

Quantile Selection Models

With an Application to Understanding Changes in Wage Inequality*

Manuel Arellano
CEMFI, Madrid

Stéphane Bonhomme
University of Chicago

Final version: September 2016

Abstract

We propose a method to correct for sample selection in quantile regression models. Selection is modelled via the cumulative distribution function, or copula, of the percentile error in the outcome equation and the error in the participation decision. Copula parameters are estimated by minimizing a method-of-moments criterion. Given these parameter estimates, the percentile levels of the outcome are re-adjusted to correct for selection, and quantile parameters are estimated by minimizing a rotated “check” function. We apply the method to correct wage percentiles for selection into employment, using data for the UK for the period 1978-2000. We also extend the method to account for the presence of equilibrium effects when performing counterfactual exercises.

JEL CODE: C13, J31.

KEYWORDS: Quantiles, sample selection, copula, wage inequality, gender wage gap.

*This paper was the basis for Arellano’s Walras-Bowley lecture given at the North American Summer Meeting of the Econometric Society in 2011. We are grateful to the editor and two anonymous referees. We thank Xiaohong Chen, David Cox, Ivan Fernández-Val, Toru Kitagawa, Roger Koenker, José Machado, Costas Meghir, Blaise Melly, Alex Torgovitsky, and seminar audiences at various venues for comments.

1 Introduction

Non-random sample selection is a major issue in empirical work. Most selection-correction approaches focus on estimating conditional mean models. In many applications, however, a flexible specification of the entire distribution of outcomes is of interest. In this paper we propose a selection correction method for quantile models.

Quantile regression is widely used to estimate conditional distributions. In a linear quantile model, each percentile is associated with a percentile-specific parameter. Conveniently, quantile parameters can be estimated by minimizing a convex (“check”) function (Koenker and Bassett, 1978). Quantile regression has proved to be a valuable tool to analyze changes in distributions, beginning with Chamberlain (1993) and Buchinsky (1994). However, to our knowledge there is yet no widely accepted quantile regression approach in the presence of sample selection.

A classic example where sample selection features prominently is the study of wages and employment (Gronau, 1974, Heckman, 1974). Only the wages of employed individuals are observed, so conventional measures of wage gaps or wage inequality may be biased. For example, in our empirical application we study the evolution of wage inequality and employment in the UK. Over the past three decades wage inequality has sharply increased. This change in the wage distribution, similar to the one experienced in the US, has motivated a large literature.¹ At the same time, employment rates have also varied during the period, especially for males. In this context, our method to correct for selection allows us to document the evolution of distributions of latent wages, by separating them from changes in employment composition. Wage inequality for those at work may provide a distorted picture of market-level wage inequality.

In regression models, correcting for sample selection involves adding a selection factor as a control. In quantile regression models, we show that selection-corrected estimates can be obtained by suitably shifting the percentile levels as a function of the amount of selection. In practice, this amounts to rotating the “check” function that is optimized in standard quantile regression. The objective function is “discordantly tilted”, since the perturbations applied to percentile levels are observation-specific and depend on the strength of selection. This

¹Gosling, Machin, and Meghir (2000) use quantile regression to study the evolution of wage inequality in the UK. Recent studies for the US are Autor, Katz and Kearney (2005) and Angrist, Chernozhukov and Fernández-Val (2006).

rotation preserves the linear programming structure, and thus the computational simplicity of quantile regression methods.

In our quantile model, sample selection is modelled via the bivariate cumulative distribution function, or *copula*, of the errors in the outcome and the selection equation. Our identification analysis covers the case where the copula is left unrestricted. However, in practice, one may wish to let the copula depend on a low-dimensional vector of parameters.² As in linear sample selection models, excluded variables (e.g., determinants of employment that do not affect wages directly) are key to achieving credible identification. We show how to estimate the parameters of the copula by minimizing a method-of-moments criterion that exploits variation in excluded regressors.

Our estimation algorithm consists of three steps: estimation of the propensity score of participation, the copula parameter, and the quantile parameters, in turn. We derive the asymptotic distribution of the estimator. We also analyze a number of extensions of the method. In particular, we propose a bounds method to assess the influence on the quantile estimates of the parametric restrictions imposed on the copula.

We apply the method to study the evolution of wage inequality in the UK in the last quarter of the twentieth century. We find that correcting for selection into employment strongly affects male wages at the bottom of the distribution, which is consistent with low-skilled males being progressively driven out of the labor market. Sample selection has smaller effects for females. As a result, correcting for sample selection accentuates the decrease in the gender wage gap at the bottom (though less at the top) of the distribution. We also perform several robustness checks, in particular regarding the specification of the copula.

Lastly, we propose a method to obtain counterfactual distributions of wages taking into account general equilibrium effects. Our approach combines the quantile selection model of wages and participation with a labor demand side in the spirit of Katz and Murphy (1992) and Card and Lemieux (2001). Because of demand responses, shifts in participation may affect latent equilibrium wage distributions. We apply the method to a counterfactual exercise where potential out-of-work income, a strong policy-based determinant of participation, is kept constant throughout the period.

²Copulas have been extensively used in statistics and financial econometrics (e.g., Joe, 1997, and Nelsen, 1999). Single-parameter copula families have been shown to yield satisfactory fit to empirical data in various contexts. For example, Bonhomme and Robin (2009) use a Plackett copula to model year-to-year earnings mobility.

Literature and outline. Our approach connects with two complementary approaches that have been used to deal with sample selection. Parametric and semiparametric versions of the Heckman (1979) sample selection model have been extensively studied. See for example Heckman and Sedlacek (1985), Heckman (1990), Ahn and Powell (1993), Donald (1995), Chen and Khan (2003), and Das, Newey, and Vella (2003). Vella (1998) provides a number of additional references. In comparison, bounds methods (Manski, 1994, Blundell, Gosling, Ichimura, and Meghir, 2007, Kitagawa, 2010) have been less studied. The sensitivity analysis in Kline and Santos (2013) is also related to our approach. However, unlike in the missing data settings that they consider, excluded variables in selection models provide information on the sign and strength of sample selection, which we exploit.

The paper also connects with the large literature on quantiles, distributions, and treatment effects. Chernozhukov and Hansen (2005, 2006) develop an instrumental variables quantile regression approach. Unlike in this paper, they rely on a rank invariance or rank similarity assumption (see also Vuong and Xu, 2014). Related models with continuous endogenous regressors are studied in Torgovitsky (2015) and D’Haultfoeuille and Février (2015). Imbens and Rubin (1997) study identification and estimation of unconditional distributions of potential outcomes in a treatment effects model with a binary instrument, and achieve identification for compliers (as in Abadie, 2003, and Abadie, Angrist and Imbens, 2002). Carneiro and Lee (2009) use the framework of Heckman and Vytlacil (2005) to identify and estimate distributions of potential outcomes on suitable “complier” subpopulations. The tools we propose could be used to provide alternative estimators in treatment effects settings. In addition, being distribution-based, our approach allows one to perform distributional decomposition exercises (as in DiNardo, Fortin and Lemieux, 1996, and Firpo, Fortin and Lemieux, 2011) while accounting for sample selection.

The literature on quantile selection models, in contrast, is scarce (see the review in Arelano and Bonhomme, 2016). Buchinsky (1998, 2001) proposes an additive approach to correct for sample selection in quantile regression. Huber and Melly (2015) consider a more general, non-additive quantile model, as we do; they focus on testing for additivity. In contrast, our focus is on providing a practical estimation method. Also related are Neal (2004), who develops imputation methods to correct the black/white wage gap among women, Olivetti and Petrongolo (2008), who apply similar methods to the gender wage gap, and Picchio and Mussida (2010), who propose a parametric model to correct the gender wage gap for selec-

tion into employment. See also Lee (1983) and Smith (2003) for parametric distributional selection-correction methods.

The rest of the paper is as follows. In Section 2 we present the quantile selection model and discuss identification. In Section 3 we describe the estimator and its asymptotic properties. In Section 4 we outline several extensions of our approach. The empirical analysis is contained in Section 5, and the counterfactual exercise in Section 6. Lastly, we conclude in Section 7. Computer codes and an appendix with additional results are provided in the supplementary material.

2 Model and identification

2.1 Model and assumptions

We consider the following sample selection model:

$$Y^* = q(U, X), \quad (1)$$

$$D = \mathbf{1}\{V \leq p(Z)\}, \quad (2)$$

$$Y = Y^* \quad \text{if} \quad D = 1, \quad (3)$$

where Y^* is the latent outcome (e.g., market wage), D is the participation indicator (employment), U and V are error terms, and $Z = (B, X)$ strictly contains X , so B are the excluded covariates. We observe (Y, D, Z) , so that potential outcomes $Y^* = Y$ are observed only when $D = 1$ (e.g., if the individual is a labor market participant).

We make four assumptions.

Assumption 1

A1 (*exclusion restriction*) (U, V) is jointly statistically independent of Z given X .

A2 (*unobservables*) The bivariate distribution of (U, V) given $X = x$ is absolutely continuous with respect to the Lebesgue measure, with standard uniform marginals and rectangular support. We denote its cumulative distribution function (cdf) as $C_x(u, v)$.

A3 (*continuous outcomes*) The conditional cdf $F_{Y^*|X}(y|x)$ and its inverse are strictly increasing. In addition, $C_x(u, v)$ is strictly increasing in u .

A4 (*propensity score*) $p(Z) \equiv \Pr(D = 1|Z) > 0$ with probability one.

Assumption A1 is satisfied if $Z = (B, X)$ strictly contains X , and (U, V) is jointly independent of B given X . In the example of wages and employment, B may measure

opportunity costs of participation in the labor market. Following Blundell *et al.* (2003), our empirical application will use a measure of potential out-of-work welfare income as exclusion restriction.

Model (1)-(3) depends on two sources of unobserved heterogeneity: the latent outcome rank U and the percentile rank V . In Assumption A2 we normalize their marginal distributions to be uniform on the unit interval, independent of Z . In particular, $\tau \mapsto q(\tau, x)$ is the conditional quantile function of Y^* given $X = x$, and it is increasing in τ by A3. A special case is the linear quantile model $Y^* = X'\beta_U$, which is widely used in applied work since Koenker and Bassett (1978). The Skorohod representation (1) is without loss of generality.³

Joint independence between (U, V) and Z given X , as stated in Assumption A2, is stronger than marginal independence. This requires the conditional cdf (that is, the *copula*) of the pair (U, V) given (B, X) to solely depend on X . The presence of dependence between U and V is the source of sample selection bias.

Lastly, A3 restricts the analysis to absolutely continuous outcomes, and A4 is a support assumption on the propensity score often made in sample selection models.

Examples. Before discussing identification of model (1)-(3) we briefly outline two special cases. A first special case is obtained when outcomes are additive in unobservables: $Y^* = g(X) + \varepsilon$, where (ε, V) is independent of Z . Note that Assumption A1 is satisfied, with $U = F_\varepsilon(\varepsilon)$, for F_ε the cdf of ε . Moreover, the following restrictions hold (as in Das *et al.*, 2003):

$$\mathbb{E}(Y | D = 1, Z) = g(X) + \mathbb{E}(\varepsilon | V \leq p(Z), Z) = g(X) + \lambda(p(Z)),$$

where $\lambda(p) \equiv \mathbb{E}(\varepsilon | V \leq p)$.

As a second special case suppose the following reservation rule:

$$D = \mathbf{1}\{Y^* \geq R(Z) + \eta\}, \tag{4}$$

where (Y^*, η) is statistically independent of Z given X . In a labor market application, (4) may represent the participation decision of an individual, who compares her potential wage Y^* with a reservation wage $R(Z) + \eta$. Note that (4) may equivalently be written as:

$$D = \mathbf{1}\{V \leq F_{\eta - Y^* | Z}(-R(Z) | Z)\},$$

³Indeed, $U = F_{Y^* | X}(Y^* | X)$, where $F_{Y^* | X}$ is the conditional cdf of Y^* given X . Moreover, U being independent of Z given X is equivalent to the potential outcome Y^* being independent of Z given X .

where $V \equiv F_{\eta-Y^*|Z}(\eta - Y^*|Z) = F_{\eta-Y^*|X}(\eta - Y^*|X)$ is uniformly distributed on the unit interval, and independent of Z . Letting $Y^* = q(U, X)$, (U, V) is independent of Z given X , so Assumption A1 is satisfied. At the same time, however, U and V are not jointly independent of X . Thus, in this reservation wage model the copula $C_x(\cdot, \cdot)$ depends on x in general.

