



Gender wage inequality: new evidence from penalized expectile regression

Marina Bonaccolto-Töpfer¹ · Giovanni Bonaccolto²

Received: 10 September 2021 / Accepted: 12 January 2023

© The Author(s) 2023

Abstract

The Machado-Mata decomposition building on quantile regression has been extensively analyzed in the literature focusing on gender wage inequality. In this study, we generalize the Machado-Mata decomposition to the expectile regression framework, which, to the best of our knowledge, has never been applied in this strand of the literature. In contrast, in recent years, expectiles have gained increasing attention in other contexts as an alternative to traditional quantiles, providing useful statistical and computational properties. We flexibly deal with high-dimensional problems by employing the Least Absolute Shrinkage and Selection Operator. The empirical analysis focuses on the gender pay gap in Germany and Italy. We find that depending on the estimation approach (i.e. expectile or quantile regression) the results substantially differ along some regions of the wage distribution, whereas they are similar for others. From a policy perspective, this finding is important as it affects conclusions about glass ceiling and sticky floors.

Keywords Expectile regression · Gender pay gap · Quantile regression · Penalized estimation

JEL Classification J31 · J16 · J45 · J51

1 Introduction

According to labor economic theory, group differences in pay for individuals with similar characteristics should not exist in competitive markets. However, we observe *ceteris paribus* wage differentials for many groups. The most prominent wage differential is probably the male wage premium or gender pay gap; see Blau and Kahn (2017) for an overview. Several

✉ Marina Bonaccolto-Töpfer
marina.toepfer@unige.it

Giovanni Bonaccolto
giovanni.bonaccolto@unikore.it

¹ Department of Economics, University of Genova, Via Francesco Vivaldi 5, 16126, Genova, Italy

² School of Economics and Law, Kore University of Enna, Cittadella Universitaria, 94100, Enna, Italy

tools can be adopted to estimate this wage gap along with its components. Many contributions in the literature estimated the gap along the wage distribution (e.g. Firpo et al. 2009; Machado and Mata 2005), building on the quantile regression model introduced by Koenker and Bassett (1978). Quantile regression is widely employed in the literature as an appealing extension of the Ordinary Least Squares (OLS) approach (Fitzenberger et al. 2013). Indeed, on the one hand, it is a semi-parametric method and, therefore, in contrast to the standard OLS approach, it does not require distributional assumptions on the error term. On the other hand, it goes beyond a simple expectation providing more information about the relationships of the involved variables along their conditional distributions.

In this study, we take into account the fact that different econometric tools may lead to substantially different results, increasing the risk of implementing wrong policy actions. Expectile regression represents an effective alternative to quantile regression to study the impact of a set of covariates on the entire distribution of a given response variable. Nevertheless, to the best of our knowledge, this method is almost absent in the labor economics literature. In contrast, expectiles have recently received increasing attention in other research areas, such as financial econometrics and operational research (see, among others, Bellini et al. 2021; Bonaccolto et al. 2022; Giacometti et al. 2021). We fill this gap by employing expectile regression to estimate and decompose gender wage inequalities into a characteristics (explained) and a coefficients (unexplained) part, along the entire wage distribution. Further, we compare these estimates to those obtained from the quantile regression method, to identify wage gaps that are not detected employing only quantiles.

Expectiles may appear unfamiliar for academics and practitioners in the field of gender wage inequality. However, they can be directly interpreted using their own inherent properties, or by way of their relationship to both quantiles and OLS (Philipps 2021). Indeed, the term ‘expectile’ has probably been suggested as a combination of ‘expectation’ and ‘quantile’ (Bellini and Di Bernardino 2017). First, we refer to the original definition given by Newey and Powell (1987), according to which expectiles are the minimizers of the asymmetric least squares loss function. Interestingly, when the expectile level $\theta \in (0, 1)$ is equal to $1/2$, the resulting expectile coincides with the expected value of the variable of interest. Therefore, expectiles can be interpreted as an asymmetric generalization of the mean (Bellini and Di Bernardino 2017). As for the connection with quantiles, there exists a functional mapping from expectiles to quantiles, which allows to estimate quantiles by least squares using expectiles. In general, expectiles correspond to the quantiles of a transformed distribution (Jones 1994). For the most common distributions, expectiles are closer to the centre of the distribution than the corresponding quantiles. Typically, the quantile and expectile curves intersect in a unique point: the centre of symmetry of a symmetric distribution (Bellini and Di Bernardino 2017).

Expectile regression provides a set of relevant advantages (see, among others, Newey and Powell 1987; Efron 1991; Jones 1994; Yao and Tong 1996; Taylor 2008; Yang and Zou 2015; Bellini and Di Bernardino 2017). First, expectile regression is computationally simple, building on an asymmetric least squares loss function, which is differentiable everywhere. In contrast, the check loss function characterizing the quantile regression model is not everywhere differentiable, so that the underlying optimization routine might require a set of restrictions affecting the computational efficiency. This issue becomes critical in high-dimensional problems. Second, it is possible to define conditional quantiles (and conditional distributions) as a function of expectiles, given the one-to-one mapping between quantiles and expectiles. Therefore, we could compute quantiles from expectiles, exploiting

the computational advantages behind the latter. Third, expectiles have a more global dependence on the form of the distribution. Altering the shape of the upper tail of the response variable's distribution does not change the quantiles of the lower tail, but it does affect all expectiles (Taylor 2008). As a result, expectiles respond more readily to extreme cases. Finally, expectile curves are typically smoother than the ones derived from quantiles.

Following Stahlschmidt et al. (2014), we apply the method of Machado and Mata (2005) to both the quantile and expectile regression frameworks. Indeed, in addition to quantile regression, expectile regression may be also employed to estimate the conditional distribution of the outcome variable. The latter derives from the fact that – as in case of quantiles – the estimated conditional expectile function also represents a consistent estimator of the population expectile function and may describe the entire conditional distribution (Newey and Powell 1987; Taylor 2008; Stahlschmidt et al. 2014). A comparison between expectile and quantile regressions often reveals that neither approach is uniformly superior compared to the other (e.g. Yang and Zou 2015). As a result, we seek to draw more robust conclusions by combining the pros and cons of each method. This exercise could be a useful tool for policy implications concerning wage differentials between specific groups and conclusions about the existence of e.g. sticky floors or glass ceiling.¹

As a second contribution, we make our expectile model flexible to be used in high-dimensional problems. For this purpose we add the ℓ_1 -norm penalty characterizing the Least Absolute Shrinkage and Selection Operator (LASSO) introduced by Tibshirani (1996) to the expectile loss function. LASSO is an effective tool to identify accurate model specifications for specific expectile or quantile levels. That is, we apply different model specifications at different points of the wage distribution based on this regularization technique. Furthermore, LASSO reduces potential omitted variable bias given the data at hand. To the best of our knowledge, this is the first study that estimates and decomposes wage gaps based on penalized expectile models. We stress the fact that our penalized model in the spirit of Machado and Mata (2005) relies on an intuitive decomposition (explained and unexplained part) and—above all—allows for an unconditional interpretation, whereas quantile and expectile treatment effects require complex double selection methods (Belloni et al. 2017; Chernozhukov et al. 2018; Kallus et al. 2019). Further, LASSO allows us to exploit the advantageous properties of machine learning when predicting sets of wages (counterfactual and empirical). Indeed, in Machado-Mata decompositions, counterfactual and empirical distributions are predicted. Note that we retrieve the coefficient estimates from post-penalization regressions in order to avoid over-shrinkage issues (Hastie et al. 2009, 2015). In order to underline the relevance of expectile- and quantile-specific model selection, as a robustness check, we compare the performance of the full model specification in Blau and Kahn (2017) with expectile- and quantile-specific model specifications. The full specification of Blau and Kahn (2017) can be considered a state-of-the art specification for augmented Mincer-type wage models.

As stated above, the existing literature on gender pay gaps mainly focused on quantile regressions since the last 20 years (e.g. Albrecht et al. 2003; Arulampalam et al. 2007; Fitzenberger et al. 2013; Castagnetti and Giorgetti 2019). This literature found evidence for glass ceiling and/or sticky floors (e.g. for Sweden or Germany, respectively Albrecht et al. 2003; Collischon 2019). So far, data-driven or machine learning methods such as the double

¹In this paper, we follow (p. 171; Arulampalam et al. 2007) in defining glass ceiling to exist “[...] if the 90th percentile wage gap exceeds the reference wage gap by at least two percentage points”. Similarly, (p. 171; Arulampalam et al. 2007) identify sticky floors “[...] if the 10th percentile wage gap exceeds the reference wage gap by at least two percentage points”.

robust LASSO procedure have rarely been used in applied economic research (exceptions are e.g. Knaus et al. 2020; Bach et al. 2018; Brunori and Neidhöfer 2021; Bonaccolto-Töpfer and Briel 2022; Wunsch and Strittmatter 2021). For instance, Bach et al. (2018) estimated individual-specific gender pay gaps, while Bonaccolto-Töpfer and Briel (2022) focused on model selection, finding that using different model specifications at different points of the distribution affects the estimated gender pay gaps. Consequently, flexible model specifications matter. Similarly, Wunsch and Strittmatter (2021) found substantially lower unexplained gender pay gaps when using more flexible specifications of the wage equation.

We employ two different datasets in our empirical analysis: i) the German Socio-Economic Panel (SOEP) 2010-2017; and ii) the Italian survey PLUS created by the Institute of Development of Vocational Training of Workers (ISFOL) 2010-2016. Both datasets include a broad set of control variables (at least 63 in our case). In such a framework, regularization techniques like LASSO turn out to be particularly useful for model selection. The results obtained suggest that, depending on the underlying estimation method (expectile or quantile regression), the coefficients effect substantially differs. As a consequence, inter-quantile gaps and, thus, policy conclusions concerning glass ceiling and sticky floors change significantly. This finding holds particularly for Germany. In case of the characteristics part, we find no marked differences and, thus, estimation results that are robust to both quantile and expectile regression. However, also in this case, we find differences in the tails what translates to different conclusions about sticky floors for Germany.

The paper is organized as follows. Section 2 describes our estimation strategy. Section 3 presents the data set used for the empirical analysis. Empirical results are given in Section 4. In Section 5, we check whether and to what extent the results change when using the full specification of Blau and Kahn (2017) as a robustness exercise. Finally, Section 6 concludes.

2 Estimation strategy

We outline in this section the estimation approach. First, we define and discuss both quantile and expectile regressions (Section 2.1). Second, we focus on variable selection in high-dimensional problems (Section 2.2). Finally, we describe the decomposition approach (Section 2.3).

