Empir Econ (2013) 45:465–482
DOI 10.1007/s00181-012-0608-y

Generalized measures of wage differentials

Philippe Van Kerm

Received: 6 September 2010 / Accepted: 30 March 2012 / Published online: 6 June 2012

Abstract This paper considers measures of wage differentials not solely determined by mean comparisons but summarizing differences across complete wage distributions. The approach builds on considerations of risk or inequality aversion and on standard expected utility concepts. In an application to the gender pay gap in Luxembourg the disadvantage of women persists according to the proposed measures: lower mean wages for women are not compensated by differences in higher moments of wage distributions (e.g., by less dispersion) at least for realistic assumptions about women preferences toward risk and inequality. The paper also illustrates an original empirical model for wage distributions in the presence of covariates and under endogenous labour market participation.

Keywords Wage differentials · Expected utility · Singh–Maddala distribution · Endogenous selection · Luxembourg

JEL Classification D63 · J31 · J70

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 often 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

P. Van Kerm (✉)
CEPS/INSTEAD, 3 Avenue de la fonte, 4364 Esch-sur-Alzette, Luxembourg
e-mail: philippe.vankerm@ceps.lu

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, mean differences in pay to be 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 adequately captured by measures of wage differentials "at the mean." With 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. Policy action aiming at reducing wage differentials would likely benefit from monitoring tools able to pick comprehensive distribution differences.

Adopting a standard expected utility framework, this paper considers two (families of) "generalized" measures of wage differentials that capture differences across complete wage distributions. These simple summary measures incorporate explicit sensitivity to dimensions of risk and inequality in the assessment of wage differentials, placing value to differences in higher moments. The indices are natural extensions of the ubiquitous Oaxaca–Blinder measures of "unexplained" wage 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 favorable configurations of higher moments—a situation that could 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 additional contribution of this paper is to illustrate estimation of a parametric but flexible model for wage distributions in the presence of covariates and under endogenous labour market participation. It 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). Its key advantage is consistency in the presence of heterogeneity where popular selectivity-corrected quantile regression with control functions fails (Huber and Melly 2011).

Section 2 sets out the statistical framework formally and briefly reviews classic measures of wage differentials. Generalized measures are described in Sect. 3. Section 4 illustrates application of these measures in an assessment of gender wage differentials in Luxembourg. Section 5 concludes.

2 Measuring wage differentials

Consider a population 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 $F_x^s : R^+ \mapsto [0, 1]$ and corresponding density function f_x^s. 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 $h^s(x)$.

In essence, all measures of wage differentials between agents of the two types take these distributions as input, namely (i) the conditional distribution of wage for any value taken by the characteristics x and (ii) the distribution of characteristics within each agent type. A common index is the (geometric) mean of the ratio of female-to-male expected wages, where expected wages are taken conditionally on x (see, e.g., Jenkins 1994):

$$\Delta_1 = \exp\left[\int_{\Xi}\left[\log\left(\frac{\mu_x^f}{\mu_x^m}\right)\right] h^f(x)\mathrm{d}x\right] \tag{1}$$

($\mu_x^s = \int x dF_x^s$). Δ_1 gives “the cents a woman makes on average, for each dollar she would make if her human capital characteristics were rewarded as men’s.” This is based on a comparison of wages in one group of interest –women, here– 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. Δ_1 is often referred to as the “unexplained” wage gap and as reflecting potential discrimination.[1] In practice, Δ_1 is typically computed by fitting log-wage regressions in samples for each group

$$\log(w_i) = x_i\beta^s + u_i \qquad i = 1, \ldots, N^s \qquad s \in \{m, f\} \tag{2}$$

and using coefficients to derive

$$\Delta_{\mathrm{OB}} = \exp\left[\int_{\Xi} x\left(\beta^f - \beta^m\right) h^f(x)dx\right]\theta \tag{3}$$

$$= \exp\left[\bar{x}^f\left(\beta^f - \beta^m\right)\right]\theta \tag{4}$$

where $\bar{x}^f$ is the vector of average characteristics in the female population. (θ is a factor reflecting differences in residual distributions in the earnings equations for men and women (see Appendix). It is often neglected in practice.) Δ_{OB} and Δ_1 are equivalent provided the log-linear specification of the wage equations is valid.[2]

[1] Notice that I take females as a group of interest and males as a 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.

