specification replaces the interactions between schooling and occupational experience with the interactions of age and occupational experience.

The use of dummies instead of continuous variables has two purposes: on one hand this helps to identify *gender differences in the supports*; on the other hand this allows us to get closer to a setup that has lower dependence on the functional form of the earnings equations (as is the case for matching). Also, in order to evaluate the effects that the failure to recognize *gender differences in the supports* have on the decomposition, this comparative exercise considers two types of BO decompositions for the linear specifications, one without recognizing the existence of such *gender differences in the supports* and hence requiring out-of-support assumptions, and another that recognizes those differences and computes the decomposition only over the common support, and hence also computes the Δ_F and Δ_M components that account for the non-overlapping parts of the supports, as defined in Section 2.²³

Before analyzing the decompositions, there is an interesting difference to note between the results obtained by matching and those obtained by regressions in the measure of the total hourly wage gap. This difference comes from the fact that while the measure of wage gap involved in the regression approach is $\ln(\overline{y_M}) - \ln(\overline{y_F})$, the measure I am using with the matching approach is $\frac{\overline{y_M}}{\overline{y_F}} - 1$. Although the former is the typical measure of the wage gap, I believe that the latter better corresponds to the definition of the concept we are addressing. For small differences in average wages (or average logarithms of the wages) the former is a good approximation of the latter. When the average differences are not so small, the approximation is poor and, in general, it is not possible to establish an order relationship between the two measures.²⁴ In

²³The variables included are:

Specification 1: age (as a continuous variable), age squared, 3 dummies for educational attainment, 11 dummies for experience at the occupation, formality and all the interactions between educational attainment and experience at the occupation.

Specification 2: 62 dummies for age, 3 dummies for educational attainment, 11 dummies for experience at the occupation, formality and all the interactions between educational attainment and experience at the occupation.

Specification 3: 62 dummies for age, 21 dummies for years of education, 11 dummies for experience at the occupation, formality and all the interactions between schooling and experience at the occupation.

Specification 4: 62 dummies for age, 21 dummies for years of education, 11 dummies for experience at the occupation, formality and all the interactions between age and experience at the occupation.

Matching variables: age, schooling, experience at the occupation and formality.

²⁴The idea of working with $\ln(\overline{y_M}) - \ln(\overline{y_F})$ as a measure of wage gap comes from the fact that $\ln(y_M) - \ln(y_F) = \ln\left(\frac{y_M}{y_F}\right) \simeq \frac{y_M}{y_F} - 1$. That is, the difference between the logarithms of wages is approximately equal to the gap between y_M and y_F in relative terms (measured as multiples of y_F). That approximation is valid whenever $\left(\frac{y_M}{y_F}\right)$ is close to 1 (the gap is small). But, even for those situations in which the gap is small enough, the measure of average gap $\frac{\overline{y_M}}{\overline{y_F}} - 1$ is approximately equal to $\ln(\overline{y_M}) - \ln(\overline{y_F})$ but not necessarily equal to $\ln(\overline{y_M}) - \ln(\overline{y_F})$, which is the measure used in the linear approach.

Figure 21: Comparison Among Different Decompositions of the Gender Wage Gap

Urban Peru 1999 Gender Wage Gap					Unexplained Component of the Gap	Explained by Individual Characteristics
Decompositions	Delta-F	Delta-0	Delta-X	Delta-M		
Matching (No specification required)	0.014	0.228	0.064	0.050	0.228	0.129
Linear Specifications						
Identifying Differences in Supports						
Specification 1	0.043	0.209	0.036	0.068	0.209	0.148
Specification 2	0.043	0.229	0.016	0.068	0.229	0.127
Specification 3	-0.010	0.216	0.064	0.087	0.216	0.141
Specification 4	-0.010	0.226	0.053	0.087	0.226	0.130
Without Identifying Differences in Supports						
Specification 1		0.224	0.132		0.224	0.132
Specification 2		0.237	0.120		0.237	0.120
Specification 3		0.277	0.080		0.277	0.080
Specification 4		0.272	0.085		0.272	0.085

this example, it is found that the “regression” (or “logarithms”) measure is substantially lower than the “matching” measure. While the former reports that males earn on average 35.7% more than females (per hour), the latter reports a figure of 62.1%.