2.1 From OLS to quantile and expectile regression

Let y_i be the log of hourly wage of individual i , for $i = 1, \dots, N$. A standard specification of the corresponding wage equation takes the following form:

$$y_i = \mathbf{x}_i \boldsymbol{\gamma}' + u_i \quad (1)$$

where $\boldsymbol{\gamma} = [\gamma_1 \ \gamma_2 \ \dots \ \gamma_k]$ is a k -dimensional row vector in which γ_1 is the intercept, whereas $\gamma_2, \dots, \gamma_k$ are slope parameters, \mathbf{x}_i is an $1 \times k$ vector which includes the value of one (as first entry) along with a set of $k - 1$ control variables observed for the i^{th} individual, and u_i is the error term.

The regression model defined in Eq. 1 allows us to estimate the impact of \mathbf{x}_i on the conditional expected value of y_i . Nevertheless, such an impact is not necessarily constant along the conditional distribution of y_i . For instance, we might observe different effects depending on whether we focus on individuals with lower or higher values of y_i . The standard linear model in Eq. 1 does not capture these potential heterogeneous effects, preventing us

from obtaining accurate results. We can overcome this shortcoming by adopting the quantile regression method introduced by Koenker and Bassett (1978). This method allows us to estimate the conditional θ^{th} quantile of y_i , that we denote as $Q_\theta(y_i|\mathbf{x}_i)$, with $\theta \in (0, 1)$. By doing so, we provide a picture about the relationships between the response variable y_i and the covariates in \mathbf{x}_i along the entire distribution of y_i . Specifically, we estimate the following model:

$$Q_\theta(y_i|\mathbf{x}_i) = \mathbf{x}_i \boldsymbol{\beta}'_\theta \quad (2)$$

using a large set of $\theta \in (0, 1)$ values.

It is important to highlight the fact that the parameters in Eq. 2 depend on θ . As a result, we obtain different estimates according to the different regions of the conditional distribution of y_i . The method introduced by Koenker and Bassett (1978) allows us to estimate $\boldsymbol{\beta}_\theta$ by minimizing the following loss function:

$$L(\boldsymbol{\beta}_\theta) = \frac{1}{N} \sum_{i=1}^N \rho_\theta(y_i - \mathbf{x}_i \boldsymbol{\beta}'_\theta) \quad (3)$$

where $\rho_\theta(u_i) = u_i(\theta - \mathbb{I}_{\{u_i < 0\}})$, $u_i = (y_i - \mathbf{x}_i \boldsymbol{\beta}'_\theta)$, and $\mathbb{I}_{\{\cdot\}}$ is an indicator function which takes the value of one if the condition into braces is true, and the value of zero otherwise.

A major contribution of this study is the application of the expectile regression method introduced by Newey and Powell (1987) to study the relationships between \mathbf{x}_i and y_i along the different regions of the conditional distribution of y_i . Expectile regression represents a relevant alternative to quantile regression to extend the standard OLS approach. Both quantile and expectile regressions have received considerable attention in the literature, that often reveals that neither approach is uniformly superior compared to the other (e.g. Yang and Zou 2015). As a result, we seek to draw more robust conclusions by combining the pros and cons of each method.

On the one hand, quantile estimates are more robust to outliers or extreme observations. On the other hand, expectile regression provides other relevant advantages, as highlighted by several contributions in the literature; see, among others, Newey and Powell (1987), Efron (1991), Jones (1994), Yao and Tong (1996), Taylor (2008), and Yang and Zou (2015), Bellini and Di Bernardino (2017) and Furno and Vistocco (2018). We summarize some of them as follows. First, expectile regression builds on an asymmetric least squares loss function that is differentiable everywhere, providing greater computational efficiency, especially in high-dimensional problems. In contrast, the loss function characterizing the quantile regression method is not differentiable everywhere. Second, it is possible to define conditional quantiles (and conditional distributions) as a function of expectiles. Indeed, there exists a one-to-one mapping between quantiles and expectiles. As a result, we can estimate quantiles from expectiles, exploiting the computational advantages provided by the latter. Third, expectiles have a more global dependence on the form of the distribution. As highlighted by Taylor (2008), changing the shape of the upper tail of the distribution of y_i does not change the quantiles of the lower tail, but affects all expectiles. Therefore, expectiles are more sensitive to extreme observations, which potentially convey important information related, for instance, to tail events. Fourth, the estimated expectile curve is smoother than the one derived from quantiles, leading to finer estimates for multiple $\theta \in (0, 1)$ values.

After discussing the motivations that prompted us to use expectile regression, we now present the model we focus on. We still focus on a linear specification, defined as:

$$\mu_\theta(y_i|\mathbf{x}_i) = \mathbf{x}_i \boldsymbol{\delta}'_\theta \quad (4)$$

where $\mu_\theta(y_i | \mathbf{x}_i)$ denotes the θ^{th} expectile of y_i conditional on \mathbf{x}_i , with $\theta \in (0, 1)$, whereas δ_θ is estimated by minimizing the following asymmetric least squares loss function (Newey and Powell 1987):

$$L(\delta_\theta) = \frac{1}{N} \sum_{i=1}^N \left| \theta - \mathbb{I}_{\{y_i < \mathbf{x}_i \delta'_\theta\}} \right| [y_i - \mathbf{x}_i \delta'_\theta]^2 \quad (5)$$

2.2 Variable selection in high-dimensional problems

A relevant point in our study concerns the selection of the control variables to insert into \mathbf{x}_i . In order to increase the informative content of our model, we take into account non-standard (e.g. past periods of unemployment or part-time experience) as well as ambiguous potential control variables (e.g. having children or the number of children), in addition to the ones which are typically used in the related literature (e.g. schooling and labor market experience). We prefer a specification containing a large set of control variables, even though some of them may turn out to be non-significant, as we employ a method, that we describe below, which automatically makes such a selection from the data at hand. Therefore, our empirical strategy does not exclude *a priori* control variables which may potentially improve the accuracy of the resulting outcome, while the coefficients of irrelevant covariates are automatically set equal to zero.

We stress that, in high-dimensional problems, where the number of parameters to estimate with standard approaches represents a challenging issue, the accumulation of estimation errors becomes a critical problem, especially when employing highly correlated regressors. On the one hand, a large number of regressors typically implies overfitting issues. On the other hand, the estimates would suffer from omitted variable bias when using a restricted subset of covariates. We deal with the curse of dimensionality using a well-known machine learning technique for model selection; that is, the Least Absolute Shrinkage and Selection Operator (LASSO) introduced by Tibshirani (1996).

In recent years, LASSO has become a widely used tool not only for standard linear regressions, but also for quantile and expectile models (e.g. Koenker 2005; Li and Zhu 2008; Belloni and Chernozhukov 2011; Liao et al. 2019). LASSO builds on the ℓ_1 -norm penalty function, that allows us to penalize the absolute size of the estimated coefficients (Hastie et al. 2009). Therefore, starting from the quantile regression model defined in Eq. 2, we estimate β_θ by minimizing the following penalized loss function:

$$L(\beta_\theta) = \frac{1}{N} \sum_{i=1}^N \rho_\theta(y_i - \mathbf{x}_i \beta'_\theta) + \lambda \sum_{j=2}^k |\beta_{j,\theta}| \quad (6)$$

where $\beta_{j,\theta}$ is the j^{th} entry of β_θ , for $j = 2, \dots, k$, so that the intercept $\beta_{1,\theta}$ is not penalized, whereas $\lambda > 0$ is the tuning parameter which governs the intensity of the penalization.

Likewise, we estimate the parameters of the expectile regression model defined in Eq. 4 by minimizing the following loss function:

$$L(\delta_\theta) = \frac{1}{N} \sum_{i=1}^N \left| \theta - \mathbb{I}_{\{y_i < \mathbf{x}_i \delta'_\theta\}} \right| [y_i - \mathbf{x}_i \delta'_\theta]^2 + \eta \sum_{j=2}^k |\delta_{j,\theta}| \quad (7)$$

where $\delta_{j,\theta}$ is the j^{th} entry of δ_θ , whereas $\eta > 0$ is the tuning parameter.

We stress the fact that the impact of LASSO in Eqs. 6 and 7 depends on the tuning parameters λ and η , respectively. The greater λ and η are, the greater the number of coefficients that approach zero. The choice of the optimal values of λ and η is then a critical point, as it determines the sparsity of the resulting solutions. We select the optimal values of λ and η in Eqs. 6 and 7, respectively, by employing the 5-fold cross-validation technique, which is commonly used in applied machine learning. Indeed, this regularization parameter selection method is flexible and easy to understand and implement, providing accurate results (Hastie et al. 2009). Furthermore, cross-validation provides the advantage of being flexible to be used for any penalized regression model, regardless of the specification of both the objective and penalty functions. This means that we select the optimal values of λ and η by employing the same method.

As we said above, LASSO is widely used because it possesses important properties. Nevertheless, it also suffers from some limitations. For instance, it typically provides biased estimates, over-shrinking the retained variables (Hastie et al. 2009, 2015, Fan and Li 2001). An effective solution to the over-shrinkage issue is the post-LASSO procedure (see, among others, Belloni and Chernozhukov 2011; Hautsch et al. 2014; Bonaccollo 2021). Starting from the expectile estimation, we implement the post-LASSO procedure as follows. In a first step, we minimize the loss function defined in Eq. 7, and discard the regressors whose coefficients are, in absolute value, sufficiently close to zero: i.e. $|\delta_{j,\theta}| \leq \eta$, for $j = 2, \dots, k$, where η is a given threshold. We then define a new $1 \times s$ vector \mathbf{x}_i^* (where $s \leq k$), which includes, in addition to the value of one as first entry (similar to \mathbf{x}_i , to take into account the role of the intercept $\delta_{1,\theta}$), only the covariates that are LASSO-selected from Eq. 7. That is, the ones which have a relevant impact on $M_\theta(y_i|\mathbf{x}_i)$, satisfying the condition $|\delta_{j,\theta}| > \eta$, for $j = 2, \dots, k$. In a second step, we estimate the coefficients corresponding to the selected covariates from the following (non-penalized) minimization problem:

$$\arg \min_{\delta_\theta^* \in \mathbb{R}^s} \frac{1}{N} \sum_{i=1}^N \left| \theta - \mathbb{I}_{\{y_i < \mathbf{x}_i^* (\delta_\theta^*)'\}} \right| \left[y_i - \mathbf{x}_i^* (\delta_\theta^*)' \right]^2 \tag{8}$$

whereas the coefficients of the regressors in \mathbf{x}_i that are excluded from \mathbf{x}_i^* (i.e. the ones which are not LASSO-selected in the first step) are set equal to zero.

