ISER Working Paper Series

Generalized measures of wage differentials

Philippe Van Kerm
CEPS/INSTEAD, Luxembourg
ISER, University of Essex

No. 2009-26
September 2009

INSTITUTE FOR SOCIAL
& ECONOMIC RESEARCH

E·S·R·C
ECONOMIC
& SOCIAL
RESEARCH
COUNCIL

NON-TECHNICAL SUMMARY

Standard statements about gender wage differentials are such as “women are paid x percent less than men on average” followed by a qualification that “controlling for differences in human capital endowments the remaining difference is y percent.” Such statements are typically inferred from regression analysis à la Oaxaca-Blinder. They mean that the difference in average wage between men and women is x percent and that the expected wage for a woman with average human capital endowments is y percent below the average wage for a man with the same characteristics.

There is growing agreement that assessments of wage differentials should go beyond such comparisons of “the average wage of the average person.” Recent studies have shown, for example, that mean differences in pay is often driven by greater differences between men and women at the top of the wage distribution -- an observation that has been interpreted as evidence of a ‘glass ceiling’ impeding progress of women to highest paid jobs. Such a phenomenon is not captured by measures of wage differentials ‘at the mean’. With narrow focus on mean wages, a hypothetical economy with three women paid \$4 each and three men paid \$5 each records the same aggregate level of wage differentials than an economy where women are also paid \$4 each but where one man is paid \$11 and two men are paid only \$2. Claiming that women experience the same degree of disadvantage in these two fictitious economies is however debatable at least.

This paper addresses this issue and considers two new ‘distributionally sensitive’ summary measures of wage differentials, not solely determined by ‘the average wage of the average person’ but by differences across complete wage distributions. Considerations of risk or inequality aversion in the assessment of wage differentials are explicitly included, transplanting expected utility concepts familiar to income distribution analysts.

The proposed indices take estimates of wage distributions conditional on human capital characteristics as input. An added contribution of this paper is to illustrate an original approach to estimate wage distributions flexibly in the presence of covariates and under endogenous labour market participation.

In an application to the gender pay gap in Luxembourg the disadvantage of women persists with these new generalized measures of wage differentials. This suggests that lower average wages for women are not “compensated” by less dispersed distributions.

Generalized measures of wage differentials

Philippe Van Kerm[*]
CEPS/INSTEAD, Luxembourg[†]
and ISER, University of Essex, UK[‡]

July 2009

Abstract

This paper considers new 'distributionally sensitive' summary measures of wage differentials, not solely determined by "the average wage of the average person" but by differences across complete wage distributions. Considerations of risk or inequality aversion in the assessment of wage differentials are explicitly included, transplanting expected utility concepts familiar to income distribution analysts. In an application to the gender pay gap in Luxembourg the disadvantage of women persists with the new generalized measures of wage differentials. This suggests that lower average wages for women are not compensated by less dispersed distributions. The paper also illustrates original estimation of wage distributions in the presence of covariates and under endogenous labour market participation.

Keywords: wage differentials ; discrimination ; expected utility ; Singh-Maddala distribution ; Luxembourg

JEL Classification: D63 ; J31 ; J70

[*]This research is part of the MeDIM project *Advances in the Measurement of Discrimination, Inequality and Mobility*) supported by the Luxembourg 'Fonds National de la Recherche' (contract FNR/06/15/08) and by core funding for CEPS/INSTEAD by the Ministry of Culture, Higher Education and Research of Luxembourg. Comments from CEPS/INSTEAD staff, seminar participants at GREQAM (Marseille) and LAMETA (Montpellier), and participants at the Warsaw International Economic Meeting 2008 and European Society for Population Economics 2009 conference are gratefully acknowledged.

[†]Centre d'Etudes de Populations, de Pauvreté et de Politiques Socio-Economiques/International Networks for Studies in Technology, Environment, Alternatives, Development. Correspondence: B.P. 48, L-4501 Differdange, Luxembourg. Contact: `philippe.vankerm@ceps.lu`.

[‡]Institute for Social and Economic Research, University of Essex, Colchester, UK.

1 Introduction

Standard statements about gender wage differentials are such as "women are paid x percent less than men on average," followed by a qualification that "controlling for differences in human capital endowments the remaining difference is y percent." Such statements are typically inferred from regression analysis à la Oaxaca-Blinder (Blinder, 1973, Oaxaca, 1973). They mean that the difference in average wage between men and women is x percent and that the expected wage for a woman with average human capital endowments is y percent below the average wage for a man with the same characteristics.[1] However, there is growing agreement that assessments of wage differentials should go beyond such comparisons of "the average wage of the average person." Recent studies have shown, for example, that mean differences in pay is often driven by greater differences between men and women at the top of the wage distribution — an observation that has been interpreted as evidence of a 'glass ceiling' impeding progress of women to highest paid jobs (see, e.g., Albrecht et al., 2003, Arulampalam et al., 2007, de la Rica et al., 2008). Such a phenomenon is not captured by measures of wage differentials 'at the mean'. With narrow focus on mean wages, a hypothetical economy with three women paid \$4 each and three men paid \$5 each records the same aggregate level of wage differentials than an economy where women are also paid \$4 each but where one man is paid \$11 and two men are paid only \$2. Claiming that women experience the same degree of disadvantage in these two fictitious economies is however debatable at least.[2]

Substantial progress has been made recently to describe wage differentials at different quantiles (or at different wage levels) taking into account differences in human capital and other observable characteristics; see, e.g., Ñopo (2008), Albrecht et al. (2009), Firpo et al. (2009) to mention only the most recent contributions. But advantage has not been taken of these new techniques to derive refined, 'distributionally sensitive' summary measures of wage differentials. Scalar summary measures are however essential to monitor progress and inform policy, alongside more flexible graphical tools.[3] This paper considers two such measures. These generalized measures of wage differentials are not solely determined by "the average wage of the average person" but depend on differences across complete wage distributions. Using a standard expected utility model, they incorporate judgments about differences in both level and spread (and higher moments), thereby adding dimensions of risk and inequality in the measurement of wage

[1] This paper is framed in terms of wage differentials by gender, but methods are directly applicable to differences across ethnic, religious, age groups, etc.

[2] Suppose in this example that everyone has the same productivity-related characteristics.

[3] One notable exception is Jenkins (1994). See also Gastwirth (1975), and, more recently, del Rìo et al. (2006) and Le Breton et al. (2008).

differentials. This is useful, for example, to examine if women are penalized twice, with both lower average wages and greater risk or inequality, or if, on the contrary, lower wages tend to be compensated by more favourable configurations of higher moments — a situation that could arguably be hypothesized as a partial explanation for the persistence of differences in mean wages.