2.2 Main restrictions and identification

We have, conditional on participation and for all $\tau \in (0, 1)$:

$$\begin{aligned} \Pr(Y^* \leq q(\tau, x) \mid D = 1, Z = z) &= \Pr(U \leq \tau \mid V \leq p(z), Z = z), \\ &= G_x(\tau, p(z)), \end{aligned} \quad (5)$$

where $G_x(\tau, p) \equiv C_x(\tau, p)/p$, and we have used Assumptions A1 to A4. The *conditional copula* $G_x(\cdot, \cdot)$ measures the dependence between U and V , which is the source of sample selection bias. As a special case, if U and V are conditionally independent given $X = x$ then $G_x(\tau, p(z)) = \tau$. More generally, (5) shows that G_x maps ranks τ in the distribution of latent outcomes (given $X = x$) to ranks $G_x(\tau, p(z))$ in the distribution of observed outcomes conditional on participation (given $Z = z$).

An implication of (5) is that, for each $\tau \in (0, 1)$, the conditional τ -quantile of Y^* coincides with the conditional $G_x(\tau, p(z))$ -quantile of Y given $D = 1$. Hence, if we knew the mapping G_x from latent to observed ranks, one could recover $q(\tau, x)$ as a quantile of observed outcomes, by suitably shifting percentile ranks.

Equation (5) is instrumental to correct quantile functions from selection. Given knowledge of the mapping G_x , latent quantiles can readily be recovered. Moreover, the exclusion restriction provides information about G_x . The intuition for this is that (5) holds for all z in the support of Z given $X = x$, thus generating restrictions on G_x .

The following result spells out the restrictions on the conditional copula G_x . We denote as \mathcal{X} the support of X , and as \mathcal{Z}_x the support of Z given $X = x$. G_x^{-1} and $F_{Y|D=1,Z}^{-1}$ denote the inverses of G_x and $F_{Y|D=1,Z}$ with respect to their first arguments, which exist by Assumption A3. Proofs are given in Appendix A.

Lemma 1 *Let $x \in \mathcal{X}$. Then, under Assumptions A1 to A4:*

$$F_{Y|D=1,Z} \left(F_{Y|D=1,Z}^{-1}(\tau|z_2) \mid z_1 \right) = G_x \left(G_x^{-1}(\tau, p(z_2)), p(z_1) \right), \quad \text{for all } (z_1, z_2) \in \mathcal{Z}_x \times \mathcal{Z}_x. \quad (6)$$

Moreover, for any G_x satisfying (6), one can find a distribution of latent outcomes $F_{Y^*|X}$ such that $G_x(F_{Y^*|X}(y|x), p(z)) = F_{Y|D=1,Z}(y|z)$ for all (z, y) in the support of (Z, Y) given $X = x$.

Note that the restrictions in (6) are uninformative in the absence of an exclusion restriction. They may become informative as soon as the conditional support of Z given $X = x$ contains two or more values. Moreover, the second part of Lemma 1 shows that these are the only restrictions on G_x , in the sense that, for any G_x satisfying (6), one can find a distribution of latent outcomes that rationalizes the data.

Nonparametric point-identification. Two simple conditions lead to nonparametric point identification of G_x , and hence to point-identification of $q(\cdot, x)$ as well. We denote as \mathcal{P}_x the conditional support of the propensity score $p(Z)$ given $X = x$.

Proposition 1 *Let Assumptions A1 to A4 hold. Let $x \in \mathcal{X}$. Suppose that one of the two following conditions holds:*

i) (identification at infinity) There exists some $z_x \in \mathcal{Z}_x$ such that $p(z_x) = 1$.

ii) (analytic extrapolation) \mathcal{P}_x contains an open interval and, for all $\tau \in (0, 1)$, the function $p \mapsto G_x(\tau, p)$ is real analytic on the unit interval.

Then the functions $(\tau, p) \mapsto G_x(\tau, p)$ and $\tau \mapsto q(\tau, x)$ are nonparametrically identified.

Both conditions in Proposition 1 allow one to point-identify the dependence mapping G_x and the quantile function $q(\cdot, x)$ using an extrapolation strategy. Under *i*), identification is achieved at the boundary of the support of the propensity score (“at infinity”). Under *ii*), extrapolation is based on the property that real analytic functions that coincide on an open neighborhood coincide everywhere. Absent conditions *i*) and *ii*) of Proposition 1, the model is nonparametrically partially identified in general.

Partial identification Let $x \in \mathcal{X}$ and $\tilde{z} \in \mathcal{Z}_x$. Using the worst-case Fréchet bounds (Fréchet, 1951, Heckman, Smith and Clements, 1997) on the copula C_x we can bound:

$$\max\left(\frac{\tau + p(\tilde{z}) - 1}{p(\tilde{z})}, 0\right) \leq G_x(\tau, p(\tilde{z})) \leq \min\left(\frac{\tau}{p(\tilde{z})}, 1\right), \quad \text{for all } \tau \in (0, 1). \quad (7)$$

Let now $z \in \mathcal{Z}_x$. Evaluating (6) at $(z_1, z_2) = (z, \tilde{z})$, and using (7) to bound $G_x(\tau, p(\tilde{z}))$, we obtain the following bounds on $G_x(\tau, p(z))$:

$$G_x(\tau, p(z)) \leq \inf_{\tilde{z} \in \mathcal{Z}_x} F_{Y|D=1,Z} \left[F_{Y|D=1,Z}^{-1} \left(\min \left(\frac{\tau}{p(\tilde{z})}, 1 \right) \mid \tilde{z} \right) \mid z \right] \quad (8)$$

$$G_x(\tau, p(z)) \geq \sup_{\tilde{z} \in \mathcal{Z}_x} F_{Y|D=1,Z} \left[F_{Y|D=1,Z}^{-1} \left(\max \left(\frac{\tau + p(\tilde{z}) - 1}{p(\tilde{z})}, 0 \right) \mid \tilde{z} \right) \mid z \right]. \quad (9)$$

Moreover, using (5) and (7) we have the following bounds on the quantiles of latent outcomes:

$$q(\tau, x) \leq \inf_{\tilde{z} \in \mathcal{Z}_x} F_{Y|D=1,Z}^{-1} \left(\min \left(\frac{\tau}{p(\tilde{z})}, 1 \right) \mid \tilde{z} \right) \quad (10)$$

$$q(\tau, x) \geq \sup_{\tilde{z} \in \mathcal{Z}_x} F_{Y|D=1,Z}^{-1} \left(\max \left(\frac{\tau + p(\tilde{z}) - 1}{p(\tilde{z})}, 0 \right) \mid \tilde{z} \right). \quad (11)$$

The quantile bounds in (10) and (11) were first derived by Manski (1994, 2003) in a slightly more general selection model. In related work, Kitagawa (2009, 2010) provides comprehensive studies of the role of independence and first-stage monotonicity restrictions in LATE and sample selection settings, respectively. The bounds in (10) and (11) coincide with the choice of the upper or lower Fréchet bounds for the copula of (U, V) . In this sense, these are worst-case bounds.⁴ In Section S1 of the supplementary appendix we show that these bounds cannot be improved upon. Importantly, in this paper we work under the maintained assumption that the model is correctly specified; that is, that (1)-(2)-(3) hold. If the threshold specification in (2) were relaxed, for example in the absence of monotonicity, it would be possible to improve over the quantile bounds (10) and (11), as shown in Kitagawa (2010).

3 Estimation

We adopt a flexible semi-parametric specification. Following a large literature on quantile regression, we assume that quantile functions are linear, that is:

$$q(\tau, x) = x' \beta_\tau, \quad \text{for all } \tau \in (0, 1) \text{ and } x \in \mathcal{X}. \quad (12)$$

Although our estimation strategy can be extended to deal with nonlinear specifications, the linear quantile model is convenient for computation. We discuss a nonparametric extension in the next section.

⁴Note, however, that the Fréchet copula bounds do not satisfy (6) in general. By (8) and (9), the bounds on G_x are generally tighter than the Fréchet bounds.

We assume that the copula function, and hence the function G_x , is indexed by a parameter vector ρ ; that is:

$$G_x(\tau, p) \equiv G(\tau, p; \rho) = \frac{C(\tau, p; \rho)}{p}.$$

The statistical literature offers a number of convenient parsimonious specifications, including the Gaussian, Frank, or Gumbel copulas. See Nelsen (1999) and Joe (1997) for comprehensive references. Flexible families may be constructed, for example by relying on the Bernstein family of polynomials (Sancetta and Satchell, 2004). In all these examples, one may let the vector ρ depend on x .⁵ For simplicity we omit the dependence of ρ on x in the following.

The parametric assumptions on the copula are substantive. Restricting the analysis to a finite-dimensional ρ allows us to focus on the case where ρ is point-identified and to propose a simple estimation method. In addition, below we propose a bounds approach to assess the influence on quantile estimates of the parametric assumptions made on the copula.

Lastly, the propensity score $p(z; \theta)$ is specified as a known function of a parameter θ . This assumption may be relaxed, at the cost of making the asymptotic analysis more involved (see the next section).

The functional form of selected quantiles. Before describing the estimator, we first comment on the form of the conditional quantiles given participation, when quantile functions of latent outcomes are linear as in (12). The τ -quantile of outcomes of participants given $z = (b, x)$ is, by (5):

$$q^d(\tau, z) \equiv F_{Y|D=1,Z}^{-1}(\tau|z) = x' \beta_{G^{-1}(\tau, p(z); \rho)}. \quad (13)$$

Equation (13) makes it clear that sample selection affects all quantiles, and that quantile functions of observed outcomes are generally non-additive in x and $p(z)$. We have the following result, where it is assumed that ρ does not depend on x .

Proposition 2 *Let $\tau \in (0, 1)$. Suppose that ρ does not depend on x . Then $z \mapsto q^d(\tau, z)$ is non-additive in x and $p(z)$, unless:*

- i) All coefficients of β_τ but the intercept are independent of τ , or*
- ii) U and V are statistically independent.*

⁵For example, for scalar $\rho \in (-1, 1)$ one may specify $\rho(x) = (e^{x'\gamma} - 1) / (e^{x'\gamma} + 1)$, where γ is a vector of parameters.

Additive specifications such as $q^d(\tau, z) = x'\beta_\tau + \lambda_\tau(p(z))$, for a smooth function $\lambda_\tau(p)$, are sometimes used in applied work (see the review in Arellano and Bonhomme, 2016). In contrast, in our framework, conditional quantiles of participants are non-additive. Huber and Melly (2015) make a related point in a testing context. Correcting for sample selection thus requires shifting the percentile ranks of individual observations. We now explain how this can be done in estimation.

3.1 Three-step estimation strategy

Let (Y_i, D_i, B_i, X_i) , $i = 1, \dots, N$, be an i.i.d. sample, with $Z_i \equiv (B_i, X_i)$. We propose to compute selection-corrected quantile regression estimates in three steps. In the first step, we compute $\hat{\theta}$, a consistent estimate of the propensity score parameter θ . In the second step, we compute a consistent estimator $\hat{\rho}$ of the copula parameter vector ρ . Lastly, given $\hat{\theta}$ and $\hat{\rho}$, for any given $\tau \in (0, 1)$ we compute $\hat{\beta}_\tau$, a consistent estimator of the τ th quantile regression coefficient.

The first step can be done using maximum likelihood. We now present the third and second steps in turn.

Rotated quantile regression (Step 3). Let us suppose that consistent estimators $\hat{\theta}$ and $\hat{\rho}$ are available. Then, for any given $\tau \in (0, 1)$ we compute:

$$\hat{\beta}_\tau = \operatorname{argmin}_{b \in \mathcal{B}} \sum_{i=1}^N D_i \left[\hat{G}_{\tau i} (Y_i - X_i' b)^+ + (1 - \hat{G}_{\tau i}) (Y_i - X_i' b)^- \right], \quad (14)$$

where \mathcal{B} is the parameter space for β_τ , $a^+ = \max(a, 0)$, $a^- = \max(-a, 0)$, and:

$$\hat{G}_{\tau i} \equiv G\left(\tau, p\left(Z_i; \hat{\theta}\right); \hat{\rho}\right).$$

Solving (14) amounts to minimizing a *rotated* check function, with individual-specific perturbed τ . As with standard quantile regression, the optimization problem takes the form of a simple linear program, and can thus be solved in a fast and reliable way. It is instructive to compare the rotated quantile regression estimate $\hat{\beta}_\tau$ with the following infeasible quantile regression estimate based on the latent outcomes:

$$\tilde{\beta}_\tau = \operatorname{argmin}_{b \in \mathcal{B}} \sum_{i=1}^N \left[\tau (Y_i^* - X_i' b)^+ + (1 - \tau) (Y_i^* - X_i' b)^- \right].$$

We see that, in order to correct for selection in (14), τ is replaced by the selection-corrected, individual-specific percentile rank $\hat{G}_{\tau i}$.