[2] See Barsky et al. (2002) for a discussion of the potential effect of mis-specification of the log-linear regression on wage differentials estimates. Racine and Green (2004) address this issue using non-parametric regression for estimating the conditional means. Nopo (2008) adopts matching techniques.

Despite its popularity, the regression-based approach can be criticized for not being "distributionally sensitive." First, as emphasized in Jenkins (1994), measures such as Δ_1 (and therefore Δ_{OB}) are not "distributionally sensitive" in the human capital dimension. It does not matter if the gap is approximately the same, say 10 %, 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 % while the other half face no wage penalty. In other words, Jenkins (1994) criticizes the averaging over all human capital characteristics Ξ in (1) and proposes alternative measures that address this issue.[3]

Second, and this is the focus of this paper, Δ_1 has been criticized 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).[4] More recently, with the growing popularity of quantile regression and non-parametric estimation techniques, analysts have been able to examine differences in wage distributions in finer detail. A number of papers have recently proposed methods to estimate complete counterfactual wage distributions if women's human capital characteristics were rewarded as men's, namely

$$F^c(w) = \int_{\Xi} F_x^m(w) h^f(x) dx, \tag{5}$$

—or its inverse, a counterfactual quantile function $Q^c(p)$—which can be compared to actual distributions of female wages $F^f(w)$ (or quantile function $Q^f(w)$, see, e.g., Fortin and Lemieux 1998; Albrecht et al. 2003; Gardeazabal and Ugidos 2005; Melly 2005; Arulampalam et al. 2007; de la Rica et al. 2008; Nopo 2008; Albrecht et al. 2009).[5] Applications of these methods have confirmed that wage distributions do not just vary in levels across gender; the magnitude of differentials can also vary substantially at different quantiles of the (unconditional) wage distribution. Restricting focus on mean wages and summarizing wage differentials with Δ_1 is therefore not fully satisfactory.

3 Two 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

Footnote 2 continued
Barsky et al. (2002) follow Di Nardo et al. (1996) and use reweighing techniques to avoid specifying the conditional mean function.

3 See also Millimet and Wang (2006); Favaro and Magrini (2008) and del Rìo et al. (2011) for further discussion and alternative treatment of this issue.

4 See also Gastwirth (1975), or, more recently, Le Breton et al. (2011).

5 Relevant econometric methods for estimating such counterfactual distribution were mostly developed for analyses of distributional changes (see Di Nardo et al. 1996; Lemieux 2002; Autor et al. 2005; Machado and Mata 2005; Chernozhukov et al. 2009; Firpo et al. 2009).

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." On the basis of this recommendation, two generalized measures of wage differentials are considered in this section. They make it straightforward to incorporate consideration of distribution differences in the construction of wage differentials indicators.

Notionally, 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 disadvantaged 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. If we represent women's preferences over lotteries by a utility functional $U : \Omega \mapsto R$ where Ω is the space of all wage distributions, the prejudice against a woman with characteristics x can then be assessed by the difference between $U\left(F_x^m\right)$ and $U\left(F_x^f\right)$.[6]

To develop the summary measures, it is assumed that U can be expressed as a standard expected utility functional:

$$U(F) = \int_0^\infty u(w)dF(w). \tag{6}$$

Unless u is linear, preferences over wage distributions will be sensitive to the mean and to higher moments of F_x^f and F_x^m. Two strategies are considered at this stage. The first makes only limited additional assumptions about u and relies on stochastic dominance criteria to derive aggregate measures 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

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

[6] I borrow from the literature on the measurement of social justice and inequality. Preferences over wage distributions are assumed formed *ex ante* or "behind a veil of ignorance" (Harsanyi 1955; Rawls 1971), that is, before women have information about their actual position within the wage distributions. This is common in the measurement of inequality (see, e.g.,Atkinson 1970). Of course, this is a notional construct used to motivate the summary measures. The objective is not to capture some "utility" individual women derive from the earnings they actually receive.