Let me now turn to the comparison of different decompositions. First, I will compare two sets of linear specifications: those that take the *gender differences in the supports* into account with those that do not. I found that differences in supports account for a significant share of the gap, particularly because there is a significant percentage of males with individual characteristics that have no female counterparts. Accounting for *gender differences in the supports* changes the unexplained differences in earnings from an estimated average of 22.4%-27.7% to an estimate of 20.9%-22.9% for 1999.

The previous comparisons —made for the same linear specifications but with different assumptions on the supports of the distributions of individual characteristics— show empirical evidence about one of the claims of this paper, namely, the failure to recognize *gender differences in the supports* implies a slight overestimation of the unexplained component of the gap (Δ_0). This point is illuminated by the matching approach.

Now let us turn to the comparison between the decomposition based on linear regressions and the decomposition based on matching. As mentioned previously, the difference between both approaches is on the weighting of the local effects. The measure of the unexplained wage gap reported by matching (22.8%) falls inside the range of estimators obtained with the different linear specifications considered in

By Jensen’s inequality it is known that $\overline{\ln(y_M)} < \ln(\overline{y_M})$ and $\overline{\ln(y_F)} < \ln(\overline{y_F})$, but it is not possible to determine an order relationship between $\overline{\ln(y_M)} - \overline{\ln(y_F)}$ and $\ln(\overline{y_M}) - \ln(\overline{y_F})$.

this comparing exercise. Restricting the modeling of the earnings equations to a specific functional form seems to have no substantial effect on the measure of unexplained differences.

It should be noted that in measuring the gender wage gap in per-hour terms (and only between working males and females), I do not take into consideration effects of participation in the labor market, either at the extensive or at the intensive margin. The gender differences in that respect have decreased during the period of analysis and this is probably related to the common perception by agents in Peru that the gender differences have decreased substantially during the last twenty years in Peru. The next section is devoted to incorporating these participation effects in the computation of the labor income gap.

6 Has there been a Decrease in the Gender Wage Gap?



According to the measure reported for unexplained gender differences in pay in the previous section, there is no evidence of a monotonic decrease of such differences. The hourly wage gap according to gender reaches its lowest levels during 1992 and 1999 while it attains peaks during 1989 and 1997, evolving in a way that seems correlated with the cycle of the Peruvian economy. That measure of gender differences does not take into consideration either labor market participation effects or unemployment rates. Interestingly, there are notorious changes in female participation and unemployment rates over the period of analysis, especially after the labor market reforms undertaken during the first half of the nineties.

While participation rates for males do not report dramatic changes over the period of analysis, that same measure for females evolved with a slight decrease at the beginning of the nineties and then an increase towards the second half of that decade, decreasing the differences in participation rates from 28% to 21%. Also, gender differences in unemployment rates decreased over the period. While the male unemployment rate increased from 4% to 7% during the fifteen-year span, the female unemployment rate evolved with substantial ups and downs, reporting a slight increase from 8% to 9% over the whole period. Interestingly, the peaks reported in this evolution of female unemployment rates coincide with the peaks found for the unexplained gender differences in hourly wages (one at the end of the eighties, another during the stabilization period running from 1992-1994 and a third peak at 1997), which in turn are correlated with the cycle of the economy. Higher gender differences in unemployment rates are linked to higher unexplained gender differences in pay.

These raw differences in participation and unemployment rates can also be controlled using the same matching procedure: re-sampling the distribution of male individuals in order to mimic the distribution of characteristics of females in the whole population. That is, generating the counterfactual: What participation (unemployment) rates would the male population have if their individual characteristics were distributed as if they were females? The gender differences, if any, in participation (unemployment) rates

obtained from the matched sample can be considered as “unexplained differences” which, as usual, can be regarded as a sign of the existence of both discrimination and unobservable characteristics determining participation (unemployment).


The next two graphs report the evolution of gender differences in unemployment and participation rates together with the controlled —by matching— versions of such rates for males, considering age, education, marital status and migratory condition as matching variables applied over the whole population.




The results suggest that gender differences in age, education, marital status and migratory condition do not explain gender differences in participation or unemployment rates. If any, the controlled unemployment rates for males are slightly smaller than the “raw” unemployment rates. There are other determinants of such differences, and discrimination may be one of those, but it also may be *choice*, in the sense that this is solely the result of differences in preferences between females and males.