The proposed summary indices take estimates of wage distributions conditional on human capital characteristics as input. An added contribution of this paper is to illustrate an original approach to estimate wage distributions flexibly in the presence of covariates and under endogenous labour market participation in a fully parametric framework which combines specification of a Singh-Maddala distribution with covariates (Biewen and Jenkins, 2005), a classic participation probability model, and a copula function to model the association between participation and wages (Smith, 2003).

Section 2 sets out the statistical framework in more details and briefly reviews classic measures of wage differentials and their critiques. Two new generalized measures are suggested in Section 3. Section 4 illustrates application of these measures in an assessment of gender wage differentials in Luxembourg. Section 5 concludes.

2 Current practice and critiques of the classic approach

2.1 Statistical framework and notation

Consider a population composed of agents of two (exogenously given) types –male or female– $s \in \{m, f\}$. Each agent is a potential labour market participant endowed with a vector of productivity-related characteristics $x \in \Xi$ reflecting, e.g., her stock of human capital. Ξ is the set of all possible combinations of characteristics. Wages of agents of type s and characteristics x are described by a random variable W_x^s with probability distribution $F_x^s : R^+ \mapsto [0, 1]$. The distribution of characteristics of type s agents in the population is described by a probability distribution $H^s : \Xi \mapsto [0, 1]$ with corresponding density function denoted $h^s(x)$.

Combining these definitions, the *unconditional* wage distribution of type s agents can be described by a random variable with probability distribution $F^s : R^+ \mapsto [0, 1]$ obtained by integrating the *conditional* wage distributions over observed characteristics among type s agents

$$F^s(w) = \int_{\Xi} F_x^s(w)h^s(x)dx. \tag{1}$$

Mean wages among type s agents are then given by

$$\mu^s \quad = \quad \int_0^{\infty} wf^s(w)dw \tag{2}$$

(where lower case f refers to the density function of F) or, equivalently,

$$\begin{aligned} \mu^s &= \int_\Xi \int_0^\infty w f_x^s(w) h^s(x) dw dx && (3) \\ &= \int_\Xi \mu_x^s h^s(x) dx. && (4) \end{aligned}$$

2.2 Common practice

A measure of wage differentials between agents of the two types takes as input (i) the conditional wage distributions for each observed characteristics and agent type and (ii) the distribution of characteristics of each agent type.

The simplest measure in use is the ratio of mean wages, which gives "the cents a woman makes for every dollar a man makes":

$$\Delta_0 = \frac{\mu^f}{\mu^m}. \quad (5)$$

The obvious problem with Δ_0 is that it compares mean wages of groups with different human capital, productivity-related characteristics. A more satisfactory way to pick up differences in wages rather than in human capital endowments is to compare wages in one group of interest –say women– with the counterfactual wages that would be observed in this same group if their human capital characteristics were rewarded as in the other, reference group. In the classic Oaxaca-Blinder approach, this is achieved by fitting log-wage regressions in each group and deriving a measure of differential based on counterfactual wage predictions when men's coefficients are used in women's equations. This leads to a wage differential measure such as

$$\begin{aligned} \Delta_{\mathrm{OB}} &= \exp\left[\int_\Xi x \left(\hat{\beta}^f - \hat{\beta}^m\right) h^f(x) dx\right] \hat{\theta} && (6) \\ &= \exp\left[\bar{x}^f \left(\hat{\beta}^f - \hat{\beta}^m\right)\right] \hat{\theta} && (7) \end{aligned}$$

where $\bar{x}^f$ is the vector of average characteristics in the female population and $\hat{\theta}$ is a factor reflecting differences in residual variance in the men and women earnings equations.[4] Δ_{OB} is interpreted as "the cents a woman makes on average for each dollar she would make if her human capital characteristics were rewarded as men's."[5] This is reflecting what is referred to as the 'unexplained' or 'discrimination' gap in Oaxaca-Blinder decompositions.

[4] Most studies mistakenly neglect the factor $\hat{\theta} = \exp\left[0.5\left(\hat{\Omega}^f - \hat{\Omega}^m\right)\right]$ with $\hat{\Omega}^s$ being the estimated residual variance in the earnings equation for group s. See Blackburn (2007) for a recent discussion.

[5] Notice that I take females as group of interest and males as reference group. The measures therefore pick up the disadvantage of women relative to men, what Jenkins (1994) calls 'discrimination against women'. See Oaxaca and Ransom (1994) for a discussion of this and alternative choices.

2.3 Critiques and alternative approaches

Despite its popularity, the regression-based approach has been criticized in several ways.

First, it rests on the validity of the log-linear specification of the wage equations (Barsky et al., 2002). Δ_{OB} is a special case of the more general measure

$$\Delta_1 = \exp\left[\int_{\Xi} [\log(\mu_x^f) - \log(\mu_x^m)]\ h_x^f dx\right]. \tag{8}$$

Δ_{OB} is equal to Δ_1 under the log-linear model specification for the conditional mean $\mu_x^m = \exp[x\hat{\beta}^m + 0.5\hat{\Omega}^m]$ and $\mu_x^f = \exp[x\hat{\beta}^f + 0.5\hat{\Omega}^f]$ where Ω^s is the homoscedastic residual variance in the earnings equation for group s. Misspecification of the earnings equations potentially translates in misleading estimates of Δ_1 by Δ_{OB}. Racine and Green (2004) address this issue by using non-parametric regression for estimating the conditional means. Ñopo (2008) adopts matching techniques. Barsky et al. (2002) follow Di Nardo et al. (1996) and use sample reweighting techniques to avoid specifying the conditional mean function altogether.

Second, as emphasized in Jenkins (1994), measures such as Δ_0 and Δ_1 (and therefore Δ_{OB}) are not 'distributionally sensitive' in the human capital dimension. It does not matter to Δ_0 or Δ_1 if the gap is approximately the same, say 10 percent, for women at all levels of human capital, or if the human capital of, say, half of women (e.g., the low-skilled, the migrants) is associated with a penalty of 20 percent while the other half face no wage penalty. In other words, Jenkins (1994) criticizes the blind averaging over all human capital characteristics Ξ in (7) or (8) and proposes alternative measures that address this issue.[6]

Third, Δ_0 and Δ_1 have been criticized, as in the Introduction, for putting narrow focus on means and discarding information about differences in higher moments in the wage distributions of men and women, that is for not being 'distributionally sensitive' in the wage dimension. Articulated critique of this kind dates back to Dolton and Makepeace (1985). However, it is only more recently, with the growing use of quantile regression and non-parametric density estimation that analysts have been able to examine differences in wage distributions in fine detail and estimate complete counterfactual wage distributions if women's human capital characteristics were rewarded as men's, namely

$$F^{(m,f)}(w) = \int_{\Xi} F_x^m(w) h^f(x) dx, \tag{9}$$

or its inverse, a counterfactual quantile function $Q^{(m,f)}(p)$ (Di Nardo et al., 1996, Fortin and Lemieux, 1998, Lemieux, 2002, Gardeazabal and Ugidos, 2005, Machado and Mata, 2005, Melly, 2005, Ñopo, 2008, Firpo et al., 2009). Applications of these methods have confirmed that wage distributions do not just vary in levels across gender; the

[6] See also del Rìo et al. (2006).

magnitude of differentials can also vary substantially at different quantiles. Restricting focus on mean wages and summarizing wage differentials with Δ_0 and Δ_1 is therefore not fully satisfactory. Yet, no alternative summary measure sensitive to finer distribution differences is in any common use. The measures considered in Section 3 attempt to fill this gap.