Estimating the copula parameter (Step 2). From (5), we obtain the following conditional moment restrictions:

$$\mathbb{E} \left[\mathbf{1} \{Y \leq X' \beta_\tau\} - G(\tau, p(Z; \theta); \rho) \mid D = 1, Z = z \right] = 0.$$

We propose to estimate the copula parameter ρ as:

$$\hat{\rho} = \underset{c \in \mathcal{C}}{\operatorname{argmin}} \left\| \sum_{i=1}^N \sum_{\ell=1}^L D_i \varphi(\tau_\ell, Z_i) \left[\mathbf{1} \{Y_i \leq X_i' \hat{\beta}_{\tau_\ell}(c)\} - G(\tau_\ell, p(Z_i; \hat{\theta}); c) \right] \right\|, \quad (15)$$

where $\tau_1 < \tau_2 < \dots < \tau_L$ is a finite grid on $(0, 1)$, $\|\cdot\|$ is the Euclidean norm, $\varphi(\tau, Z_i)$ are instrument functions with $\dim \varphi \geq \dim \rho$, and:

$$\begin{aligned} \hat{\beta}_\tau(c) \equiv \underset{b \in \mathcal{B}}{\operatorname{argmin}} \sum_{i=1}^N D_i \left[G(\tau, p(Z_i; \hat{\theta}); c) (Y_i - X_i' b)^+ \right. \\ \left. + \left(1 - G(\tau, p(Z_i; \hat{\theta}); c) \right) (Y_i - X_i' b)^- \right]. \end{aligned} \quad (16)$$

Effectively, in this step we are estimating ρ together with $\beta_{\tau_1}, \dots, \beta_{\tau_L}$. Hence, if the researcher is only interested in β_τ for $\tau \in \{\tau_1, \dots, \tau_L\}$, Step 3 is not necessary.

This step is computationally more demanding. In particular, the objective function in (15) is not continuous, due to the presence of the indicator functions, and generally non-convex. In practice, for low-dimensional ρ one may use grid search, as in our application. For higher-dimensional ρ , simulation-based methods such as simulated annealing (see, e.g., Judd, 1998), or the pseudo-Bayesian approach of Chernozhukov and Hong (2003), could be used. Importantly, evaluating the objective function is usually fast and straightforward. The reason is that (16) is a linear programming problem, for which there exist fast algorithms.⁶

In addition, in experiments we observed that using a large number of percentile values τ_ℓ in (15) tends to smooth the objective function. In Section S2 of the supplementary appendix we consider a nonparametric quantile specification with discrete covariates, and show that in this case an integrated version of the objective function in (15), with a continuum of τ values, is differentiable with respect to the copula parameter c under weak conditions.

Finally, solving (15) is only one possibility to estimate the copula parameter. In Section S3 of the supplementary appendix we describe an alternative estimator of ρ that relies on the copula restrictions (6). The method provides a fast and straightforward way to obtain good starting values to minimize the objective function in (15). Another possibility would

⁶For example, the Matlab version of Morillo, Koenker and Eilers is directly applicable to the problem at hand. Available at: <http://www.econ.uiuc.edu/~roger/research/rq/rq.m>

be to estimate ρ using a likelihood approach, based on the semi-parametric structure of the model. An interesting question, which we do not address in this paper, would be to construct a semi-parametric efficient estimator for ρ by exploiting the continuum of moment restrictions in (6).

Remark: unconditional quantiles. Once θ and ρ have been estimated, the parameters β_τ are estimated by simple quantile regression using the rescaled percentile levels $\widehat{G}_{\tau i} = G\left(\tau, p\left(Z_i; \widehat{\theta}\right); \widehat{\rho}\right)$ in place of τ . So, the techniques developed in the context of ordinary quantile regression can be used in the presence of sample selection. As an example, counterfactual distributions may be constructed as explained in Machado and Mata (2005) and Chernozhukov, Fernández-Val and Melly (2013). Specifically, the unconditional cdf of Y^* may be estimated as a discretized or simulated version of:

$$\widehat{F}_{Y^*}(y) = \frac{1}{N} \sum_{i=1}^N \int_0^1 \mathbf{1}\left\{X_i' \widehat{\beta}_\tau \leq y\right\} d\tau,$$

and unconditional quantiles can be estimated as $\widehat{q}(\tau) = \inf\left\{y, \widehat{F}_{Y^*}(y) \geq \tau\right\}$. Also, a pervasive problem in quantile regression is that estimated quantile curves may cross each other because of sampling error. The approach proposed by Chernozhukov, Fernández-Val and Galichon (2010), based on quantiles rearrangement, may also be applied in our context.⁷

3.2 Asymptotic properties

In Section S4 of the supplementary appendix we derive the asymptotic distributions of $\widehat{\rho}$ and $\widehat{\beta}_\tau$ for given τ . Under standard conditions for quantile regression estimators (as in Koenker, 2005), an identification condition to be discussed below, and suitable differentiability conditions on G , the estimators satisfy:

$$\sqrt{N} \begin{pmatrix} \widehat{\beta}_\tau - \beta_\tau \\ \widehat{\rho} - \rho \end{pmatrix} \xrightarrow{d} \mathcal{N}(0, V_\tau), \quad (17)$$

where ρ and β_τ denote true parameter values. We provide an explicit expression for the asymptotic variance V_τ , which can be estimated using an approach similar to the one in Powell (1986). These results can be easily generalized to derive the asymptotic distribution for a finite number of quantiles $(\widehat{\beta}_{\tau_1}, \dots, \widehat{\beta}_{\tau_L})$. An interesting extension is to derive the large

⁷A difference with standard quantile regression concerns inference, as one needs to take into account that ρ and θ have already been estimated when computing asymptotic confidence intervals.

sample theory of the quantile process $\tau \mapsto \sqrt{N} \left(\widehat{\beta}_\tau - \beta_\tau \right)$, which can be done along the lines of Koenker and Xiao (2002) or Chernozhukov and Hansen (2006). Confidence bands for unconditional effects may be derived using the results in Chernozhukov, Fernández-Val and Melly (2013). Alternatively, given the distributional characterization in (17), confidence intervals may be estimated using subsampling (Politis, Romano and Wolf, 1999). In our empirical application, given the large sample sizes, subsampling is computationally attractive relative to other methods such as the conventional nonparametric bootstrap.

An important condition in the asymptotic analysis is the identification of ρ based on the following unconditional moment restrictions:

$$\sum_{\ell=1}^L \mathbb{E} \left[D\varphi(\tau_\ell, Z) \left(\mathbf{1} \{Y \leq X' \overline{\beta}_{\tau_\ell}(\rho)\} - G(\tau_\ell, p(Z; \theta); \rho) \right) \right] = 0, \quad (18)$$

where $\overline{\beta}_\tau(c)$ solves the population counterpart to (16). A rank condition for local identification is readily obtained.⁸ Identification intuitively requires that the propensity score vary sufficiently conditionally on X , and that both φ and the ρ -derivative of G depend on it.

3.3 Estimating bounds

The above method to estimate the copula parameter ρ relies on the assumption that the copula, and hence the quantile functions, are point-identified. In the absence of functional form assumptions on the copula, both G and $q(\tau, x)$ are partially identified in general. In particular, the quantiles of latent outcomes are bounded by (10) and (11).⁹ In practice, a simple way to informally assess the influence of functional form assumptions on the results is to compute estimates of the bounds in (10) and (11), obtained from the semi-parametric model.

Denoting $\overline{p}_x = \sup_b p(x, b)$ the supremum of the support of the excluded variable B for

⁸For example, when $L = 1$ and $\tau_1 = \tau$, it suffices that the following matrix be full column rank:

$$\mathbb{E} [D\varphi(\tau, Z) X' f_Z(X' \beta_\tau)] \mathbb{E} [D X X' f_Z(X' \beta_\tau)]^{-1} \mathbb{E} [D X \nabla G'_Z] - \mathbb{E} [D\varphi(\tau, Z) \nabla G'_Z],$$

where f_Z denotes the conditional density of Y given $D = 1$ and Z , and $\nabla G_Z = \frac{\partial G(\tau, p(Z; \theta); \rho)}{\partial c}$.

⁹Note that (10) and (11) do not impose a linear representation of the quantile functions as in $q(\tau, X) = X' \beta_\tau$. Under linearity one could in principle derive tighter bounds, although such bounds would not be valid under misspecification of the quantile functions.

given $X = x$, the model implies the following bounds:¹⁰

$$\underline{q}(\tau, x) \equiv x' \beta_{G^{-1}(\max(\frac{\tau + \bar{p}_x - 1}{\bar{p}_x}, 0), \bar{p}_x; \rho)} \leq q(\tau, x) \leq x' \beta_{G^{-1}(\min(\frac{\tau}{\bar{p}_x}, 1), \bar{p}_x; \rho)} \equiv \bar{q}(\tau, x). \quad (19)$$

Under the assumption that the support of B given $X = x$ is independent of x , \bar{p}_x can be consistently estimated by $\hat{p}_x = \sup_{i \in \{1, \dots, N\}} p(x, B_i; \hat{\theta})$. As these estimates may be sensitive to outliers, in the application we will also consider alternative estimates based on a trimming approach. Consistent estimates of $\underline{q}(\tau, x)$ and $\bar{q}(\tau, x)$ are then obtained by replacing \bar{p}_x , β_τ , and ρ , by \hat{p}_x , $\hat{\beta}_\tau$, and $\hat{\rho}$, respectively.

We are thus using our model as a semi-parametric specification for the self-selected conditional quantiles, and therefore for the bounds, which themselves are nonparametrically identified. An alternative, fully nonparametric strategy, robust to violation of the parametric assumptions on the copula, would be to construct estimators and confidence sets for the identified sets of the copula and quantile functions. We will return to this possibility in the conclusion.

4 Extensions

In this section we briefly discuss several extensions of our approach. More details are given in Section S5 of the supplementary appendix.

Nonparametric quantile regression. Consistency of the estimator described in Section 3 requires quantile linearity (12) to hold, at least at all τ values of interest.¹¹ Nonparametric estimators could be used instead. As an example, denoting X_i net of the constant as \tilde{X}_i , one might consider replacing (15)-(16) using the following local linear approach:

$$\hat{\rho} = \operatorname{argmin}_{c \in \mathcal{C}} \left\| \sum_{i=1}^N \sum_{\ell=1}^L D_i \varphi(\tau_\ell, Z_i) \left[\mathbf{1} \left\{ Y_i \leq \hat{q}_{\tau_\ell}(c, \tilde{X}_i) \right\} - G(\tau_\ell, p(Z_i; \hat{\theta}); c) \right] \right\|,$$

where:

$$\begin{aligned} \hat{q}_\tau(c, x) \equiv \operatorname{argmin}_{b_0 \in \mathcal{B}_0} \min_{b_1 \in \mathcal{B}_1} \sum_{i=1}^N D_i \kappa \left(\frac{\tilde{X}_i - x}{h} \right) & \left[G(\tau, p(Z_i; \hat{\theta}); c) \left(Y_i - b_0 - (\tilde{X}_i - x)' b_1 \right)^+ \right. \\ & \left. + \left(1 - G(\tau, p(Z_i; \hat{\theta}); c) \right) \left(Y_i - b_0 - (\tilde{X}_i - x)' b_1 \right)^- \right], \end{aligned}$$

¹⁰One can show that, given that $G(\cdot, \cdot; \rho)$ is a conditional copula, $p \mapsto G^{-1}(\min(\frac{\tau}{p}, 1), p)$ is non-increasing, and $p \mapsto G^{-1}(\max(\frac{\tau + p - 1}{p}, 0), p)$ is non-decreasing, for all $\tau \in (0, 1)$.

¹¹For example, if one is only interested in the median $\beta_{1/2}$, when using Step 2 of the algorithm with $L = 1$ and $\tau_1 = 1/2$, consistency will only require a linearity assumption on the conditional median; that is, $q(1/2, x) = x' \beta_{1/2}$.

where h is a vanishing bandwidth and κ is a kernel function (e.g., Chaudhuri, 1991).

Treatment effects with selection on unobservables. As a direct extension of model (1)-(3), consider the following system of equations:

$$Y_0^* = q(U_0, X), \quad Y_1^* = q(U_1, X), \quad Y = (1 - D)Y_0^* + DY_1^*, \quad (20)$$

where, in the spirit of Assumption A1, (U_0, U_1, V) is assumed independent of Z given X . This model coincides with the standard potential outcomes framework in the treatment effects literature (Vytlacil, 2002). In the context of the empirical application, $Y_0^* = 0$, and Y_1^* is the partial equilibrium causal effect of working. In this framework, the quantile IV method of Chernozhukov and Hansen (2005) relies on an assumption of rank invariance or rank similarity which restricts the dependence between U_0 and U_1 . Specifically, rank invariance (respectively, similarity) requires the comonotonicity of U_0 and U_1 (resp., given V), thus ruling out most patterns of sample selection. In contrast, in the identification analysis our approach leaves the joint distribution of U_0 , U_1 and V given X unrestricted.

