March 2026

WORKING PAPER SERIES

2026-EQM-01

Glass Ceilings and Sticky Floors in France: Quantile and Frontier Hedonic Wage Analysis of the Gender Pay Gap

Audrey DUMAS
Université de Perpignan, CDED, 52 Avenue Villeneuve, Perpignan F-66000, France, audrey.dumas@univ-perp.fr

Kristiaan KERSTENS
IESEG School of Management, Univ. Lille, CNRS, UMR 9221 - LEM - Lille Économie Management, F-59000 Lille, France, k.kerstens@ieseg.fr

Poleth VEGA RUALES
Université de Perpignan, CDED, 52 Avenue Villeneuve, Perpignan F-66000, France, Corresponding author, poleth.vegaruales@univ-perp.fr

IÉSEG School of Management Lille Catholic University 3, rue de la Digue F-59000 Lille Tel: 33(0)3 20 54 58 92
www.ieseg.fr

Glass Ceilings and Sticky Floors in France: Quantile and Frontier Hedonic Wage Analysis of the Gender Pay Gap

Audrey Dumas[*], Kristiaan Kerstens[†], Poleth Vega Ruales[‡]

March 17, 2026

Abstract

This paper examines the gender pay gap by focusing on the glass ceiling and sticky floor effects. First, we use parametric methods based on the Blinder-Oaxaca decomposition which measure the adjusted average wage difference between men and women. Second, for comparison, we propose an original approach based on a nonparametric double hedonic wage frontier. This method evaluates how far each individual's wage is from both the maximum (upper frontier) and minimum (lower frontier) potential salary they could obtain given their characteristics, thereby capturing the glass ceiling and sticky floor effects, respectively. Using 2022 French labour-market data, the nonparametric analysis shows that the gender pay gap is larger than what parametric methods suggest. We confirm that the glass ceiling effect is the more severe issue. Financial and insurance activities display the strongest glass ceiling and sticky floor effects, while employment in the public sector offers women better protection from these inequalities.

Keywords: Gender pay gap, Luenberger indicator, glass ceiling, sticky floor
JEL Classification: C14, J31, J71

[*]Université de Perpignan, CDED, 52 Avenue Villeneuve, Perpignan F-66000, France, `audrey.dumas@univ-perp.fr`

[†]Univ. Lille, CNRS, IESEG School of Management, UMR 9221 - LEM - Lille Économie Management, F-59000 Lille, France. `k.kerstens@ieseg.fr`

[‡]Corresponding author: Université de Perpignan, CDED, 52 Avenue Villeneuve, Perpignan F-66000, France, `poleth.vegaruales@univ-perp.fr`

1 Introduction

The gender pay gap (GPG) remains a persistent and deeply researched phenomenon in labour economics (see Olivetti et al. (2024) for a recent overview). Despite significant advancements in educational attainment, labour force participation, and legal protections, empirical evidence continues to demonstrate that women earn on average less than men across most sectors and countries. The drivers of this gap are complex and multifaceted, including factors such as occupational segregation, differences in human capital, work experience, institutional constraints, and, importantly, discriminatory practices. To inform evidence-based policy interventions aimed at achieving gender equity in the labour market, it is essential to understand both the size of the GPG as well as the mechanisms behind it.

Occupational segregation further exacerbates these disparities, with women often concentrated in lower-paid positions and facing barriers such as "glass ceilings" and "sticky floors", which limit career mobility (Cotter et al., 2001; International Labour Organization (ILO), 2018). Traditionally, the glass ceiling metaphor describes the unseen obstacles that hinder women from reaching senior leadership roles, despite having the requisite qualifications and experience. These barriers are often systemic, rooted in organizational cultures and practices that favour male advancement. Conversely, the sticky floor metaphor captures the difficulties women encounter at the lower end of the occupational hierarchy. It refers to the tendency of women to be concentrated in low-paying jobs with minimal prospects for career advancement (Espinosa and Ferreira, 2022).

One of the most influential tools used in the analysis of the GPG is the Blinder-Oaxaca (BO) decomposition (Blinder, 1973; Oaxaca, 1973), a method grounded in the neoclassical human capital model of wage determination. This approach decomposes the observed average wage differential between male and female workers into two components: one that can be explained by differences in observable characteristics (such as education, experience, or economic activity) and an "unexplained" component, often interpreted as a proxy for discrimination. The method has become a standard in labour economics due to its intuitive interpretability and compatibility with linear regression frameworks. However, it is inherently limited by its reliance on parametric assumptions and its focus on mean outcomes, which may fail to account for the full distributional structure of wages.

To address this limitation and provide a more nuanced analysis, recent extensions of the BO framework apply quantile regression techniques (see Machado and Mata (2005), and Firpo et al. (2009)). A quantile-based BO decomposition allows us to estimate the gender wage gap at different points of the wage distribution (e.g., at the 10th, 50th, and 90th percentiles), thereby capturing distributional heterogeneity. This adaptation has been especially relevant in labour markets characterized by glass ceilings and sticky floors, since it allows to analyse both the upper and lower ends of the wage distribution. The glass ceiling (sticky floor) phenomenon is established if the unexplained effect is largest or significantly wider at the upper (lower) quantiles of the wage distribution (Christofides et al., 2013).

Both concepts may oversimplify the complex nature of workplace inequality by segmenting the labour market into discrete zones of disadvantage, one at the top (i.e., the glass ceiling) and one at the bottom (i.e., the sticky floor). Espinosa and Ferreira (2022) challenge the traditional understanding of this glass ceiling by demonstrating that gender-based barriers are not confined to the uppermost levels of organizational hierarchies. In their study, they model candidate promotion processes using a Markov chain framework to examine how implicit gender biases can influence career advancement over time. Their findings reveal that even subtle asymmetries in decision-making processes can accumulate into significant disparities in representation and outcomes for men and women. Crucially, they argue that the glass ceiling should not be viewed solely as an obstacle to reaching top executive positions, but rather as a systemic phenomenon that manifests itself throughout the entire wage and occupational structure. This cumulative disadvantage results in persistent gender gaps not only at the top, but also in middle and lower segments of the wage distribution. The authors introduce a formal characterization of various types of glass ceilings, illustrating that gender disparities are a structural feature of biased systems rather than isolated to a specific echelon (Espinosa and Ferreira, 2022).

Blackaby et al. (2002) find out that while women are promoted at rates comparable to men, they often receive smaller wage increases upon promotion. This suggests that women are more likely to remain at the lower end of the wage distribution, even as they advance in their careers (i.e., a clear illustration of the sticky floor effect). However, the study reveals that this disparity in wage progression highlights a subtler but persistent mechanism of inequality: even when women advance within organizations, their economic rewards do not match those received by their male counterparts. This suggests that women's human capital, whether in the form of experience, education, or job performance, is not valued equally in the labour market. The authors interpret this as a key feature of the sticky floor effect. It is not merely about being "stuck" in entry-level roles, but about the systematic undervaluation of women's capabilities and contributions throughout the lower and middle tiers of the wage distribution.

Carli and Eagly (2016) propose that both metaphors may not adequately represent the multifaceted challenges women face in leadership roles. They introduce the "labyrinth" metaphor to depict the complex and often non-linear paths women navigate in their careers, marked by various obstacles and dead ends. This framework acknowledges that barriers exist at multiple levels and stages, not just at the top or bottom of the hierarchy.

In this study, we propose to examine the two phenomena across the entire wage distribution, arguing that these should be reconceptualized not as isolated or static barriers, but as evolving mechanisms that operate throughout the occupational and wage distribution. To do so, we develop a methodology based on hedonic wage frontiers.

Based on Rosen (1974) who is the first to study market equilibria for differentiated commodities differing along multiple characteristics, it has become rather common to estimate hedonic price functions relating prices (wages) to a range of characteristics (see Nesheim (2008)). A traditional estimation strategy to measure the GPG and discrimination is by considering a relationship between the average wage attainable by an individual given their human capital and other work characteristics. Polachek and Yoon (1987) introduce a highly innovative stochastic parametric frontier methodology that simultaneously looks at both extremes of the wage distribution. On the one hand, a relationship is established between the maximum wage attainable by an individual and their human capital and other work characteristics. On the other hand, an alternative relationship is estimated between the minimum wage obtained by an individual given its human capital and work characteristics. This estimation of an upper and a lower hedonic wage frontier seems to offer an adequate complementary perspective to study both the glass ceiling and the sticky floor effects, respectively.[1] In addition to parametric stochastic frontiers, also alternative nonparametric deterministic hedonic wage frontiers have become popular since the seminal work of Hadley and Ruggiero (2006).

For this purpose, the hedonic wage frontier methodology can contribute to an alternative estimation of the GPG and discrimination because it establishes a relationship between the maximum (minimum) wage attainable by an individual given their human capital and other work characteristics, instead of considering an average wage obtained by the estimation of some reduced wage equation. Then, the earnings function represents the relation between the human capital variables (inputs) and the maximum (minimum) wage attainable (output) and allows to compare the observed wage obtained by a worker with its potential or theoretical wage. Thus, both parametric stochastic and nonparametric deterministic frontiers assume that human capital characteristics are ideally rewarded the same way for men and women. Any differences between observed wages and theoretical wages at the frontier simply reveal informational inefficiencies (see Polachek and Robst (1998)), that may be labeled discrimination. This approach based on maximum or minimum wages is particularly relevant for capturing very exclusive high-salary “clubs” or low wage traps linked to glass ceiling and sticky floor effect.

Each of these parametric stochastic (see Kumbhakar and Lovell (2000)) and nonparametric deterministic (see Hackman (2008)) frontier methods have advantages and disadvantages. The main reason to privilege the nonparametric deterministic frontiers for this topic is that they do not impose any assumption on the functional form of the wage structure and therefore better respect the actual structure of wages. Indeed, the labour market has institutional rigidities such as minimum wages or collective bargaining agreements that imply irregularities in the wage structure, which may be particularly high at the top or the bottom of the wage distribution. Accordingly, and following some studies on hedonic price functions, we aim to test the convexity of the estimated hedonic wage frontiers (see, e.g., Chumpitaz et al. (2010) and Kerstens et al. (2024) for single and double frontier cases) in order to select the best model for computing the hedonic wage frontier. In the hedonic pricing literature it is rather known that (a) the implicit pricing functions for characteristics are not linear as commonly assumed by empirical studies and (b) in hedonic equilibrium the market does not need to offer a continuum of products. To our knowledge convexity has never been tested in a hedonic wage estimation context.

It is important to remark in this context that the frontier methodology traditionally assesses performance at the organizational level. However, more recently several studies have adapted this approach to the individual level: see, e.g., Ramos and Silber (2005) for measuring human development; Deutsch and Silber (2005) for evaluating poverty; and Johnes (2006) for an application in education.

Amado et al. (2018) develop an enhanced method to measure and decompose the gender pay gap based on nonparametric deterministic frontiers and the Malmquist productivity index under a constant returns to scale assumption and by measuring using an output orientation. The authors estimate a pay frontier representing the maximum pay that could be achieved for certain characteristics and measure the gender pay gap with a Malmquist index that compares male and female pay data. However, the use of a constant returns to scale

[1] An overview of this two-tier frontier model is found in Papadopoulos and Parmeter (2025).

assumption is unrealistic in the context of salary structures (see Chumpitaz et al. (2010)). Therefore, the objective of our work is to identify a methodology capable of capturing both effects (i.e., sticky floors and glass ceilings) within a more flexible and realistic framework.