3 Two new generalized measures of wage differentials

Dolton and Makepeace (1985, pp.391–392) describe the 'ideal' way to measure wage differentials as follows: "In principle, the amount of sex discrimination should be deduced from a comparison of the distribution of earnings actually paid to females and the distribution when there is no discrimination. Ideally a utility function should be used to rank the distributions, but, unless the function is linear, higher moments of the earnings distributions than the mean will affect the choice between them." Consider measures of wage differentials that match this description.

In the stochastic model used thus far, the wage of an agent with characteristics x and type s can be interpreted as a realization from a lottery with payoff structure given by the (conditional) wage distribution F_x^s. As basic principle, consider that an agent of type s is prejudiced against if, given her characteristics, she would prefer to be paid a wage according to the other type's wage distribution. In other words, a woman is "discriminated" if she would prefer her wage to be drawn from men's lottery (F_x^m) rather than from women's lottery (F_x^f), given her characteristics.

Represent agents' preferences over lotteries by a utility functional $U : \Omega \mapsto R$ where Ω is the space of all wage distributions. $U(F)$ assigns a utility level to distribution F. The prejudice against a woman with characteristics x is then assessed by the difference between $U(F_x^m)$ and $U(F_x^f)$. Let us further assume that U can be expressed as a standard expected utility functional:

$$U(F) = \int_0^\infty u(w)dF(w).$$

Unless u is linear, preferences over wage distributions will be sensitive to the mean and to higher moments of F.

Two strategies are considered at this stage. The first makes only limited additional assumptions about u and relies on first-, second-, and third-order stochastic dominance criteria to derive an aggregate measure of wage differentials. The second proceeds with further assumptions about u.

3.1 Stochastic dominance

It is well-known that first-order stochastic dominance of distribution function F over G, that is $F(w) \leq G(w)$ for all w, is equivalent to

$$\int u(w)dF(w) \geq \int u(w)dG(w)$$

implying $U(F) \geq U(G)$, for every non-decreasing function u (see, e.g., Hadar and Russell, 1969). So, any woman with characteristics x would prefer being paid as a man whenever F_x^m first-order stochastically dominates F_x^f if her preferences can be represented with an expected utility functional with non-decreasing u, which merely reflects that she values receiving higher wages. Similarly, second-order stochastic dominance of the distribution function F over G, that is $\int_0^w F(s)ds \leq \int_0^w G(s)ds$ for all w, leads to $U(F) \geq U(G)$ for every non-decreasing, concave function u. A woman with characteristics x would prefer being paid as a man whenever F_x^m second-order stochastically dominates F_x^f if her preferences can be represented with an expected utility functional with nondecreasing and concave u, reflecting that she values receiving higher wages but that the marginal value of an extra dollar is decreasing with the wage level. Higher order dominance is associated with comparisons of utility with additional restrictions on higher order derivatives of u. Third-order stochastic dominance of F over G implies $U(F) \geq U(G)$ if u is nondecreasing ($u'(w) \geq 0$), concave ($u''(w) \leq 0$) and $u'''(w) \geq 0$, so that the marginal value of an extra dollar is decreasing at an increasing rate with the wage level. These results are well-known and are routinely used to rank income distributions (Saposnik, 1981, Davidson, 2008).

Stochastic dominance comparisons between F_x^m and F_x^f lead to a first way of identifying the wage prejudice of a woman with characteristics x compared to the reference distribution of males. Let $D(F_x^m, F_x^f; p)$ be a binary function that takes value one if F_x^m stochastically dominates F_x^f at the order p, and zero otherwise. A first set of summary measures of wage differentials is then given by

$$\Delta_2(p) = \int_\Xi D(F_x^m, F_x^f; p)\, h_x^f dx \qquad p \in \{1, 2, \ldots\}. \tag{10}$$

$\Delta_2(p)$ is the proportion of women in the population that would unambiguously prefer to be paid according to men's wage structure as captured by stochastic dominance at order p.[7] First-order stochastic dominance is unambiguous –it is hard to motivate a case for decreasing u. Assumptions about risk/inequality aversion associated with higher order may admittedly be more debatable (see *supra* for a measure allowing preference for risk).

[7]Note that the measure is not symmetric. That a proportion $\Delta_2(p)$ of women would prefer to be paid like men does not imply that a proportion $1 - \Delta_2(p)$ would prefer to be paid like women because dominance orderings are not complete.

Notice that it is wage distributions conditional on x that are compared, not aggregate, unconditional wage distributions. This allows for some types of women to be prejudiced, given their human capital endowment, and not others. For instance, this accommodates the observation that the wage gap tends to increase with age because of "glass ceilings" (Cobb-Clark, 2001) or career interruptions (Manning and Robinson, 2004). More fundamentally, this approach rests on the view that prejudice is first identified by comparing the wage distributions of observationally equivalent men and women. How the assessment of these differentials at given human capital endowments are then aggregated into an overall, summary measure when pooling agents with different levels of human capital is a distinct question. This is entirely consistent with Jenkins' (1994) discussion of the identification and aggregation issues in the measurement of wage differentials. Similarly, how conditional wage distribution differences combine to shape differences in the unconditional wage distribution across gender is not directly considered.[8]

3.2 Constant relative risk aversion

Stochastic dominance-based measures do not require many assumptions on the shape of the expected utility functional U. However comparisons of F_x^m and F_x^f are only ordinal. The resulting measures therefore tell us nothing about 'how much' distributions differ. To do so, the function u needs to be specified completely. A familiar specification is

$$u(w) = \begin{cases} \frac{w^{1-\epsilon}}{1-\epsilon} & \text{if } \epsilon \neq 1 \\ \ln(w) & \text{if } \epsilon = 1 \end{cases} \tag{11}$$

which implies constant relative risk aversion. This is the form of the well-known Atkinson social welfare function for income distribution analysis (Atkinson, 1970). The parameter ϵ determines the degree of risk (or inequality) aversion. To work with a monetary metric, it is convenient to determine the 'certainty equivalent' wage, that is the amount $C(F)$ that, if received with certainty, would lead to the same utility as the uncertain outcome described by F. $C(F)$ is defined implicitly as the solution of $U(F) = U(\tilde{F})$ where $\tilde{F}$ is a distribution concentrated on a point-mass at the value $C(F)$. With the chosen specification of u, the certainty equivalent is

$$C(F;\epsilon) = \left(\int_0^\infty w^{1-\epsilon} dF(w) \right)^{\frac{1}{1-\epsilon}}$$

for $\epsilon \neq 1$ and $C(F;1) = \exp\left[\int_0^\infty \ln(w) dF(w)\right]$. $C(F;\epsilon)$ is in a money metric and can be conveniently interpreted in dollar terms. $C(F;\epsilon)$ is homogeneous of degree one;