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. 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 in comparison to the reference distribution of observationally equivalent 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 given by

$$\Delta_2(p) = \int_{\Xi} D(F_x^m, F_x^f; p)\, h^f(x)dx \qquad p \in \{1, 2, \ldots\}. \tag{7}$$

$\Delta_2(p)$ 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]

Notice that it is wage distributions *conditional on* x (F_x^s) that are fed into the utility functionals and compared through stochastic dominance checks, not aggregate unconditional wage distributions (F^s). This fundamentally distinguishes this approach from the related literature cited earlier that has focused on comparisons of *unconditional* wage distributions. The approach adopted here rests crucially on the view that a measure of distributional prejudice is identified by comparing the wage distributions of observationally equivalent men and women. Consequently, women may be prejudiced at some levels of human capital endowment, and not others. For instance, this accommodates the observation that the wage gap may increase with age because of "glass ceilings" (Cobb-Clark 2001) or career interruptions (Manning and Robinson 2004). Aggregation of differentials captured at given human capital endowments into an overall population measure when pooling agents with different levels of human capital is then addressed separately without reference to any utility function. This is

[7] Note that observing that a proportion $\Delta_2(p)$ of women would prefer being paid like men according to the stochastic dominance comparisons at order p does not imply that, in reverse, a proportion $1 - \Delta_2(p)$ prefers being paid like women. The stochastic dominance ordering is not complete so the latter proportion also includes undeterminate cases at order p.

consistent with Jenkins (1994) discussion of the identification and aggregation issues in the measurement of wage differentials.

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 do not provide orders of magnitude of the wage differential, but only identifies proportions of prejudiced women. To achieve this, the function u needs to be specified completely. A familiar specification used for example by Atkinson (1970) in deriving a measure of inequality assumes constant relative risk aversion:

$$u(w) = \begin{cases} \frac{w^{1-\epsilon}}{1-\epsilon} & \text{if } \epsilon \neq 1 \\ \ln(w) & \text{if } \epsilon = 1 \end{cases} \tag{8}$$

where 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 $C(F; 1) = \exp\left[\int_0^\infty \ln(w) dF(w)\right]$. As is well-known, $\epsilon = 0$ reflects risk neutrality 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. 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 an uncertain outcome rather than a certain outcome.

Summarizing wage differentials by comparing certainty equivalent wages, rather than mean wages, is a straightforward extension to Δ_1 and to the classic Oaxaca–Blinder approach, and one that allows taking differences in higher moments into account. $C(F; \epsilon)$ has the form of a "general mean" which has been characterized as an appealing income standard *per se*, independently on any reference to utility functions by Foster and Székely (2008).[8] This suggests a family of "distributionally sensitive" measures of wage differential (indexed by ϵ), directly comparable to Δ_1,

[8] See Cruces and Wodon (2007) for a recent application of this approach in the context of poverty measurement. See also Makdissi and Wodon (2003).

$$\Delta_3(\epsilon) = \exp\left[\int_{\Xi} \left[\log\left(C(F_x^f;\epsilon)\right) - \log\left(C(F_x^m;\epsilon)\right)\right] h^f(x)dx\right]. \tag{9}$$

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 that $\Delta_3(0) = \Delta_1$.