The first contribution of this paper is to propose a hedonic wage frontier analysis by using a Luenberger productivity indicator (introduced by Chambers et al. (1996)) using an output oriented measurement, but dropping the restrictive constant returns to scale assumption. Irrespective of the returns to scale assumption, Boussemart et al. (2003, Proposition 2) convincingly show that an output-oriented Malmquist index overestimates productivity change compared to an output-oriented Luenberger indicator. Furthermore, the arithmetic structure of the Luenberger indicator permits a comparison with the additive BO decomposition. This Luenberger indicator approach allows us to evaluate wages relative to a benchmark defined by a maximum attainable salary. We then compare the results derived from this efficiency framework with those obtained from a quantile regression-based BO decomposition, which relies on mean-based counterfactuals. By contrasting these two methodologies, one frontier-based and the other rooted in conventional econometric decomposition, we aim to generate complementary insights into the structure of the GPG. This comparison also serves to assess the coherence and robustness of the hedonic wage frontier approach relative to a widely established technique in the analysis of the GPG.

Considering the Luenberger productivity indicator and output-oriented distances, our second contribution is to identify and measure glass ceilings. Therefore, we reinterpret the output-oriented distance from an individual to the efficiency frontier as a measure of the glass ceiling effect.

The third contribution is to use a double frontier approach (see Hadley and Ruggiero (2006)). This framework positions each individual's observed wage between a lower frontier (the minimum acceptable wage from the worker's perspective) and an upper frontier (the maximum wage an employer is willing to pay for this worker's characteristics). In a gender context, this enables to test whether for equivalent characteristics women systematically earn wages closer to the lower bound, which indicates a lower valuation of their qualifications, or a segregation into less competitive labour market segments or the presence of a sticky floor. Consequently, the output-oriented distance from an individual to the lower frontier is interpreted as a measure of the sticky floor effect.

A fourth contribution of this study lies in the explicit statistical testing of the convexity of the hedonic wage frontier. By formally testing whether the wage structure satisfies convexity, we in principle allow for a departure from the assumption adopted by Amado et al. (2018) who imposed a convex frontier under constant returns to scale. Using this statistical test allows us to develop a framework that potentially more accurately reflects the reality of hedonic wage formation.

In this study, we use harmonized micro data from the 2022 French Employment Survey, desegregated by economic activity. By examining all economic activities in the French labour market, we aim to demonstrate the applicability of our methodologies and contribute to a broader understanding of gender pay disparities.

The article is structured as follows. Section 2 explains the methodology. Section 3 describes the data and variables. Section 4 provides descriptive statistics. Section 5 presents the empirical findings. Finally, Section 6 offers the concluding remarks.

2 Methodology

Our method consists of two stages. First, we compute the BO decomposition, both at the mean and across quantiles, to capture the glass ceiling and sticky floor effects using a traditional counterfactual framework. Second, we focus on a double frontier nonparametric hedonic wage approach. We construct the hedonic wage frontiers under variable returns to scale. On the one hand, the upper frontier represents the maximum wage an employer is willing to pay for a worker's characteristics. On the other hand, the lower frontier represents the minimum acceptable wage from the worker's perspective given its characteristics. Moreover, we test for the convexity of these double hedonic wage frontiers to determine the shape of the wage distribution for adapting our methodology. Then, we calculate the three output-oriented Luenberger indicators: (i) to the upper frontier, using only men as the reference group in order to make our results directly comparable with those from the BO decomposition; (ii) to the upper frontier using both male and female as the reference group to measure the glass ceiling effect; and (iii) to the lower frontier using both male and female as the reference group to measure the sticky floor effect.

2.1 Blinder-Oaxaca decomposition

2.1.1 Original method

The most straightforward and widely used indicator of the gender wage gap is the difference in earnings between men and women measured as the extent to which women's wages are lower (or higher) than men's wages. This measure is referred to as the unadjusted GPG. As a metric of the wage distribution, both the mean (i.e., average wage) and the median (i.e., the midpoint of the wage distribution) are highly sensitive to how wages are defined, to the subpopulation examined, and to its overall composition. Consequently, this approach can overlook important differences across sectors, education levels, and other worker characteristics.

Due to the limitations of relying solely on unadjusted average wages to evaluate gender pay equality, a commonly used alternative is the BO decomposition. This method divides the observed GPG into two parts: (i) an explained component reflecting differences in observable characteristics (such as education and work experience), and (ii) an unexplained component which may be associated with discrimination or other unmeasured factors (Oaxaca, 1973). A key feature of this approach is the construction of counterfactual wages. For example, estimating what an average woman would earn if she were remunerated according to the male wage structure allows the decomposition to distinguish the part of the gap due to differences in characteristics from the part arising from differences in returns to those characteristics.

Figure 1 visually represents the two counterfactual scenarios, where Y denotes the wage, and X represents the explanatory variables for workers used to account for the variation in Y. The subscripts M and F refer to the male and female, respectively, allowing to measure the difference between the means of the earnings of men $\bar{Y}^M$ and women $\bar{Y}^F$. The counterfactual female outcome $Y_{C,F}$ is the outcome a female would receive if she had the male returns, but her own mean endowments $(\bar{X}_F)$. The decomposition is written as: $\bar{Y}_M - \bar{Y}_F = (\bar{Y}_M - Y_{C,F}) + (Y_{C,F} - \bar{Y}_F)$. The first term $(\bar{Y}_F - Y_{C,F})$ is the explained gap (endowment effect) and it captures the portion of the difference in Y attributable to the difference in mean characteristics $(\bar{X}_M - \bar{X}_F)$ when valued using the male coefficients. The second term $(Y_{C,F} - \bar{Y}_F)$ is the unexplained gap. This segment represents the variation caused by the differing returns to the same characteristic. Economically, this second component is often interpreted as a measure of discrimination.

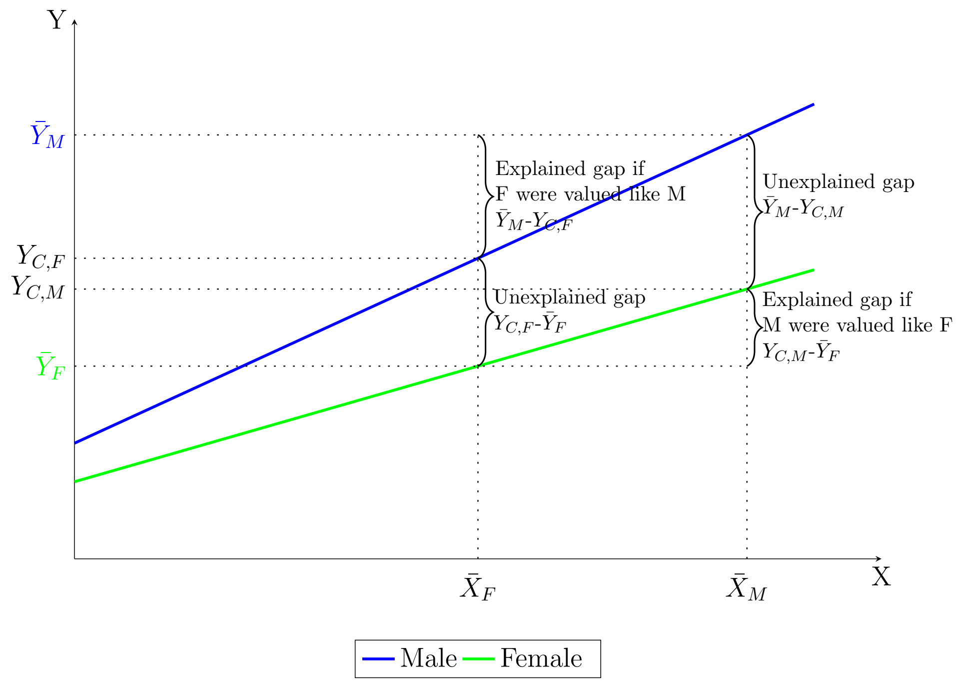

Figure 1: Blinder-Oaxaca decomposition of the wage gap between men and women

When using the female wage structure as the reference, the counterfactual male outcome $Y_{C,M}$ is the outcome a male would receive if he had the female returns, but his own mean endowments. The decomposition is written as $\bar{Y}_M - \bar{Y}_F = (\bar{Y}_M - Y_{C,M}) + (Y_{C,M} - \bar{Y}_F)$. The first term $(\bar{Y}_M - Y_{C,M})$ is the unexplained gap and captures the difference due to the application of differing returns to the male mean endowment $(\bar{X}_M)$. The second term $(Y_{C,M} - \bar{Y}_F)$ is the explained gap. It captures the portion of the difference attributable to the difference in characteristics $(\bar{X}_M - \bar{X}_F)$ when valued using the female coefficients.

The BO decomposition is based on two regression equations:

$$Y_i^M = \beta_0^M + \sum_{j=1}^{n} \beta_j^M X_{ij}^M + u_i^M \tag{1}$$

$$Y_i^F = \beta_0^F + \sum_{j=1}^{n} \beta_j^F X_{ij}^F + u_i^F \tag{2}$$

where Y_i^M and Y_i^F denote the natural logarithm of male and female wages and X_{ij}^M and X_{ij}^F represent the explanatory variables for male workers i and female workers i, respectively. Subsequently, the difference between the means of earnings of men $\bar{Y}^M$ and women $\bar{Y}^F$ is computed:

$$\bar{Y}^M - \bar{Y}^F = \hat{\beta}_0^M + \sum_{j=1}^{n} \hat{\beta}_j^M \bar{X}_j^M - \hat{\beta}_0^F - \sum_{j=1}^{n} \hat{\beta}_j^F \bar{X}_j^F \tag{3}$$

This equation shows the relationship between the mean of log earnings and the observed average characteristics for men and women ($\bar{X}_j^M$ and $\bar{X}_j^F$, respectively). In this relationship, the estimated constant $\hat{\beta}_0^M$ and coefficients $\hat{\beta}_j^M$ measure the financial returns to the characteristics of male employees, whereas $\hat{\beta}_0^F$ and $\hat{\beta}_j^F$ measure the returns for female employees.