[8] This discussion is empirically irrelevant with classic measures that compare average wages such as Δ_0 or Δ_1 since the unconditional average is equal to the average of conditional averages because of the implicit linearity in u in these measures. Note also that recent papers that have looked at gender differentials 'at quantiles' have been concerned with the *unconditional* wage distributions of men and women directly (see, e.g., Gardeazabal and Ugidos, 2005, Firpo et al., 2009).

multiplying all wages by a constant increases $C(F;\epsilon)$ by the same constant. As is well-known, for $\epsilon = 0$ there is no risk aversion in U and $C(F;0)$ is equal to the expected wage μ. Increasing ϵ leads to $C(F;\epsilon) < \mu$: risk aversion makes people ready to accept lower expected wages for less uncertainty and lower dispersion in the wage distribution, there is a trade-off. On the contrary, $\epsilon < 0$ represents preference for risk or inequality. Greater risk is perceived positively, so $C(F;\epsilon) > \mu$; people are ready to pay a premium for facing the uncertain outcome rather than the certain outcome.

Summarizing wage differentials by comparing certainty equivalent wages, rather than mean wages, is a straightforward alternative to the classic Oaxaca-Blinder, regression-based approach, and one that allows taking differences in higher moments into account. Additionally, $C(F;\epsilon)$ has the form of a "general mean" which has been recently characterized as an appealing income standard *per se*, independently on any reference to utility functions by Foster and Székely (2008). This suggests a family of 'distributionally sensitive' measures of wage differential (indexed by ϵ), directly comparable to Δ_1,

$$\Delta_3(\epsilon) = \exp\left[\int_\Xi \left[\log(C(F_x^f;\epsilon)) - \log(C(F_x^m;\epsilon))\right] h_x^f dx\right]. \tag{12}$$

The interpretation of $\Delta_3(\epsilon)$ is now as the "certainty-equivalent cents a woman makes for every certainty-equivalent dollar an observationally equivalent men makes." Notice also that $\Delta_3(0) = \Delta_1$. Ideally a distinct ϵ should be assigned to each women to reflect individual attitudes to risk and inequality. In the absence of such information, the second-best approach is to select one ϵ representative of attitudes among women. There is however little guidance about setting one particular, representative value for ϵ.[9] The best strategy is therefore to compute $\Delta_3(\epsilon)$ over a range of ϵ where increasing ϵ reflects greater risk aversion, thereby giving greater importance to differences in the conditional wage distributions at low wage levels.

4 A re-assessment of the gender pay gap in Luxembourg

I illustrate use of $\Delta_2(p)$ and $\Delta_3(\epsilon)$ in an assessment of the gender pay gap in Luxembourg. STATEC, the national statistical institute of Luxembourg, recently reported that the average monthly gross wage is 20 percent lower for women than for men, and that half of this gap can be explained by differences in human capital and job characteristics (STATEC, 2007). The report also pointed out that unemployment risk is twice as high for women and that their participation rate is much lower (55 percent against 73 percent for the age group 15-64). These observations are similar to those found elsewhere in Europe despite peculiarities of the labour market in Luxembourg (e.g., high

[9] A body of evidence has shown that women tend to be more risk averse than men in general (Hartog et al., 2002, Agnew et al., 2008), but that is not helpful here since only attitudes towards risk among women matter, not men's attitudes.

fraction of immigrant and cross-border workers, prevalence of banking industry and comparatively high wage rates). Using data from the *Panel Socio-Economique Liewen zu Lëtzebuerg* (PSELL-3/EU-SILC) survey, I check how much the wage gap appears aggravated or lessened once generalized measures of wage differentials are considered.

4.1 Estimation

The measures of wage differential considered in this paper require estimates of conditional wage distributions F_x^m and F_x^f for all x observed in the sample. Flexible estimators are desirable to avoid restrictive parametric assumptions on the shape of these functions. This complicates estimation compared to the classic Oaxaca-Blinder approach. Several non-parametric or semi-parametric estimators can be chosen from; see, *inter alia*, Hall et al. (1999), Donald et al. (2000), Peracchi (2002). More standard linear quantile regression models estimated on a fine grid of points and inverted where necessary as in Machado and Mata (2005) or Melly (2005) form a simple alternative option. In this application, I opt instead for a fully parametric, yet flexible, approach which allows estimation of the distribution functions under endogenous labour market participation, unlike most semi- or non-parametric approaches.[10] Conditional wage distributions are assumed to follow a Singh-Maddala distribution with each of the three parameters allowed to vary log-linearly with a set of covariates:

$$F_x^s(w) = 1 - \left[1 + \left(\frac{w}{b^s(x)}\right)^{a^s(x)}\right]^{-q^s(x)} \tag{13}$$

where $b^s(x) = \exp(x\theta_b^s)$ is a scale parameter, $q^s(x) = \exp(x\theta_q^s)$ is a shape parameter affecting the right tail, and $a^s(x) = \exp(x\theta_a^s)$ is a shape parameter affecting both tails (Singh and Maddala, 1976, Kleiber and Kotz, 2003). This specification is discussed in Biewen and Jenkins (2005) to model income distributions in the presence of covariates. While this is a fully parametric specification, the Singh-Maddala distribution is a flexible model for unimodal distributions allowing varying degrees of skewness and kurtosis and dealing with the heavy tails typical of income and earnings distributions.

A parametric specification permits to consider the effect of endogenous labour market participation decision on the estimation of the wage distribution functions. Let z denote a binary employment indicator for a given agent. Her wage w is only observed if $z = 1$. Consider z^* to be a continuous, latent propensity to participate in the labour market with $z = 1$ if $z^* > 0$ and $z = 0$ otherwise. Assume the pair (w, z^*) is jointly

[10] Buchinsky's (1998) selectivity-corrected quantile regression model is the only alternative that I am aware of for dealing with self-selection in a semi-parametric setting. See Albrecht et al. (2009) for an application to gender pay gap analysis. Identification of the constant term –which is crucial in the context of this paper– is however difficult in these models. This is avoided here, albeit at the cost of imposing stronger parametric assumptions. See also Huber (2009) for a recent critique of the identification assumptions of the sample selection quantile regression model.