As there are gender differences in participation and unemployment rates (extensive margin), there are also differences in the number of hours worked (intensive margin). On average, males worked 48 hours while females worked 41 hours per week during the period of analysis, this represents an approximate 16% difference in the average number of hours worked by men versus women. Also, these differences have decreased over the period 1986-2000. While males worked 21% more hours than females in 1986, they did work for 13% more hours than females during 2000. It would be interesting to analyze the impact on income that those changes in participation had.

The next graph shows the evolution of gender differences in the Peruvian labor market using a measure of earnings that incorporates participation at both margins (intensive and extensive). For this purpose, the measure to analyze is the Fraction of Total Labor Income Generated by Males, computed at the individual level. In a world in which there were no gender differences in participation, employment and pay in the labor markets, this would take the value 50%. In a world in which only males generate labor income such a measure would be 100%. The evolution of this measure is reported below.




According to this measure of gender differences, there is a monotonic evolution that falls from an average that was around 75% at the end of the eighties and beginning of the nineties to an average of 61% at the end of the nineties, with a tendency to decrease. Furthermore, in analyzing separately this measure for different age groups, it is possible to find interesting differences. Among the young these measures of the difference in the generation of labor income are smaller and, especially for the last two years of the period of analysis, this measure reveals almost no differences between males and females. Labor income is equally generated by gender among that cohort. Among the older people the fraction of labor income generated by males is higher and seems to show the same monotonically decreasing pattern starting in the second half of the nineties.



This measure can be considered as raw in the sense that it does not take into account gender differences in the individual characteristics that determine participation, employment and wages. Again, the generation of a counterfactual is required in order to answer the question: What fraction of total labor income would be generated by males in case their individual characteristics are distributed in the population according to the empirical distribution of individual characteristics of females?

The matching algorithm is applied now, not only to the working population reporting positive wages as was done for the hourly wage gap analysis, but also to the non-working, economically active population. The measure that matters for the purpose of this new exercise is the total labor income generated by females and males in the matched sample. The variables considered for this matching exercise belong to a set of variables that can be considered as “exogenous” to the labor market: age, education, marital status and migratory condition. The results are shown next.



Interestingly, there are almost no differences among the “raw” gender differences in total labor income and its various controlled versions considering combinations of the variables mentioned above. That is, these gender differences in total labor income can not be attributed to gender differences in age, education, marital status and migratory condition; there are other determinants of such differences in generation of labor income (among which we can consider discrimination). Even more, for the last years of our analysis the controlled fraction exceeds the raw fraction. This is explained by an increase in years of education for females which is not accompanied by a corresponding increase in labor income. Females are acquiring more education but they are getting neither more jobs nor higher pay.

This section concludes by summarizing an interesting finding. The set of variables that helped us to (partially) explain gender differences in hourly wages does not have explanatory power for gender differences in participation (both at the extensive and intensive margins). Moreover, for a measure that combines participation and pay —namely, the fraction of total labor income by gender— this set of “exogenous” variables has no explanatory power either. Or equivalently, if any, I find more unexplained gender differences in access to the labor market than in pay (conditional on working), although those gender differences in access have been decreasing during the last fifteen years.

7 Conclusions

This paper introduced a new non-parametric technique to decompose gaps in terms of explained and unexplained components, paying attention to the problem of *gender differences in the supports*. One of the purposes of the paper was to challenge the linear specifications that involve the estimation of earnings equations and to propose matching as an alternative for a world in which the relationships that govern the co-movement of wages and individual characteristics are not necessarily linear. In that respect, I have found empirically that there are no substantial differences that matter for the wage gap decomposition. In the end, the linearity assumption does not make a great difference for the wage gap decomposition —provided it is estimated only over the common support.

There is a substantial new issue that this matching approach raises: the importance of recognizing the problem of *gender differences in the supports*. That is, “not all males are comparable to all females”. The failure to recognize this problem typically implies a slight overestimation of the unexplained component of the wage gap. Also, there is an important share of the wage gap that can be attributed to the fact that in the labor market a significant number of males exhibit a set of characteristics that have no females counterparts, and these characteristics are highly rewarded in the market.

Besides the issue of *gender differences in the supports* that the matching approach raises, there is also an advantage of matching over the traditional linear regressions in terms of the information it allow us to

take directly. By means of matching, instead of obtaining only *average* measures of unexplained differences in pay, it is possible to also obtain a *distribution* for such unexplained differences in pay.