The treatment effects literature has characterized quantities of economic interest, which may be identified in model (20) absent rank invariance. Related to this paper, Carneiro and Lee (2009) extend the analysis in Heckman and Vytlacil (2005) to identify and estimate conditional distributions of potential outcomes. They provide conditions under which conditional cdfs and quantiles of Y_0^* and Y_1^* are identified given $V = p$ and $X = x$, for p in the support of $p(Z)$ given $X = x$. In estimation, Carneiro and Lee specify potential outcomes as additive in X and an unobservable independent of Z . Interestingly, our approach may be used to estimate such conditional quantiles (or cdfs), while allowing observables X and unobservables (U_0, U_1) to interact.¹² At the same time, as pointed out in Section 2, identification of the *unconditional* distributions of potential outcomes in a nonseparable setup would require either identification at infinity or analytic extrapolation.¹³ In the absence of such conditions, unconditional quantiles may only be bounded in general.

Other extensions. In Section S5 of the supplementary appendix we outline several additional extensions of the framework. The first one is to allow for a nonparametric propensity

¹²Specifically, when applying our approach one could parametrically specify the copulas of (U_0, V) and (U_1, V) given X , or alternatively specify the trivariate copula of (U_0, U_1, V) given X .

¹³A related though different extrapolation strategy is introduced in Brinch, Mogstad and Wiswall (2015), who rely on parametric restrictions on the marginal treatment effects functions.

score, instead of a parametric specification. The second one is the construction of a test statistic to test for the absence of sample selection. We also outline how to adapt the method to allow for some regressors to be endogenous (as in Chernozhukov and Hansen, 2005, 2006), and for outcomes to be partially censored (as in Powell, 1986).

5 Wages and labor market participation in the UK

In this section, we apply our method to measure market-level changes in wage inequality in the UK. Moreover, we compare wages of males and females in the UK at different quantiles, correcting for selection into work. Due to changes in employment rates, wage inequality for those at work may provide a distorted picture of market-level inequality. Our exercise decomposes actual changes in the aggregate wage distribution into different interpretable sources (selection and non-selection components). Our procedure could be standardized into building economic statistics, similar to other decomposition-based statistics such as price indices adjusted for changes in quality.

In this application, the latent variable Y^* represents the opportunity cost of working for each person, whether employed or not, *at* given employment rates. It is not a potential outcome in the conventional treatment-effect sense, because Y^* depends on the market price of skill, which may be affected by changes in participation rates. In order to account for equilibrium effects on skill prices, in Section 6 we also propose an extension of the method and we apply it to a counterfactual exercise.

5.1 Data and methodology

We use data from the Family Expenditure Survey (FES) from 1978 to 2000. To construct the sample, we closely follow previous work using these data: Gosling *et al.* (2000) and Blundell *et al.* (2003), who focus on males, and Blundell *et al.* (2007), who consider both males and females. We select individuals aged 23 to 59 who are not in full-time education, and drop observations for which education is not reported, or for which wages are missing but the individual is working. Hourly wages are constructed by dividing usual weekly pre-tax earnings by usual weekly hours worked. In addition, we drop the self-employed from the sample. We end up with 77,630 observations for males, and 89,848 observations for females.

During the period of analysis, wage inequality increased sharply in the UK. For example, in our sample, the logarithm of the 90/10 percentile ratio of male hourly wages increased

Table 1: Descriptive statistics (conditional on employment)

	Mean	Min	Max	q10	q50	q90
Males						
Married						
Log-wage	2.10	.172	4.30	1.56	2.06	2.71
Propensity score	.879	.021	1.00	.766	.893	.979
Single						
Log-wage	1.99	.319	4.28	1.45	1.95	2.58
Propensity score	.753	.259	1.00	.574	.765	.916
Females						
Married						
Log-wage	1.64	-.378	3.59	1.11	1.57	2.32
Propensity score	.681	.006	.998	.512	.699	.844
Single						
Log-wage	1.78	-.465	3.58	1.20	1.76	2.42
Propensity score	.718	.019	1.00	.475	.735	.933

Source: Family Expenditure Survey, 1978-2000.

Note: The propensity score is estimated using a probit model.

from .90 in 1978 to 1.34 in 2000. This is in line with previous evidence on wage inequality (Gosling *et al.*, 2000). Moreover, a comparison of mean log-wages between males and females shows a mean log-wage gap of .44 in 1978, and a mean gap of .30 in 2000. During the same period the overall employment rate of males fell from 92% to 80%. The mean employment rate of females also changed over the period, though not in a monotone way. This suggests that correcting for selection into employment might be important. We now use our approach to provide selection-corrected measures of wage inequality and gender wage gaps.

We use the quantile selection model to model log-hourly wages Y and employment status D . Our controls X include linear, quadratic, and cubic time trends, four cohort dummies (born in 1919-34, 1935-44, 1955-64, and 1965-77, the baseline category being 1945-54), two education dummies (end of schooling at 17 or 18, and end of schooling after 18), and 11 regional dummies. In addition, we include as regressors the marital status and the number of kids split by age categories (six dummies, from 1 year old to 17-18 years old). Our sample contains 75% of married men and 74% of married women.

We follow Blundell *et al.* (2003) and use their measure of potential out-of-work (welfare) income, interacted with marital status, as our excluded regressor B . This variable is con-

structed for each individual in the sample using the Institute of Fiscal Studies (IFS) tax and welfare-benefit simulation model. We estimate the propensity score using a probit model. In Table 1 we report several descriptive statistics on the distribution of log-wages, and on the distribution of the estimated propensity score, by gender and marital status. Out-of work income is a strong determinant of labor market participation. For example, in the sample of married (respectively, single) males the log-likelihood of the probit model of participation increases from $-21,454$ to $-20,438$ (resp., $-10,480$ to $-10,275$) when out-of-work income is added.

The main sources of variation in out-of-work income are the demographic composition of households (age, household size) and the housing costs that households face, as well as changes in policy over time. Our maintained assumption is that those determinants are exogenous to the latent wage equation, and the participation equation satisfies a monotonicity condition. Though not uncontroversial,¹⁴ out-of-work income provides a natural choice for an excluded variable in this context. Moreover, variations in out-of-work income over time are partly due to changes in policy, motivating the counterfactual analysis that we will present at the end of this section.

Implementation. We specify the copula $C(\cdot, \cdot; \rho)$ as a member of the one-parameter Frank family (Frank, 1979). We provide details on Frank copulas in Section S6 of the supplementary appendix. We let the copula parameter be gender- and marital-status specific, as both dimensions play an important role in potential out-of-work income. We will return to the choice of the copula below. In addition, to compute $\hat{\rho}$ in (15) we take $\tau_\ell = \ell/10$ for $\ell = 1, \dots, 9$, and $\varphi(\tau_\ell, Z_i) = \varphi(Z_i) = p(Z_i; \hat{\theta})$.¹⁵ Finally, we use grid search for computation of $\hat{\rho}$, and take 200 grid points.

¹⁴For example, as argued by Blundell *et al.* (2007), the way the out-of-work income variable operates may imply a positive correlation with potential wages, if individuals who earn more on the labor market have better housing, hence a higher out-of-work income. Kitagawa (2010) tests the validity of independence assumptions based on a discretization of X and B , and finds a rejection in 5 out of 16 covariates cells.

¹⁵When considering a two-parameter copula we take $p(Z_i; \hat{\theta})$ and $p(Z_i; \hat{\theta})^2$ as instrument functions. We also estimated the model with $\varphi(\tau_\ell, Z_i) = \sqrt{\tau_\ell(1 - \tau_\ell)} p(Z_i; \hat{\theta})$, in order to give more weight to central quantiles, and obtained very similar results. As already mentioned, here we do not attempt to address the question of efficient estimation of ρ .

5.2 Selection-corrected wage distributions

On the nine panels of Figure 1 we plot the evolution of the log-wage deciles for men (thick lines), and women (thin lines). The solid lines show the deciles of observed log-wages, conditional on employment. The dashed lines show the selection-corrected deciles, by gender. To compute the latter, we estimated the selection-corrected quantile regression coefficients using our method, and we then simulated the wage distribution using the method of Machado and Mata (2005), re-adjusting the percentile levels in order to correct for sample selection.¹⁶

Focusing first on male wages, we see that correcting for sample selection makes a strong difference at the bottom of the wage distribution. For example, at the 10% percentile male wages increased by 10% conditional on employment, while latent wages remained broadly flat. We also see sizable differences between latent and observed wages at the 20% and 30% percentiles. There are smaller differences in the middle and at the top of the distribution. In addition, differences across quantiles illustrate the sharp increase in male wage inequality in the UK over the period.

The results for male wages are consistent with low-skilled individuals being progressively driven out of the labor market. Our estimated copula has a rank correlation of $-.24$ for married males, and of $-.79$ for singles,¹⁷ which means that individuals with higher wages (higher U) tend to participate more (lower V). Thus, associated with the fall in participation over time, positive selection into employment implies that individuals at the bottom of the latent wage distribution tend to become increasingly non-employed. Selection into employment is stronger for singles than for married males. The 95% confidence intervals for the rank correlation coefficients are $(-.35, -.06)$ for married males, and $(-.84, -.42)$ for singles, respectively.¹⁸

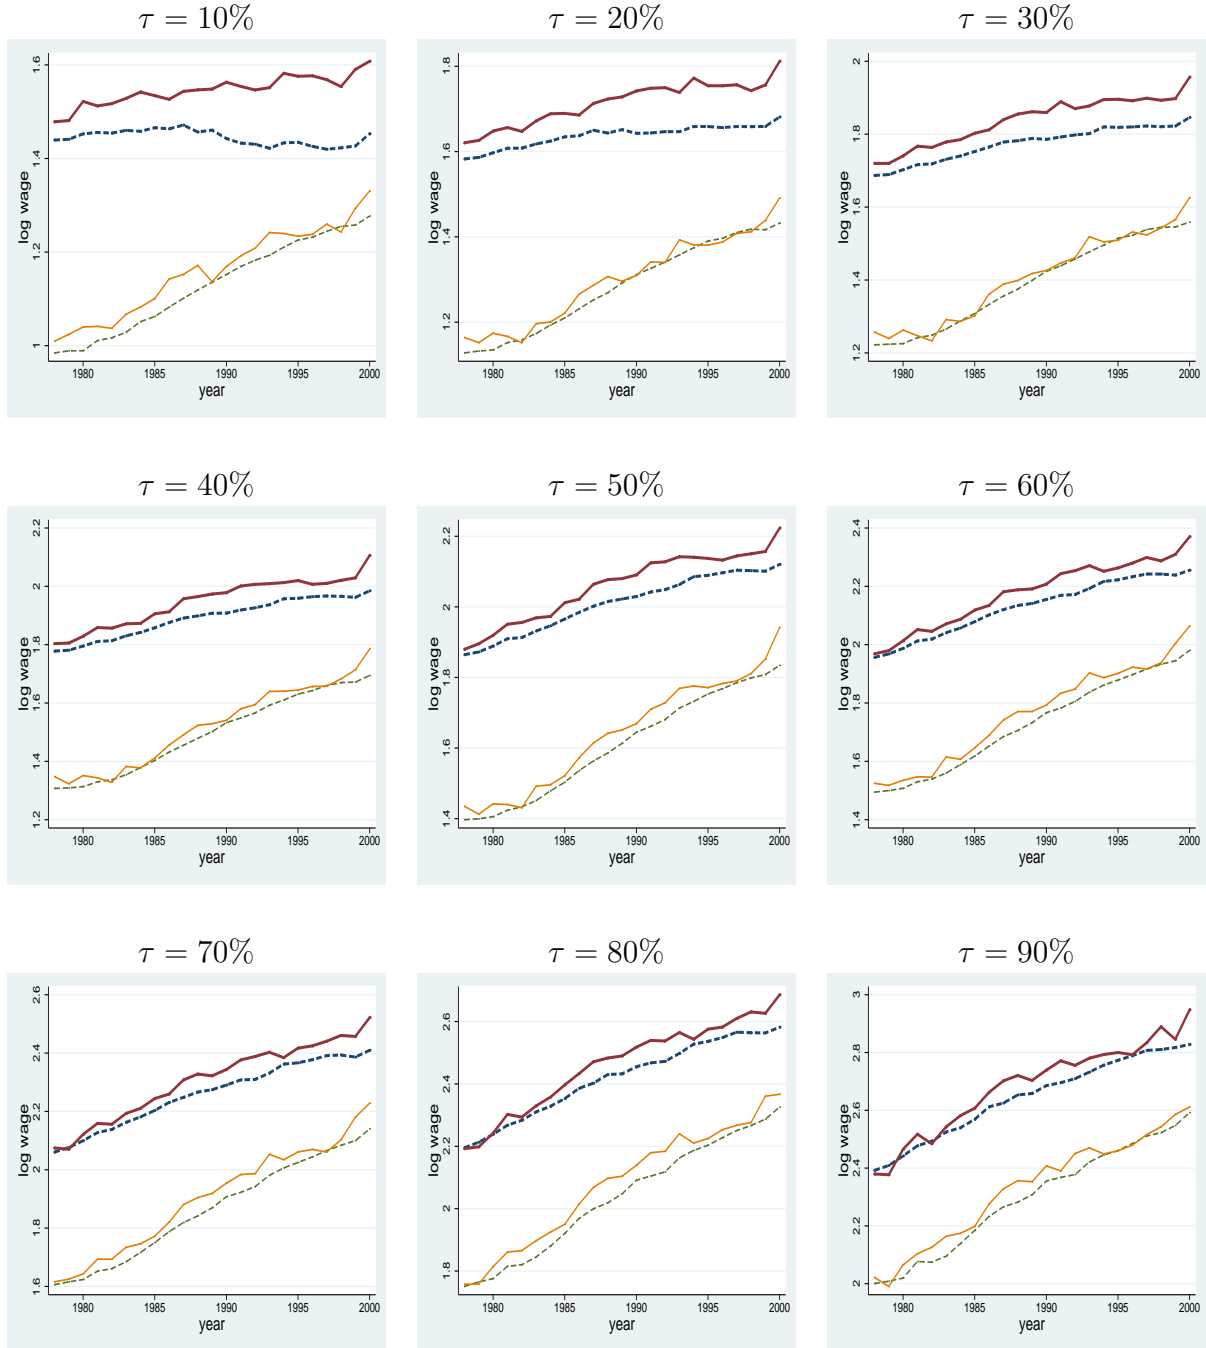
Looking now at female wages, we observe less difference between wages conditional on employment and latent wages. Indeed, we estimate a copula with rank correlation of $-.17$ for married females, and of $-.08$ for singles, suggesting that there is less positive selection into