distributed conditionally on human capital characteristics with cumulative distribution H_x and express H_x using its copula and the marginal distributions for w (namely F_x, assumed to be Singh-Maddala distributed) and z^* (denoted G_x):

$$H_x(w, z^*) = C(F_x^{-1}(w), G_x^{-1}(z^*)). \tag{14}$$

(Superscript s for agent type is dropped for exposition clarity. All model parameters are allowed to vary by agent type.) I make the standard parametric assumption that G_x is normal with mean $x\delta$ and unit variance and take C to be a Clayton copula.[11] This copula-based approach to specifying a parametric self-selection model is detailed in Smith (2003); see also Trivedi and Zimmer (2007). Smith (2003) demonstrates how standard maximum likelihood methods can be used to estimate this fully parametric model which combines the standard participation model assumptions with a flexible three-parameters Singh-Maddala specification for the wage distribution.[12]

It is straightforward to compute the proposed generalized measures of wage differentials once maximum likelihood estimates $(\hat{\theta}_a^f, \hat{\theta}_b^f, \hat{\theta}_q^f)$ and $(\hat{\theta}_a^m, \hat{\theta}_b^m, \hat{\theta}_q^m)$ are available. $\hat{F}_x^s$ is estimated by plugging $(\hat{\theta}_a^s, \hat{\theta}_b^s, \hat{\theta}_q^s)$ in (13) and the formula for the general means of a Singh-Maddala distribution can be used to compute $C(\hat{F}_x^s; \epsilon)$ (Kleiber and Kotz, 2003):

$$C(\hat{F}_x^s; \epsilon) = \hat{b}^s(x) \left(\frac{\Gamma\left(1 + (1-\epsilon)/\hat{a}^s(x)\right) \Gamma\left(\hat{q}^s(x) - (1-\epsilon)/\hat{a}^s(x)\right)}{\Gamma\left(\hat{q}^s(x)\right)} \right)^{\frac{1}{1-\epsilon}} \tag{15}$$

where $\Gamma(\cdot)$ is the Gamma function. These estimates are then plugged into (10) and (12) where average is taken over all females in the sample (in place of integration over h_x^f). While the delta method could be used to estimate the sampling variability the final estimates from the maximum likelihood-based estimate of the covariance matrix of $(\hat{\theta}_a^f, \hat{\theta}_b^f, \hat{\theta}_q^f)$ and $(\hat{\theta}_a^m, \hat{\theta}_b^m, \hat{\theta}_q^m)$, I estimate all standard errors using the repeated half-sample bootstrap of Saigo et al. (2001) that allows me to take into account complex survey design features with some small stratum sizes.

4.2 Data

The analysis uses data from the *Panel Socio-Economique Liewen zu Lëtzebuerg* (PSELL-3/EU-SILC). PSELL-3/EU-SILC is a general purpose panel survey carried out annually since 2003. More than 3,500 private households residing in Luxembourg are surveyed and all adult members of sampled households are interviewed. The questionnaire covers

[11] The Clayton copula is $C(u, v; \theta) = \left(u^{-\theta} + v^{-\theta} - 1\right)^{-1/\theta}$ where θ is an association parameter to be estimated. This specification provided a better fit to the data than the Frank copula that was also considered.

[12] Maximum likelihood estimations for this paper were all done using the standard built-in Newton-Raphson optimizer of Stata$^{\text{TM}}$(StataCorp, 2007).

topics such as income and living conditions, employment, education, health. Covering both employed and non-employed respondents, PSELL-3/EU-SILC makes it possible to account for differential labour market participation between men and women.

Sample data from 2003 to 2007 are pooled in this analysis.[13] Only 25- to 55-year-old respondents are considered to avoid issues related to gender differences in retirement and labour market entry. I consider differences in gross hourly wages of full-time workers (private and public sectors pooled).[14] Gross hourly wage is computed as gross monthly salary in current job (including paid overtime) divided by 4.32 times work hours in a normal week on the job. Wages are expressed in constant January 2007 prices. In the vector x of human capital characteristics, I take age, educational attainment, nationality and actual work experience into account.[15] I do not control for tenure or contract and job characteristics. This parsimonious specification avoids including variables that are largely determined by gender and, possibly, by wage differentials themselves as advised, e.g., in Neal and Johnson (1996). It leads to measures of wage differentials that capture the total effect of gender on wages rather than its partial effect after controlling for job-related confounding variables.

4.3 Results

To fix ideas, Figure 1 shows estimates of the conditional wage distributions f_x^m and f_x^f for a subset of combinations of human capital characteristics observed in the data. All densities follow Singh-Mandala distributions with estimated parameters $(a^s(x), b^s(x), q^s(x))$. Unsurprizingly, male distributions (solid lines) are generally more densely concentrated towards higher wages than female distributions (dashed lines). There are also noticeable

[13]The bootstrap resampling algorithm takes into account the repetition of individuals in the pooled dataset, the correlation of individuals from the same household, as well as the correlation of respondents from different household but related to the same original household at wave 1.

[14]Self-employed and international civil servants are excluded from the sample, as well as workers with hourly wages below €3 or above €60. The final sample includes 9,168 observations for men, among which 7,919 are full-time workers with mean wage of €21, and 10,015 observations for women among which 3,543 are full-time workers with mean wage of €18.

[15]Specifically, equation for $b^s(x)$ includes three age group dummies (25–34, 35–44, 45–55) fully interacted with four dummies for educational attainment (lower secondary or below, upper secondary, BAC+2/3, and BAC+4 or above) and four nationality groups (Luxembourg nationals, Portuguese, other Europeans, non-Europeans). Equation for $a^s(x)$ includes age group and nationality dummies. Equation for $q^s(x)$ includes work experience. In selecting this model, I considered the convergence and stability of the model across bootstrap replications, the precision and significance of estimated coefficients as well as the Akaike information criterion. While it is standard in copula-based models of self-selection to identify all parameters through functional form assumptions and not by exclusion restrictions (Smith, 2003), in order to strengthen identification of the model, I have augmented the participation equation (which includes all covariates used in the Singh-Maddala equations) with a dummy variable indicating whether the sample observation filled the questionnaire directly or whether it was filled by another household member ('proxy interview'). Absence from home at the visit of the interviewer is correlated with employment but assumed independent on wage.

differences in the variance of the distribution and female distributions tend to be more right skewed. Taking selection into account (gray lines) shifts distributions towards lower wages.

Figure 1. Examples of conditional probability distribution function estimates f_x^m (solid lines, men) and f_x^f (dashed lines, women) for illustrative combinations of education level, age and nationality (all at 6–10 years of work experience).