To explore the distribution of unexplained differences in pay provided interesting insights. The average gender wage gap is mainly driven by gender differences in pay at the top percentiles of the wages distributions. Wages at the highest quintile of the distributions of wages for females and males explain more than one half of the average wage gap in Peru for the period of analysis. At the poorest percentiles of earnings, the wage gap in absolute terms is small and does not contribute substantially to the average wage gap of the population; but the same wage gap in the poorest percentiles, measured in relative terms is the highest among all percentiles (around 94%). Also, I found that there is more dispersion of the unexplained gender differences in pay among the married than among singles. There is also slight evidence of an increase in the gender wage gap with age for married individuals and a substantially higher (and more disperse) unexplained gender wage gap among the highly educated (more than college degree).

These two advantages of matching over linear regressions are not a free lunch. There is a cost to pay: the curse of dimensionality. The inclusion of many explanatory variables —that is, the use of many matching characteristics— may reduce the chances of obtaining an adequate number of matched observations, limiting as a consequence the possibility of exploring the distribution of unexplained differences in pay.

The attempt to explain gender differences in participation and unemployment rates in terms of observable out-of-the-labor-market variables (age, education, marital status and migratory condition) fails considerably. This is also true for gender differences in the generation of total labor income. That lack of explanatory power has two interpretations. On the one hand, it may be the case that the discriminatory practices according to gender are more severe in hiring and work load than in the determination of hourly payments. On the other hand, it may be also the case that these gender differences in participation are explained by differences in other non-observable individual characteristics (among which we can include “preferences” or “social roles”).

In general, this matching approach can also be used to control for observed characteristics in any other measure for which it is expected to find some sort of explained and unexplained components. It only requires the generation of the adequate counterfactuals.

References

- [1] Abadie, Alberto and Guido Imbens (2002). “Simple and Bias-Corrected Matching Estimators for Average Treatment Effects.” Unpublished.

- [2] Angrist, Joshua. "Estimating The Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants." *Econometrica*, Vol. 66, No. 2 (March 1998), 249-288.
- [3] Barsky Robert, John Bound, Kerwin Charles and Joseph Lupton (2001). "Accounting for the Black-White Wealth Gap: A Nonparametric Approach." *NBER Working Paper* No. 8466, National Bureau of Economic Research, Cambridge, MA.
- [4] Blau, Francine and Marianne Ferber (1992). "The Economics of Women, Men and Work." 2nd Ed. Englewood Cliffs, NJ: Prentice-Hall.
- [5] Blau, Francine and Lawrence Khan (1996). "International Differences in Male Wage Inequality: Institutions versus Market Forces." *Journal of Political Economy*, 104(4), 791-837.
- [6] Blinder, Allan (1973). "Wage Discrimination: Reduced Form and Structural Estimates." *The Journal of Human Resources*, VII, 4, pp. 436-55.
- [7] Bourguignon, Francois, Francisco Ferreira and Phillippe Leite (2002). "Beyond Oaxaca-Blinder: Accounting for Differences in Household Income Distributions." Unpublished.
- [8] Buchinsky, Moshe (1994). "Changes in U.S. Wage Structure 1963 - 87: An application of quantile regression." *Econometrica*, 62, 405-458
- [9] Dehejia, Rajeev and Sadek Wahba (1998). "Propensity Score Matching Methods for Non-Experimental Causal Studies." *NBER Working Paper* No. 6829, National Bureau of Economic Research, Cambridge, MA.
- [10] Deutsch, Ruthanne, Andrew Morrison, Claudia Piras and Hugo Nopo (2002). "Working Within Confines: Occupational Segregation by Gender in Three Latin American Countries." *Inter-American Development Bank. Sustainable Development Department. Technical Papers Series #126*. Washington D.C.
- [11] DiNardo, John; Nicole Fortin and Thomas Lemieux (1996). "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach." *Econometrica*, Volume 64, No. 5, 1001-1044.
- [12] Dolton, P.J. and G.H. Makepeace (1987). "Marital Status, Child Rearing and Earnings Differentials in the Graduate Labour Market." *Economic Journal*, 97, 897-922.
- [13] Donald, Stephen, David Green and Harry Paarsch (2000). "Differences in Wage Distributions between Canada and the United States: An Application of a Flexible Estimator of Distribution Functions in the Presence of Covariates." *Review of Economic Studies*, 67, 609-633.