By doing so, we use LASSO as a selection variable tool in a first step, whereas the final coefficients are computed, in a second step, from a non-penalized minimization problem. The post-LASSO procedure provides relevant improvements. For instance, as highlighted by Belloni and Chernozhukov (2011) and Hautsch et al. (2014), the post-LASSO method outperforms both the standard LASSO and the standard quantile regression, which suffer from over-shrinkage and over-identification problems, respectively.² We implement the same procedure described above to estimate the quantile model. We then minimize the loss function in Eq. 6 and discard the regressors whose absolute coefficients approach zero: $|\beta_{j,\theta}| \leq \eta$, for $j = 2, \dots, k$. The coefficients of the covariates which are not LASSO-selected are set equal to zero. In contrast, we compute the coefficients of the LASSO-selected regressors (that we still denote as \mathbf{x}_i^* to simplify the notation) from the following minimization problem:

$$\arg \min_{\beta_\theta^* \in \mathbb{R}^s} \frac{1}{N} \sum_{i=1}^N \rho_\theta \left[y_i - \mathbf{x}_i^* (\beta_\theta^*)' \right] \tag{9}$$

²We show in a simulation study that this holds also in our framework (see Appendix C).

We conclude this section providing some details about the empirical setup. We estimate the expectile regression model (with and without the ℓ_1 -norm penalty) by employing the R package ‘gcdnet’. In contrast, we use the R package ‘quantreg’ to estimate the quantile regression model (with and without the ℓ_1 -norm penalty). As for the LASSO-selection, we set the η threshold equal to $1e-08$ for both expectile and quantile estimations. By doing so, we exclude from the post-LASSO estimation those covariates whose impact is null or negligible.

2.3 Decomposition approach

We describe in this section the decomposition procedure for both expectile and quantile estimations. The intuition behind this approach builds on the Machado-Mata decomposition; see Machado and Mata (2005) and Fortin et al. (2011) for additional details. A main advantage of this decomposition is that it allows for an unconditional interpretation. Our extension consists in estimating the expectile- and quantile-based wage equations separately for men and women:

$$\begin{aligned} Q_{\theta} (y_i^g | \mathbf{x}_i^g) &= \mathbf{x}_i^g (\boldsymbol{\beta}_{\theta}^g)' \\ \mu_{\theta} (y_i^g | \mathbf{x}_i^g) &= \mathbf{x}_i^g (\boldsymbol{\delta}_{\theta}^g)' \end{aligned}$$

with $g = \{f, m\}$, where ‘ f ’ and ‘ m ’ stand for women and men, respectively.

Specifically, the pairs of parameters $(\boldsymbol{\beta}_{\theta}^f, \boldsymbol{\delta}_{\theta}^f)$ and $(\boldsymbol{\beta}_{\theta}^m, \boldsymbol{\delta}_{\theta}^m)$ are separately estimated from the data observed, respectively, for women—i.e. y_i^f and \mathbf{x}_i^f —and men—i.e. y_i^m and \mathbf{x}_i^m —by employing the post-LASSO procedure described in Section 2.2. By doing so, we circumvent the usage of complex double robust LASSO estimation, as we do not require to identify a treatment effect (e.g. gender).

The Machado and Mata (2005) decomposition combined with our machine-learning technique can be summarized as follows:

1. Randomly draw θ_j , with $j = 1, \dots, 5000$, from the uniform distribution $U[0, 1]$;
2. For each θ_j , estimate the coefficients $\boldsymbol{\delta}_{\theta_j}^g$ and $\boldsymbol{\beta}_{\theta_j}^g$ of the expectile and quantile models defined in Eqs. 4 and 2, respectively, employing the post-LASSO approach using the original female and male datasets, respectively;
3. Randomly draw, with replacement, 5000 individuals from group f and get their characteristics, that we insert into the $5000 \times k$ matrix $\mathbf{x}^{r,f}$. Likewise, randomly draw, with replacement, 5000 individuals from group m and get their characteristics, that we insert into the $5000 \times k$ matrix $\mathbf{x}^{r,m}$;
4. By employing: i) the coefficients $\hat{\boldsymbol{\delta}}_{\theta_j}^g$ and $\hat{\boldsymbol{\beta}}_{\theta_j}^g$ obtained from step 2; and ii) the characteristics $\mathbf{x}^{r,g}$ obtained from step 3, generate the following 5000×1 vectors of predicted wages:

$$\begin{aligned} &\bullet \mu_{\theta_j} (\hat{\mathbf{y}}^g | \mathbf{x}^{r,g}) = \mathbf{x}^{r,g} (\hat{\boldsymbol{\delta}}_{\theta_j}^g)' \\ &\bullet \mu_{\theta_j} (\tilde{\mathbf{y}}^f | \mathbf{x}^{r,f}) = \mathbf{x}^{r,f} (\hat{\boldsymbol{\delta}}_{\theta_j}^m)' \\ &\bullet Q_{\theta_j} (\hat{\mathbf{y}}^g | \mathbf{x}^{r,g}) = \mathbf{x}^{r,g} (\hat{\boldsymbol{\beta}}_{\theta_j}^g)' \\ &\bullet Q_{\theta_j} (\tilde{\mathbf{y}}^f | \mathbf{x}^{r,f}) = \mathbf{x}^{r,f} (\hat{\boldsymbol{\beta}}_{\theta_j}^m)' \end{aligned}$$

The empirical distribution functions of $\mu_{\theta_j} (\tilde{\mathbf{y}}^f | \mathbf{x}^{r,f})$ and $Q_{\theta_j} (\tilde{\mathbf{y}}^f | \mathbf{x}^{r,f})$ are the counterfactual distributions—i.e. what women would have earned if they were paid like men

(Castagnetti and Giorgetti 2019)—that we estimate using post-penalized expectile and quantile regressions, respectively. Therefore, following Machado and Mata (2005) and Fortin et al. (2011), Castagnetti and Giorgetti (2019) and Bonaccollo-Töpfer et al. (2022), we adopt the male wage structure as reference category. Note that the above estimated expectile and quantile functions are consistent for their respective population counterparts (Bassett and Koenker 1982; 1986; Newey and Powell 1987; Machado and Mata 2005; Stahlschmidt et al. 2014). Similar to Machado and Mata (2005), we use usual summary statistics to measure the resulting changes. Let $\alpha(\cdot)$ be one such statistic (we use the mean in our empirical analysis), we compute, for each $j = 1 \dots, 5000$, the following quantities derived from the expectile estimation:

1. $\alpha [\mu_{\theta_j} (\hat{\mathbf{y}}^m | \mathbf{x}^{r,m})] - \alpha [\mu_{\theta_j} (\hat{\mathbf{y}}^f | \mathbf{x}^{r,f})]$, which measures the total gap;
2. $\alpha [\mu_{\theta_j} (\hat{\mathbf{y}}^m | \mathbf{x}^{r,m})] - \alpha [\mu_{\theta_j} (\tilde{\mathbf{y}}^f | \mathbf{x}^{r,f})]$, which measures the characteristics effect;
3. $\alpha [\mu_{\theta_j} (\tilde{\mathbf{y}}^f | \mathbf{x}^{r,f})] - \alpha [\mu_{\theta_j} (\hat{\mathbf{y}}^f | \mathbf{x}^{r,f})]$, which measures the coefficients effect.

Likewise, we compute the following quantities: i) $\alpha [Q_{\theta_j} (\hat{\mathbf{y}}^m | \mathbf{x}^{r,m})] - \alpha [Q_{\theta_j} (\hat{\mathbf{y}}^f | \mathbf{x}^{r,f})]$; ii) $\alpha [Q_{\theta_j} (\hat{\mathbf{y}}^m | \mathbf{x}^{r,m})] - \alpha [Q_{\theta_j} (\tilde{\mathbf{y}}^f | \mathbf{x}^{r,f})]$; and iii) $\alpha [Q_{\theta_j} (\tilde{\mathbf{y}}^f | \mathbf{x}^{r,f})] - \alpha [Q_{\theta_j} (\hat{\mathbf{y}}^f | \mathbf{x}^{r,f})]$ from the quantile estimation. We repeat the steps described above 100 times to get the bootstrapped standard errors of the gaps along with the corresponding confidence intervals.

3 Data description

We implement the methods described in Section 2 on two different countries: Germany and Italy. When focusing on Germany, the empirical analysis builds on the survey waves 2010–2017 of the German Socio-Economic Panel (SOEP) (Wagner et al. 2008; SOEP 2019). We use the period 2010–2017 in order to reduce heterogeneity in the pay gaps over time (see Fig. A.1 in the Appendix). We consider only full-time employees in West Germany in order to form a more homogeneous sample of employees. In total, we use 78 potential control variables.

As for Italy, we use the survey-years 2010, 2011, 2014, 2016 from the Participation, Labor, Unemployment Survey (PLUS) from ISFOL (Corsetti et al. 2014). Again, we use the period 2010–2016 in order to reduce heterogeneity in the pay gaps over time (see Fig. A.1 in the Appendix). In line with the German sample, we restrict the sample to full-time employees and use 63 potential control variables. Due to few observations, we exclude armed forces and activities of extraterritorial organizations and bodies in both samples.

As we aim at obtaining robust estimates of the wage gaps and its components, it is important to control for all relevant control variables to predict male and female wages, respectively. Therefore, the potential set of control variables from which we LASSO-select the relevant regressors is pivotal.