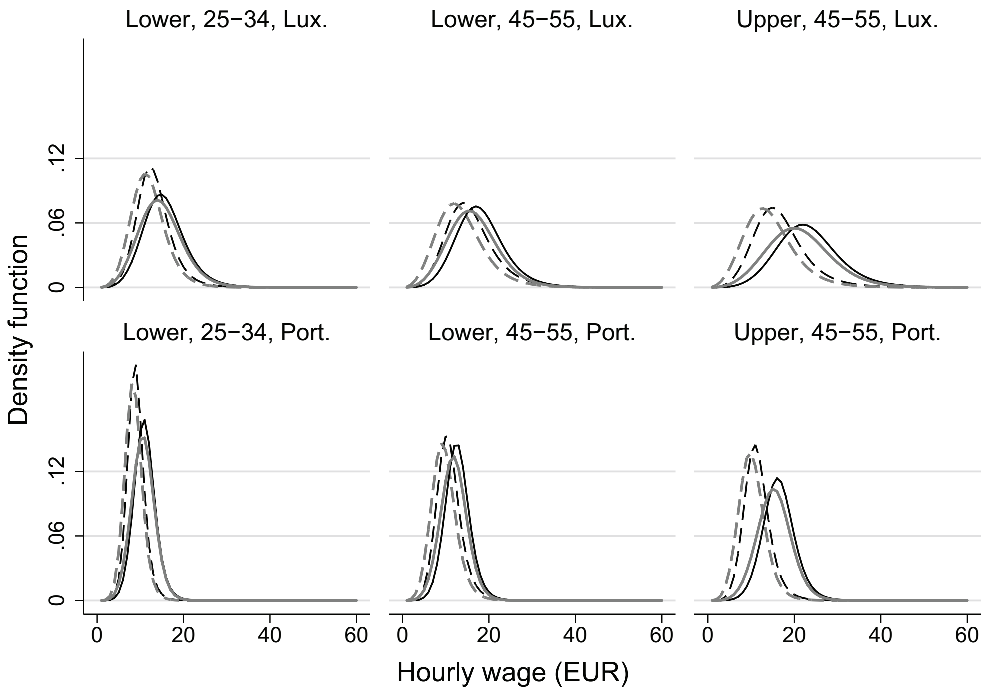

Notes: Black lines are estimates unadjusted for labour market participation, gray lines are adjusted for self-selection. Examples are combinations of lower secondary (or less) and tertiary education, 25–34 and 45–55 age groups and Luxembourg vs. Portuguese nationality.

Table 1 reports estimates of $\Delta_2(p)$ for $p = 1$ (first-order dominance) and $p = 2$ (second-order dominance). Disregarding endogenous selection issues, 10 percent of all women in Luxembourg face a wage distribution dominated at the first order by the wage distribution of observationally equivalent men. The proportion increases up to 38 percent at the second order. These proportions increase to 25 and 63 percent once endogenous participation is taken into account. They also vary substantially by population subgroups. They are substantially smaller for younger women (25–34) and those with higher, tertiary-level education. The sampling variability of these estimates is however large, in fact extremely large for the models that account for sample selection.

Estimates of $\Delta_3(\epsilon)$ for ϵ in the range –4 to 4 are reported in Figure 2. Estimates for $\epsilon < 0$ assume preference for risk/inequality, thereby rewarding a more dispersed distribution. On the contrary, $\epsilon > 0$ assume risk/inequality aversion thereby penalizing

Table 1. Dominance-based measures of wage differentials, $\Delta_2(p)$

	No selection correction		With selection correction	
	$p = 1$	$p = 2$	$p = 1$	$p = 2$
All	0.10	0.38	0.25	0.63
	[0.03,0.46]	[0.18,0.69]	[0.03,0.87]	[0.03,1.00]
Luxembourg	0.10	0.35	0.37	0.70
	[0.03,0.56]	[0.16,0.83]	[0.03,0.90]	[0.03,1.00]
Portuguese	0.03	0.53	0.02	0.80
	[0.02,0.53]	[0.02,0.99]	[0.01,0.96]	[0.02,1.00]
Other EU	0.19	0.32	0.18	0.33
	[0.03,0.49]	[0.03,0.67]	[0.03,1.00]	[0.03,1.00]
25–34	0.01	0.01	0.01	0.16
	[0.01,0.35]	[0.01,0.56]	[0.01,0.85]	[0.01,1.00]
35–44	0.01	0.18	0.32	0.68
	[0.00,0.60]	[0.00,0.85]	[0.00,0.98]	[0.00,1.00]
45–55	0.24	0.92	0.35	0.99
	[0.01,0.62]	[0.26,0.99]	[0.00,0.87]	[0.01,1.00]
Lower secondary ed.	0.04	0.57	0.08	0.86
	[0.04,0.54]	[0.14,0.89]	[0.04,0.90]	[0.04,1.00]
Secondary ed.	0.20	0.35	0.55	0.66
	[0.02,0.63]	[0.23,0.80]	[0.02,0.97]	[0.02,1.00]
Tertiary ed.	0.04	0.10	0.04	0.19
	[0.02,0.22]	[0.02,0.36]	[0.02,0.74]	[0.02,1.00]

Notes: Figures in brackets are 90% percentile boostrap confidence intervals based on 999 repeated half-sample bootstrap replications.

a more dispersed distribution. The dividing case of $\epsilon = 0$ is risk/inequality neutrality and corresponds to the classic measure Δ_1 focusing on conditional means only (see (8)). In this latter case, the measure evaluates to 0.86 in this application (0.82 if self-selection is taken into account) suggesting that women suffer a 14 percent (18 percent) 'unexplained' mean wage disadvantage.

The clear message from Figure 2 is that there is little impact of considering the generalized measure for small and positive ϵ; women's disadvantage subsists. The gap is substantially reduced (and becomes statistically insignificant) only for large negative ϵ. According to these results, the argument that women's lower average wages are compensated by differences in higher order moments would only hold if women have substantial taste for risk or inequality reflected in large, negative ϵ. Taking endogenous selection into account leads to a larger gap, but to a moderate degree.

A smaller set of results for specific population subgroups are reported in Table 2. Observations from Table 1 are largely confirmed. The wage gap for women with higher, tertiary level education is lower than for other women and does not appear to be statistically significant (i.e., $\Delta_3(\epsilon)$ is not significantly different from 1). That holds true for all levels of ϵ. Younger women also tend to face a smaller gap but that is reversed for large, negative ϵ. The pay gap is also slightly smaller for women of Luxembourg

Figure 2. Certainty-equivalent-based indices, $\Delta_3(\epsilon)$, for $\epsilon \in [-4, 4]$.