The male wage structure is used as a reference wage. Therefore, the estimated constant and coefficients in the men's equation are treated as the non-discriminatory benchmarks for the financial returns to characteristics of employees. A counterfactual equation is constructed, where the coefficients from the male model are applied to the female's equation: $\hat{\beta}_0^M + \sum_{j=1}^{n} \hat{\beta}_j^M \bar{X}_j^F$, and inserted in the difference of means:

$$\begin{aligned}
\bar{Y}_i^M - \bar{Y}_i^F &= \left(\hat{\beta}_0^M + \sum_{j=1}^{n} \hat{\beta}_j^M \bar{X}_j^M - \left(\hat{\beta}_0^M + \sum_{j=1}^{n} \hat{\beta}_j^M \bar{X}_j^F\right)\right) \\
&\quad + \left(\left(\hat{\beta}_0^M + \sum_{j=1}^{n} \hat{\beta}_j^M \bar{X}_j^F\right) - \hat{\beta}_0^F - \sum_{j=1}^{n} \hat{\beta}_j^F \bar{X}_j^F\right) \\
&= \underbrace{\sum_{j=1}^{n} \hat{\beta}_j^M \left(\bar{X}_j^M - \bar{X}_j^F\right)}_{\text{Explained part}} + \underbrace{\left(\hat{\beta}_0^M - \hat{\beta}_0^F\right) + \sum_{j=1}^{n} \bar{X}_j^F \left(\hat{\beta}_j^M - \hat{\beta}_j^F\right)}_{\text{Unexplained part}}
\end{aligned} \tag{4}$$

This equation can be interpreted as what the average female worker would have earned if she had been paid on the same basis as an equivalent male worker (Leythienne and Ronkowski, 2018). The first component is the endowment effect $\sum_{j=1}^{n} \hat{\beta}_j^M \left(\bar{X}_j^M - \bar{X}_j^F\right)$ or the explained part. This term captures the difference between the mean male outcome and the counterfactual female outcome, representing the portion of the wage differential that is attributable to differences in mean characteristics ($\bar{X}_j^M - \bar{X}_j^F$) between the groups, valued at the male returns ($\hat{\beta}_j^M$). This explains the wage gap that arises due to observable differences in human capital factors.

The second component is the coefficient effect $\left(\hat{\beta}_0^M - \hat{\beta}_0^F\right) + \sum_{j=1}^{n} \bar{X}_j^F \left(\hat{\beta}_j^M - \hat{\beta}_j^F\right)$ or the unexplained part. This component isolates the difference between the counterfactual female outcome and the actual mean female outcome. It accounts for differences in the estimated returns (coefficients, $\hat{\beta}$) paid to men versus women for their characteristics. This unexplained component is commonly used to estimate the extent of wage discrimination or the effect of unobserved productivity differences (see Bazen and Charni (2023)).

2.1.2 Blinder-Oaxaca decomposition based on quantile regression

For this study, we implement an extended version of the BO decomposition, based on quantile regression. This approach enables the decomposition of wage differences at various points in the wage distribution, thereby providing insight into the glass ceiling and the sticky floor effects.

Therefore, to analyze the GPG beyond the mean and capture potential glass ceilings across the wage distribution, we apply a reweighting-based quantile decomposition following the approach of DiNardo *et al.* (1996). Let Q_τ^M and Q_τ^F represent the τ-th quantiles of male and female wage distributions, respectively. The wage gap at quantile τ is decomposed as:

$$Q_\tau^M - Q_\tau^F = \underbrace{\left(Q_\tau^M - Q_\tau^{(M|F)}\right)}_{\text{Unexplained part}} + \underbrace{\left(Q_\tau^{(M|F)} - Q_\tau^F\right)}_{\text{Explained part}} \tag{5}$$

The term $Q_\tau^{M|F}$ represents the counterfactual wage distribution: the wage that a female worker would earn at quantile τ if she faced the same wage structure as her male counterpart, but retained her own characteristics.

On the one hand, the explained component $(Q_\tau^{M|F} - Q_\tau^F)$ captures the part of the wage gap that can be attributed to observable differences in characteristics such as education, experience, and occupation. On the other hand, the unexplained component $(Q_\tau^M - Q_\tau^{M|F})$ reflects differences in returns to these characteristics. This residual is typically interpreted as indicative of discrimination or other structural inequalities affecting wage determination.

To construct the counterfactual distribution $Q_\tau^{M|F}$, we apply a reweighting function that adjusts the distribution of female characteristics to resemble that of males. The reweighting function is defined as:

$$w(X) = \frac{f(X \mid D = M)}{f(X \mid D = F)}, \tag{6}$$

where X denotes the vector of covariates and D is a binary indicator of group membership. In empirical applications, this density ratio is typically estimated by modelling group membership using a logit or probit specification. By applying these estimated weights to the female sample, we obtain a reweighted distribution that mirrors the male distribution of characteristics. This procedure isolates the contribution of the wage structure itself to the observed wage disparities. Consequently, the method enables a more detailed and distribution-sensitive examination of wage differentials, which is especially valuable when potential discrimination varies across the wage distribution rather than occurring uniformly (see Fortin et al. (2011)).

This method has been applied for the analysis of wage inequality by many authors. Machado and Mata (2005) provide one of the first explicit quantile regression-based decomposition applications to analyze the GPG in Portugal. Their work establishes the necessity of moving beyond the mean, showing that the GPG's evolution is non-uniform across quantiles and confirming that the role of unobserved skills is not constant. This finding has been largely supported across multiple countries by Arulampalam et al. (2006), who analyze ten European nations. These authors establish the glass ceiling as a widespread phenomenon in countries like the UK and Netherlands, while also identifying evidence of a sticky floor in other countries such as Germany and Portugal.

In their decomposition of unconditional quantiles, Firpo et al. (2009) show that distributional patterns missed by mean regressions are essential to understanding GPG dynamics in the United States. Subsequent applications include Redmond and McGuinness (2017) who analyse the GPG across 28 European Union countries and identify a U-shaped GPG across the distribution, and Blau and Kahn (2017) who show that -despite a narrowing average GPG- a substantial upper-tail gap persists in the United States.

2.2 Nonparametric deterministic hedonic wage frontiers

The inclusion of other forms of measurement of the GPG is very relevant because traditional measures may be limited since they depend heavily on the assumptions of linearity and the specific parametric specification selected. The nonparametric deterministic frontier approach is a methodology usually applied to position a firms' production relative to a frontier that represents, for instance, the maximum amount of outputs attainable with a given level of inputs.

Hofler and Polachek (1985) is the seminal contribution of a similar parametric stochastic frontier estimating a hedonic wage relation to measure the amount of inefficiency suffered by workers due to ignorance of informational inefficiencies in the labour market. Polachek and Robst (1998) offer independent confirmation that the inefficiency measured by a parametric stochastic frontier correlates very well with survey evidence on the knowledge of workers on the functioning of the labour market.

Our nonparametric deterministic hedonic wage frontier methodology can contribute to estimating the wage gap and discrimination. We adapt the methodology as if a worker is a producer that transforms its available resources (such as education and work experience, or human capital characteristics) into outputs (e.g., wage). Then, the earnings or hedonic wage function represents the relationship between the human capital variables (inputs) and the maximum wage attainable or the minimum wage acceptable from the worker's perspective (output) if one adopts a double frontier perspective. This allows for a comparison between the wage actually obtained by a worker and their potential or theoretical maximum or minimum wage offering a flexible and insightful approach to identify glass ceilings and sticky floors, respectively.

2.2.1 Construction of the nonparametric deterministic frontiers

Following Hadley and Ruggiero (2006), we apply a double hedonic wage frontier approach slightly adapted here to the analysis of the GPG. The strength of this framework lies in its ability to disentangle supply-side mechanisms, stemming from workers' wage expectations, from demand-side mechanisms, rooted in firms' willingness to pay. In this setting, each individual's observed wage is conceptualised as lying between two frontier functions: an upper frontier representing the maximum wage an employer would be willing to offer for that worker's characteristics, and a lower frontier representing the minimum wage the worker would accept. On the one hand, if men are found to earn closer to the upper frontier, this would indicate more favourable wage opportunities or contexts in which their human capital is more fully rewarded, consistent with glass ceiling dynamics. On the other hand, when for equivalent observable characteristics women tend to receive wages closer to the lower frontier then this signals a lower valuation of their qualifications or the presence of sticky floor effects.

We consider n individuals that use q inputs to produce m outputs. $X \in \mathbb{R}_+^q$ is the vector of inputs and $Y \in \mathbb{R}_+^m$ is the vector of outputs. The production possibility set, T, describes how inputs $x = (x_1, \ldots, x_q) \in \mathbb{R}_+^q$ are transformed into outputs $y = (y_1, \ldots, y_m) \in \mathbb{R}_+^m$.

The set of feasible production possibilities of an individual T is the set of all input-output vectors achievable for a given gender. It is defined as follows:

$$T = \{(x, y) \in \mathbb{R}_{+}^{q+m} : x \text{ can produce } y\} \tag{7}$$

The technology satisfies the following axioms:

A1: If $(0, 0) \in T$, $(0, y) \in T$, then $y = 0$, i.e., an input equal to zero cannot produce a non-zero output.

A2: The set $A(x) = \{(u, v) \in T : u \leq x, v \geq y\}$ is bounded for all $(x, y) \in T$, i.e., we cannot obtain infinite outputs from a finite quantity of inputs.

A3: T is closed.

A4: $\forall (x, y) \in T$, $u \geq x$, $0 \leq v \leq y$ implies that $(u, v) \in T$, i.e., fewer outputs can always be produced with more inputs and vice versa.

A5: T is convex.

The technology satisfies axioms $A1$-$A5$ and is implicitly subject to a variable returns to scale assumption and is defined as:

$$T_G = \left\{(x, y) : x \geq \sum_{l=1}^{n_G} t_l x_G^l, y \leq \sum_{l=1}^{n_G} t_l y_G^l, \sum_{l=1}^{n_G} t_l = 1, t_l \in \Lambda\right\} \tag{8}$$

where $G \in \mathcal{G} := \{F, M\}$ is a generic notation that can describe the female and male technologies, and:
(i) $\Lambda \equiv \Lambda^{\mathrm{C}} = \left\{t : \sum_{k=1}^{K} t_l = 1 \text{ and } t_l \geq 0\right\}$; (ii) $\Lambda \equiv \Lambda^{\mathrm{NC}} = \left\{t : \sum_{k=1}^{K} t_l = 1 \text{ and } t_l \in \{0, 1\}\right\}$.
Note that $A5$ is not always maintained: the activity vector t operates under constraints of either convexity Λ^{C} or nonconvexity Λ^{NC}. We explicitly test for convexity in the specification of the double nonparametric hedonic wage frontiers.

2.2.2 Luenberger productivity indicator

We apply the Luenberger indicator within a nonparametric frontier analysis framework to study the GPG. First, we construct a Luenberger indicator that is directly comparable to the BO decomposition. Since BO conventionally treats men as the non-discriminatory benchmark, we mirror this structure by evaluating both men and women against the male wage frontier. This captures the GPG in terms of the average distance of women, and of men, from the male hedonic wage frontier.

Second, we depart from this restriction. In order to analyse glass ceiling and sticky floor effects, we apply an approach that does not treat men as the privileged reference group. Instead, we estimate two frontiers jointly from the full sample of workers, irrespective of gender. The upper hedonic wage frontier represents the maximal wage that employers are willing to pay for a given set of characteristics, while the lower wage hedonic frontier represents the minimum acceptable wage for those same characteristics. Every individual's wage lies between these two bounds, thereby defining a feasible wage interval for workers with similar characteristics. If women with equivalent characteristics are found to lie systematically closer to the lower frontier, this would indicate a sticky floor mechanism: women receive wages nearer to the minimum feasible valuation of their skills. Conversely, if men tend to lie closer to the upper frontier, this is consistent with glass ceiling dynamics: men capture wage offers that approach the maximal valuation of their human capital. Output-oriented distances composing the Luenberger indicator evaluated to the upper and lower boundaries provide a quantitative measure of these patterns.