In principle, a distinct ϵ could be assigned to each women to reflect heterogeneity in individual attitudes toward risk and inequality. In the absence of such information, one constant ϵ is selected. 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

The national statistical institute of Luxembourg recently reported a 20 % difference in average monthly gross wage between men and women, with half of this gap being accounted for by differences in human capital and job characteristics (STATEC 2007). The report also pointed out a lower employment rate for women (55 % against 73 % 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 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, and, thereby, check whether it is realistic to claim that differences in mean wages are compensated by more favourable configurations of higher moments for women.

4.1 Estimation

The measures of wage differential $\Delta_2(p)$ and $\Delta_3(\epsilon)$ 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. Several non-parametric or semi-parametric estimators can be chosen from; see, *inter alia*, Hall et al. (1999), Donald et al. (2000), and Peracchi (2002). 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 parametric, yet flexible, model which allows estimation of the distribution functions under endogenous labor market participation, unlike most semi- or non-parametric approaches.

[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 toward risk among women matter in computing $\Delta_2(p)$ and $\Delta_3(\epsilon)$, not men's attitudes.

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:

$$F_x^s(w) = 1 - \left[1 + \left(\frac{w}{b^s(x)}\right)^{a^s(x)}\right]^{-q^s(x)} \tag{10}$$

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 used 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 fully parametric specification permits to account for the effect of endogenous labor market participation on the estimation of the wage distribution functions. Buchinsky (1998) selectivity-corrected quantile regression model has become popular for dealing with self-selection in a semi-parametric setting; see Albrecht et al. (2009) for an application to gender pay gap analysis. However, this estimator has been shown to be consistent only when conditional quantiles are parallel (Huber and Melly 2011). This conflicts with our objective to capture variations in conditional distributions and therefore rules out this approach here. The model adopted does not suffer from the limitation of the control functions approach, albeit at the cost of imposing stronger parametric assumptions.

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 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):

$$H_x(w, z^*) = \Psi(F_x(w), G_x(z^*)). \tag{11}$$

(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 Ψ to be a Clayton copula.[10] 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.[11]

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

[11] Maximum likelihood estimations for this paper were all done using the built-in Newton–Raphson optimizer of Stata™ (StataCorp 2007).

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 (10), dominance checks are derived from comparisons of $\hat{F}_x^f$ and $\hat{F}_x^m$, and a closed-form 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):

$$C(\hat{F}_x^s; \epsilon) = \hat{b}^s(x) \left(\frac{\Gamma\left(1 + (1-\epsilon)/\hat{a}^s(x)\right) \Gamma\left(\hat{q}^s(x) - (1-\epsilon)/\hat{a}^s(x)\right)}{\Gamma\left(\hat{q}^s(x)\right)} \right)^{\frac{1}{1-\epsilon}} \tag{12}$$

where $\Gamma(\cdot)$ is the Gamma function. These estimates are then plugged into (7) and (9) where average is taken over all females in the sample (in place of integration over h^f).

While the delta method could be used to estimate the sampling variability of the wage differential measures 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 taking into account complex survey design features including 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 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 labor market participation between men and women.

Sample data from 2003 to 2007 are pooled.[12] 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. I consider differences in gross hourly wages of full-time workers (private and public sectors pooled). Only 25- to 55-year-old respondents are considered to avoid issues related to gender differences in retirement and labor market entry. 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.

[12] 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.

In the vector x of human capital characteristics, I take age, educational attainment, nationality, and actual work experience into account. 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). Differences in job characteristics across gender therefore contribute here to differences in the distributions of wage, conditional on human capital endowments.

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), the participation equation (which includes all covariates used in the Singh–Maddala equations) is augmented to strengthen identification of the model 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.