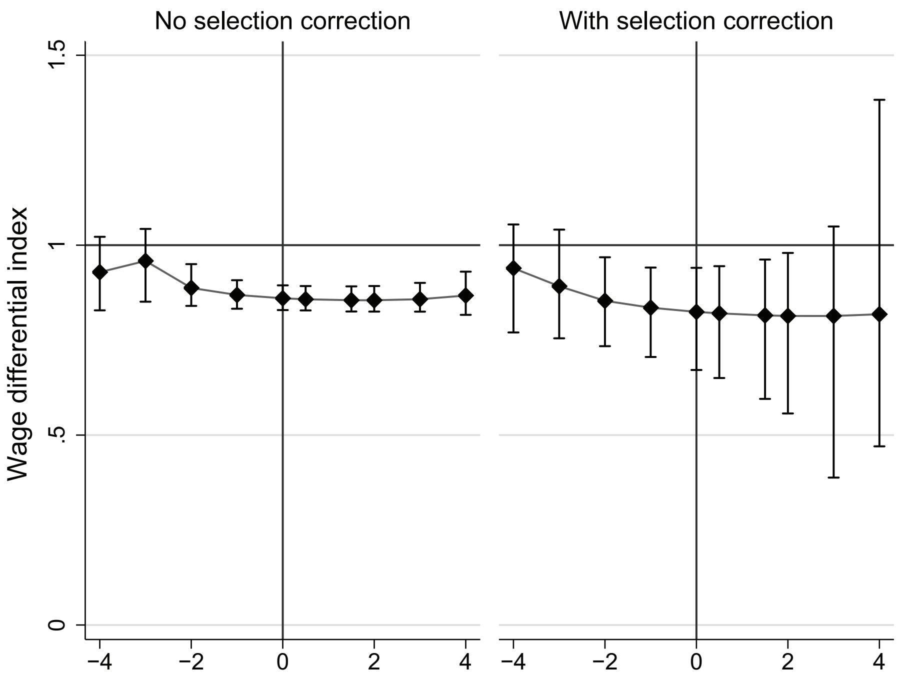

Note: Vertical capped bars show 90 percent percentile-based bootstrap variability bands (based on 999 repeated half-sample bootstrap replications).

nationality.

Table 2. Certainty-equivalent-based measures of wage differentials, $\Delta_3(\epsilon)$

	No selection correction			With selection correction		
	$\epsilon=-3$	$\epsilon=0$	$\epsilon=3$	$\epsilon=-3$	$\epsilon=0$	$\epsilon=3$
All	0.96	0.86	0.86	0.89	0.82	0.81
	[0.85,1.04]	[0.83,0.89]	[0.82,0.90]	[0.75,1.04]	[0.67,0.94]	[0.39,1.05]
Luxembourg	1.02	0.88	0.86	0.93	0.83	0.81
	[0.87,1.14]	[0.83,0.92]	[0.81,0.93]	[0.76,1.13]	[0.66,0.97]	[0.34,1.08]
Portuguese	0.86	0.83	0.83	0.83	0.80	0.80
	[0.78,0.94]	[0.79,0.88]	[0.79,0.87]	[0.74,0.92]	[0.70,0.89]	[0.59,0.93]
Other EU	0.87	0.82	0.85	0.83	0.79	0.82
	[0.75,1.00]	[0.76,0.87]	[0.78,0.92]	[0.69,1.01]	[0.65,0.91]	[0.42,1.07]
25–34	0.91	0.90	0.93	0.88	0.88	0.91
	[0.86,0.99]	[0.88,0.94]	[0.89,0.98]	[0.81,0.98]	[0.76,0.96]	[0.56,1.09]
35–44	0.88	0.84	0.86	0.84	0.80	0.82
	[0.81,1.02]	[0.80,0.89]	[0.81,0.92]	[0.72,1.00]	[0.65,0.92]	[0.42,1.06]
45–55	1.13	0.84	0.79	0.97	0.80	0.73
	[0.84,1.30]	[0.77,0.91]	[0.73,0.87]	[0.71,1.28]	[0.61,0.97]	[0.26,1.04]
Lower secondary ed.	0.95	0.85	0.83	0.89	0.82	0.79
	[0.82,1.04]	[0.79,0.90]	[0.78,0.89]	[0.73,1.04]	[0.65,0.94]	[0.38,1.02]
Secondary ed.	0.92	0.81	0.81	0.85	0.77	0.76
	[0.81,1.02]	[0.76,0.87]	[0.76,0.87]	[0.70,1.02]	[0.62,0.91]	[0.35,1.00]
Tertiary ed.	1.04	0.97	0.99	0.98	0.94	0.95
	[0.94,1.13]	[0.93,1.02]	[0.94,1.06]	[0.87,1.13]	[0.79,1.05]	[0.49,1.20]

Notes: Figures in brackets are 90% percentile boostrap confidence intervals based on 999 repeated half-sample bootstrap replications.

5 Summary and concluding remarks

This paper is about computing summary measures of wage differential that capture potentially complex distribution differences between men and women (or any two population groups). It is indeed nowadays largely accepted that there is interest in going beyond comparisons "at the mean" and researchers have now described distribution differences between men and women in a lot more details than in earlier decades. However, no 'distributionally sensitive' index measure has correspondingly been in frequent use. I consider two sets of simple generalized measures of wage differentials transplanting concepts familiar to income distribution analysts. The measures take into account higher-order differences in conditional wage distributions rather than just mean differences using a simple expected utility framework.

Note that this is a different, complementary perspective from recent examinations of the 'glass ceiling' as in, e.g., Albrecht et al. (2003) or Arulampalam et al. (2007). These studies focus on aggregate, unconditional distribution differences whereas the key justification for the measures adopted here is that women with given human capital endowments evaluate their wage disadvantage from the wage distributions of men with similar human capital endowments, that is, by comparing *conditional* wage distributions.

The methods adopted in this paper are also closely complementary to the distributionally sensitive approach of Jenkins (1994) who is primarily concerned with the aggregation of individual-level discrimination across people of different levels of human capital (and other characteristics). Discrimination for a woman with given human capital is however simply captured by differences in expected wages in a log-linear regression framework. This is the direct complement to the approach followed here since I consider a richer identification of differentials, but my aggregation across human capital levels is agnostic about distributional issues in the human capital dimension. Combining the two approaches is an avenue for future research.

Empirical results from survey data for Luxembourg seem to invalidate the hypothesis that lower wages for women are in fact 'compensated' by lower risk/uncertainty in pay (an hypothesis that might be advanced to explain the persistence of wage gaps in mean wages). They would rather suggest that women are penalized twice: wages are lower on average *and* their distributions tend to be more dispersed. One would therefore require substantial taste for inequality for holding to the compensation argument. Results also suggest that taking self-selection into account is relevant and increases estimates of the gender pay gap, but by a moderate amount. There is however interest in confronting these results to estimates from other labour markets.

References

Agnew, J. R., Anderson, L. R., Gerlach, J. R. and Szykman, L. R. (2008), 'Who chooses annuities? An experimental investigation of the role of gender, framing, and defaults', *American Economic Review* **98**(2), 418–22.

Albrecht, J., Björklund, A. and Vroman, S. (2003), 'Is there a glass ceiling in Sweden?', *Journal of Labor Economics* **21**(1), 145–177.

Albrecht, J., van Vuuren, A. and Vroman, S. (2009), 'Counterfactual distributions with sample selection adjustments: Econometric theory and an application to the Netherlands', *Labour Economics* **16**, 383–396.