¹⁶We jointly simulate wages and participation decisions as follows. For every individual, we draw $(U_i^{(m)}, V_i^{(m)})$, $m = 1, \dots, M$, from the relevant copula. Then we compute $Y_i^{(m)} = X_i \hat{\beta}_{U_i^{(m)}}$, and $D_i^{(m)} = \mathbf{1}\{V_i^{(m)} \leq \hat{p}(Z_i)\}$. Finally, we compute unconditional quantiles, either latent or conditional or participation, as empirical quantiles from the simulated data $(Y_i^{(m)}, D_i^{(m)})$. In practice we take $M = 20$, and we round τ in $\hat{\beta}_\tau$ to the closest percentile.

¹⁷The rank (or ‘‘Spearman’’) correlation of a copula C is given by: $12 \int_0^1 \int_0^1 uv dC(u, v) - 3$.

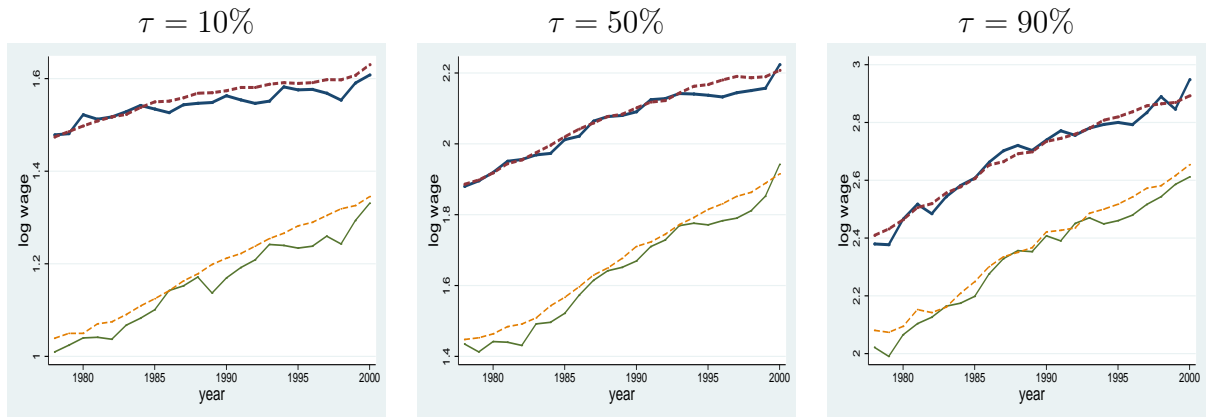
¹⁸We computed the confidence intervals using subsampling. Following Chernozhukov and Fernández-Val (2005) we chose the subsample size as a constant plus the square-root of the sample size, where the constant (≈ 1000) was taken to ensure reasonable finite sample performance of the estimator.

Figure 1: Wage quantiles, by gender



Note: FES data for 1978-2000. Quantiles of log-hourly wages, conditional on employment (solid lines) and corrected for selection (dashed). Male wages are plotted in thick lines (top lines in each graph), while female wages are in thin lines (bottom lines).

Figure 2: Fit to wage quantiles, by gender (employed individuals)



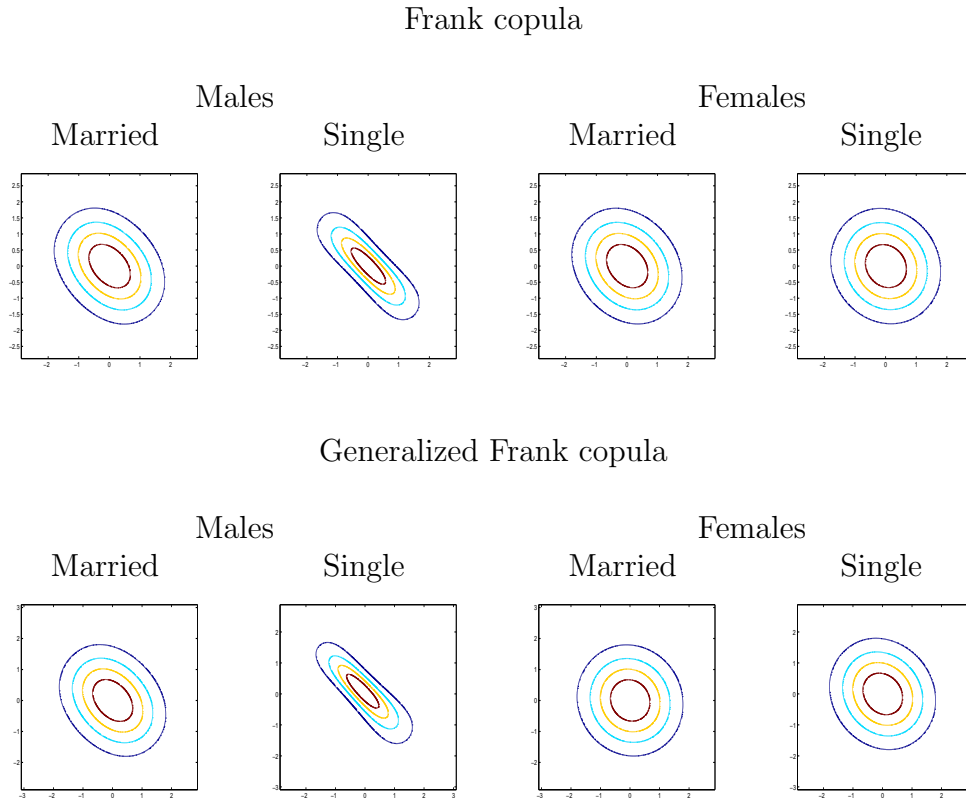
Note: FES data for 1978-2000. Quantiles of log-hourly wages conditional on employment, data (solid lines) and model fit (dashed). Male wages are plotted in thick lines (top lines in each graph), while female wages are in thin lines (bottom lines).

employment for women than for men. A tentative explanation could be that for females non-economic factors play a bigger role in participation decisions. The confidence intervals for the correlation coefficients are $(-.30, -.01)$ for married females, and $(-.24, .16)$ for singles.

As a result of this evolution, the selection-corrected gender wage gap tends to decrease over time. This is especially true at the bottom of the wage distribution. For example, at the 10% percentile, the difference in log wages between men and women decreases from 45% at the beginning of the period to 18% at the end. A comparable decrease can be seen at the 20% and 30% percentiles. Hence, correcting for sample selection magnifies the reduction in the wage gap in this part of the distribution. At the top of the distribution the gap seems to decrease less, from 39% to 24% at the 90% percentile.

Model fit. Figure 2 shows the model fit to the wage percentiles of employed workers. To predict wage percentiles, we simulated wages using our parameter estimates. The results show that the fit to wage quantiles is accurate at the top of the distribution for both genders. At the bottom of the distribution we observe some discrepancies, particularly for females. In addition, we estimated the model allowing the Frank copula parameter to vary with calendar time, on subsamples before and after 1990, in addition to gender and marital status (not reported). We found some evidence of increasingly positive selection into employment

Figure 3: Contour plots of the copula



Note: FES data for 1978-2000. Contour plots of the estimated copula. Negative correlation indicates positive selection into employment. The first row shows the Frank copula, while the second row shows the generalized Frank copula; see Section S6 of the supplementary appendix.

for females.¹⁹ The fit to the selected wage quantiles improved slightly. At the same time, quantiles of latent wages were comparable to the ones in Figure 1.

Choice of copula. We then investigate the robustness of our results to the choice of the copula. The symmetry properties of the Frank copula are apparent in the first two rows of Figure 3, which shows the contour plots of the copula densities that we estimated on the FES data.²⁰ As a specification check, we consider an encompassing two-parameter family,

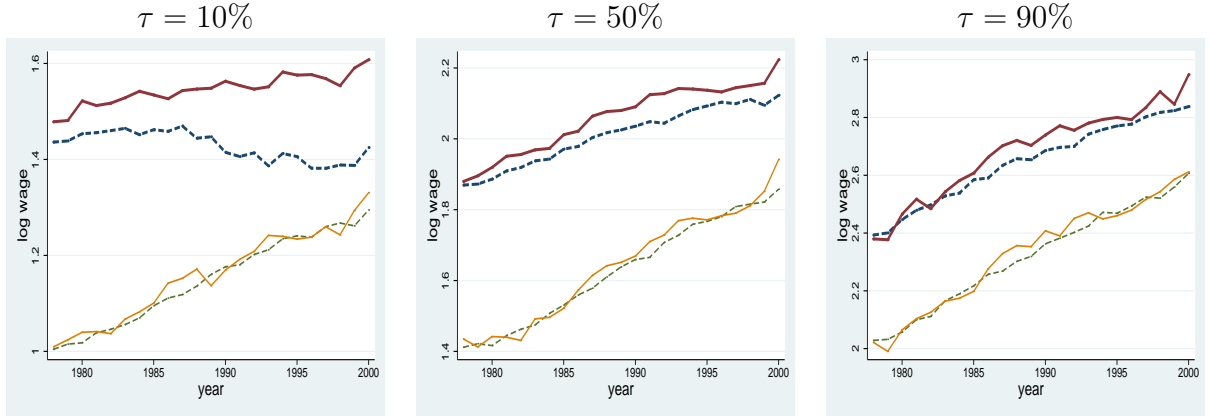
¹⁹On US data, Mulligan and Rubinstein (2008) document that women’s selection into participation shifted from being negative in the 1970s to being positive in the 1990s.

²⁰As a graphical convention (common in the literature on copulas), we plot the copula density by rescaling the margins so that they are standard normal. That is, if $C(u, v)$ denotes the copula, we plot the contours of:

$$(x, y) \mapsto \phi(x)\phi(y) \frac{\partial^2 C}{\partial u \partial v}(\Phi(x), \Phi(y)),$$

where ϕ and Φ denote the standard normal density and cdf, respectively.

Figure 4: Wage quantiles, by gender (generalized Frank copula)



Note: FES data for 1978-2000. Percentiles of log-hourly wages, conditional on employment (solid lines) and corrected for selection (dashed). Male wages are plotted in thick lines (top lines in each graph), while female wages are in thin lines (bottom lines).

which we call the “generalized Frank copula”. This family may capture different degrees of dependence in different regions of the (U, V) plane, as we explain in Section S6 of the supplementary appendix. The estimated copula densities in the generalized Frank family are shown in the last two rows of Figure 3. We see that, for both males and females, the differences between the estimated Frank and generalized Frank copulas are relatively small. Moreover, as shown by Figure 4, the quantiles of latent wages are quite similar for both genders when using a Frank or a generalized Frank copula.