4.3 Results

To fix ideas, Fig. 1 shows estimates of the conditional wage distributions f_x^m and f_x^f for a subset of six combinations of human capital characteristics observed in the data (out of 131 possible combinations). All densities follow Singh-Mandala distributions with estimated parameters $(\hat{a}^s(x), \hat{b}^s(x), \hat{q}^s(x))$. Unsurprisingly, male distributions (solid lines) are generally more densely concentrated toward higher wages than female distributions (dashed lines). There are also noticeable differences in spread and female distributions tend to be more right-skewed. Taking selection into account (gray lines) shifts distributions toward lower wages.

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 % 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 % at the second order. These proportions increase to 25 and 63 % 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 Fig. 2. Estimates for $\epsilon < 0$ assume preference for risk/inequality, thereby rewarding a more dispersed dis-

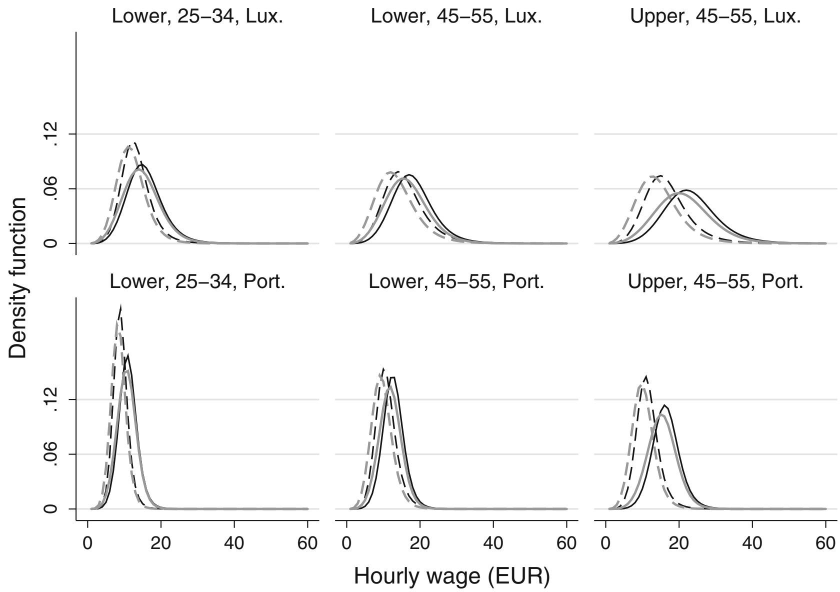

Fig. 1 Examples of conditional probability distribution function estimates f_x^m (*solid lines*, men) and f_x^f (*dashed lines*, women) for illustrative combinations of education level, age and nationality (all at 6–10 years of work experience). *Notes*: *Black lines* are estimates unadjusted for labour 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 versus Portuguese nationality, all with between 6 and 10 years of work experience

tribution. On the contrary, $\epsilon > 0$ assumes risk/inequality aversion thereby penalizing 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 (1)). 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 % (18 %) "unexplained" mean wage disadvantage after controlling for human capital differences.

The message from Fig. 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 nationality.

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]

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 wage distribution differences between men and women (or any two population groups). It is nowadays largely accepted that there is interest in going beyond comparisons "at the mean" and researchers now frequently describe distribution differences between men and women in detail. However, even though scalar summary measures are key to monitor progress and inform policy alongside more flexible graphical tools, no "distributionally sensitive" index measure has been in frequent use. I consider two sets of such 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 and 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