The output-oriented directional distance function (see Chambers et al. (1998)) is the map:

$$D : \mathbb{R}_+^{q+m} \times \mathbb{R}_+^{q+m} \to \mathbb{R} \cup \{-\infty, +\infty\} \tag{9}$$

which is defined by:

$$D(z, g) = \begin{cases} \sup\{\delta : (x, y + \delta g) \in T\} & \text{if } (x, y + \delta g) \in T \text{ for some } \delta \in \mathbb{R}, \\ -\infty & \text{otherwise .} \end{cases} \tag{10}$$

In this expression, δ is a non-negative scalar that quantifies the maximum feasible expansion of outputs y in the specified direction $g = (0, k)$, where k denotes the output expansion direction, determining the path along which efficiency is measured. The value of δ represents the distance to the frontier: a zero value indicates that the unit is efficient, while a positive δ signals inefficiency because of a positioning away from the frontier. A larger δ implies greater potential for improvement. In the empirical analysis we put the direction vector $g = (0, k)$ equal to the output values of the evaluated observation. This leads to a proportional interpretation for the output-oriented directional distance function that satisfies strong commensurability, an important property for productivity measurement (see Briec et al. (2022)). Figure 2 presents an example of the output-oriented distances relative to both an upper ($D_U(.)$) and a lower ($D_L(.)$) frontier.

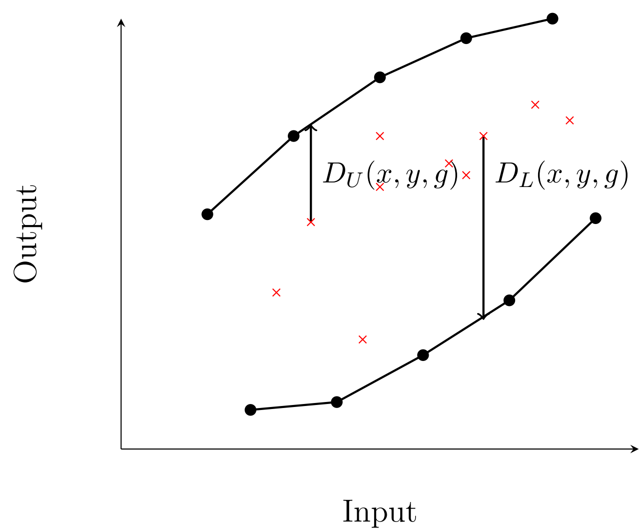

Figure 2: Output-oriented distances to the upper and lower hedonic wage frontiers

The traditional Luenberger indicator considers two periods, t and $t+1$, a distance given by the vector g, and one firm i, and is defined as follows:

$$L\big(x_t^i, y_t^i, x_{t+1}^i, y_{t+1}^i; g\big) = \frac{1}{2}\Bigg[\ D_{t+1}^i\big(x_t^i, y_t^i; g\big) - D_{t+1}^i\big(x_{t+1}^i, y_{t+1}^i; g\big) \\ + D_t^i\big(x_t^i, y_t^i; g\big) - D_t^i\big(x_{t+1}^i, y_{t+1}^i; g\big)\Bigg] \tag{11}$$

This indicator has an additive structure and takes the arithmetic average of two pairs of differences in distances. The first pair of distances evaluates observations in t and $t+1$ relative to a technology in $t+1$ represented by $D_{t+1}^i(., g)$. The second pair of distances evaluates observations in t and $t+1$ relative to a a technology in t represented by $D_t^i(., g)$. The arithmetic average taken with respect to both technologies reflects the indifference between the choice of base period. Note that the arithmetic structure of the Luenberger indicator (11) allows for a straightforward comparison with the equally additive BO decomposition (4).

To adapt this framework to the study of gender wage gaps, inspired by Färe et al. (2004) and Zaim (2004) we redefine the time dimension of the traditional Luenberger indicator to instead compare two groups: men (M), and women (F). Each observation of a woman is indexed by j, and every observation of a man is indexed by i. With $g = (0, k)$ representing the output-oriented direction, the static gender-based Luenberger indicator is given by:

$$L^{(F,M)}\big(x_F^{(j)}, y_F^{(j)}, x_M^{(i)}, y_M^{(i)}; g\big) = \frac{1}{2}\Bigg[\ D_M\big(x_F^{(j)}, y_F^{(j)}; g\big) - D_M\big(x_M^{(i)}, y_M^{(i)}; g\big) \\ + D_F\big(x_F^{(j)}, y_F^{(j)}; g\big) - D_F\big(x_M^{(i)}, y_M^{(i)}; g\big)\Bigg] \tag{12}$$

Each term measures an output-oriented distance. The term $D_M(x_F^{(j)}, y_F^{(j)}, g)$ represents the distance of a woman with characteristics $(x_F^{(j)}, y_F^{(j)})$ to a male frontier represented by $D_M(., g)$. The term $D_M(x_M^{(i)}, y_M^{(i)}, g)$ similarly measures the distance of a man to a male frontier. Similarly, $D_F(x_F^{(j)}, y_F^{(j)}, g)$ denotes the distance of a woman to a female frontier, while $D_F(x_M^{(i)}, y_M^{(i)}, g)$ represents the distance of a man with characteristics $(x_M^{(j)}, y_M^{(j)})$ to a female frontier represented by $D_F(., g)$.

Aggregating over all individuals in the sample, where the number of women (men) is represented by n_F (n_M), the static Luenberger indicator becomes:

$$L^{(F,M)} = \frac{1}{2}\Bigg[\ \sum_{j=1}^{n_F} D_M\left(x_F^{(j)}, y_F^{(j)}, g\right) - \sum_{i=1}^{n_M} D_M\left(x_M^{(i)}, y_M^{(i)}, g\right) \\ + \sum_{j=1}^{n_F} D_F\left(x_F^{(j)}, y_F^{(j)}, g\right) - \sum_{i=1}^{n_M} D_F\left(x_M^{(i)}, y_M^{(i)}, g\right)\Bigg] \tag{13}$$

The interpretation of these indicators is straightforward. If the value is zero, it indicates the absence of gender inequalities. A positive (negative) value suggests that gender inequalities favour men (women). The higher the absolute value of the indicator, the greater the degree of inequality.

To ensure comparability with the BO decomposition, where men are conventionally treated as the non-discriminatory benchmark, we construct an output-oriented Luenberger indicator using the male wage frontier as the sole reference technology. This modified indicator quantifies the GPG between women and men by evaluating both groups relative to the male hedonic wage frontier. Formally, it is defined as:

$$L_{BO}^{(F,M)} = \frac{1}{2}\left[\sum_{j=1}^{n_F} D_M\left(x_F^{(j)}, y_F^{(j)}; 0, k\right) - \sum_{i=1}^{n_M} D_M\left(x_M^{(i)}, y_M^{(i)}; 0, k\right)\right] \tag{14}$$

A positive value of $L_{BO}^{(F,M)}$ indicates that, on average, women lie farther from the male frontier, signaling a disadvantage in wage-setting conditions consistent with the presence of glass ceilings.

To measure glass ceilings and sticky floors, we extend the Luenberger framework by replacing the male-specific frontier with two wage frontiers estimated from the pooled sample of workers in each sector.[2] In this specification, men and women jointly determine the upper and lower hedonic wage frontiers, denoted by P. The upper pooled frontier represents the maximum wage employers are willing to pay for a given bundle of characteristics, while the lower pooled frontier captures the minimum acceptable wage from the worker's perspective. All observed wages lie within this interval; together, these two pooled frontiers define the feasible wage range for individuals with comparable characteristics. Systematic gender differences in the distances to these frontiers indicate whether women are disproportionately located near the lower bound (sticky floors) or whether men are disproportionately situated near the upper bound (glass ceilings).

Since the two boundaries represent opposite directions of adjustment, the vectors in the distance functions differ. For the upper frontier (U), the relevant transformation increases wages upward, and we define the output direction as $g^U = (0, k)$ with $k > 0$. The output-oriented Luenberger indicator for the upper frontier which captures barriers to reaching high-wage opportunities is computed as:

$$L_U^{(F,M)} = \frac{1}{2}\left[\sum_{j=1}^{n_F} D_U^P(x_F^{(j)}, y_F^{(j)}; 0, k) - \sum_{i=1}^{n_M} D_U^P(x_M^{(i)}, y_M^{(i)}; 0, k)\right]. \tag{15}$$

[2]The use of a pooled frontier estimated from the combined sample of men and women follows the nonparametric frontier literature when the objective is to evaluate outcomes relative to a common benchmark. Bowlin et al. (2003) construct a single compensation frontier encompassing both male and female executives in order to measure pay efficiency relative to a unified "best-practice" standard. Amado et al. (2018) also emphasize that comparisons of gender pay outcomes depend critically on the choice of reference technology. A frontier based on the entire sample can be interpreted as representing the maximum attainable wage conditional on characteristics irrespective of gender. In the context of the GPG, such a pooled frontier provides a non-discriminatory benchmark against which both groups can be evaluated. This framework allows one to assess whether women are systematically located farther from the upper bound or closer to the lower bound of the feasible wage set than men, corresponding to glass ceiling and sticky floor effects, respectively. Alternatively, it is eventually possible to compare the GPG with respect to a male frontier solely (as in (14)), or with respect to a separate male frontier and a separate female frontier separately and then take the arithmetic average of both measurements (as in (12) and (13)). Both of the latter options are kept for eventual future work.

Each term of this equation measures an output-oriented distance to the upper frontier. $D_U^P(x_F^{(j)}, y_F^{(j)}; 0, k)$ represents the distance of a woman to the upper pooled frontier, capturing how much her wage would need to increase to reach the maximal attainable wage given her characteristics. Similarly, $D_U^P(x_M^{(i)}, y_M^{(i)}; 0, k)$ measures the distance of a man to the upper pooled frontier. Positive values of this indicator $L_U^{(F,M)}$ indicate that women are, on average, farther away from the maximal wage frontier than men, consistent with a glass ceiling effect.

The lower frontier (L) instead forms a lower envelope, so the relevant transformation moves wages downward. The direction vector becomes $g^L = (0, -k)$ and the Luenberger indicator for the lower frontier, capturing undervaluation of women's human capital, is computed as:

$$L_L^{(F,M)} = \frac{1}{2}\left[\sum_{i=1}^{n_M} D_L^P(x_M^{(i)}, y_M^{(i)}; 0, -k) - \sum_{j=1}^{n_F} D_L^P(x_F^{(j)}, y_F^{(j)}; 0, -k)\right], \tag{16}$$

where $D_L^P(x_F^{(j)}, y_F^{(j)}; 0, -k)$ denotes the distance of a woman to the lower pooled frontier, reflecting how close her wage is to the minimum acceptable level, while $D_L^P(x_M^{(i)}, y_M^{(i)}; 0, -k)$ represents the distance of a man to the same lower pooled frontier. Positive values this indicator $L_L^{(F,M)}$ indicate that women are systematically closer to the lower frontier than men, highlighting the presence of sticky floors.