Table 1 shows descriptive statistics for selected potential control variables by gender for Germany, while Table 2 presents the corresponding statistics for Italy. Over the last seven years, we find a statistically significant GPG of 25% (log approximation) in Germany. According to Eurostat (2018), the German gender pay gap amounts to 22% in 2017, while the corresponding EU average lies at 16%. Individuals in our sample are on average 45 years old. Women outperform men in terms of schooling. Men have, on average, more

Table 1 Descriptive statistics by gender (selected controls) – Germany

Variable	(1)	(2)	(3)	(4)	(5)
	Men		Women		
	Mean	Std.Dev.	Mean	Std.Dev.	Difference
<i>Panel (a): Dependent Variable</i>					
Log hourly wage	2.992	0.459	2.747	0.435	0.245***
<i>Panel (b): Selected Potential Control Variables</i>					
Age (in years)	44.63	9.955	44.72	9.771	- 0.090
Education (in years)	12.75	2.838	12.95	2.684	- 0.200***
Academic degree (dummy)	0.285	0.451	0.260	0.439	0.025***
Labor market experience (in years)	20.86	10.75	11.47	9.522	9.390***
Tenure (in years)	13.43	10.92	11.20	9.932	2.230***
Small firm	0.156	0.363	0.232	0.422	- 0.076***
Medium firm	0.261	0.439	0.264	0.441	- 0.003
Permanent contract	0.906	0.292	0.873	0.332	0.033***
Metropolitan area (dummy)	0.0375	0.190	0.0385	0.192	- 0.001***
Public sector (dummy)	0.223	0.417	0.348	0.476	- 0.125***
Married (dummy)	0.730	0.444	0.572	0.495	0.158***
Child (dummy)	0.354	0.478	0.374	0.484	- 0.020***
Child \leq 6 years (dummy)	0.188	0.391	0.112	0.316	0.076***
Income from assets (dummy)	0.256	0.436	0.248	0.432	0.008**
Father college (dummy)	0.462	0.499	0.484	0.500	- 0.022***
Mother college (dummy)	0.482	0.500	0.502	0.500	- 0.020***
Part-time experience (in years)	0.882	2.493	7.538	7.510	- 6.656***
Past years of unemployment (in years)	0.483	1.510	0.570	1.517	- 0.087***
Observations	30,277		30,091		60,368

Notes: ‘Small Firm’ equals one if the firm has at most 19 employees, zero otherwise. ‘Medium Firm’ equals one if the firm has between 20 and 199 employees, zero otherwise. Reported differences are based on a regression of a male dummy on the respective selected variable. *, ** and *** denote significance at the 10%-, 5%- and 1%-level, respectively. Robust standard errors (clustered at the individual level) are used. Further potential controls include federal state, survey-year, sectoral and occupational dummies. For classification of sectors, we use NACE (level 1), while for classification of occupations, we use ISCO88 (1-digit). Moreover, we add dummies for parents education (*Realschule* and *Abitur*), for having a migration background, dependent children (≤ 18), studied Science, Technology, Engineering or Mathematics (STEM), being employed in a big firm and quadratic polynomials of age, experience and tenure. Further, we add an interaction term between being married and having income from assets. Source: SOEP data v34

than nine years more labor market experience and stay more than two years longer with the same employer and have more often a permanent contract. We observe more women than men in smaller firms or living in metropolitan areas. Men are more often married but we observe more mothers than fathers in our sample. Yet, we have more fathers with young children than mothers. Further, men have, on average, more income from assets, less part-time experience as well as less unemployment spells compared to women.

Moving to the Italian sample (see Table 2), we also find a statistically significant, though smaller, wage gap amounting to 7% (log approximation). In Italy the gender wage gap is

Table 2 Descriptive statistics by gender (selected controls) – Italy

Variable	(1)	(2)	(3)	(4)	(5)
	Men		Women		
	Mean	Std.Dev.	Mean	Std.Dev.	Difference
<i>Panel (a): Dependent Variable</i>					
Log hourly wage	2.150	0.330	2.081	0.280	0.069***
<i>Panel (b): Selected Potential Control Variables</i>					
Age (in years)	43.84	13.12	41.69	12.09	2.150***
Education (in years)	12.81	2.717	13.66	2.455	- 0.850***
Honors degree (dummy)	0.0288	0.167	0.0619	0.241	- 0.033***
Labor market experience (in years)	22.56	13.24	19.35	12.52	3.210***
Tenure (in years)	16.72	12.59	14.43	11.74	2.290***
Small firm (dummy)	0.570	0.495	0.708	0.455	- 0.138***
Medium firm (dummy)	0.257	0.437	0.190	0.393	0.067***
Permanent contract (dummy)	0.889	0.314	0.861	0.346	0.028***
Metropolitan area (dummy)	0.188	0.391	0.212	0.409	- 0.024***
Public sector (dummy)	0.345	0.476	0.478	0.500	- 0.133***
Married (dummy)	0.518	0.500	0.512	0.500	0.006
Child (dummy)	0.559	0.497	0.572	0.495	- 0.013**
Child \leq 6 years (dummy)	0.097	0.295	0.149	0.356	- 0.052***
Homeowner (dummy)	0.896	0.305	0.892	0.311	0.004
Mother college (dummy)	0.0454	0.208	0.0538	0.226	- 0.008***
Father college (dummy)	0.0693	0.254	0.0794	0.270	- 0.010***
Past years out of employment (in years)	8.470	4.508	8.679	5.121	- 0.209***
Observations	18,964		14,097		33,061

Notes: ‘Small Firm’ equals one if the firm has at most 19 employees, zero otherwise. ‘Medium Firm’ equals one if the firm has between 20 and 199 employees, zero otherwise. Reported differences are based on a regression of a male dummy on the respective selected variable. *, ** and *** denote significance at the 10%-, 5%- and 1%-level, respectively. Robust standard errors (clustered at the individual level) are used. Further potential controls include sector, occupation (ISCO88 (1-digit)) and survey-year dummies as well as dummies for living in the South, North or Centre of Italy, being employed in a big firm and quadratic polynomials of age, experience and tenure. Further, we add dummies for parental education (*Diploma*, *Medie inferiori*, *Elementare*), an interaction term between the dummies homeowner and being married. *Source*: ISFOL PLUS 2010, 2011, 2014, 2016

generally well below the EU-average of 16% (the Italian gender pay gap was 5% in 2017 according to Eurostat (2018)). On average, women are two years younger than their male counterparts, have more years of education and hold more often an honors degree. Men have more than three years more labor market experience and more than two years more tenure compared to their female colleagues. Italian women are also more often employed in smaller firms. Men and women are about equally often homeowners or married. We observe both more women with children and more women with young children in our sample compared to men. Further, women stay on average longer out of employment compared to men.

Note that the potential set of control variables we consider is much broader (at least 63 raw controls) than the variables represented here. We provide a full list of the set of potential controls for both countries in Table B.1 in the Appendix. All in all, the descriptive

statistics show that men and women differ substantially in several observable characteristics. These control variables may thus be relevant for explaining the corresponding wage gaps. As we do not know *a priori* which controls are pivotal for estimating the gaps, we use LASSO for model selection. For example, controls such as past years of unemployment or out of employment are generally not included in augmented Mincer-type wage models (potentially due to data restrictions). Moreover, the literature is ambiguous concerning the relevance of having children or the number of children in wage equations Castagnetti et al. (2020). Further, as we are interested in wage gaps at different points of the wage distribution, expectile- or quantile-specific model specifications may matter. In fact, the literature highlighted recently that for different groups (Wilde et al. 2010; Gensowski 2018; Juhn and McCue 2017) or points of the wage distribution (Bonaccolto-Töpfer and Briel 2022), different sets of covariates are required.

Overall, penalized estimation offers a convenient tool for model selection and thus an answer to uncertainties concerning whether or not to include past years of unemployment or the number of children in the model. Beyond that, it provides an answer on whether different model specifications are required across the distribution.

4 Empirical findings

We analyze in this section the empirical results. We compare the coefficients and characteristics effects obtained from post-penalized expectile (PE) and quantile (PQ) regressions, highlighting the statistically significant differences between these two estimation approaches. We also study common inter-quantile wage gaps for both competing methods. Both quantiles and expectiles show the relation of the covariates to the response variable across the distribution (Sobotka et al. 2013). As in case of the mean, the coefficient estimates can be interpreted as the effect of the controls \mathbf{x} on the conditional quantile or expectile of the response variable \mathbf{y} given \mathbf{x} (Newey and Powell 1987; Fortin et al. 2011).

Panel (a) of Fig. 1 shows the decomposition resulting from the German data along the wage distribution. PQ provides a lower coefficients effect up to the 70th percentile. In this

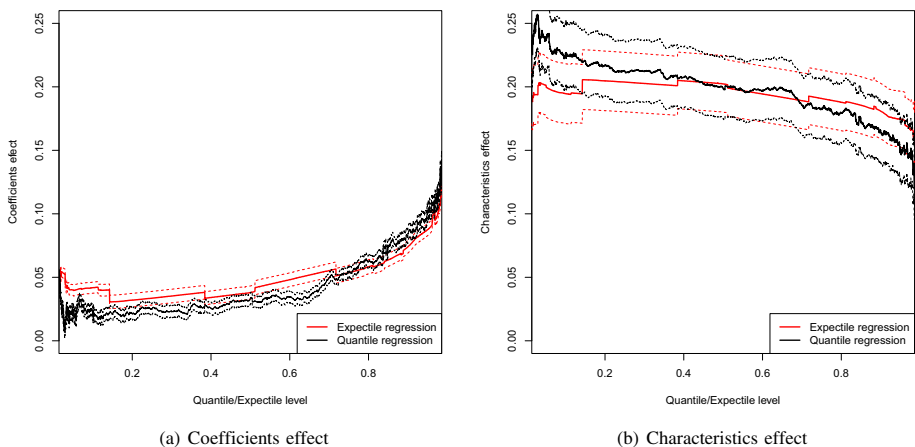


Fig. 1 Decomposition of the gender pay gap along the wage distribution – Germany. *Notes:* 95%-bootstrapped confidence bands presented (100 replications). *Source:* SOEP v34

region of the wage distribution, we often find significant differences between PE and PQ. In contrast, the estimates do not substantially differ from each other in the right tail of the same distribution.³ Overall, the coefficients effect takes low values and is a positive function of $\theta \in (0, 1)$. We stress the fact that the coefficients or unexplained part is related to gender differences in prices. Since the unexplained part reflects differences in pay of individuals with identical observable characteristics, some authors consider it as a proxy of discrimination (Goldin 2014). Yet, the unexplained part may also include unmeasured productivity differences or compensating wage differentials.

The empirical evidence suggests no significant discrepancies between PE and PQ for the characteristics effect, which is often referred to as explained part, arising from gender differences in control variables. This effect is also relatively stable along the wage distribution, ranging between 0.18 and 0.25. Consider for example the estimates at the 10th percentile of the wage distribution for Germany (Fig. 1). Using expectile regression yields a coefficients effect of 19 percentage points. Given an aggregate 10th-percentile wage gap of 23% (log approximation), 82.6% of the total wage gap at the 10th percentile can be explained based on differences in observable characteristics such as education and labor market experience. The corresponding coefficients effect amounts to four percentage points. Thus, at the 10th percentile, men earn four percentage points more compared to women with identical characteristics. This finding implies that 17.4% of the wage gap remain unexplained at this part of the distribution.

When looking at the quantile regression results, the results suggest that 8% of the 10th-percentile wage gap remain unexplained, while 92% can be explained based on differences in endowments.⁴ Thus, both estimation approaches suggest that most part of the 10th percentile wage gap can be explained based on the selected controls. However, the magnitude differs by about ten percentage points across the two approaches for each component. That is, even though, the main conclusions persist and the estimates in case of the characteristics effect do not differ markedly, they may be economically different. Empirical research typically detects a substantial unexplained fraction (i.e. coefficients effect); see, among others Goldin (2006), Mandel and Semyonov (2014) and Blau and Kahn (2017). In contrast, our estimation approach explains most of the gender pay gap in Germany, as we can see by comparing Panels (a) and (b) of Fig. 1, where it is clear that the characteristics effect plays a more relevant role. A possible reason might be related to the selection process made by our regularization technique (i.e. LASSO), which filters the relevant control variables from a large set of covariates for both PE and PQ. Interestingly, the coefficients effect is highest at the top, while the characteristics effect is smallest at the top. The latter underlines that it is important to consider gender wage gaps at different points of the wage distribution. This finding is robust to both PE and PQ regression.