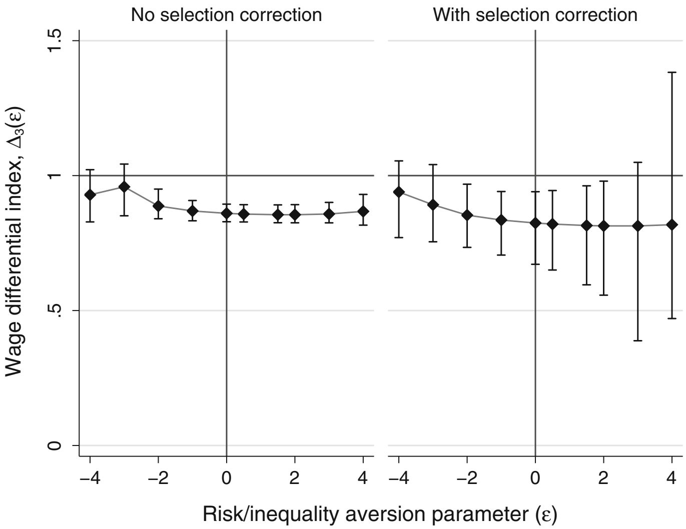

Fig. 2 Certainty-equivalent-based indices, $\Delta_3(\epsilon)$, for $\epsilon \in [-4, 4]$. *Note*: Vertical capped bars show 90 % percentile-based bootstrap variability bands (based on 999 repeated half-sample bootstrap replications)

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. Also, this literature is typically interested in assessing wage gaps at specific (unconditional) quantiles directly, while distribution function estimates *per se* are not looked at here but serve as input to aggregate, summary index measures based on particular utility functionals.

The methods adopted in this paper are also 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). For Jenkins (1994) however, discrimination for a woman with given human capital is 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. It would practically be easy to combine the two approaches. This 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 risk (or 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 labor markets.

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]

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

Acknowledgments 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 Higher Education and Research of Luxembourg. Comments from Stephen Jenkins, CEPS/INSTEAD staff, seminar participants at GREQAM (Marseille), ISER (Uni. Essex), LAMETA (Montpellier), OEI (Regensburg), STATEC (Luxembourg), KU Leuven, and University of Verona and participants at the WIEM 2008, ESPE 2009 and IARIW 2010 conferences are gratefully acknowledged.

Appendix: Derivation of Δ_{OB} from log-wage regressions

This appendix details the derivation of Δ_{OB} given in equation (4).

A log-wage regression among type s agents assumes

$$\begin{aligned} \log(w_i) &= x_i\beta^s + u_i \\ \Leftrightarrow w_i &= \exp(x_i\beta^s)\exp(u_i). \end{aligned}$$

Under this model, the conditional expectation of wages is

$$\begin{aligned}\mu_x^s &= \int w f_x^s(w)dw \\ &= \int \exp(x_i\beta^s)\exp(u) f_u^s(u)du \\ &= \exp(x\beta^s)\int \exp(u) f_u^s(u)du\end{aligned}$$

(where f_x^s is the conditional distribution of wage given x in group s and f_u^s is the distribution of residuals) and the ratio of female-to-male conditional expectations is

$$\begin{aligned}\frac{\mu_x^f}{\mu_x^m} &= \frac{\exp(x\beta^f)}{\exp(x\beta^m)}\ \frac{\int \exp(u) f_u^f(u)du}{\int \exp(u) f_u^m(u)du} \\ &= \exp(x(\beta^f-\beta^m))\ \frac{\int \exp(u) f_u^f(u)du}{\int \exp(u) f_u^m(u)du}.\end{aligned}$$

Setting, for clarity,

$$\theta = \frac{\int \exp(u) f_u^f(u)du}{\int \exp(u) f_u^m(u)du},$$

the measure Δ_{OB} is derived from expression (1) as

$$\begin{aligned}\Delta_{\mathrm{OB}} &= \exp\left(\int_{\Xi} \log\left(\frac{\mu_x^f}{\mu_x^m}\right) h^f(x)dx\right) \\ &= \exp\left(\int_{\Xi} \left(\log(\exp(x(\beta^f-\beta^m))) + \log(\theta)\right) h^f(x)dx\right) \\ &= \exp\left(\left(\int_{\Xi} x(\beta^f-\beta^m)\, h^f(x)dx\right) + \left(\int_{\Xi} \log(\theta) h^f(x)dx\right)\right) \\ &= \exp\left(\bar{x}(\beta^f-\beta^m) + \log(\theta)\right) \\ &= \exp\left(\bar{x}(\beta^f-\beta^m)\right)\theta.\end{aligned}$$

Under normality of the residual distribution, $\theta = \exp[0.5\left(\Omega^f - \Omega^m\right)]$ where Ω^s is the residual variance in the earnings equation for group s (see, e.g., Mood et al. 1987; Kleiber and Kotz 2003). These expressions can be adapted easily to a case with heteroscedasticity in x (that is, when f_u^s varies with x). See Blackburn (2007) for a detailed discussion.