Arulampalam, W., Booth, A. L. and Bryan, M. L. (2007), 'Is there a glass ceiling over Europe? Exploring the gender pay gap across the wage distribution', *Industrial and Labor Relations Review* **60**(2), 163–186.

Atkinson, A. B. (1970), 'On the measurement of inequality', *Journal of Economic Theory* **2**, 244–263.

Barsky, R., Bound, J., Charles, K. K. and Lupton, J. P. (2002), 'Accounting for the black-white wealth gap: A nonparametric approach', *Journal of the American Statistical Association* **97**(459), 663–673.

Biewen, M. and Jenkins, S. P. (2005), 'Accounting for differences in poverty between the USA, Britain and Germany', *Empirical Economics* **30**(2), 331–358.

Blackburn, M. L. (2007), 'Estimating wage differentials without logarithms', *Labour Economics* **14**(1), 73–98.

Blinder, A. (1973), 'Wage discrimination: Reduced form and structural estimates', *Journal of Human Resources* **8**, 436–455.

Buchinsky, M. (1998), 'The dynamics of changes in the female wage distribution in the USA: A quantile regression approach', *Journal of Applied Econometrics* **13**(1), 1–30.

Cobb-Clark, D. (2001), 'Getting ahead: The determinants of and payoffs to internal promotion for young men and women', *Research in Labor Economics* **20**, 339–372.

Davidson, R. (2008), Stochastic dominance, *in* S. N. Durlauf and L. E. Blume, eds, 'The New Palgrave Dictionary of Economics', Palgrave Macmillan, Basingstoke.

de la Rica, S., Dolado, J. and Llorens, V. (2008), 'Ceilings or floors? Gender wage gaps by education in Spain', *Journal of Population Economics* **21**(3), 751–776.

del Rìo, C., Gradin, C. and Cantò, O. (2006), The measurement of gender wage discrimination: The distributional approach revisited, ECINEQ Working Paper 25, Society for the Study of Economic Inequality.

Di Nardo, J., Fortin, N. M. and Lemieux, T. (1996), 'Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach', *Econometrica* **64**(5), 1001–1044.

Dolton, P. J. and Makepeace, G. (1985), 'The statistical measurement of discrimination', *Economics Letters* **18**, 391–5.

Donald, S. G., Green, D. A. and Paarsch, H. J. (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**(4), 609–633.

Firpo, S., Fortin, N. M. and Lemieux, T. (2009), 'Unconditional quantile regressions', *Econometrica* **77**(3), 953–973.

Fortin, N. M. and Lemieux, T. (1998), 'Rank regressions, wage distributions, and the gender gap', *Journal of Human Resources* **33**(3), 610–643.

Foster, J. E. and Székely, M. (2008), 'Is economic growth good for the poor? Tracking low incomes using general means', *International Economic Review* **49**(4), 1143–1172.

Gardeazabal, J. and Ugidos, A. (2005), 'Gender wage discrimination at quantiles', *Journal of Population Economics* **18**(1), 165–179.

Gastwirth, J. L. (1975), 'Statistical measures of earnings differentials', *The American Statistician* **29**, 32–35.

Hadar, J. and Russell, W. R. (1969), 'Rules for ordering uncertain prospects', *American Economic Review* **59**(1), 97–122.

Hall, P., Wolff, R. C. L. and Yao, Q. (1999), 'Methods for estimating a conditional distribution function', *Journal of the American Statistical Association* **94**(445), 154–63.

Hartog, J., Ferrer-i-Carbonell, A. and Jonker, N. (2002), 'Linking measured risk aversion to individual characteristics', *Kyklos* **55**(1), 3–26.

Huber, M. (2009), Treatment evaluation in the presence of sample selection, Working Paper 2009-07, Department of Economics, University of St. Gallen.

Jenkins, S. P. (1994), 'Earnings discrimination measurement: A distributional approach', *Journal of Econometrics* **61**, 81–102.

Kleiber, C. and Kotz, S. (2003), *Statistical size distributions in Economics and Actuarial Sciences*, Wiley series in probability and statistics, John Wiley and Sons, Inc.

Le Breton, M., Michelangeli, A. and Peluso, E. (2008), Wage discrimination measurement: In defense of a simple but informative statistical tool, IDEI Working Papers 516, Institut d'Économie Industrielle, Toulouse.

Lemieux, T. (2002), 'Decomposing changes in wage distributions: a unified approach', *Canadian Journal of Economics* **35**(4), 646–688.

Machado, J. A. F. and Mata, J. (2005), 'Counterfactual decomposition of changes in wage distributions using quantile regression', *Journal of Applied Econometrics* **20**, 445–465.

Manning, A. and Robinson, H. (2004), 'Something in the way she moves: A fresh look at an old gap', *Oxford Economic Papers* **56**(2), 169–188.

Melly, B. (2005), 'Decomposition of differences in distribution using quantile regression', *Labour Economics* **12**, 577–90.

Neal, D. A. and Johnson, W. R. (1996), 'The role of premarket factors in black-white wage differences', *The Journal of Political Economy* **104**(5), 869–895.

Ñopo, H. R. (2008), 'Matching as a tool to decompose wage gaps', *The Review of Economics and Statistics* **90**(2), 290–299.

Oaxaca, R. L. (1973), 'Male-female wage differentials in urban labor markets', *International Economic Review* **14**, 673–709.

Oaxaca, R. L. and Ransom, M. R. (1994), 'On discrimination and the decompostion of wage differentials', *Journal of Econometrics* **61**, 5–21.

Peracchi, F. (2002), 'On estimating conditional quantiles and distribution functions', *Computational Statistics and Data Analysis* **38**(4), 433–447.

Racine, J. S. and Green, C. A. (2004), Robust estimation of union and gender wage gaps, Working Paper COBA-ECON-04-01, College of Business Administration, University of South Florida.

Saigo, H., Shao, J. and Sitter, R. R. (2001), 'A repeated half-sample bootstrap and balanced repeated replications for randomly imputed data', *Survey Methodology* **27**(2), 189–196.

Saposnik, R. (1981), 'Rank-dominance in income distributions', *Public Choice* **36**(1), 147–151.

Singh, S. K. and Maddala, G. S. (1976), 'A function for size distribution of incomes', *Econometrica* **44**(5), 963–970.

Smith, M. D. (2003), 'Modelling sample selection using Archimedean copulas', *Econometrics Journal* **6**(1), 99–123.

StataCorp (2007), *Stata Statistical Software: Release 10*, StataCorp LP, College Station.

STATEC (2007), Egalité hommes-femmes, mythe ou réalité?, Cahier économique 105, STATEC, Luxembourg.

Trivedi, P. K. and Zimmer, D. M. (2007), 'Copula modeling: An introduction for practitioners', *Foundations and Trends in Econometrics* **1**(1), 1–111.