- [14] Gardeazabal, J. and A. Ugidos (2000). "Measuring the Wage Distribution Gender Gap at Quantiles." paper presented at the *ESPE2000* conference, Bonn.
- [15] Hansen, Jorgen and Roger Wahlberg (1999). "Endogenous Schooling and the Distribution of the Gender Wage Gap." Discussion Paper No. 78. Institute for the Study of Labor (IZA), Bonn.
- [16] Heckman, James J. and Jeffrey Smith (1995). "Assessing the Case for Social Experiments." *Journal of Economic Perspectives*, 9, No.2, 85-110.
- [17] Heckman, James J. and Jeffrey A. Smith (1996). "Experimental and Nonexperimental Evaluation." Schmid, Gunther O'Reilly, Jacqueline Schomann, Klaus, eds. *International Handbook of Labour Market Policy and Evaluation*. Cheltenham, U.K. and Lyme, N.H.: Elgar, 37-88.
- [18] Heckman, James J., Hidehiko Ichimura and Petra E. Todd (1997). "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme." *Review of Economic Studies*, 64, No.4, 605-54.
- [19] Heckman, James and Carmen Pages (2000). "The Cost of Job Security Regulation: Evidence from Latin American Labor Markets." *NBER Working Paper* No. 7773, National Bureau of Economic Research, Cambridge, MA.
- [20] Heckman, James, Lance Lochner and Petra Todd (2001). "Fifty Years of Mincer Earnings Regressions." Unpublished.
- [21] Jenkins, S. P. (1994). "Earnings Discrimination Measurement—A Distributional Approach." *Journal of Econometrics*, 61, 81-102.
- [22] Juhn, Chinhui; Kevin Murphy and Brooks Pierce (1993). "Wage Inequality And The Rise In Returns To Skill." *The Journal of Political Economy*, Vol. 101, Issue 3, 410-442.
- [23] LaLonde, Robert (1986). "Evaluating the Econometric Evaluations of Training Programs With Experimental Data" *American Economic Review*, 76, 604-20.
- [24] Machado, Jose and Jose Mata (2001). "Earning Functions in Portugal 1982-1994: Evidence from Quantile Regressions." *Empirical Economics*, 26(1), 115-134.
- [25] Meyer, Bruce (1995). "Natural and Quasi-Experiments in Economics." *Journal of Business & Economic Statistics*, Vol. 13, No. 2.
- [26] Montenegro, Claudio (1999). "Wage Distribution in Chile: Does Gender Matter? A Quantile Regressions Approach." Mimeo, The World Bank

- [27] Montenegro, Claudio and Ricardo Paredes (1999). "Gender Wage Gap and Discrimination: A Long Term View Using Quantile Regressions." Mimeo, Universidad de Chile.
- [28] Montes, Jose and Hugo Nopo (2002). "Latin America and the Gender Wage Gap During the Nineties". Unpublished.
- [29] Mora, Ricardo (2000). "A Nonparametric Decomposition of the Mexican American Average Wage Gap." Mimeo, Universidad Carlos III de Madrid.
- [30] Munroe, A. (1988). "The Measurement of Racial and Other Forms of Discrimination." *Discussion Paper in Economics No. 148*, University of Stirling, Stirling.
- [31] Oaxaca, Ronald (1973). "Male-Female Wage Differentials in Urban Labor Market." *International Economic Review*, Vol.14, No.3, 693-709.
- [32] Oaxaca, R. & M. Ransom (1994). "On Discrimination and the Decomposition of Wage Differentials." *Journal of Econometrics*, 61 (1), 5-21.
- [33] Pratap, Sangeeta and Erwan Quintin (2002). "Are Labor Markets Segmented in Argentina? A Semi-parametric Approach." *ITAM Discussion Paper 02-02*, Instituto Tecnológico Autónomo de México, México.
- [34] Rubin, Donald and Paul Rosenbaum (1983). "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika*, Vol. 70, No. 1, 41-55.
- [35] Rubin, Donald and Paul Rosenbaum (1985). "The Bias Due To Incomplete Matching." *Biometrics*, Vol. 41, No. 1, 103-116.
- [36] Rubin, Donald and Neal Thomas (1996). "Matching Using Estimated Propensity Scores: Relating Theory to Practice." *Biometrics*, 52(1), 249-264.
- [37] Saavedra, Jaime and Eduardo Maruyama (1999). "Los Retornos a la Educación y a la Experiencia en el Perú 1985-1997." In "*Pobreza y Economía Social: Análisis de una encuesta (ENNIV-1997)*". Webb, Richard Ed; Ventocilla, Moisés ed, Instituto Cuánto, Lima.
- [38] Saavedra, Jaime (2000). "La Flexibilización del Mercado Laboral." In "*La Reforma Incompleta: Rescatando los Noventa.*" Abusada, Roberto Ed. Universidad del Pacífico, Lima, 379-428
- [39] Saavedra, Jaime and Maximo Torero (2000). "Labor Market Reforms and Their Impact on Formal Labor Demand and Job Market Turnover: The Case of Peru." *Inter-American Development Bank, Research Department, Research Network Working Paper #394*. Washington D.C.