We enrich our analysis by implementing the two-sample Kolmogorov-Smirnov test on the empirical distribution functions of the coefficients and characteristics effects obtained from PE and PQ, respectively. By doing so, we test the null hypothesis that the distributions provided by these two competing methods are equal. We report the results of the two-sample Kolmogorov-Smirnov test implemented on the German data in Panel (a) of Table 3. Here, we can see that the differences between PE and PQ are highly statistically significant for

³Note that overlapping 95% confidence bands do not necessarily imply statistical insignificant differences between point estimates (see for details ; Knol et al. 2011).

⁴Note that the 10th percentile wage gap in case of quantiles amounts to 25% (log approximation), the coefficients effect to two and the characteristics effect to 23 percentage points.

Table 3 Two-sample Kolmogorov-Smirnov test

Effect	(1) Difference	(2) P-value
<i>Panel (a): Germany</i>		
Coefficients	0.449	0.000
Characteristics	0.4212	0.000
<i>Panel (b): Italy</i>		
Coefficients	0.086	0.000
Characteristics	0.807	0.000

Notes: The null hypothesis that the true distribution function of expectile regression is equal to the distribution function of quantile regression. This is a comparison of cumulative distribution functions, and the test statistic is the maximum difference in value. *Source*: SOEP v34 for Germany and ISFOL PLUS 2010, 2011, 2014, 2016 for Italy

both the coefficients and characteristics effects. This evidence is also clear in Panels (a) and (b) of Fig. A.2 reported in the Appendix, where we display the empirical distribution functions based on predicted wages.

We now analyze the decomposition derived from the Italian data, displaying the results in Fig. 2. Starting from the coefficients effect, we find that PE and PQ provide significantly different estimates in many regions of the wage distribution, from the left to the right tail. PE almost always leads to greater coefficients effects with respect to PQ. However, a clear exception is observed in correspondence of extremely lower values of θ (approaching zero), where the confidence interval associated to PE is below, and does not overlap, the one resulting from PQ. Therefore, policy implications may differ depending on the estimation approach (PE or PQ) as well as on the different regions of the wage distribution. The coefficients effect always takes positive values for both PE and PQ and is a positive function of θ . This means that, given the same set of characteristics, women earn substantially less than men, and this phenomenon is more pronounced when increasing the wage levels, similar to the German analysis.

Moving to the characteristics part, we do not find significant differences between PE and PQ (see Panel (b) of Fig. 2), similar to the German case. This finding implies that the catch-up of women in terms of observable characteristics (Goldin 2014) is robust to both quantile and expectile regression. Again, the characteristics effect is quite stable along the different values of $\theta \in (0, 1)$, exhibiting a slightly decreasing trend. However, in contrast to Germany, we find a negative characteristics effect in Italy along the entire wage distribution, from the left to the right tail. This finding implies that the catch-up of women in terms of observable characteristics (Goldin 2014) is robust to both quantile and expectile regression. Furthermore, differently to Germany, the coefficients effect is the main driver of wage gaps in Italy. This result is in line with Arulampalam et al. (2007) and Castagnetti et al. (2020), who found a greater unexplained fraction of the gap compared to the explained one.

Again, we implement the two-sample Kolmogorov-Smirnov test, which suggests the rejection of the null hypothesis of equal distribution also for Italy; see Panel (b) of Table 3. Therefore, depending on the estimation approach (PE or PQ), we obtain significantly different empirical distribution functions (see Panels (c) and (d) of Fig. A.2 reported in the Appendix).

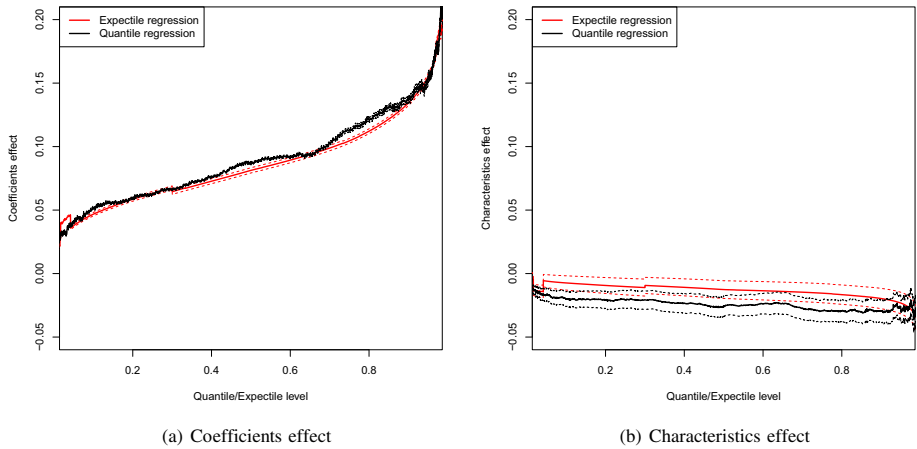


Fig. 2 Decomposition of the gender pay gap along the wage distribution – Italy. *Notes:* 95-% bootstrapped confidence bands presented (100 replications). *Source:* ISFOL PLUS 2010, 2011, 2014, 2016

All in all, we find significant differences in the decompositions derived from PE and PQ. This evidence is clearer when analyzing the coefficients effects for both Germany and Italy, with statistically significant differences along the entire distribution of the wage distribution. As a result, policy implications may change depending on the method chosen. For effective policy implications, a careful selection of the adequate estimation approach and a consideration of potential confounders is pivotal. Indeed, the set of relevant controls may differ across the distribution as well as for men and women. For instance, Wilde et al. (2010) found that having children penalizes in particular high-skilled women’s wages. The results of Bonaccolto-Töpfer and Briel (2022) suggest that control variables related to labor market experience are dominant at the top of the distribution. Thus, using different sets of controls at different parts of the distribution and for different groups (men or women) might be important in order to obtain valid estimation results. However, detecting the most suitable variable set among a large set of potential controls and at several points (each 5th quantile or expectile in our case) of the distribution is cumbersome for a human researcher. Using a machine learning approach offers a convenient way to identify the most adequate set of controls for the data set at hand.

Moreover, for inequality studies including studies on gender pay gaps, or related wage gaps (e.g. public-private wage gaps), expectile estimation represents an appealing estimation tool, being more sensitive to extreme observations. In order to detect patterns of inequality that are often concentrated in the lower or upper tails, expectile regression may thus be more adequate compared to traditional quantile regression. In fact, we find relevant differences in the tails of the coefficients effect distribution for Germany. This issue may be particularly important for research on labor market differentials as well as for policy evaluation. Depending on the estimation method (PE or PQ) and the effect analyzed (coefficients or characteristics), we may find sticky floors or glass ceiling or not.

In order to better understand the implications of the different estimates derived from PE and PQ, we calculate common inter-quantile (or -expectile) wage gaps. Table 4 shows the inter-quantile and inter-expectile wage gaps between: i) the top and bottom (90-10); ii) the top and median (90-50); and iii) the median and bottom (50-10). These cases represent inter-quantile (or inter-expectile) gaps that are mostly analyzed in the literature (e.g. Albrecht

Table 4 Inter-quantile and - expectile gaps, with penalization

	(1)	(2)	(3)	(4)
	Quantile Regression		Expectile Regression	
Inter-wage gap ↓ /Effect →	Characteristics	Coefficients	Characteristics	Coefficients
<i>Panel (a): Germany</i>				
90-10	-0.062*** (0.017)	0.064*** (0.004)	-0.014 (0.016)	0.032*** (0.003)
90-50	-0.034** (0.017)	0.051*** (0.003)	-0.022 (0.016)	0.036*** (0.004)
50-10	-0.027 (0.017)	0.013*** (0.003)	0.007 (0.016)	-0.004 (0.003)
<i>Panel (b): Italy</i>				
90-10	-0.009* (0.005)	0.086*** (0.001)	-0.012** (0.005)	0.088*** (0.001)
90-50	-0.004 (0.006)	0.051*** (0.001)	-0.007 (0.005)	0.053*** (0.001)
50-10	-0.005 (0.005)	0.034*** (0.001)	-0.005 (0.004)	0.035*** (0.001)

Notes: *, ** and *** denote significance at the 10%-, 5%- and 1%-level, respectively. Bootstrapped standard errors (clustered at the individual level, 100 replications) in parentheses. Source: SOEP data v34 for Germany and ISFOL PLUS 2010, 2011, 2014, 2016 for Italy

et al. 2003; Arulampalam et al. 2007; Fortin et al. 2011). Recall that glass ceiling exists when the 90th quantile- (or expectile-) level wage gap exceeds the corresponding 50th or 10th gap by at least two percentage points. Sticky floors are defined as the differences between the 10th and 50th or 90th percentile wage gap.

Panel (a) of Table 4 shows the results for Germany. We often find statistically significant differences, especially when contrasting the top ($\theta = 0.9$) quantile or expectile level with both the median ($\theta = 0.5$) and the bottom ($\theta = 0.1$) ones. This evidence holds when looking at both the explained and unexplained parts for PQ, whereas it is clear only for the coefficients effect when adopting PE. In contrast, we find that the 50-10 difference is significant only for the unexplained part by employing PQ. The differences in the coefficients effects take a positive value, whereas the opposite holds for the characteristics effects, given the different trends, increasing and decreasing, of these two components, respectively, displayed in Fig. 1. As a result, we find evidence of glass ceiling in the coefficients effect by implementing both PE and PQ. In contrast, we highlight sticky floors in the characteristics effects when employing PQ only. The relative pronounced difference in the 90-10 gender wage gap between PE and PQ in Germany is driven by the lower part of the distribution. Figure 1 shows that the pattern of PQ across the distribution is steeper compared to PE yielding to non-negligible differences at the lower part of the wage distribution between the two estimation approaches. To be precise, differences in characteristics and coefficients across the distribution are higher when using PQ. Thus, German men outperform German women especially at the upper part of the distribution leading to marked 90-10 and 90-50 gaps in case of PQ.