2.2.3 Selective review of empirical hedonic wage frontiers related to GPG

The existing empirical work applying hedonic wage frontiers to GPG can be broadly divided into two methodological strands: parametric stochastic and nonparametric deterministic frontier estimation. Both strands typically estimate the distance between observed wages and an earnings frontier that reflects the maximum attainable wage conditional on observable characteristics. In this study, we extend this perspective by also considering the minimum wage observed, or feasibly attainable, for comparable workers.

The parametric strand consists primarily of stochastic frontier analysis (SFA) models. These approaches specify a functional form for the wage–characteristics relationship and impose distributional assumptions. Empirical studies in this strand estimate gender-specific frontiers, derive scores for each worker, and compare the distributions of these scores across male and female workers. Early and influential empirical implementations established the use of frontier methods for discrimination analysis. Hofler and Polachek (1985) use this method for the first time to study the existence of imperfect information and to measure its effect on workers' actual wages. Hofler and Murphy (1992) also employ this technique to talk about "underpayment" in the job market. Robinson and Wunnava (1989) also represent one of the first applications of a SFA approach to measure direct discrimination by estimating a female earnings frontier and treating systematic shortfalls as evidence of underpayment relative to potential earnings. Subsequent work extends this approach across countries, refines selection corrections, and compares group-specific frontiers.

The most recent study using SFA is Garcia-Prieto and Gomez-Costilla (2017) who measure wage discrimination against women in Spain and identify three sources of gender wage differentials: differences in human capital, differences in the probability of reaching the high wage class (associated with a glass ceiling), and differences in the ability to reach the potential wage. These results are consistent with other studies (Díaz and Sánchez, 2011; Bishop et al., 2007) using the same method, showing that a significant part of the GPG in all the countries analysed is not attributable to differences in human capital endowment or personal and job-related characteristics.

The nonparametric strand includes deterministic frontier estimators. This approach relaxes functional-form and distributional assumptions and constructs the hedonic wage frontier directly from the upper envelope of observed earnings. Mohan and Ruggiero (2003, 2007), Bowlin et al. (2003) and Bowlin and Renner (2008) initiate the use of nonparametric hedonic wage frontiers to estimate differences in pay between male and female workers. They find out that female executives generally appear less efficient relative to the pooled frontier, implying lower observed pay relative to potential pay.

Amado et al. (2018) use this approach to measure and decompose the GPG. In their study, the authors established gender-specific frontiers representing the maximum achievable pay given a set of inputs or productive traits. They compare the position of male and female workers in relation to these hedonic wage frontiers, identifying whether women are systematically less efficient in converting their productive attributes into pay, or if they face pay penalties despite similar efficiency levels. This method captures heterogeneity in individual performance and helps isolate the effect of gender on pay disparities, avoiding some of the parametric assumptions that standard econometric models require. The empirical analysis focuses specifically on business and administrative associate professionals working in the financial and insurance industry across 20 European countries. Their results reveal the existence of a GPG in all countries, while the value of the gap, and of its components, varies considerably among countries.

Furthermore, their study finds that France exhibits a relatively low Malmquist-based GPG of 9.3%, placing it among the countries with the smallest adjusted GPG in their sample. However, the decomposition of the gap reveals that this low value arises from an atypical configuration of hedonic wage frontiers rather than from strong pay outcomes for either gender. In France, female workers are closer to their gender-specific pay frontier than male workers are to their male frontier. This indicates that women's pay, given their education and tenure, is relatively close to the maximum levels observed for women across Europe. By contrast, French male workers receive only 56.8% of the pay observed in the best-paying countries for comparable male workers, positioning them substantially below the male hedonic wage frontier. Therefore, this study emphasizes that France's low GPG reflects depressed male wages rather than equitable gender pay outcomes. This characteristic must be acknowledged by policymakers: lowering male pay can mechanically reduce the measured GPG without improving women's earnings. Thus, the decomposition highlights France as a case where the GPG appears small, but only because men are unusually far removed from their highest attainable hedonic wage levels, and not because women fare particularly well.

A key difference between Bowlin et al. (2003) and Amado et al. (2018) lies in how the frontiers are constructed and interpreted. Bowlin et al. (2003) build a single hedonic wage frontier that encompasses both male and female executives, using firm-level inputs such as company size and performance to benchmark compensation. This common frontier represents the best practice in translating firm characteristics into executive pay, allowing for a direct comparison of pay efficiency across genders relative to the same standard.

The approach of Amado et al. (2018) recognizes that men and women may operate under different conditions or labour market constraints and thus allows for a comparison of individual pay efficiency relative to their own gender's best practice. Subsequently, differences between these gender-specific hedonic wage frontiers highlight structural disparities or discrimination affecting pay outcomes beyond individual productivity.

3 Data source and variables

3.1 Labour Force Survey

This study uses micro data from the 2022 French Labour Force Survey (Enquête Emploi), conducted by the National Statistics Office (INSEE). The Labour Force Survey is a nationally representative household survey that provides detailed individual-level data on employment, income, education, and demographic characteristics. It is the primary source for analysing labour market dynamics in France. This survey is carried out quarterly and contains approximately 92,000 households. Sampling is random within strata defined by region. The gross response rate is higher than 80%. The data contains a rich set of wage determinants that are considered important in the GPG literature such as age, education, hours usually worked and economic activity. The 2022 dataset includes 348,964 individuals, with 166,128 individuals participating actively in the labour market.

We analyse gender inequalities by examining the distribution of individuals across different economic activities, further distinguishing whether they work with a public or private status within key industries such as education, health, and public administration. These distinctions are essential for two main reasons.

Firstly, the ways in which inputs, such as skills, qualifications, and experience, are transformed into outputs like wages, promotions, and job security vary considerably depending on the economic activity and institutional sector. For example, activities such as professional and scientific may place greater emphasis on formal qualifications and specialised training, whereas sectors like administration may prioritise seniority and standardised career progression. Additionally, public status salaries are often governed by specific regulations, including pay scales and recruitment practices, which differ markedly from the more market-driven mechanisms typical of the private sector. These institutional differences can either alleviate or intensify gender inequalities depending on how they interact with any eventual systemic biases and labour market dynamics.

Secondly, understanding the role of economic activity is crucial when analysing gender inequalities, particularly in relation to gender pay gaps. Economic activity largely determines the types of roles individuals occupy, the skills required, and the working conditions encountered—all of which have a direct impact on earnings disparities. Research, such as that conducted by Eurostat (2022), shows that differences in economic activity are a primary factor explaining gender pay gaps in France. Their findings highlight that sectors dominated by women often feature lower average wages and fewer opportunities for advancement, which contribute significantly to persistent pay disparities. Furthermore, the segregation of men and women into different industries and sectors, known as occupational segregation, reinforces these gaps by channeling women into lower-paid roles, often within the public sector's distinctive frameworks.

3.2 Hourly earnings and explanatory variables

The analysis employs a unified set of variables across the two methodological frameworks to ensure comparability of insights. While the quantile regression framework can only handle a single dependent variable, the nonparametric hedonic frontier framework can in principle combine multiple inputs (variables for which less is better) and outputs (variables for which more is better). However, in our data set we only have a single output or outcome, namely hourly wages. The dependent variable is the hourly earnings, calculated by dividing reported net monthly earnings by usual monthly hours worked.[3] Hourly earnings are a standard measure of labour market returns and are widely used in wage gap studies (Mincer and Polachek, 1974; Blau and Kahn, 2017). This measure avoids biases associated with variations in working time and better captures individual productivity on a per-hour basis.

The explanatory variables used to characterize human capital and employment history of workers, both of which are theoretically grounded in human capital theory (Becker, 1964) and empirically validated in wage determination literature (Mincer and Polachek, 1974; Altonji and Shakotko, 1987), are years of education, work experience and tenure.

Years of education serve as a proxy for formal human capital investment and are a key determinant of wage differentials across individuals. This variable is derived from the individual's highest educational attainment, using standard mappings of French degrees to years of schooling (for instance, twelve years for a high school degree ("baccalauréat") and seventeen years for a master's degree). Numerous studies (e.g., Blau and Kahn (2006); Van Ophem and Mazza (2024); OECD (2025); Bar-Haim et al. (2023)) have shown that education significantly raises productivity and earnings, although the magnitude of its return often differs by gender and occupation.

Job tenure represents the number of years that the individual has been working in his or her current position. Work experience (excluding current position) is approximated by subtracting years of education and the typical school-starting age of six from the respondent's current age. It captures the accumulation of general labour market knowledge and skills over time. This proxy is widely used in empirical labour economics (Mincer, 1974) and remains one of the strongest predictors of earnings, especially in the context of GPG decompositions. Although not a perfect measure, this widely used proxy provides a reasonable estimate of total labour market exposure.

3.3 Data cleaning and sample selection

The data cleaning process for the labour force survey aims to construct a reliable and analytically consistent sample of employed individuals. The initial stage involves removing observations with missing responses in key areas, including earnings, job characteristics, education, and employment history. Entries with logically inconsistent or implausible values, such as future-dated events or placeholder codes, are also excluded.

[3]The usual monthly hours worked are calculated by multiplying the usual weekly hours worked by 4.33, which is the legal number of weeks.

The sample is then restricted to individuals who were employed at the time of the survey, excluding those unemployed or out of the labour force. To align with conventional definitions of the working-age population in France, only individuals between the ages of 15 and 64 are retained.

To ensure comparability across observations, the sample is further refined to include only respondents with a single job. Occupations not classified or related to military service are excluded, as were individuals working in sectors that are either marginal or not representative of the broader economy. The analysis focuses exclusively on full-time civilian employees, thereby removing the self-employed, unpaid workers, and part-time workers.

Additional sector-specific adjustments are applied to the variable indicating usual weekly working hours. Respondents employed in the public education sector and reporting fewer than 35 hours are re-coded to 35 hours, to reflect the standard legal hours in that sector. More generally, observations with implausibly low or high usual working hours (i.e., < 35 or > 60 hours per week) are excluded to ensure data quality and comparability.

Earnings data underwent extensive cleaning to remove outliers and potential misreports. Monthly earnings are trimmed by excluding observations below the 5th percentile (less than €985) and above the 99th percentile (greater than €7,000). A more conservative threshold is applied to the lower end of the wage distribution, reflecting the fact that wages in France are typically bounded below by the statutory minimum wage (SMIC), and thus implausibly low values are more likely to result from reporting errors. After these restrictions, hourly wages are computed by dividing monthly income by usual hours worked. Further filtering is applied to this derived variable by removing the bottom 1% of observations (hourly wage $<$ €7.65), accounting for likely inconsistencies or exceptional contractual arrangements (e.g., apprentices, or minors legally paid below the SMIC).

Final consistency checks are performed to eliminate any remaining anomalies, such as unrealistically high working hours. The resulting data set comprises a clean, targeted sample of full-time employees of working age, suitable for the descriptive statistics, econometric and hedonic wage frontier analysis of the GPG. The final analytical sample includes 14,300 individuals, of whom 6,661 are women and 7,639 are men, ensuring sufficient representation across genders for this study.