References

Agnew JR, Anderson LR, Gerlach JR, Szykman LR (2008) Who chooses annuities? An experimental investigation of the role of gender, framing, and defaults. Am Econ Rev 98(2):418–422
Albrecht J, Björklund A, Vroman S (2003) Is there a glass ceiling in Sweden?J Labor Econ 21(1):145–177
Albrecht J, van Vuuren A, Vroman S (2009) Counterfactual distributions with sample selection adjustments: Econometric theory and an application to the Netherlands. Labour Econ 16:383–396
Arulampalam W, Booth AL, Bryan ML (2007) Is there a glass ceiling over Europe? Exploring the gender pay gap across the wage distribution. Ind Labor Relat Rev 60(2):163–186
Atkinson AB (1970) On the measurement of inequality. J Econ Theory 2:244–263
Autor DH, Katz LF, Kearney MS (2005) Rising wage inequality: the role of composition and prices NBER working paper 11628. National Bureau of Economic Research, Cambridge
Barsky R, Bound J, Charles KK, Lupton JP (2002) Accounting for the black-white wealth gap: a nonparametric approach. J Am Stat Assoc 97(459):663–673
Biewen M, Jenkins SP (2005) Accounting for differences in poverty between the USA, Britain and Germany. Empir Econ 30(2):331–358
Blackburn ML (2007) Estimating wage differentials without logarithms. Labour Econ 14(1):73–98
Blinder A (1973) Wage discrimination: Reduced form and structural estimates. J Hum Resour 8:436–455
Buchinsky M (1998) The dynamics of changes in the female wage distribution in the USA: A quantile regression approach. J Appl Econ 13(1):1–30
Chernozhukov V, Fernández-Val I, Melly B (2009) Inference on counterfactual distributions. CEMMAP working paper CWP09/09. Centre for Microdata Methods and Practice, Institute for Fiscal Studies, University College of London
Cobb-Clark D (2001) Getting ahead: The determinants of and payoffs to internal promotion for young men and women. Res Labor Econ 20:339–372
Cruces G, Wodon Q (2007) Risk-adjusted poverty in Argentina: Measurement and determinants. J Dev Stud 43(7):1189–1214
Davidson R (2008) Stochastic dominance. In: Durlauf SN, Blume LE (eds) The new palgrave dictionary of economics. Palgrave Macmillan, Basingstoke
de la Rica S, Dolado J, Llorens V (2008) Ceilings or floors? Gender wage gaps by education in Spain. J Popul Econ 21(3):751–776
del Rìo C, Gradin C, Cantò O (2011) The measurement of gender wage discrimination: the distributional approach revisited. J Econ Inequal 9(1):57–86
Di Nardo J, Fortin NM, Lemieux T (1996) Labor market institutions and the distribution of wages, 1973–1992: a semiparametric approach. Econometrica 64(5):1001–1044
Dolton PJ, Makepeace G (1985) The statistical measurement of discrimination. Econ Lett 18:391–395
Donald SG, Green DA, Paarsch HJ (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. Rev Econ Stud 67(4):609–633
Favaro D, Magrini S (2008) Group versus individual discrimination among young workers: A distributional approach. J Socio-Econ 37(5):1856–1879
Firpo S, Fortin NM, Lemieux T (2009) Unconditional quantile regressions. Econometrica 77(3):953–973
Fortin NM, Lemieux T (1998) Rank regressions, wage distributions, and the gender gap. J Hum Resour 33(3):610–643
Foster JE, Székely M (2008) Is economic growth good for the poor? Tracking low incomes using general means. Int Econ Rev 49(4):1143–1172
Gardeazabal J, Ugidos A (2005) Gender wage discrimination at quantiles. J Popul Econ 18(1):165–179
Gastwirth JL (1975) Statistical measures of earnings differentials. Am Stat 29:32–35
Hadar J, Russell WR (1969) Rules for ordering uncertain prospects. Am Econ Rev 59(1):97–122
Hall P, Wolff RCL, Yao Q (1999) Methods for estimating a conditional distribution function. J Am Stat Assoc 94(445):154–163
Harsanyi JC (1955) Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. J Polit Econ 63(4):309–321
Hartog J, Ferrer-i-Carbonell A, Jonker N (2002) Linking measured risk aversion to individual characteristics. Kyklos 55(1):3–26
Huber M, Melly B (2011) Quantile regression in the presence of sample selection. Economics working paper series 1109. University of St. Gallen, School of Economics and Political Science