Panel (b) of Table 4 shows the results obtained for Italy. Here, the differences in the coefficients effect between expectile and quantile inter-wage gaps are less pronounced with respect to the German case. The inter-wage gaps are generally larger and are always highly significant. This is then a clear evidence of glass ceiling for both PE and PQ, whose magnitude is similar between these two competing estimation approaches. The inter-wage gaps become less evident when looking at the characteristics effects, while both estimation approaches still yield similar results. Therefore, the estimation outcome for Italy in terms of inter-expectile or -quantile wage gaps is robust to these two different econometric tools (expectiles and quantiles). For Italy, the difference between PQ and PE across the distribution is small (e.g. both curves are either quite steep (Panel (a)) or flat (Panel (b)), Fig. 2). As a consequence, the difference in inter-wage gaps for PE and PQ is not pronounced for Italy.

To summarize, the results analyzed above suggest that, despite of using the same data set over the same time period, policy implications concerning labor market differentials may change depending on the estimation approach used. These issues may significantly distort policy implications. On the one hand, it is then crucial to carefully choose the estimation method based on the problem at hand. On the other hand, the resulting empirical findings may be considered as more reliable when confirmed by both methods. The results provided by PE and PQ are similar for Italy when looking at inter-wage gaps and, thus, at implications for wage inequality across the distribution (such as sticky floors or glass ceiling). However, this finding does not hold when looking at the components at a specific part of the distribution (see Fig. 2). In case of Germany, significant differences were mainly located in the tails, while, for Italy, we find them at various points of the wage distribution. Further, given that expectiles depend more global on the form of the distribution and are thus sensitive to outliers (Taylor 2008), they present an attractive alternative in the estimation of gender pay gaps across the distribution.

5 Robustness analysis

We compare here the estimates derived from expectile- and quantile-specific model specification with those obtained using pre-defined regressors. We use the control variables of the full specification in Blau and Kahn (2017) for each value of $\theta = (0.05, 0.1, 0.15, \dots, 0.95)$ along the conditional distribution of y_i , for both expectiles and quantiles. We stress that the latter (i.e., using one set of selected controls for estimation at all points of the distribution) is the main approach in applied labor economics. We use the specification suggested by Blau and Kahn (2017) as their paper presents a thorough and up-to-date review of the literature on gender differences in pay. The full specification of Blau and Kahn (2017) includes apart from human capital, labor market and background characteristics (such as migration information), controls for the sector, the occupation, the survey year and the federal state.

Figure 3 suggests that expectile- and quantile-specific model specifications, respectively, are important when decomposing gender pay gaps for Germany. Using the full specification of Blau and Kahn (2017) for all points of the response variable's distribution and for both approaches instead of quantile- or expectile-specific specifications, yields substantial differences in the estimates. Generally both expectile- and quantile-specific model specifications (i.e. with penalization) explain a larger fraction of the wage gap (characteristics part). Analogously, the coefficients part is lower in case of penalization compared to the model without penalization. The point estimates of the coefficients part differ substantially from each other when using penalization compared to using one set of controls in case of

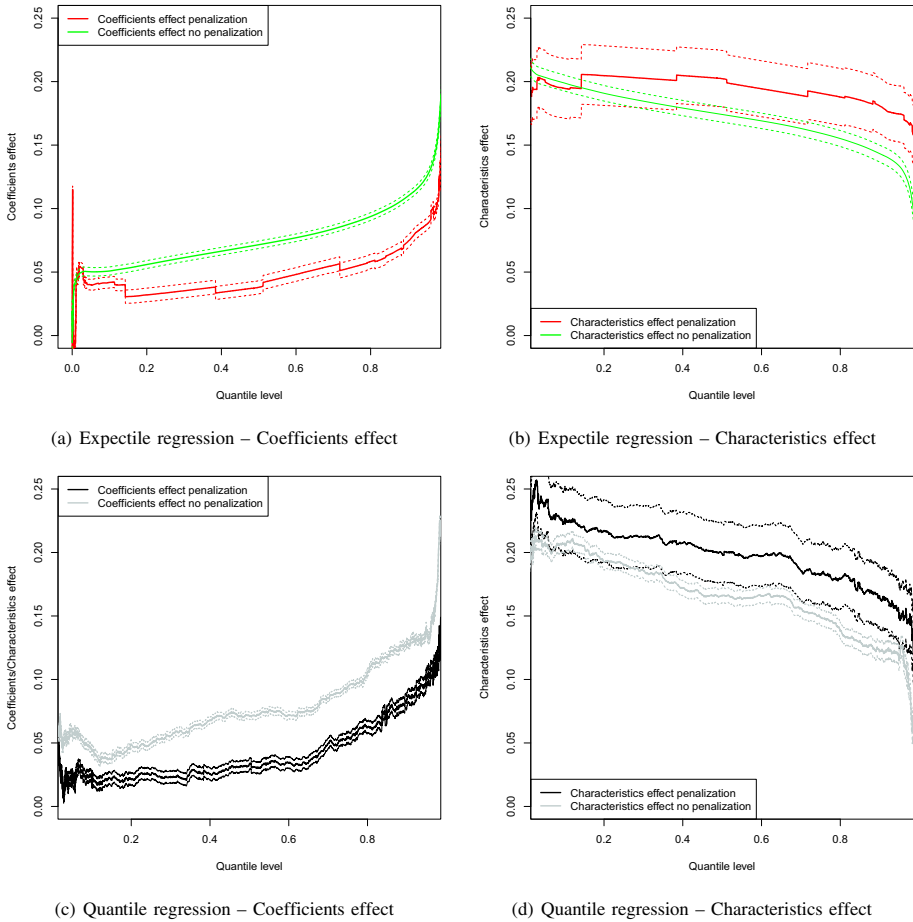


Fig. 3 Decomposition of the gender pay gap along the wage distribution (no penalization, full specification of Blau & Kahn, 2017) – Germany. *Notes:* 95-% bootstrapped confidence bands presented (100 replications). Figure shows the estimates from Fig. 1 (penalization) in comparison with the corresponding estimates based on the full specification in Blau and Kahn (2017) (no penalization). The full specification includes: years of education, labor market experience, age, migration-background dummy, metropolitan-region dummy as well as federal-state dummies, survey-year dummies, sector dummies and occupation dummies. Source: SOEP v34

expectile regression. The latter holds for most points of the distribution and the economic difference is relatively pronounced (except at the very top). This finding holds for both expectile and quantile regressions. In contrast, differences in the characteristics part are less evident.

Figure 4 displays the robustness results for Italy. Similar to Germany, we find relevant differences between the full specification of Blau and Kahn (2017) and the penalized expectile model in the coefficients effect. The point estimates of the coefficients part are significantly different. This result holds along the entire wage distribution, except at the right tail, when θ approaches one, where the two competing methods provide similar results (see Panel (a) of Fig. 4). Therefore, once again, our regularized method yields lower coefficients effects compared to the full specification of Blau and Kahn (2017), which does not

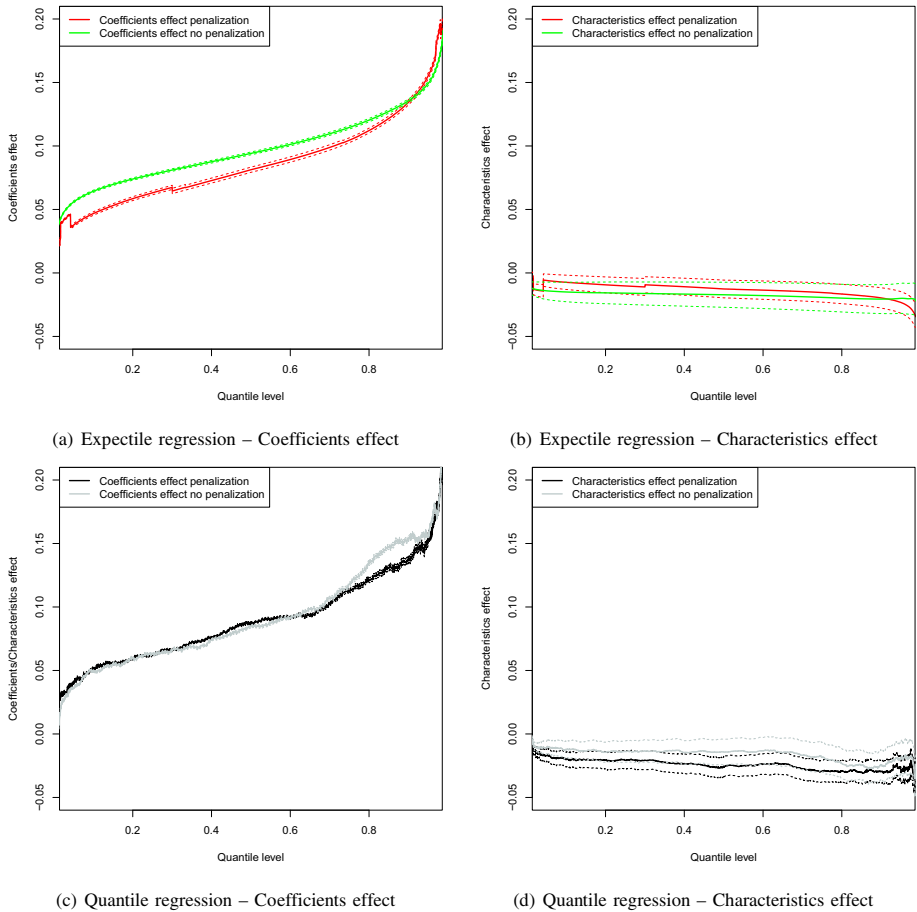


Fig. 4 Decomposition of the gender pay gap along the wage distribution (no penalization, full specification of Blau & Kahn, 2017) – Italy. *Notes:* 95-% bootstrapped confidence bands presented (100 replications). Figure shows the conditional estimates from Figure 2 (penalization) in comparison with the corresponding estimates based on the full specification in Blau and Kahn (2017) (no penalization). The full specification includes: years of education, labor market experience, age, metropolitan-region dummy as well as dummies for living in the North, South or Centre of Italy, survey-year dummies, sector dummies and occupation dummies. Source: ISFOL PLUS 2010, 2011, 2014, 2016

take into account the potential bias typically observed in large-dimensional problems. The correction made by our model may then affect policy implications. In contrast to Germany, the differences in the coefficients effect are now less marked when comparing the penalized and non-penalized quantile regressions (see Panel (c) of Fig. 4). However, we still detect non-negligible differences from $\theta = 0.75$ to $\theta = 0.90$. Similar to Germany, we do not find significant differences in the characteristics part for Italy, by employing both expectile and quantile methods (see Panels (b) and (d) of Fig. 4).

We finally report the inter-quantile and inter-expectile wage gaps obtained from the full specification of Blau and Kahn (2017) in Table 5, distinguishing between Germany (Panel (a)) and Italy (Panel (b)). By comparing Tables 4 and 5, we can see that the full specification of Blau and Kahn (2017) leads to more significant sticky floors and glass ceilings.