4 Descriptive statistics

This section presents an overview of the sample used in the analysis. In particular, we focus on key descriptive statistics related to the key variables both for the total sample and disaggregated by economic activity.

Table 1 shows the descriptive statistics for the variables involved in our analysis of the entire sample by gender. The hourly earnings are expressed in euros, while the other variables involve the number of years. First, educational attainment is slightly higher among women who average 13.6 years of schooling compared with 12.8 years for men. The dispersion of education as measured by the coefficient of variation is identical across genders (0.2), suggesting a comparable heterogeneity in educational backgrounds.

Second, regarding labour market experience, both women and men report the same average level of work experience (22.9 years), with very similar medians and variability. This

Table 1: Descriptive statistics of the variables by gender

	Female				Male			
	Education	**Work experience**	**Tenure**	**Net hourly wage**	**Education**	**Work experience**	**Tenure**	**Net hourly wage**
Mean	13.6	22.9	12.8	13.1	12.8	22.9	12.0	13.9
Median	12	23	9	11.9	12	23	9	12.5
Coefficient of variation	0.2	0.5	0.9	0.3	0.2	0.5	0.9	0.4
Minimum	7	0.1	0.1	7.7	7	0.1	0.1	7.7
Maximum	20	49	46	40.4	20	49	45	46.2

indicates that the sample includes men and women at comparable career stages. However, tenure exhibits a modest gender gap, with women averaging 12.8 years in their current job compared to 12.0 years for men. This suggests slightly greater job stability or longer employment relationships among women in the sample.

Third, the wage statistics reveal more substantial differences. Men earn a higher mean net hourly wage (€13.9) than women (€13.1), and the gap is also visible in the medians (€12.5 vs. €11.9), implying an unadjusted gender wage gap of approximately 6.2% in favour of men. While the coefficient of variation is somewhat higher for men (0.4 vs. 0.3) indicating greater wage dispersion among male workers, the minimum and maximum values show that the highest wages in the sample are concentrated among men.

Turning to disaggregated statistics by economic activity and institutional status in Table 2, gender distributions and characteristics vary substantially across sectors. A first striking feature is the pronounced gender segregation across industries. Women are heavily concentrated in sectors such as education (68.4%) and public human health and social work (79.0%), while they remain a minority in traditionally male-dominated sectors like construction (13.5%), transportation (28.4%), manufacturing (28.6%) and information and communication (33.4%).

Women generally exhibit equal or higher levels of schooling than men, with a few exceptions in professional and technical sectors, such as information and communication or financial and insurance activities, where men slightly surpass women by only 0.4-0.5 years on average. The gender gap is more pronounced in male-dominated sectors, such as construction, where women average 13.3 years of schooling compared with 11.3 years for men.

Differences in work experience appear generally smaller than those in education, but women tend to have slightly lower average experience in several industries, such as administrative and support services (women: 21.2 years vs. men: 22.9 years) and human health and social work in the private sector (women: 21.8 years vs. men: 22.1 years). These gaps may reflect career interruptions, part-time work, or different labour market trajectories. By contrast, public sector activities, such as public administration, education, and human health and social work, exhibit high levels of experience for both genders, often exceeding 23 years, and the gender gap largely narrows or disappears, indicating more stable and continuous employment for women in these sectors.

Tenure patterns also reveal sectoral segmentation. Sectors with traditionally stable or career-oriented employment structures, such as public administration, education, and health, display the longest average tenure, often exceeding 15 years for both genders. This reflects the prevalence of long-term employment relationships in public sector occupations. In con-

Table 2: Descriptive statistics of the variables by gender and economic activity

Economic activity	Number of obs.	Sex	Number of obs.	% of women	Mean years of education	Mean years of work experience	Mean years of tenure	Mean net hourly wage
Total	14300	Female	6661	46.6	13.6	22.9	12.8	13.1
		Male	7639		12.8	22.9	12.0	13.9
Manufacturing	2198	Female	628	28.6	12.9	23.6	11.2	12.9
		Male	1570		12.5	24.0	13.2	14.1
Construction	1031	Female	139	13.5	13.3	22.4	9.2	12.2
		Male	892		11.3	22.8	8.9	12.6
Wholesale and retail trade	1990	Female	804	40.4	12.9	21.2	10.1	12.0
		Male	1186		12.2	21.9	10.0	13.0
Transportation and storage	894	Female	254	28.4	12.5	25.1	14.5	12.7
		Male	640		11.6	24.9	12.3	12.9
Accommodation and food service	426	Female	201	47.2	12.1	20.4	8.0	10.8
		Male	225		11.7	20.8	7.1	11.5
Information and communication	626	Female	209	33.4	14.9	20.3	10.2	15.4
		Male	417		15.4	18.2	9.1	17.5
Financial and insurance	598	Female	350	58.5	14.5	21.6	12.9	14.2
		Male	248		14.9	21.2	13.0	16.8
Professional, scientific and technical	988	Female	497	50.3	15.2	19.3	8.5	14.8
		Male	491		15.2	18.7	9.1	16.1
Administrative and support service	566	Female	240	42.4	13.1	21.2	8.6	11.8
		Male	326		11.9	22.9	9.2	12.3
Public status								
Public administration	1395	Female	683	49.0	13.4	26.5	18.1	13.2
		Male	712		12.8	26.7	18.2	14.0
Education	1287	Female	880	68.4	14.9	25.0	18.8	14.2
		Male	407		15.4	24.4	18.6	15.3
Human health and social work	984	Female	777	79.0	13.1	23.1	15.3	13.2
		Male	207		13.4	23.7	16.3	14.1
Private status								
Public administration	200	Female	143	71.5	13.9	25.8	14.6	12.7
		Male	57		13.8	23.4	16.0	14.2
Education	157	Female	92	58.6	14.2	22.6	8.4	12.8
		Male	65		13.6	23.4	9.0	13.0
Human health and social work	960	Female	764	79.6	12.8	21.8	9.3	11.8
		Male	196		12.8	22.1	8.6	12.6

trast, sectors such as accommodation and food service, administrative support, and certain commercial activities exhibit substantially lower levels of tenure, frequently below 10 years. These shorter tenure profiles are consistent with higher mobility, more fixed-term or seasonal contracts, and generally less stable employment conditions. Across most sectors, men tend to have slightly longer tenure than women, although the differences are modest and vary by activity.

The largest gender gaps are observed in net hourly wages. Men earn higher wages than women across nearly all sectors, even in those sectors where women have equal or higher levels of education. For example, in professional, scientific, and technical activities, where women's qualifications and experience are comparable to men's, male average hourly wage is still 1.3 euros higher. The pay gap is particularly pronounced in male-dominated sectors such as manufacturing, construction, and information and communication. Even in low-wage service industries, such as accommodation and food services, where overall pay is relatively low for both genders, the unadjusted gender wage gap persists.

5 Empirical Results

This section presents a comprehensive analysis of the GPG in France in 2022. We compare the results of the Luenberger productivity indicator with a BO decomposition by quantiles, and then we apply a double-frontier approach to detect potential glass ceiling and sticky floor effects. Prior to conducting the main analysis, we formally test the axiomatic assumption of convexity. The results are further disaggregated by economic activity to highlight sector-specific patterns.

5.1 Testing convexity of the nonparametric hedonic wage frontiers

Table 3: Convexity test results for double hedonic wage frontiers

Economic Activity	Frontier	T_{KSW} (p-value)	T_{SW} (p-value)	Conclusion
Manufacturing	U	0.112	0.239	Convex
	L	0.605	0.814	Convex
Construction	U	0.931	0.501	Convex
	L	0.367	0.648	Convex
Wholesale and retail trade	U	0.190	0.380	Convex
	L	0.275	0.472	Convex
Transportation and storage	U	0.518	0.706	Convex
	L	0.694	0.835	Convex
Accommodation and food service	U	0.145	0.203	Convex
	L	0.311	0.109	Convex
Information and communication	U	0.770	0.903	Convex
	L	0.612	0.738	Convex
Financial and insurance	U	0.401	0.592	Convex
	L	0.322	0.449	Convex
Professional, scientific and technical	U	0.850	0.661	Convex
	L	0.719	0.940	Convex
Administrative and support service	U	0.298	0.495	Convex
	L	0.171	0.344	Convex
Public status				
Public administration	U	0.573	0.671	Convex
	L	0.468	0.555	Convex
Education	U	0.750	0.842	Convex
	L	0.801	0.910	Convex
Human health and social work activities	U	0.395	0.528	Convex
	L	0.240	0.337	Convex
Private status				
Public administration	U	0.621	0.700	Convex
	L	0.540	0.655	Convex
Education	U	0.138	0.211	Convex
	L	0.305	0.408	Convex
Human health and social work activities	U	0.899	0.925	Convex
	L	0.785	0.830	Convex

Our approach utilizes the Luenberger indicator and we perform a rigorous check of the convexity axiom for both the estimated upper and lower hedonic wage frontiers. When the relationship between human capital inputs and the maximum (minimum) attainable wage satisfies convexity, then its boundary can be adequately represented by a convex piecewise linear upper (lower) frontier. When this relationship fails convexity, then this boundary can only be adequately represented by a monotone upper (lower) frontier.

We employ two distinct bootstrap-based procedures to test the null hypothesis of convexity (H_0: T is convex) against the alternative hypothesis of nonconvexity. The first test proposed by Kneip et al. (2016) T_{KSW} uses a smooth bootstrap procedure and is known for its robustness. The second test is based on the refined methodology developed by Simar and Wilson (2020) (T_{SW}) and provides a complementary perspective on structural nonconvexity by implementing multiple (instead of a single) sample splits. For each economic activity and institutional status (public/private), we test the convexity assumption separately for the upper frontier (U) and the lower frontier (L).

The results of the convexity tests are summarized in Table 3. Across all tested economic activities, the null hypothesis of convexity cannot be rejected for both the upper and lower hedonic wage frontiers by both the T_{KSW} and T_{SW} tests. Note that this table summarizes the p-values from these two nonparametric tests (from Kneip et al. (2016) and Simar and Wilson (2020a)) whereby p-values > 0.01 indicate failure to reject the null hypothesis of convexity. This robust finding suggests that the structure of wage formation within the French labour market, as modeled by our specification of the hedonic wage frontier, naturally satisfies the convexity axiom, perhaps because of the large sample size and the low amount of dimensions (mitigating the curse of dimensionality typical for nonparametric methods). Therefore, the use of the convexity assumption in the construction of our double hedonic wage frontier is statistically justified for all subsequent calculations.

5.2 Results of the three output-oriented Luenberger indicators

This subsection presents the results comparing the BO decomposition using quantile regression and the three output-oriented Luenberger indicators derived from the hedonic wage frontier analysis in Table 4. The first column indicates the economic activity sector. The next two columns indicate the unadjusted and the BO-adjusted GPG. Thereafter, we present quantile regression results for nine deciles of the wage distribution. Finally, the outcomes across all three Luenberger indicators, disaggregated by economic activity, are summarized in the final three columns of Table 4. At a glance, these results reveal a consistent pattern of gender pay disparities across various economic sectors, with notable differences in the magnitude and distribution of the gaps. Generally, the unadjusted GPGs tend to increase after applying the adjusted measures, indicating that observed worker characteristics do not fully explain the wage differences.