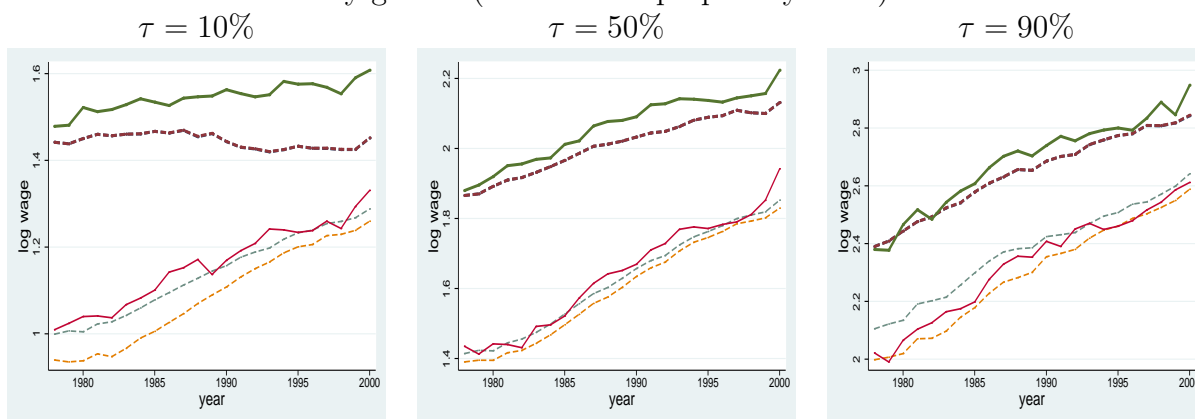
Lastly, we also estimated the model based on a Gaussian copula. With a Gaussian copula and Gaussian marginals the quantile selection model boils down to the Heckman (1979) model. Our approach makes it possible to combine a Gaussian copula with a non-Gaussian outcome distribution given by (12). The results of this specification (not reported) are very similar to the ones based on the Frank copula. In particular, the Spearman correlation coefficients of the estimated copulas are almost identical.²¹

Bounds estimates. As a further check of the influence of functional forms on the estimates, in Figure 5 we report estimates of the bounds derived in equation (19). On the top panel we plot bounds estimates by gender. We see that bounds on wage quantiles for males

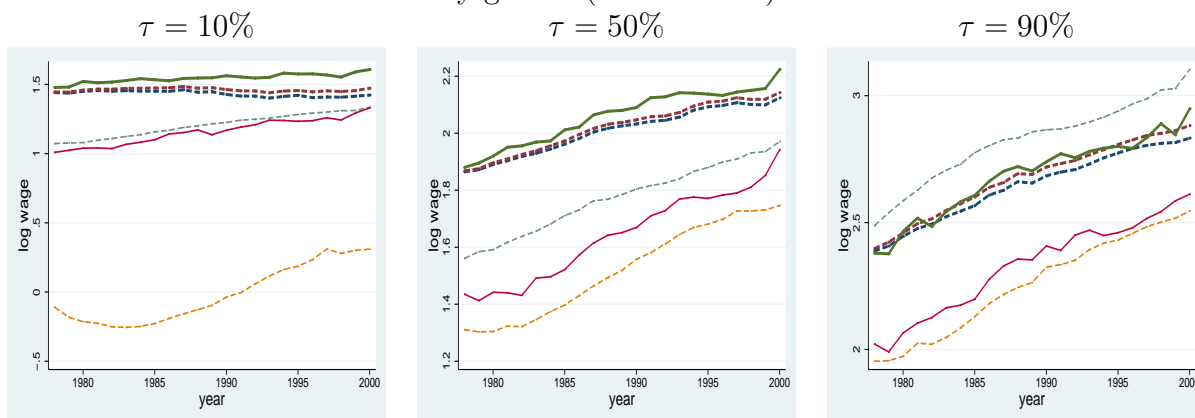
²¹Dependence of the copula on additional covariates could also be relevant. In unreported results we found that higher education, conditional on gender and marital status, tends to be associated with more positive selection into employment, particularly for females.

Figure 5: Estimated bounds on latent wage quantiles

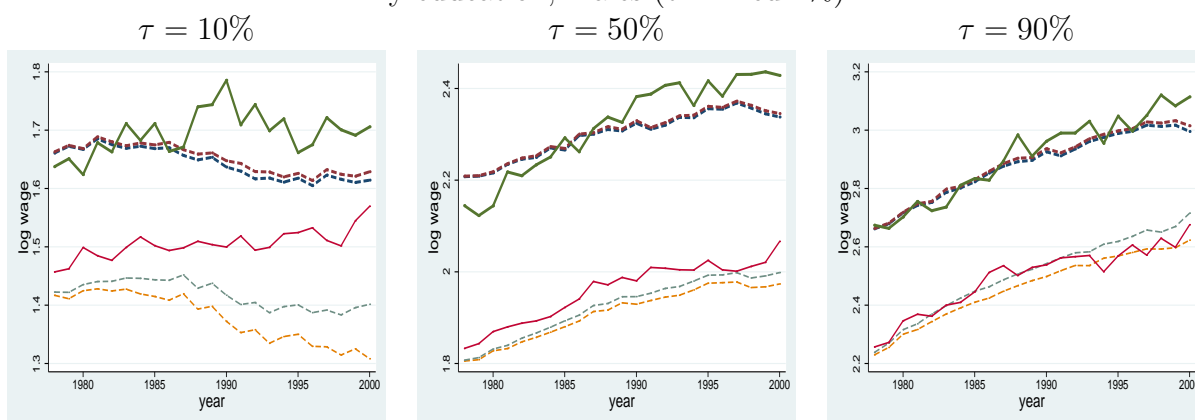
By gender (un-trimmed propensity score)



By gender (trimmed 1%)



By education, males (trimmed 1%)



Note: FES data for 1978-2000. Estimated bounds on quantiles of log-hourly wages (dashed). The solid lines show the quantiles conditional on employment. Top two panels: male wages are plotted in thick lines, female wages are in thin lines. Bottom panel: wages for high-school and college are plotted in thick lines, wages for statutory schooling are in thin lines.

(in dashed lines) are very close to each other. The bounds for females are wider, though still informative. However, the results for females are sensitive to the estimator of the supremum of the propensity score (\bar{p}_x) that we use. Larger participation rates are associated with smaller values of out-of-work income. On the middle panel of Figure 5 we report estimates of the bounds when trimming 1% of extreme observations in out-of-work income. We see that, while the results for males are very stable, those for females are very different, showing extremely wide bounds throughout the wage distribution. This reflects the fact that the selection problem is more severe for females, as their employment rates are lower.

Lastly, on the bottom panel of Figure 5 we compare the bounds, for males, for two education groups: statutory schooling (71% of the sample, in thin lines) and high-school and college (29%, in thick lines). We use a trimmed estimator of the supremum of the propensity score. We see that the bounds are narrow for more educated individuals, and that they are wider for the low educated whose employment rates are lower. We observe some evidence of an increase in the education gap over time, particularly at the median, although the evidence after correcting for selection is more mixed. The graphs also show evidence of an inequality increase within the education groups that we consider (similarly as in Blundell *et al.*, 2007).

6 Counterfactuals in the presence of equilibrium effects

In this last section we consider a simple equilibrium model of wage quantile functions and nonrandom selection into work as a flexible tool for examining changes in the distribution of wages over time. We show how the simplicity of linear quantiles can be essentially preserved while embedding wage functions in a model of human capital, employment decisions, and labor demand. We then use the model to recompute wage and employment distributions in a counterfactual scenario where potential out-of-work income is kept at its 1978 value.

6.1 Model and computation

We abstract from hours of work and dynamics. Let r_t^s be the skill price of a worker of education level s in time period t . Let also $h(s, x, u)$ be the amount of human capital of a worker with education (or “skill level”) s , observed characteristics x (such as cohort and gender), and unobserved ability u . The wage rate for an individual i of schooling level S_i in

period t is:

$$W_{it} = r_t^{S_i} \cdot h(S_i, X_{it}, U_{it}),$$

where there are two skill levels ($S_i \in \{1, 2\}$). Note that the human capital function h is time-invariant.²²

Letting $Z_{it} = (B_{it}, X_{it})$, the individual work decision is:

$$D_{it} = 1 \{r_t^{S_i} h(S_i, X_{it}, U_{it}) \geq W^R(S_i, Z_{it}, \eta_{it})\}.$$

Let $\bar{X}_{it} \equiv (S_i, X_{it})$. The log-human capital function and log-reservation wage are specified as: $\ln h(S_i, X_{it}, U_{it}) \equiv \bar{X}'_{it} \beta(U_{it})$, and: $\ln W^R(S_i, Z_{it}, \eta_{it}) \equiv \bar{X}'_{it} \gamma(\eta_{it}) + B'_{it} \varphi$, so that the participation decision is:

$$D_{it} = 1 \left\{ \bar{X}'_{it} \gamma(\eta_{it}) - \bar{X}'_{it} \beta(U_{it}) \leq \ln r_t^{S_i} - B'_{it} \varphi \right\} = 1 \{V_{it} \leq \bar{F}(\ln r_t^{S_i} - B'_{it} \varphi, \bar{X}_{it})\},$$

where the composite error $\bar{X}'_{it} (\gamma(\eta_{it}) - \beta(U_{it}))$ is assumed independent of Z_{it} given $\bar{X}_{it} = \bar{x}$ with cdf $\bar{F}(\cdot, \bar{x})$, and V_{it} is its uniform transformation. In practice, we approximate the propensity score by a single-index (probit) model of the form $F(\ln r_t^{S_i} - Z'_{it} \psi)$.

Using wage and participation equations, our quantile selection approach allows one to perform partial equilibrium counterfactual exercises where skill prices r_t^s are kept constant. In order to allow for equilibrium responses in skill prices, we now introduce a model for labor demand. See Heckman, Lochner and Taber (1998) and Lee and Wolpin (2006) for related approaches in dynamic structural settings.

Labor demand. Consider a one-sector economy with one physical capital input (which we assume fixed) and two types of human capital. We assume a standard aggregate production function: $F_t(L_t, K_t) = A_t L_t^\alpha K_t^{1-\alpha}$, where L_t is a CES aggregator of the human capital inputs: $L_t = \left[a_t H_{1t}^\phi + (1 - a_t) H_{2t}^\phi \right]^{1/\phi}$. If $\phi = 1$ the two labor skills are perfect substitutes, in which case an increase in the supply of one type of human capital does not affect the relative skill prices. The scope for equilibrium effects critically depends on the structure of production.

From the first-order conditions we obtain:

$$\ln \left(\frac{r_t^1}{r_t^2} \right) = \ln \left(\frac{a_t}{1 - a_t} \right) + (\phi - 1) \ln \left(\frac{H_{1t}}{H_{2t}} \right). \quad (21)$$

In Appendix B we discuss how to recover estimates of H_{1t} , H_{2t} , ϕ , and a_t from micro-data based on (21). In practice, due to weak identification from our time-series, we calibrate

²²This assumption is called the ‘‘proportionality hypothesis’’ in Heckman and Sedlacek (1985).

$\phi = .4$ using Card and Lemieux (2001)'s estimate on UK data. We take $\alpha = .6$ in the results below. We varied α between .4 and .8 and found small effects on the results.

Counterfactual equilibrium skill prices. Suppose we are interested in estimating the counterfactual equilibrium skill prices, $\ln \tilde{r}_t^s$ say, that would have prevailed under technology conditions in period t and the labor force composition or the welfare policy in some other period.

Equilibrium log skill prices satisfy the equations:

$$\ln r_t^s = \ln A_t + \ln \alpha + (1 - \alpha) \ln \left(\frac{K_t}{L_t} \right) + \ln a_{st} + (\phi - 1) \ln \left(\frac{H_{st}}{L_t} \right),$$

where $a_{1t} \equiv a_t$ and $a_{2t} \equiv 1 - a_t$. In addition, the labor supply equations imply:

$$H_{st}(r_t^s) = \sum_{S_i=s} F(\ln r_t^s - Z'_{it}\psi) \int_0^1 e^{\bar{X}'_{it}\beta(u)} dG[u, F(\ln r_t^s - Z'_{it}\psi); \rho], \quad s = 1, 2, \quad (22)$$

$$L_t = \left(a_{1t} [H_{1t}(r_t^1)]^\phi + a_{2t} [H_{2t}(r_t^2)]^\phi \right)^{1/\phi}. \quad (23)$$

The log-difference between observed and counterfactual skill prices is given by:

$$\ln \tilde{r}_t^s - \ln r_t^s = (1 - \alpha) \ln \left(\frac{L_t}{\tilde{L}_t} \right) + (\phi - 1) \left[\ln \left(\frac{\tilde{H}_{st}}{\tilde{L}_t} \right) - \ln \left(\frac{H_{st}}{L_t} \right) \right], \quad s = 1, 2, \quad (24)$$

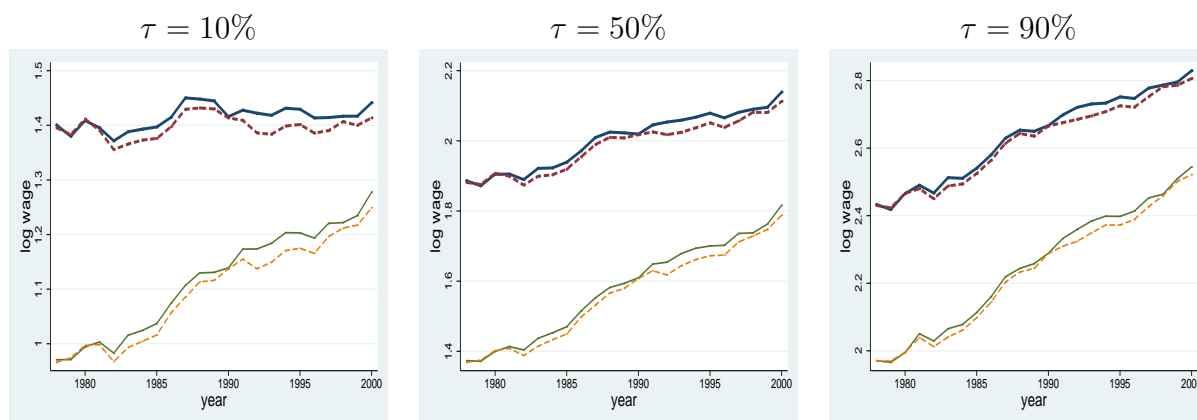
where the counterfactual skill aggregates \tilde{H}_{st} and \tilde{L}_t satisfy (22)-(23) at prices $(\tilde{r}_t^1, \tilde{r}_t^2)$. Note that capital (which is fixed) and neutral technical progress are common to both sets of prices and thus cancel out in (24).