Table 5 Inter-quantile and - expectile gaps, no penalization

Inter-wage gap ↓ /Effect →	(1)	(2)	(3)	(4)
	Quantile Regression		Expectile Regression	
	Characteristics	Coefficients	Characteristics	Coefficients
<i>Panel (a): Germany</i>				
90-10	-0.085*** (0.005)	0.086*** (0.002)	-0.056*** (0.005)	0.059*** (0.002)
90-50	-0.043*** (0.005)	0.058*** (0.002)	-0.031*** (0.004)	0.038*** (0.002)
50-10	-0.041*** (0.004)	0.028*** (0.002)	-0.024*** (0.004)	0.021*** (0.002)
<i>Panel (b): Italy</i>				
90-10	-0.012*** (0.007)	0.104*** (0.001)	-0.007 (0.007)	0.071*** (0.001)
90-50	-0.011 (0.008)	0.069*** (0.001)	-0.003 (0.008)	0.041*** (0.001)
50-10	-0.002 (0.006)	0.035*** (0.001)	-0.002 (0.006)	0.030*** (0.001)

Notes: *, ** and *** denote significance at the 10%-, 5%- and 1%-level, respectively. Bootstrapped standard errors (clustered at the individual level, 100 replications) in parentheses. Source: SOEP data v34 for Germany and ISFOL PLUS 2010, 2011, 2014, 2016 for Italy

This result is again more evident for Germany. Indeed, all inter-quantile and inter-expectile wage gaps reported in Panel (a) of Table 5 are highly significant. In contrast, the values given in the fourth column of Table 4, as well as the 50-10 gaps in the second and fifth columns of the same table, are not significant. Again, we find non negligible differences between expectile and quantile estimates. For instance, we find in the characteristics part an inter-quantile 50-10 gap of approximately -4 percentage points, but an inter-expectile gap of approximately -2 percentage points. Therefore, even in the full specification of Blau and Kahn (2017), where we do not employ penalty functions to select the relevant regressors, the comparison between expectile and quantile estimates reflects significant differences. As for Italy, the inter-quantile and inter-expectile gaps do not substantially change by moving from Table 4 to Table 5. Policy conclusions about glass ceiling or sticky floors are then not affected. However, we still detect non-negligible differences in the magnitude of these gaps. For instance, the estimated 90-10 coefficients inter-quantile gap of 8.6 percentage points in Panel (b) of Table 4 increases to 10.4 percentage points by employing the standard specification of Blau and Kahn (2017). Furthermore, the resulting estimates also differ between expectile and quantile models (e.g. approximately 7 versus 4 percentage points in the 90-50 coefficients gap). The analysis reported in this section allows us to verify that a standard non-penalized model, such as the full specification proposed by Blau and Kahn (2017), provides different results compared to more accurate regularized methods, where we LASSO-select the relevant variables. In fact, the former approach typically suffers from bias estimates when dealing with high-dimensional problems, implying the risk of leading to misleading policy conclusions about gender wage inequalities. Moreover, this risk might significantly change according to the different regions of the wage distribution.

Summing up, the findings analyzed in this section suggest that a regularization technique such as LASSO significantly impacts the decomposition of wage gaps, with heterogeneous effects along the wage distribution. We find marked and significant differences in point estimates in both expectile and quantile methods. As a consequence, inter-quantile wage gaps and policy implications change. A penalized or data-driven approach is typically an efficient tool for model selection at specific expectile or quantile levels; in this study, we confirm its appealing properties also in the estimation of gender pay gaps. As an additional robustness check, we show in Appendix D the estimation results when using linear unconditional quantile regression or RIF-OLS (Firpo et al. 2009) for the analysis. Again, we find that results differ based on the estimation approach.

6 Conclusion, discussion and future research

In this study, we decompose the gender pay gap in Germany and Italy over the periods 2010-2017 and 2010-2016, respectively. We adapt the Machado and Mata (2005) decomposition for post-penalized expectile and quantile regression. Even though expectile regression may be particularly well-suited for research on wage differentials, it has not yet been used in applied labor economics. Depending on the estimation approach (i.e. expectile or quantile regression), the decomposition components (characteristics and coefficients part) differ substantially along the wage distribution, for both Germany and Italy. This result becomes particularly evident when analyzing coefficients effects. Therefore, policy implications may differ depending on the estimation approach. However, we also find cases in which the estimates derived from the expectile and quantile methods are not significantly different, mainly in the characteristics part. The corresponding estimates may then be considered more reliable, as they are robust when contrasting different estimation approaches (i.e. expectile and quantile methods).

The original Machado and Mata (2005) decomposition builds on quantile regression for each possible quantile along the wage distribution, and uses a simulation procedure. This method predicts wages and, thus, is particularly appealing to exploit the machine-learning property of prediction. Our study extends the approach proposed by Machado and Mata (2005) by introducing regularization techniques. Moreover, we apply post-penalized expectile regression to this decomposition. Expectile regression may represent an interesting alternative approach to the more popular quantile regression framework when analyzing heterogeneous effects along the wage distribution. Indeed, expectiles are more sensitive to changes in the tails of the wage distribution and, therefore, respond more readily to extreme observations. To the best of our knowledge, this is the first study that uses (post-) penalized expectile regression for decomposing gender pay gaps.

Our approach is suitable to deal with high-dimensional problems, building on the LASSO penalty function to identify the relevant control variables. Thus, given data restrictions, the approach minimizes potential omitted variable bias. As previous literature suggested that different sets of covariates for different groups (Wilde et al. 2010; Juhn and McCue 2017; Gensowski 2018) or points of the distribution (Bonaccollo-Töpfer and Briel 2022) should be used, the data-driven penalization offers a convenient and efficient way for model selection. Further, as we have rich data sets (at least potential 63 control variables), machine learning helps to conduct model selection, dealing with the curse of dimensionality.

We find, in line with the literature (e.g. Blau and Kahn 2017), strictly positive gender pay gaps along the wage distribution. For Germany, most part of the gap can be attributed to the characteristics part, while the main driver of the gap in Italy is the unexplained part. The

explained fraction changes depending on the estimation approach that is adopted (expectile or quantile regression), affecting the conclusions about glass ceiling and sticky floors. We find weak evidence in case of expectile but rather strong evidence in case of quantile regressions for the characteristics part. Likewise, inter-expectile and inter-quantile coefficients effects differ substantially by contrasting expectile and quantile regression (up to three percentage points).

As a robustness exercise, we decompose the gender pay gaps using a pre-defined set of controls that does not change across the wage distribution; that is, the full specification proposed by Blau and Kahn (2017). The estimated components across the wage distribution as well as inter-wage gaps are again sensitive to the estimation approach that is employed. Moreover, these non-penalized estimates significantly differ from those obtained by adopting the LASSO method. Thus, using a data-driven approach for model selection has a relevant impact.

Building on a large set of potential controls, penalization techniques may also select variables that are affected by gender (wage) discrimination or feedback effects. For instance, gender division of labor in specific industries, jobs or child-rearing drives gender differences in labor market outcomes and vice versa (Blau and Winkler 2017). Controlling for these variables may thus lead to underestimation of gender inequality or discrimination attributed typically to the coefficients component. Note, however, that we do not interpret the coefficients effect as gender discrimination but consider it (at most) a proxy for discrimination as it incorporates effects of group differences in unobserved predictors (Blau and Kahn 2006). Further, standard specifications without penalization generally include these kind of variables (e.g. the full specification of Blau and Kahn 2017, contains occupational and industrial controls). Finally, previous research found that using penalization techniques yields estimation results that are robust to selection on unobservables (Bonaccolto-Töpfer and Briel 2022). The latter suggests that the penalized approach leads to more adequate estimates of *ceteris paribus* wage differentials.

A caveat of the Machado and Mata (2005) method is that we cannot compute detailed decompositions for both components (Fortin et al. 2011). Another important point is related to sample selection correction; that is, estimated wage gaps may be biased and inconsistent in case of nonrandom selection into the labor market (Heckman 1979; Arellano and Bonhomme 2017). However, the main focus of this study is to check whether and to what extent results change when implementing expectile regression as an alternative method to quantile regression. Therefore, in this study, we do not take sample selection into account (as many other studies, e.g.; Albrecht et al. 2003; Melly 2005; Arulampalam et al. 2007; Firpo et al. 2009; Depalo et al. 2015).

All in all, our study underlines the importance of using appropriate statistical tools. As expectile regression is more sensitive to extreme cases, it may be particularly well-suited for policy analysis focused on heterogeneous effects that we detect along different regions of the wage distribution. Further, penalized regressions substantially affect the results and offer an interesting tool for model selection in applied research.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10888-023-09565-x>.

Acknowledgements We thank the co-editor and two anonymous referees of this journal as well as Carolina Castagnetti, Michael Knaus, Sascha Satlukal, Claus Schnabel and participants of the workshop on machine learning in economics 2020 at FAU and IAB for helpful comments and suggestions.

Funding Open access funding provided by Università degli Studi di Genova within the CRUI-CARE Agreement. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data Availability Access to SOEP and PLUS data is subject to prior registration of researchers and their institution but is free of charge for academic purposes (see https://www.diw.de/en/diw_02.c.238223.en/soep_conditions.html and <https://www.isfol.it/open-data-delle-ricerche/isfol-microdati> for details). Yet, as data access is restricted to registered users, we cannot deposit the files in an openly accessible data repository.

Declarations

Conflict of Interests none

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>; <http://creativecommons.org/licenses/by/4.0/>; <http://creativecommons.org/licenses/by/4.0/>; <http://creativecommons.org/licenses/by/4.0/>.