5.2.1 Luenberger indicator comparable with the Blinder–Oaxaca decomposition

First, we comment on the conventional BO adjusted mean gap and the adjusted mean gaps by quantile, which represent the average unexplained difference in returns. We observe that the adjusted BO gap exceeds the unadjusted gap for most activities, with the exception of information and communication, financial and insurance, education, and human health and

Table 4: GPG by economic activity: Blinder-Oaxaca decomposition and Luenberger Productivity Indicators (%)

			BO-Adjusted GPG by percentile									Luenberger Indicator		
Economic activity	**Unadjusted GPG**	**BO-Adjusted GPG**	**10%**	**20%**	**30%**	**40%**	**50%**	**60%**	**70%**	**80%**	**90%**	**Comparable with BO-Adjusted GPG**	**Measure of the glass ceiling**	**Measure of the sticky floor**
Manufacturing	8.9	10.2	5.3	10.9	11.8	12.5	12.7	11.5	11.9	10.3	11.3	12.7	13.2	7.2
Construction	3.1	9.1	6.6	10.4	10.5	11.3	10.6	12.1	14.5	19.7	24.7	15.0	15.0	7.5
Wholesale and retail trade	7.7	11.8	3.5	6.1	9.8	10.6	11.3	11.9	14.6	16.8	20.1	14.1	14.1	6.7
Transportation and storage	1.3	8.0	4.0	3.3	3.5	6.7	8.6	11.5	13.2	16.3	12.8	9.3	9.3	3.1
Accommodation and food service	6.3	9.5	-2.0	0.4	1.4	4.0	5.6	5.2	2.5	10.7	25.9	6.7	6.5	3.3
Information and communication	11.6	11.1	8.9	10.8	9.8	9.6	7.9	10.9	9.3	7.0	11.2	13.7	12.6	9.5
Financial and insurance	15.3	13.3	5.6	10.1	10.7	14.9	12.8	14.4	13.5	16.1	20.8	16.1	19.6	11.4
Professional, scientific and technical	8.1	8.8	5.4	8.6	9.7	9.5	10.0	10.2	11.6	12.8	12.9	10.2	10.4	6.6
Administrative and support service	4.6	12.2	1.0	1.0	4.0	6.5	6.0	11.8	15.2	19.5	30.6	12.8	11.5	4.5
Public status														
Public administration	5.9	9.5	8.5	8.3	8.2	10.1	8.1	8.4	9.1	11.6	12.2	11.5	10.8	5.2
Education	7.3	5.5	2.8	1.5	4.0	6.5	6.6	4.7	3.8	5.2	6.1	8.9	7.3	4.2
Human health and social work	6.4	3.9	6.4	4.7	3.1	0.3	0.4	0.0	2.6	0.0	6.7	2.6	2.9	2.6
Private status														
Public administration	10.2	11.3	12.6	14.3	10.7	9.1	13.7	13.0	11.7	6.1	22.2	12.9	9.5	11.2
Education	1.7	1.7	2.9	-0.6	1.4	4.3	8.8	9.3	6.0	-0.6	5.6	8.3	7.1	3.1
Human health and social work	6.1	6.8	-0.1	5.0	0.7	1.3	4.4	4.5	7.6	12.0	18.5	6.3	8.0	3.3

social work under public status. This pattern reflects the fact that women tend to have higher educational attainment than men.

When turning to the BO results by quantile, particularly at the 80–90 percent range, we find that the GPG is substantially larger at the upper end of the wage distribution, confirming a glass ceiling effect across nearly all economic activities. Exceptions include manufacturing, public administration (public status), and education. At the lower end of the wage distribution, the GPG is generally similar or smaller, providing little evidence of a sticky floor effect, except in human health and social work with public status. Overall, traditional GPG metrics in France highlight a pronounced glass ceiling effect and no meaningful sticky floor effect. These findings contrast with those of Arulampalam et al. (2006), who report evidence of both effects for France using European Union household panel data.

Second, we compare these findings with the first Luenberger indicator, which uses the male wage frontier as the non-discriminatory reference technology. This latter indicator, by quantifying the average output-oriented distance women are situated from the maximum attainable wage, is interpreted as a robust measure of the overall glass ceiling effect in the two-group context.

The most salient finding is the systematic difference in the magnitude of the measured unexplained gap across the two methods. For instance, while the BO-adjusted GPG ranges from 1.7% (private education) to 13.3% (financial and insurance), the comparable Luenberger indicator ranges from 2.6% (public health) to a higher 16.1% (financial and insurance). This disparity suggests that the conventional BO approach, constrained by its reliance on mean regression estimates, may understate the economic cost of structural wage inequality.

As noted in the previous paragraph, the sector with the highest BO-adjusted gap, financial and insurance (13.3%), exhibits a glass ceiling effect of 16.1% when evaluated against the nonparametric hedonic wage frontier. This result is strongly supported by (Amado et al., 2018), who measure the glass ceiling primarily via the cross-gender frontier gap. Their study indicates a 12.4% gap in the financial and insurance industry, reflecting a persistent difference between the maximum attainable pay for male and female workers.

In sectors characterised by tightly regulated wage structures, such as human health and social work under public status, the BO-adjusted GPG (3.9%) and the corresponding Luenberger value (2.6%) exhibit a high degree of convergence. This alignment reflects the strong institutional constraints that compress the wage distribution, thereby limiting the scope for discretionary, high-end discriminatory returns and reducing the measurable glass ceiling effect. An exception to this pattern is private education, where the BO-adjusted gap is relatively small (1.7%) while the Luenberger indicator is considerably higher (8.3%), suggesting that despite modest average unexplained differences, substantial unrealised wage potential persists at the hedonic wage frontier.

The quantile BO decomposition helps to explain why the Luenberger indicator reaches such high levels in some economic activities. For instance, in construction, the BO-adjusted GPG rises from 6.6% at the 10th percentile to 24.7% at the 90th percentile. This extreme widening at the top of the distribution is the main source of the large frontier-based estimate (15.0%). Likewise, financial and insurance activities show BO gaps rising from 5.6% (10th) to 20.8% (90th), underpinning the large Luenberger value of 16.1%. Overall, while the BO decomposition measures average differences in returns, the Luenberger indicator quantifies unrealised wage potential across the entire wage structure, capturing structural inequality.

5.2.2 Luenberger indicator to measure the glass ceiling effect

The second Luenberger indicator isolates the distance to the upper wage frontier consisting of both women and men, thus measuring the severity of the glass ceiling across sectors. A higher value reflects a larger structural constraint preventing women from achieving the highest feasible wages given their characteristics.

The results show that glass ceiling effects are strongest in private market-driven sectors. Financial and insurance activities exhibit the highest glass ceiling effect at 19.6%, followed by construction (15.0%), wholesale and retail trade (14.1%), manufacturing (13.2%), and administrative and support services (11.5%).

In sectors where wage determination is more constrained, such as public health (2.9%), the glass ceiling effect is attenuated. This is consistent with institutional mechanisms that reduce the scope for gender disparities at higher earnings levels.

Taken together, the upper hedonic wage frontier results confirm that the French labour market is characterised by rather marked glass ceiling effects, particularly in industries where wage-setting is decentralised and rewards for career advancement are more unevenly distributed.

5.2.3 Luenberger indicator to measure the sticky floor effect

The third Luenberger indicator measures the distance to the lower wage frontier again consisting of both women and men, capturing the extent to which women are disproportionately positioned near the lowest boundary of feasible wages, i.e., the sticky floor effect.

Sticky floor patterns are evident across multiple sectors, though generally less pronounced than glass ceiling effects, with the exception of public administration (private status). The strongest effects appear in financial and insurance (11.4%), private public administration (11.2%), and information and communication (9.5%), showing that even in higher-paying sectors, women are more likely to occupy positions closer to the minimum attainable wages.

Moderate sticky floor effects are observed in administrative and support services (4.5%) and public education (4.2%), while the lowest effects are found in public health (2.6%), accommodation and food services (3.3%), private education (3.1%), and transportation (3.1%).

6 Conclusions

This study contributes to the labour economics literature on the GPG by proposing a new methodology to analyse wage disparities. Our primary methodological advancement lies in pioneering the application of nonparametric double hedonic wage frontiers and the use of the Luenberger indicator within a wage discrimination context. This approach moves beyond the focus on mean outcomes of the conventional BO decomposition to quantify gender disparities across the entire wage distribution, with a specific focus on the coexisting phenomena of the glass ceiling and sticky floor effects.

By applying our approach to the French labour market, using harmonized micro data from the 2022 French Employment Survey, we propose original results that identify the activities in which the GPG is highest and clarify whether these inequalities are driven primarily by glass ceiling or by sticky floor effects.

Empirically, the hedonic wage frontier analysis consistently shows a larger unexplained GPG than the unadjusted and adjusted mean gaps derived from the traditional methods. In every sector, the Luenberger indicator, measuring the distance to the maximum attainable wage frontier, reveals a greater loss than the BO results, indicating that standard approaches may understate the economic penalty associated with gender inequality. For example, in financial and insurance activities, the BO-adjusted gap is 13.3%, while the comparable Luenberger value reaches 16.1%.

The nonparametric hedonic wage frontier approach further confirms the coexistence and variation of the two major mechanisms of inequality along the wage distribution. The glass ceiling effect, measured by the distance to the upper frontier, is generally more severe. This effect is most pronounced, reaching 19.6% in financial and insurance, 15.0% in construction, and 14.1% in wholesale and retail trade. This points to significant barriers preventing women from achieving the highest feasible rewards for their human capital. While generally smaller, sticky floor effects remain present in most sectors. Financial and insurance activities exhibit a clear "double jeopardy" pattern, showing both the strongest glass ceiling (19.6%) and sticky floor (11.4%) effects. This pattern indicates both the systematic undervaluation of women's skills (wages closer to the minimum acceptable level) as well as persistent barriers to upward wage progression (wages farther from the maximum attainable level). Conversely, public status activities, such as human health and social work, display notably smaller gaps across all Luenberger measures, highlighting the protective role of standardized pay structures.

These results have direct implications for evidence-based policy interventions. Firstly, interventions should shift from monitoring unadjusted average GPGs to addressing the substantial unrealised wage potential revealed by adjusted measures. Secondly, the presence of sticky floor effects even in high-skilled sectors underscores the need for policies targeting not only increasing remunerations, but also the systemic undervaluation of female skills. Finally, the protective effect of public sector wage standardization suggests that expanding pay transparency requirements and limiting discretionary wage-setting may be effective tools for reducing this variety of structural inequalities.

This study has certain limitations since it controls only for education, experience, and tenure. A promising direction is to embed the Luenberger indicators in second-stage regressions incorporating richer qualitative factors such as origin, contract type, firm size, or place of residence, to better isolate the eventual further drivers of the current observed GPG. Future research could also deepen the analysis by integrating other dimensions of inequality, such as employment access (including unemployed individuals) and working time (including part-time workers). Indeed, the main advantage of this nonparametric hedonic wage frontier approach is that it allows researchers to consider several "variables to explain", rather than focusing solely on hourly wages. In this sense, it has the potential to provide a richer and more comprehensive measure of gender inequalities.