Counterfactual log-skill prices $\ln \tilde{r}_t^1$ and $\ln \tilde{r}_t^2$ are then obtained as the solution to the two nonlinear equations in (24), subject to (22)-(23). This fixed-point problem depends on the following inputs: the parameters β , ψ , ρ , and r_t^s (estimated using our quantile selection method), the aggregate quantities H_{st} and L_t and the technological shocks a_t (estimated as explained in Appendix B), and the parameters ϕ and α (which we take from the literature). As starting value for the counterfactual \tilde{r}_t^s we take the estimated r_t^s , and we solve for the fixed point iteratively.

6.2 Results

Figure 6 shows the estimates of latent wage quantiles in two scenarios: when out-of-work income is as in the data (solid lines), and in a counterfactual scenario where out-of-work

Figure 6: Latent wage quantiles and counterfactual equilibrium latent wage quantiles, by gender



Note: FES data for 1978-2000. Quantiles of log-hourly wages corrected for selection. Latent wage quantiles (solid lines) and counterfactual general equilibrium latent wage quantiles (dashed). Male wages are plotted in thick lines (top lines in each graph), while female wages are in thin lines (bottom lines).

income is kept at its 1978 value (dashed). The specification that we use has some differences compared to the one in Figure 1. In particular, here the two education groups are college and non-college, the specification is pooled across genders, and controls are interacted with gender.²³ We present the results by gender.

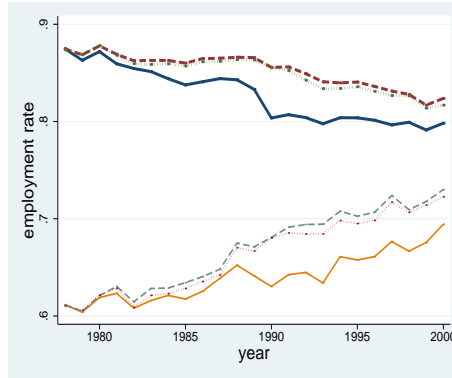
We see that accounting for general equilibrium responses tends to lower latent counterfactual quantiles throughout the distribution. This is due to the fact that in the counterfactual scenario out-of-work income is lower, thus increasing employment rates, and as a result pushing skill prices down. General equilibrium effects appear to be relatively small for both genders, although they seem more sizable at the bottom of the distribution.

Figure 7 shows actual employment rates (as predicted by the model), and employment rates in the partial equilibrium and general equilibrium counterfactuals. We see that in the counterfactual scenario employment rates tend to increase (dashed lines). The dampening effect on employment that comes from the general equilibrium response of skill prices is quantitatively small (dotted lines).

Lastly, Figure 8 shows the actual evolution of wages conditional on employment as predicted by the model (solid lines), and the evolution in the counterfactual scenario where

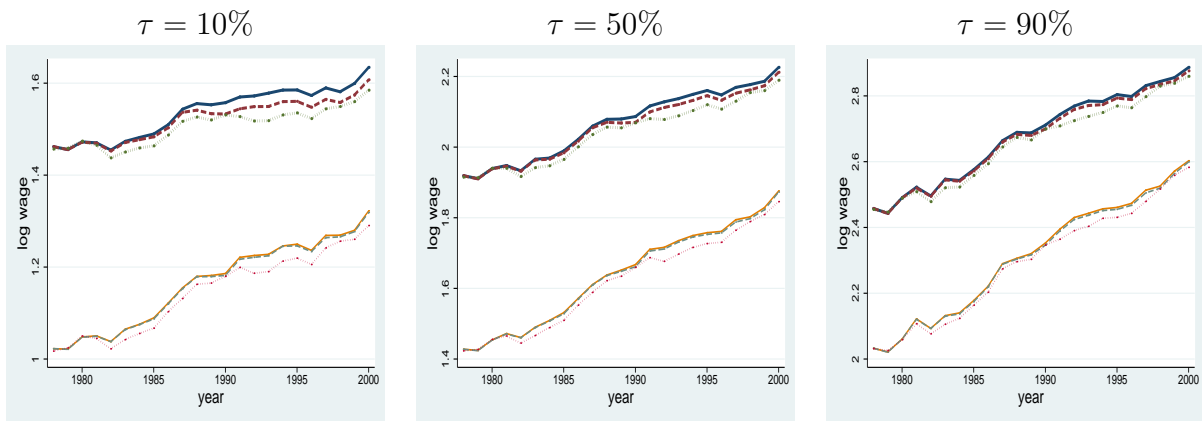
²³The fit of the model used in this subsection is shown in Section S7 of the supplementary appendix.

Figure 7: Employment (actual and counterfactual), by gender



Note: FES data for 1978-2000. Actual employment rate predicted by the model (solid lines), counterfactual employment rate at constant prices (dashed), and counterfactual employment rate at equilibrium prices (dotted). Male employment is plotted in thick lines (top lines), while female employment is in thin lines (bottom lines).

Figure 8: Wage quantiles conditional on employment (actual and counterfactual), by gender



Note: FES data for 1978-2000. Quantiles of log-hourly wages conditional on employment. Actual quantiles predicted by the model (solid lines), counterfactual quantiles in partial equilibrium (dashed), and counterfactual quantiles in general equilibrium (dotted). Male wages are plotted in thick lines (top lines in each graph), while female wages are in thin lines (bottom lines).

out-of-work income is kept at its 1978 value, with skill prices fixed (dashed) and with skill prices adjusting through general equilibrium (dotted). We see that, in the partial equilibrium counterfactual, wages of male workers tend to be lower at the bottom of the distribution, due to positive selection into employment. In addition, general equilibrium responses imply further reduction in wages. In the middle and at the top of the distribution, and for females, differences between actual and counterfactual evolution appear to be smaller.

7 Conclusion

We have presented a three-step method to correct quantile regression estimates for sample selection. In a first step, the parameters of the participation equation are estimated. In a second step, the parameters of the copula linking the percentile error of the outcome equation to the participation error are computed by minimizing a method-of-moments objective function. In a third step, quantile parameters are computed by minimizing a weighted check function, using a fast linear programming routine. The method provides a simple and intuitive way to compute selection-adjusted quantile parameters. Moreover, our application shows that such selection corrections for quantiles may be as empirically relevant as in the standard regression context of the popular Heckman (1979) sample selection model.

An important issue is the choice of the copula. An approach that treats the copula non-parametrically is conceptually attractive, for example a sieve approach based on conditional moment restrictions as in Chen and Pouzo (2009, 2012). It would be desirable to allow the copula to be unspecified, and to conduct inference on the identified set of quantile functions. The empirical application suggests that nonparametric bounds might be informative when selection is not too severe (as in the case of men in our application).

References

- [1] Abadie, A. (2003): “Semiparametric Instrumental Variable Estimation of Treatment Response Models,” *Journal of Econometrics*, 113(2), 231–263.
- [2] Abadie, A., J. Angrist, and G. Imbens (2002): “Instrumental Variables Estimates of the Effect of Subsidized Training on the Quantiles of Trainee Earnings,” *Econometrica*, 70(1), 91–117.
- [3] Ahn, H., and J. L. Powell (1993): “Semiparametric Estimation of Censored Selection Models with a Nonparametric Selection Mechanism,” *Journal of Econometrics*, 58, 3–29.
- [4] Albrecht, J., A. van Vuuren, and S. Vroman (2009): “Counterfactual Distributions with Sample Selection Adjustments: Econometric Theory and an Application to the Netherlands,” *Labour Economics*, 16(4), 383–396.
- [5] Angrist, J., V. Chernozhukov, and I. Fernández-Val (2006): “Quantile Regression Under Misspecification, with an Application to the U.S. Wage Structure,” *Econometrica*, 74, 539–563.
- [6] Arellano, M., and S. Bonhomme (2016): “Sample Selection in Quantile Regression: A Survey,” to appear in the *Handbook of Quantile Regression*.
- [7] Autor, D., L. Katz, and M. Kearney (2005): “Residual Wage Inequality: The Role of Composition and Prices,” NBER Working Paper 11628.
- [8] Blundell, R., H. Reed, and T. M. Stoker (2003): “Interpreting Aggregate Wage Growth: The Role of Labor Market Participation,” *American Economic Review*, 93(4), 1114–1131.
- [9] Blundell, R., A. Gosling, H. Ichimura, and C. Meghir (2007): “Changes in the Distribution of Male and Female Wages Accounting for Employment Composition Using Bounds,” *Econometrica*, 75, 323–364.
- [10] Bonhomme, S., and J. M. Robin (2009): “Assessing the Equalizing Force of Mobility Using Short Panels: France, 1990–2000,” *Review of Economic Studies*, 76(1), 63–92.
- [11] Brinch, C. N., M. Mogstad, and M. Wiswall (2015): “Beyond LATE with a Discrete Instrument,” to appear in the *Journal of Political Economy*.
- [12] Buchinsky, M. (1994): “Changes in the U.S. Wage Structure 1963 to 1987; An Application of Quantile Regressions,” *Econometrica*, 62, 405–458.
- [13] Buchinsky, M. (1998): “The dynamics of changes in the female wage distribution in the USA: a quantile regression approach,” *Journal of Applied Econometrics*, 13, 1–30.
- [14] Buchinsky, M. (2001): “Quantile Regression with Sample Selection: Estimating Women’s Return to Education in the US,” *Empirical Economics*, 26, 87–113.
- [15] Card, D., and T. Lemieux (2001): “Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis,” *Quarterly Journal of Economics*, 116, 705–746.
- [16] Carneiro, P., and S. Lee (2009): “Estimating Distributions of Potential Outcomes using Local Instrumental Variables with an Application to Changes in College Enrollment and Wage Inequality,” *Journal of Econometrics*, 149, 191–208.
- [17] Chamberlain, G. (1993): “Quantile Regressions, Censoring and the Structure of Wages,” in C. Sims (ed.), *Advances in Econometrics: Proceedings of the 6th World Congress in Barcelona*, Vol. I (Cambridge: Cambridge University Press).

- [18] Chaudhuri, P. (1991): “Nonparametric Estimates of Regression Quantiles and Their Local Bahadur Representation”, *Annals of Statistics*, 19(2), 760–777.
- [19] Chen, S., and S. Khan (2003): “Semiparametric Estimation of a Heteroskedastic Sample Selection Model,” *Econometric Theory*, 19(6), 1040–1064.
- [20] Chen, X., and D. Pouzo (2009): “Efficient Estimation of Semiparametric Conditional Moment Models with Possibly Nonsmooth Residuals,” *Journal of Econometrics*, 152, 46–60.
- [21] Chen, X., and D. Pouzo (2012): “Estimation of Nonparametric Conditional Moment Models with Possibly Nonsmooth Moments,” *Econometrica*, 80(1), 277–322.
- [22] Chernozhukov, V., and I. Fernández-Val (2005): “Subsampling Inference on Quantile Regression Processes,” *Sankhya: The Indian Journal of Statistics*, 67(2), 253–276.
- [23] Chernozhukov, V., I. Fernández-Val, and A. Galichon (2010): “Quantile and Probability Curves without Crossing,” *Econometrica*, 78(3), 1093–1125.
- [24] Chernozhukov, V., I. Fernández-Val, and B. Melly (2013): “Inference on Counterfactual Distributions,” *Econometrica*, 81(6), 2205–68.
- [25] Chernozhukov, V., and C. Hansen (2005): “An IV model of Quantile Treatment Effects,” *Econometrica*, 73, 245–262.
- [26] Chernozhukov, V., and C. Hansen (2006): “Instrumental Quantile Regression Inference for Structural and Treatment Effect Models,” *Journal of Econometrics*, 132, 491–525.
- [27] Chernozhukov, V., and H. Hong (2003): “An MCMC Approach to Classical Estimation”, *Journal of Econometrics*, 115, 293–346.
- [28] Das, M., W. K. Newey, and F. Vella (2003): “Nonparametric Estimation of Sample Selection Models,” *Review of Economic Studies*, 70, 33–58.
- [29] D’Haultfoeuille, X., and P. Février (2015): “Identification of Nonseparable Triangular Models With Discrete Instruments,” *Econometrica*, 3, 1199–1210.
- [30] DiNardo, J., N. Fortin, and T. Lemieux (1996): “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach,” *Econometrica*, 64, 1001–1044.
- [31] Donald, S. G. (1995): “Two-Step Estimation of Heteroskedastic Sample Selection Models,” *Journal of Econometrics*, 65(2), 347–380.
- [32] Firpo, S., N. Fortin, and T. Lemieux (2011): “Decomposition Methods in Economics,” in O. Ashenfelter and D. Card, eds., *Handbook of Economics*, Amsterdam: North-Holland, Vol. IV.A, 1–102.
- [33] Frank, M. J. (1979): “On the Simultaneous Associativity of $F(x, y)$ and $x + y - F(x, y)$,” *Aequationes Mathematicae*, 19, 194–226.
- [34] Fréchet, M. (1951): “Sur les Tableaux de Corrélacion dont les Marges sont Données,” *Ann. Univ. Lyon Sér. 3*, 14, 53–77.
- [35] Gosling, A., S. Machin, and C. Meghir (2000): “The Changing Distribution of Male Wages in the UK,” *Review of Economic Studies*, 67, 635–666.