References

- Albrecht, J., Björklund, A., Vroman, S.: Is there a glass ceiling in Sweden? *J. Labor Econ.* **21**(1), 145–177 (2003)
- Arellano, M., Bonhomme, S.: Quantile selection models with an application to understanding changes in wage inequality. *Econometrica* **85**(1), 1–28 (2017)
- Arulampalam, W., Booth, A., Bryan, M.: Is there a glass ceiling over Europe? Exploring the gender pay gap across the wages distribution. *Industr. Labor Relations Rev.* **60**(2), 163–186 (2007)
- Bach, P., Chernozhukov, V., Spindler, M.: Closing the US gender wage gap requires understanding its heterogeneity. arXiv:1812.04345 (2018)
- Bassett, G., Koenker, R.: An empirical quantile function for linear models with iid errors. *J. Amer. Stat. Assoc.* **77**(378), 407–415 (1982)
- Bassett, G., Koenker, R.: Strong consistency of regression quantiles and related empirical processes. *Economet. Theor.* **2**(2), 191–201 (1986)
- Bellini, F., Cesarone, F., Colombo, C., Tardella, F.: Risk parity with expectiles. *European J. Oper. Res.* **291**(3), 1149–1163 (2021)
- Bellini, F., Di Bernardino, E.: Risk management with expectiles. *European J. Finance* **23**(6), 487–506 (2017)
- Belloni, A., Chen, D., Chernozhukov, V., Hansen, C.: Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica* **80**(6), 2369–2429 (2012)
- Belloni, A., Chernozhukov, V.: L1-Penalized quantile regression in high-dimensional sparse models. *Annals Stat.* **39**(1), 82–130 (2011)
- Belloni, A., Chernozhukov, V., Fernández-Val, I., Hansen, C.: Program evaluation and causal inference with high-dimensional data. *Econometrica* **85**(1), 233–298 (2017)
- Blau, F.D., Kahn, L.: The US gender pay gap in the 1990s: slowing convergence. *Industr. Labor Relations Rev.* **60**, 45–66 (2006)
- Blau, F.D., Kahn, L.M.: The gender wage gap: extent, trends, and explanations. *J. Econ. Lit.* **55**(3), 789–865 (2017)
- Blau, F.D., Winkler, A.E.: The economics of women, men, and work (vol. 8). New York: Oxford University Press (2017)
- Bonaccollo, G.: Quantile-based portfolios: post-model-selection estimation with alternative specifications. *Computat. Manag. Sci.* **18**(3), 355–383 (2021)
- Bonaccollo-Töpfer, M., Briel, S.: The gender pay gap revisited: does machine learning offer new insights? *Labour Econom.* **78**, 102223 (2022). Elsevier

- Bonaccolto, G., Caporin, M., Maillat, B.B.: Dynamic large financial networks via conditional expected shortfalls. *European J. Oper. Res.* **298**(1), 322–336 (2022)
- Bonaccolto-Töpfer, M., Castagnetti, C., Prümer, S.: Understanding the public-private sector wage gap in Germany: new evidence from a fixed effects quantile approach. *Econom. Model.* **116**, 106037 (2022). Elsevier
- Brunori, P., Neidhöfer, G.: The evolution of inequality of opportunity in Germany: a machine learning approach. *Rev. Income Wealth* **67**(4), 900–927 (2021)
- Castagnetti, C., Giorgetti, M.L.: Understanding the gender wage-gap differential between the public and private sectors in Italy: a quantile approach. *Econ. Model.* **78**(1), 240–261 (2019)
- Castagnetti, C., Rosti, L., Töpfer, M.: Discriminate me – if you can! The disappearance of the gender pay gap among public-contest selected employees in Italy. *Gender, Work Organization* **27**(6), 1040–1076 (2020)
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., Robins, J.: Double/debiased machine learning for treatment and structural parameters. *Econometrics J.* **21**(1), C1–C68 (2018)
- Collischon, M.: Is there a glass ceiling over Germany? *German Economic Rev.* **20**(4), 329–359 (2019)
- Corsetti, G., Mandrone, E., Spizzichino, A.: L'indagine Isfol-Plus. *Rivista Italiana di Economia Demografia e Statistica.* *Rivista Italiana di Economia Demografia e Statistica* **48**(3/4), 1–8 (2014)
- Davino, C., Romano, R., Vistocco, D.: Handling multicollinearity in quantile regression through the use of principal component regression, vol. 1–22. <https://doi.org/10.1007/s40300-022-00230-3> (2022)
- Depalo, D., Giordano, R., Papapetrou, E.: Public–private wage differentials in euro-area countries: evidence from quantile decomposition analysis. *Empirical Econ.* **49**(3), 985–1015 (2015)
- Efron, B.: Regression percentiles using asymmetric squared error loss. *Stat. Sin.* **1**(1), 93–125 (1991)
- Eurostat: Gender pay gap statistics. http://ec.europa.eu/eurostat/statistics-explained/index.php/Gender_pay_gap_statistics. Accessed 10 Oct 2018 (2018)
- Fan, J., Li, R.: Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Stat. Assoc.* **96**(456), 1348–1360 (2001)
- Firpo, S., Fortin, N.M., Lemieux, T.: Unconditional quantile regressions. *Econometrica* **77**(3), 953–973 (2009)
- Fitzenberger, B., Koenker, R., Machado, J.A.: *Economic applications of quantile regression.* Springer Science & Business Media (2013)
- Fortin, N., Lemieux, T., Firpo, S.: Decomposition methods in economics. In: Ashenfelter, O., Card, D. (eds.) *Handbook of Labor Economics*, vol. 4, pp. 1–102. Elsevier (2011)
- Furno, M., Vistocco, D.: Quantile regression: estimation and simulation, vol 2. Wiley (2018)
- Genowski, M.: Personality, IQ, and lifetime earnings. *Labour Econ.* **51**, 170–183 (2018)
- Giacometti, R., Torri, G., Paterlini, S.: Tail risks in large portfolio selection: penalized quantile and expectile minimum deviation models. *Quantitative Finance* **21**(2), 243–261 (2021)
- Goldin, C.: The quiet revolution that transformed women's employment, education, and family. *Amer. Econ. Rev.* **96**(2), 1–21 (2006)
- Goldin, C.: A grand gender convergence: its last chapter. *Amer. Econ. Rev.* **104**(4), 1091–1119 (2014)
- Hastie, T., Tibshirani, R., Friedman, J.: *Linear methods for regression.* In: *The elements of statistical learning*, pp. 43–100. Springer, New York (2009)
- Hastie, T., Tibshirani, R., Wainwright, M.: *Statistical learning with sparsity.* Chapman and Hall/CRC, <https://doi.org/10.1201/b18401> (2015)
- Hautsch, N., Schaumburg, J., Schienle, M.: Financial network systemic risk contributions. *Rev. Finance* **19**(2), 685–738 (2014)
- Heckman, J.: Sample selection bias as a specification error. *Econometrica* **47**(1), 153–161 (1979)
- Jones, M.: Expectiles and m-quantiles are quantiles. *Stat. Probab. Lett.* **20**(2), 149–153 (1994)
- Juhn, C., McCue, K.: Specialization then and now: marriage, children, and the gender earnings gap across cohorts. *J. Econ. Perspectives* **31**(1), 183–204 (2017)
- Kallus, N., Mao, X., Uehara, M.: Localized debiased machine learning: efficient inference on quantile treatment effects and beyond. arXiv:1912.12945 (2019)
- Knaus, M.C., Lechner, M., Strittmatter, A.: Heterogeneous employment effects of job search programmes: a machine learning approach. *J. Human Resources*, published online (2020)
- Knol, M.J., VanderWeele, T.J., Groenwold, R.H., Klungel, O.H., Rovers, M.M., Grobbee, D.E.: Estimating measures of interaction on an additive scale for preventive exposures. *European J. Epidemiology* **26**(6), 433–438 (2011)
- Koenker, R.: *Quantile regression.* Cambridge: Cambridge University Press (2005)
- Koenker, R., Bassett, G.: Regression quantiles. *Econometrica: J. Econometric Society* **1**, 33–50 (1978)
- Li, Y., Zhu, J.: L1-norm quantile regression. *J. Computat. Graph. Stat.* **17**(1), 163–185 (2008)

- Liao, L., Park, C., Choi, H.: Penalized expectile regression: an alternative to penalized quantile regression. *Annals Institute Stat. Math.* **71**, 409–438 (2019)
- Machado, J.A., Mata, J.: Counterfactual decomposition of changes in wage distributions using quantile regression. *J. Appl. Econ.* **20**(4), 445–465 (2005)
- Mandel, H., Semyonov, M.: Gender pay gap and employment sector: sources of earnings disparities in the United States, 1970–2010. *Demography* **51**(5), 1597–1618 (2014)
- Melly, B.: Public-private sector wage differentials in Germany: evidence from quantile regression. *Empirical Econ.* **30**(2), 505–520 (2005)
- Newey, W.K., Powell, J.L.: Asymmetric least squares estimation and testing. *Econometrica* **77**(1), 819–847 (1987)
- Philipps, C.: Interpreting expectiles. <https://doi.org/10.48550/arXiv.2204.10403> (2021)
- Riccobello, R., Bogdan, M., Bonaccolto, G., Kremer, P., Paterlini, S., Sobczyk, P.: Sparse graphical modelling via the sorted l_1 -norm. arXiv:2204.10403, <https://doi.org/10.48550> (2022)
- Sæbø, S., Almøy, T., Helland, I.S.: simrel — a versatile tool for linear model data simulation based on the concept of a relevant subspace and relevant predictors. *Chemometrics Intell. Laboratory Syst.* **146**, 128–135 (2015)
- Sobotka, F., Kauermann, G., Schulze Waltrup, L., Kneib, T.: On confidence intervals for semiparametric expectile regression. *Stat. Comput.* **23**, 135–148 (2013)
- SOEP: Data for years 1984–2017, version 34. <https://doi.org/10.5684/soep.v34> (2019)
- Stahlschmidt, S., Eckardt, M., Härdle, W.K.: Expectile treatment effects: an efficient alternative to compute the distribution of treatment effects (Tech. Rep.). SFB 649 Discussion Paper No. 2014–059 (2014)
- Taylor, J.W.: Estimating value at risk and expected shortfall using expectiles. *J. Financial Econ.* **6**(2), 231–252 (2008)
- Tibshirani, R.: Regression shrinkage and selection via the lasso. *J. Royal Stat. Society. Series B (Methodological)* **58**(1), 267–288 (1996)
- Wagner, G.G., Göbel, J., Krause, P., Pischner, R., Sieber, I.: Das sozio-oekonomische panel (SOEP): Multidisziplinäres Haushaltspanel und Kohortenstudie für Deutschland—Eine Einführung (für neue Datennutzer) mit einem Ausblick (für erfahrene Anwender). *ASTA Wirtschafts- und Sozialstatistisches Archiv* **2**(4), 301–328 (2008)
- Wilde, E.T., Batchelder, L., Ellwood, D.T.: The mommy track divides: the impact of childbearing on wages of women of differing skill levels (Tech. Rep.). National bureau of economic research working paper No. 16582 (2010)
- Wunsch, C., Strittmatter, A.: The gender pay gap revisited with big data: do methodological choices matter? (Tech. Rep.). CEPR Discussion Paper No. DP15840 (2021)
- Yang, Y., Zou, H.: Nonparametric multiple expectile regression via ER-Boost. *J. Stat. Computat. Simul.* **85**(7), 1442–1458 (2015)
- Yao, Q., Tong, H.: Asymmetric least squares regression estimation: a nonparametric approach. *J. Nonparametric Stat.* **6**(2-3), 273–292 (1996)