Jenkins SP (1994) Earnings discrimination measurement: a distributional approach. J Econ 61:81–102
Kleiber C, Kotz S (2003) Statistical size distributions in economics and actuarial sciences. Wiley series in probability and statistics. John Wiley and Sons Inc, New York
Le Breton M, Michelangeli A, Peluso E (2011) A stochastic dominance approach to the measurement of discrimination. Forthcoming in J Econ. Theory
Lemieux T (2002) Decomposing changes in wage distributions: a unified approach. Can J Econ 35(4):646–688
Machado JAF, Mata J (2005) Counterfactual decomposition of changes in wage distributions using quantile regression. J Appl Econ 20:445–465
Makdissi P, Wodon Q (2003) Risk-adjusted measures of wage inequality and safety nets. Econ Bull 9(1):1–10
Manning A, Robinson H (2004) Something in the way she moves: a fresh look at an old gap. Oxford Econ Papers 56(2):169–188
Melly B (2005) Decomposition of differences in distribution using quantile regression. Labour Econ 12:577–590
Millimet DL, Wang L (2006) A distributional analysis of the gender earnings gap in urban China. Contribut Econ Anal Policy 5(1):1–48
Mood AM, Graybill FA, Boes DC (1987) Introduction to the theory of statistics, 3rd edn. McGraw-Hill series in probability and statistics. McGraw-Hill, Auckland
Neal DA, Johnson WR (1996) The role of premarket factors in black-white wage differences. J Polit Econ 104(5):869–895
Ñopo HR (2008) Matching as a tool to decompose wage gaps. Rev Econ Stat 90(2):290–299
Oaxaca RL (1973) Male-female wage differentials in urban labor markets. Int Econ Rev 14:673–709
Oaxaca RL, Ransom MR (1994) On discrimination and the decompostion of wage differentials. J Econ 61:5–21
Peracchi F (2002) On estimating conditional quantiles and distribution functions. Comput Stat Data Anal 38(4):433–447
Racine JS, Green CA (2004) Robust estimation of union and gender wage gaps. Working Paper COBA-ECON-04-01, College of Business Administration, University of South Florida
Rawls J (1971) A theory of justice. Belknap Press, Cambridge
Saigo H, Shao J, Sitter RR (2001) A repeated half-sample bootstrap and balanced repeated replications for randomly imputed data. Surv Methodol 27(2):189–196
Saposnik R (1981) Rank-dominance in income distributions. Public Choice 36(1):147–151
Singh SK, Maddala GS (1976) A function for size distribution of incomes. Econometrica 44(5):963–970
Smith MD (2003) Modelling sample selection using Archimedean copulas. Econometr J 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 PK, Zimmer DM (2007) Copula modeling: an introduction for practitioners. Found Trends Econometr 1(1):1–111