- [36] Gronau, R. (1974): “Wage Comparison - A Selectivity Bias,” *Journal of Political Economy*, 82, 1119–1143.
- [37] Heckman, J. J. (1974): “Shadow Prices, Market Wages and Labour Supply,” *Econometrica*, 42, 679–694.
- [38] Heckman, J. J. (1979): “Sample Selection Bias as a Specification Error,” *Econometrica*, 47, 153–161.
- [39] Heckman, J. J. (1990): “Varieties of Selection Bias,” *The American Economic Review*, 80, 313–318. (Papers and Proceedings of the Hundred and Second Annual Meeting of the American Economic Association.)
- [40] Heckman, J. J., L. Lochner, and C. Taber (1998): “Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Earnings with Heterogeneous Agents,” *Review of Economic Dynamics*, 1, 1–58.
- [41] Heckman, J. J., and G. Sedlacek (1985): “Heterogeneity, Aggregation and Market Wage Functions; an Empirical Model of Self Selection into the Labor Market,” *Journal of Political Economy*, 98, 1077–1125.
- [42] Heckman, J. J., J. Smith, and N. Clements (1997): “Making the Most out of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts,” *Review of Economic Studies*, 64, 487–535.
- [43] Heckman, J. J., and E. Vytlacil (2005): “Structural Equations, Treatment Effects, and Econometric Policy Evaluation,” *Econometrica*, 73(3), 669–738.
- [44] Huber, M., and B. Melly (2015): “A Test of the Conditional Independence Assumption in Sample Selection Models,” *Journal of Applied Econometrics*, 30(7), 1144–1168.
- [45] Imbens, G., and J. Angrist (1994): “Identification and Estimation of Local Average Treatment Effects,” *Econometrica*, 62, 467–476.
- [46] Imbens, G., and D. Rubin (1997): “Estimating Outcome Distributions for Compliers in Instrumental Variables Models,” *Review of Economic Studies*, 64, 555–574.
- [47] Joe, H. (1997): *Multivariate Models and Dependence Concepts*. London: Chapman & Hall.
- [48] Judd, K. (1998): *Numerical Methods in Economics*, MIT Press. Cambridge, London.
- [49] Katz, L. F., and K. M. Murphy (1992): “Changes in the Wage Structure, 1963-1987: Supply and Demand Factors,” *Quarterly Journal of Economics*, 107, 35–78.
- [50] Kitagawa, T. (2009): “Identification Region of the Potential Outcome Distributions under Instrument Independence,” Cemmap Working Paper CWP30/09.
- [51] Kitagawa, T. (2010): “Testing for Instrument Independence in the Selection Model,” Unpublished Manuscript.
- [52] Kline, P., and A. Santos (2013): “Sensitivity to Missing Data Assumptions: Theory and an Evaluation of the US Wage Structure,” *Quantitative Economics*, 4(2), 231–267.
- [53] Koenker, R. (2005): *Quantile Regression*, Econometric Society Monograph Series, Cambridge: Cambridge University Press.

- [54] Koenker, R., and G. Bassett (1978): “Regression Quantiles,” *Econometrica*, 46, 33–50.
- [55] Koenker, R., and Z. Xiao (2002): “Inference on the quantile regression process,” *Econometrica*, 70, 1583–1612.
- [56] Lee, D., and K. I. Wolpin (2006): “Intersectoral Labor Mobility and the Growth of the Services Sector,” *Econometrica*, 74(1), 1–46.
- [57] Lee, L.-F. (1983): “Generalized Econometric Models with Selectivity,” *Econometrica*, 51(2), 507–512.
- [58] Machado, J. A. F., and J. Mata (2005): “Counterfactual Decomposition of Changes in Wage Distributions using Quantile Regression,” *Journal of Applied Econometrics*, 20, 445–465.
- [59] Manski, C. F. (1994): “The Selection Problem” in *Advances in Econometrics*, Sixth World Congress, Vol 1, ed. by C. Sims. Cambridge, U.K.: Cambridge University Press, 143–170.
- [60] Manski, C. F. (2003): *Partial Identification of Probability Distributions*, Springer Series in Statistics, Vol. 12. Berlin: Springer Verlag.
- [61] Mulligan, C., and Y. Rubinstein (2008): “Selection, Investment, and Women’s Relative Wages Over Time,” *Quarterly Journal of Economics*, 123(3), 1061–1110.
- [62] Neal, D. (2004): “The Measured Black-White Wage Gap among Women Is too Small,” *Journal of Political Economy*, 112, 1–28.
- [63] Nelsen, R. B. (1999): *An Introduction to Copulas*. New-York: Springer Verlag.
- [64] Olivetti, C., and B. Petrongolo (2008): “Unequal Pay or Unequal Employment? A Cross-Country Analysis of Gender Gaps,” *Journal of Labor Economics*, 26(4), 621–654.
- [65] Picchio, M., and C. Mussida (2010): “Gender Wage Gap: a Semi-Parametric Approach with Sample Selection Correction,” IZA Discussion Paper No. 4783. Bonn: IZA.
- [66] Politis, D. N., J. P. Romano, and M. Wolf (1999): *Subsampling*. Springer, New York.
- [67] Powell, J. L. (1986): “Censored Regression Quantiles,” *Journal of Econometrics*, 32, 143–155.
- [68] Sancetta, A., and S. Satchell (2004): “The Bernstein Copula and its Applications to Modeling and Approximations of Multivariate Distributions,” *Econometric Theory*, 20, 535–562.
- [69] Smith, M. D. (2003): “Modelling Sample Selection Using Archimedean Copulas,” *The Econometrics Journal*, 6, 99–123.
- [70] Torgovitsky, A. (2015): “Identification of Nonseparable Models Using Instruments with Small Support,” *Econometrica*, 83(3), 1185–1197.
- [71] Vella, F. (1998): “Estimating Models with Sample Selection Bias: A Survey,” 33(1), 127–169.
- [72] Vuong, Q., and H. Xu (2014): “Counterfactual Mapping and Individual Treatment Effects in Nonseparable Models with Discrete Endogeneity,” Unpublished Manuscript.
- [73] Vytlacil, E. (2002): “Independence, Monotonicity, and Latent Index Models: An Equivalence Result,” *Econometrica*, 70, 331–341.

APPENDIX

A Proofs on identification

Proof of Lemma 1. Equation (6) is a direct application of (5), using the fact that by A3 both G_x and $F_{Y|D=1,Z}$ are strictly increasing in their first argument.

To show the second part, let $x \in \mathcal{X}$ and let G_x satisfy (6). Pick a $z_x \in \mathcal{Z}_x$, and define:

$$F_{Y^*|X}(y|x) \equiv G_x^{-1} (F_{Y|D=1,Z}(y|z_x), p(z_x)).$$

For all (z, y) in the support of (Z, Y) given $X = x$ we have:

$$\begin{aligned} G_x (F_{Y^*|X}(y|x), p(z)) &= G_x (G_x^{-1} (F_{Y|D=1,Z}(y|z_x), p(z_x)), p(z)) \\ &= F_{Y|D=1,Z} (F_{Y|D=1,Z}^{-1} (F_{Y|D=1,Z}(y|z_x) | z_x) | z) \\ &= F_{Y|D=1,Z} (y|z), \end{aligned}$$

where we have used (6) to obtain the second equality.

Proof of Proposition 1. Let us start with *i*). Evaluating (6) at $z_1 = z$ and $z_2 = z_x$, and noting that $G_x^{-1}(\tau, 1) = \tau$, we have that $G_x(\tau, p(z)) = F_{Y|D=1,Z} (F_{Y|D=1,Z}^{-1}(\tau | z_x) | z)$. Hence G_x is identified. The identification of q then comes from (5) and Assumption A3.

Let us now suppose *ii*). Let G_x and \tilde{G}_x satisfy model (1)-(3), and let Assumptions A1 to A4 hold. Then, by (6) we have:

$$G_x [G_x^{-1}(\tau, p_2), p_1] - \tilde{G}_x [\tilde{G}_x^{-1}(\tau, p_2), p_1] = 0, \quad \text{for all } (p_1, p_2) \in \mathcal{P}_x \times \mathcal{P}_x.$$

Hence, for each $\tau \in (0, 1)$, the function:

$$(p_1, p_2) \mapsto G_x [G_x^{-1}(\tau, p_2), p_1] - \tilde{G}_x [\tilde{G}_x^{-1}(\tau, p_2), p_1],$$

which is real analytic, is zero on a product of two open neighborhoods. As a result it is zero everywhere on $(0, 1) \times (0, 1)$, and evaluating it at $p_2 = 1$ leads to:

$$G_x(\tau, p_1) - \tilde{G}_x(\tau, p_1) = 0, \quad \text{for all } p_1 \in (0, 1).$$

Hence G_x and \tilde{G}_x coincide on $(0, 1) \times (0, 1)$. This implies that G_x , and hence q (as in the first part of the proof), are identified.

Proof of Proposition 2. For clarity here we denote $x = (\tilde{x}, 1)$, where \tilde{x} contains all covariates but the constant term. Let also $\tilde{\beta}$ contain all β coefficients except the intercept. Finally, let $\tilde{q}^d(x, p) = x' \beta_{G^{-1}(\tau, p; \rho)}$. For $\tilde{q}^d(x, p)$ to be additive in \tilde{x} and p , it is necessary and sufficient that $\tilde{\beta}_{G^{-1}(\tau, p; \rho)}$ does not depend on p . This happens only if $\tilde{\beta}_\tau$ does not depend on τ , or if $G^{-1}(\tau, p; \rho)$ does not depend on p . In the second case, U and V are independent on the relevant support. For example, if the conditional support of $p(Z)$ contains 1, taking $p = 1$ implies that $G^{-1}(\tau, p; \rho) = \tau$ for all (τ, p) , so U and V are independent.

B Estimating the elasticity of substitution

The estimation of equation (21) is based on time series aggregate data. We use the microdata to construct time series of the relevant aggregates. The time series of the log-relative price of skill $\ln(\hat{r}_t^1/\hat{r}_t^2)$ is obtained from the estimation of the wage functions. Time series of relative aggregate labor supplies can be estimated by aggregation of individual units of human capital of employed workers:

$$\ln\left(\frac{\widehat{H_{1t}}}{\widehat{H_{2t}}}\right) = \ln\sum_{S_i=1}\frac{W_{it}}{\hat{r}_t^1} - \ln\sum_{S_i=2}\frac{W_{it}}{\hat{r}_t^2} = \ln\left(\frac{\sum_{S_i=1}W_{it}}{\sum_{S_i=2}W_{it}}\right) - \ln(\hat{r}_t^1/\hat{r}_t^2).$$

The log ratio of factor-specific productivities $\ln\left(\frac{a_t}{1-a_t}\right)$ is allowed to vary over time to capture skill-biased technical change. It is specified as a trend $\lambda(t)$ plus an unobservable shock ε_t . The equation to be estimated is therefore:

$$\ln(\hat{r}_t^1/\hat{r}_t^2) = \lambda(t) + (\phi - 1)\ln(\hat{H}_{1t}/\hat{H}_{2t}) + \varepsilon_t. \quad (\text{B1})$$

This equation was estimated on aggregate US data by Katz and Murphy (1992), who obtained $\hat{\phi} = 0.3$. A comparable estimate on UK data in Card and Lemieux (2001) is $\hat{\phi} = 0.4$. We then estimate a_t as:

$$\hat{a}_t \equiv \Lambda\left(\ln(\hat{r}_t^1/\hat{r}_t^2) - (\hat{\phi} - 1)\ln(\hat{H}_{1t}/\hat{H}_{2t})\right),$$

where $\Lambda(r) = \exp(r)/(1 + \exp(r))$.

Finally, note that the explanatory variable $\ln(\hat{H}_{1t}/\hat{H}_{2t})$ is likely to be correlated with ε_t in (B1), in which case OLS estimates are inconsistent. Natural instrumental variables would be aggregates (by skill) of labor supply shifters such as potential out-of-work welfare income.