- [40] Smith, Jeffrey and Petra Todd (2000). "Does Matching Overcome Lalonde's Critique of Nonexperimental Estimators?" Unpublished.

Appendix

In this appendix I will show the application of the δ – *method* for the computation of the standard errors of the first component of the right-hand side of (9). I will start introducing some additional notation. K denotes the number of values that the set of characteristics x can attain, \widehat{W}^F is the K -dimensional vector whose elements are the weights $\widehat{\omega}^F(x)$ and \overline{Y}^M is the K -dimensional vector whose elements are the conditional means $\overline{y}^M(x)$. I will also denote by f_i the population proportion of females that exhibit the set of characteristics x_i and by σ_i^2 the population variance of wages of males exhibiting the set of characteristics x_i ($i = 1, \dots, K$).

Using this notation, the first component of the right-hand side of (9) can be expressed as $\widehat{P} = \widehat{W}^F \cdot \overline{Y}^M$.

The asymptotic distributions of \widehat{W}^F and \overline{Y}^M are also well know,

$$\begin{aligned} \sqrt{\frac{1}{n_F}} \left(\widehat{W}^F - W^F \right) &\xrightarrow[n_F \rightarrow \infty]{} N(0, V_{w^F}), \\ \sqrt{\frac{1}{n_M}} \left(\overline{Y}^M - Y^M \right) &\xrightarrow[n_M \rightarrow \infty]{} N(0, V_{Y^M}), \end{aligned}$$

$$\text{where } V_{w^F} = \begin{bmatrix} f_1(1-f_1) & -f_1f_2 & \cdots & -f_1f_K \\ -f_2f_1 & f_2(1-f_2) & \cdots & -f_2f_K \\ \vdots & \vdots & \ddots & \vdots \\ -f_Kf_1 & -f_Kf_2 & \cdots & f_K(1-f_K) \end{bmatrix}$$

$$\text{and } V_{Y^M} = \begin{bmatrix} \sigma_1^2 & 0 & \cdots & 0 \\ 0 & \sigma_2^2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_K^2 \end{bmatrix}.$$

Assuming further that n_F and n_M grow at the same rate ($n_f = \alpha n_m$) and that the sample of females is independent of the sample of males, we have

$$\sqrt{\frac{1}{n_M}} \begin{bmatrix} \widehat{W}^F - W^F \\ \overline{Y}^M - Y^M \end{bmatrix} \xrightarrow[n_M \rightarrow \infty]{} N(0, V),$$

$$\text{where } V \equiv \begin{bmatrix} \frac{V_{w^F}}{\alpha^2} & 0 \\ 0 & V_{Y^M} \end{bmatrix}$$

Applying the delta method ($\frac{\partial P}{\partial W} = Y^M$ and $\frac{\partial P}{\partial Y^M} = W^F$), the limiting distribution of the product \hat{P} can be approximated as

$$\sqrt{\frac{1}{n_M}} (\hat{P} - P) \sim N \left(0, \begin{bmatrix} Y^{M^T} & W^{F^T} \end{bmatrix} V \begin{bmatrix} Y^M \\ W^F \end{bmatrix} \right).$$

In such a way, the asymptotic variance of the first component of the right-hand side of (9) can be computed as

$$\sum_{i=1}^K \left[\frac{\hat{\omega}^F(x_i) (1 - \hat{\omega}^F(x_i))}{\alpha^2} (\bar{y}^M(x_i))^2 + \hat{\sigma}_i^2 (\hat{\omega}^F(x_i))^2 \right] - 2 \sum_{i=1}^K \sum_{j=1}^{i-1} \frac{\hat{\omega}^F(x_i) \hat{\omega}^F(x_j)}{\alpha^2} \bar{y}^M(x_i) \bar{y}^M(x_j).$$