On the methodological side, one can sketch the following avenues for future research. First, instead of taking the male hedonic wage frontier as a basis to construct a Luenberger indicator as in (14), it would be possible to take a female hedonic wage frontier as a basis to define an alternative Luenberger indicator to investigate eventual wage gaps of men compared to women. This possibility has so far not been implemented. Furthermore, one could also take an arithmetic average of both such bases to average out the effect of a specific choice of basis. Second, while quite some work is available on statistical inference in the

Malmquist index (see Zelenyuk and Zhao (2025)), to our knowledge it would be useful to obtain confidence intervals for our estimates of the GPG.

References

Altonji, J. G. and R. A. Shakotko (1987): "Do Wages Rise with Job Seniority?" *The Review of Economic Studies*, 54, 437–459.

Amado, C. A., S. P. Santos, and J. M. São José (2018): "Measuring and Decomposing the Gender Pay Gap: A New Frontier Approach," *European Journal of Operational Research*, 271, 357–373.

Arulampalam, W., A. Booth, and M. L. Bryan (2006): "Is There a Glass Ceiling over Europe? Exploring the Gender Pay Gap across the Wages Distribution," CEPR Discussion Paper 510, Centre for Economic Policy Research, Research School of Economics, Australian National University.

Bar-Haim, E., L. Chauvel, J. C. Gornick, M. Raitano, and E. Yacobovich (2023): "The Persistence of the Gender Earnings Gap: Cohort Trends and the Role of Education in Twelve Countries," *Social Indicators Research*, 165, 821–841.

Bazen, S. and K. Charni (2023): "Gender Differences in the Early Career Earnings of Economics Graduates," IZA Discussion Paper 15954, IZA.

Becker, G. S. (1964): *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education*, New York: National Bureau of Economic Research.

Bishop, J. A., A. Grodner, H. Liu, and J.-R. Chiou (2007): "Gender Earnings Differentials in Taiwan: A Stochastic Frontier Approach," *Journal of Asian Economics*, 18, 934–945.

Blackaby, D., A. L. Booth, and J. Frank (2002): "Outside Offers and the Gender Pay Gap: Empirical Evidence from the UK Academic Labour Market," *Economic Journal*, 112, 659–676.

Blau, F. and L. Kahn (2017): "The Gender Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, 55, 789–865.

Blau, F. and L. M. Kahn (2006): "The U.S. Gender Pay Gap in the 1990S: Slowing Convergence," *ILR Review*, 60, 45–66.

Blinder, A. S. (1973): "Wage Discrimination: Reduced Form and Structural Estimates," *Journal of Human Resources*, 8, 436–455.

Boussemart, J.-P., W. Briec, K. Kerstens, and J.-C. Poutineau (2003): "Luenberger and Malmquist Productivity Indices: Theoretical Comparisons and Empirical Illustration," *Bulletin of Economic Research*, 55, 391–405.

Bowlin, W. and C. Renner (2008): "Assessing Gender and Top-Management-Team Pay in the S&P Mid-Cap and Small-Cap Companies Using Data Envelopment Analysis," *European Journal of Operational Research*, 185, 430–437.

Bowlin, W., C. Renner, and J. Rives (2003): "A DEA Study of Gender Equity in Executive Compensation," *Journal of the Operational Research Society*, 54, 751–757.

Briec, W., A. Dumas, K. Kerstens, and A. Stenger (2022): "Generalised Commensurability Properties of Efficiency Measures: Implications for Productivity Indicators," *European Journal of Operational Research*, 303, 1481–1492.

Carli, L. L. and A. H. Eagly (2016): "Women Face a Labyrinth: An Examination of Metaphors for Women Leaders," *Gender in Management: An International Journal*, 31, 514–527.

Chambers, R., Y. Chung, and R. Färe (1998): "Profit, Directional Distance Functions, and Nerlovian Efficiency," *Journal of Optimization Theory and Applications*, 98, 351–364.

Chambers, R., R. Färe, and S. Grosskopf (1996): "Productivity Growth in APEC Countries," *Pacific Economic Review*, 1, 181–190.

Christofides, L. N., A. Polycarpou, and K. Vrachimis (2013): "Gender wage gaps, sticky floors and glass ceilings in Europe," *Labour Economics*, 21, 86–102.

Chumpitaz, R., K. Kerstens, N. Paparoidamis, and M. Staat (2010): "Hedonic Price FunÇction Estimation in Economics and Marketing: Revisiting Lancaster's Issue of "Noncombinable" Goods," *Annals of Operations Research*, 173, 145–161.

Cotter, D. A., J. M. Hermsen, S. Ovadia, and R. Vanneman (2001): "The Glass Ceiling Effect," *Social Forces*, 80, 655–681.

Deutsch, J. and J. Silber (2005): "Measuring Multidimensional Poverty: An Empirical Comparison of Various Approaches," *Review of Income and Wealth*, 51, 145–174.

Díaz, M. A. and R. Sánchez (2011): "Gender and Potential Wage in Europe: A Stochastic Frontier Approach," *International Journal of Manpower*, 32, 410–425.

Espinosa, M. P. and E. Ferreira (2022): "Gender Implicit Bias and Glass Ceiling Effects," *Journal of Applied Economics*, 25, 37–57.

Eurostat (2022): "Gender pay gaps in the European Union – A statistical analysis – 2021 edition – Revision 1," Tech. rep., Publications Office of the European Union.

Färe, R., S. Grosskopf, and F. Hernandez-Sancho (2004): "Environmental Performance: An Index Number Approach," *Resource and Energy Economics*, 26, 343–352.

Firpo, S., N. Fortin, and T. Lemieux (2009): "Unconditional Quantile Regressions," *Econometrica*, 77, 953–973.

Fortin, N., T. Lemieux, and S. Firpo (2011): "Decomposition methods in economics," in *Handbook of Labour Economics*, ed. by D. Card and O. Ashenfelter, Amsetrdam: Elsevier, vol. 4, 1–102.

Garcia-Prieto, C. and P. Gomez-Costilla (2017): "Gender Wage Gap and Education: A Stochastic Frontier Approach," *International Journal of Manpower*, 38.

Hackman, S. (2008): *Production Economics: Integrating the Microeconomic and Engineering Perspectives*, Berlin: Springer.

Hadley, L. and J. Ruggiero (2006): "Final-Offer Arbitration in Major League Baseball: A Nonparametric Analysis," *Annals of Operations Research*, 145, 201–209.

Hofler, R. A. and K. J. Murphy (1992): "Underpaid and Overworked: Measuring the Effect of Imperfect Information on Wages," *Economic Inquiry*, 30, 511–529.

Hofler, R. A. and S. W. Polachek (1985): "A New Approach for Measuring Wage Ignorance in the Labor Market," *Journal of Economics and Business*, 37, 267–276.

International Labour Organization (ILO) (2018): *Global Wage Report 2018/19: What lies behind gender pay gaps*, Geneva: ILO.

Johnes, J. (2006): "Measuring Teaching Efficiency in Higher Education: An Application of Data Envelopment Analysis to Economics Graduates from UK Universities 1993," *European Journal of Operational Research*, 174, 443–456.

Kerstens, K., M. Azadi, R. K. Matin, and R. F. Saen (2024): "Double hedonic price-characteristics frontier estimation for IoT service providers in the industry 5.0 era: A nonconvex perspective accommodating ratios," *European Journal of Operational Research*, 319, 222–233.

Kneip, A., L. Simar, and P. W. Wilson (2016): "Testing Hypotheses in Nonparametric Models of Production," *Journal of Business & Economic Statistics*, 34, 435–456.

Kumbhakar, S. and C. Lovell (2000): *Stochastic Frontier Analysis*, Cambridge: Cambridge University Press.

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.

Mincer, J. (1974): *Schooling, Experience and Earnings*, New York: National Bureau of Economic Research.

Mincer, J. and S. Polachek (1974): "Family Investment in Human Capital: Earnings of Women," *Journal of Political Economy*, 82, S76–S108.

Mohan, N. and J. Ruggiero (2003): "Compensation Differences between Male and Female CEOs for Publicly Traded Firms: A Nonparametric Analysis," *Journal of the Operational Research Society*, 54, 1242–1248.

——— (2007): "Influence of Firm Performance and Gender on CEO Compensation," *Applied Economics*, 39, 1107–1113.

Nesheim, L. (2008): "Hedonic Price Functions," in *The New Palgrave Dictionary of Economics*, ed. by S. Durlauf and L. Blume, New York: Palgrave Macmillan, 2705–2711, 2 ed.

Oaxaca, R. (1973): "Male-Female Wage Differentials in Urban Labour Markets," *International Economic Review*, 14, 693–709.

OECD (2025): *Education at a Glance 2025: OECD Indicators*, Paris: OECD Publishing.

Olivetti, C., J. Pan, and B. Petrongolo (2024): "The Evolution of Gender in the Labor Market," in *Handbook of Labor Economics*, ed. by C. Dustmann and T. Lemieux, Amsterdam: Elsevier, vol. 5, 619–677.

Papadopoulos, A. and C. F. Parmeter (2025): *Two-Tier Stochastic Frontier Analysis for the Social Sciences*, Cham: Springer.

Polachek, S. and J. Robst (1998): "Employee Labor Market Information: Comparing Direct World of Work Measures of Workers' Knowledge to Stochastic Frontier Estimates," *Labour Economics*, 5, 231–242.

Polachek, S. and B. Yoon (1987): "A Two-tiered Earnings Frontier Estimation of Employer and Employee Information in the Labor Market," *Review of Economics and Statistics*, 69, 296–302.

Ramos, X. and J. Silber (2005): "On the Application of Efficiency Analysis to the Study of the Dimensions of Human Development," *Review of Income and Wealth*, 51, 285–309.

Redmond, P. and S. McGuinness (2017): "The Gender Wage Gap in Europe: Job Preferences, Gender Convergence and Distributional Effects," IZA Discussion Paper 10933, IZA Institute of Labor Economics.

Robinson, M. D. and P. V. Wunnava (1989): "Measuring Direct Discrimination in Labor Markets Using a Frontier Approach: Evidence from CPS Female Earnings Data," *Southern Economic Journal*, 56, 212–218.

Rosen, S. (1974): "Hedonic Prices and Implicit Markets: Production Differentiation in Pure Competition," *Journal of Political Economy*, 82, 34–55.

Simar, L. and P. W. Wilson (2020): "Hypothesis Testing in Nonparametric Models of Production Using Multiple Sample Splits," *Journal of Productivity Analysis*, 53, 287–303.

Van Ophem, H. and J. Mazza (2024): "Educational Choice, Initial Wage and Wage Growth," *Empirical Economics*, 67, 1235–1274.

Zaim, O. (2004): "Measuring Environmental Performance of State Manufacturing through Changes in Pollution Intensities: A DEA Framework," *Ecological Economics*, 48, 37–47.

Zelenyuk, V. and S. Zhao (2025): "Inference for DEA estimators of Malmquist productivity indices: An overview, further improvements, and a guide for practitioners," *Macroeconomic Dynamics*, 29, e81.