Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

How Cyclical Is a Gap? Composition-Adjusted Evidence on the Gender Wage Gap

Pedro Martins

pedro.martins@novasbe.pt

Universidade Nova de Lisboa

João Pereira

University of Évora

Raul Ramos

University of Barcelona

Research Article

Keywords: Wage decompositions, Gender gaps, Wage cyclicality, Discrimination

Posted Date: June 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9544432/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

How Cyclical Is a Gap?

Composition-Adjusted Evidence on the Gender Wage Gap*

João Pereira[1]

Raul Ramos[2]

Pedro S. Martins[3]

ABSTRACT

We propose a simple and novel approach to estimate the cyclicality of gaps using cross-sectional data. Our method controls for composition effects in contexts in which panel data cannot be leveraged. Our method also provides insights on assessing common trends in difference-in-differences analyses. We illustrate the approach studying the gender wag gap (GWG) in Portugal, finding that it is generally procyclical, except at the top of the distribution. This procyclicality stems from the unexplained component of the GWG rather than from differences in observed characteristics. The results also do not support the hypothesis that GWG cyclicality is driven by discrimination.

JEL classification: J31, E32, J16, J71, J82

Keywords: Wage decompositions, Gender gaps, Wage cyclicality, Discrimination.

* We thank comments from Robert Inklaar and participants in the 44th Meeting of the Association of Southern European Economic Theorists and the 17th Labor Economics Conference, Spanish Association of Labor Economics. We also thank funding from FCT (Portugal). All errors are our own.

[1] University of Évora. Email: jpereira@uevora.pt.

[2] University of Barcelona. Email: rramos@ub.edu.

[3] New University of Lisbon (Universidade Nova de Lisboa). Corresponding author. Email: pedro.martins@novasbe.pt. Web: https://pmrsmartins.wixsite.com/website. Address: Nova School of Business & Economics, Rua da Holanda, 1, 2775-405 Carcavelos, Portugal.

1. INTRODUCTION

While the gender wage gap (GWG) has been decreasing over the last decades (Blau and Kahn 2017), there may be a business cycle dimension to the GWG which has been largely ignored. This cyclicality can shed light on the drivers of GWG, including the role of wage discrimination. For instance, men's wages may benefit more from booms because of stronger bargaining power and rent sharing (Card et al. 2018; Biasi and Sarsons, 2022), leading to higher GWG during expansions. Also, if the cycle is not gender neutral (Razzu and Singleton 2016), identifying the most affected group is important in order to set-up adequate economic policies (Kovalenko and Töpfer 2021).

Moreover, GWG *cyclicality*, if any, can distort GWG-related policy evaluations and discrimination comparisons over time and across countries. An example is the common trends assumption in difference-in-differences evaluations (Baker et al. 2023; Biasi and Sarsons 2022; Bennedsen et al. 2023; Roth et al. 2023). Indeed, if the wages of men and of women tend to vary differently over the business cycle, time differences in the two wages series cannot necessarily be used to estimate the impact of a GWG policy, even under common trends before the intervention. Furthermore, the low power of pre-trends tests may make it difficult to identify these cases, mainly when the pre-trend differences are small (Freyaldenhoven et al. 2019; Roth 2022). Therefore, knowing how the GWG behaves over the business cycle is a way of "bringing economic knowledge to bear on how the parallel trends might be plausibly violated…", in (Roth 2022, p. 319).

This is a general issue that concerns any analysis of a gap over time, including other outcome variables than wages and other dichotomies than gender. This issue can emerge in any difference-in-difference evaluation concerning a gap between groups that is sensitive to the business cycle. A concrete example is the impact of transparency policies on the gender wage gap (GWG), as these are subject to potential bias introduced by the cyclical behaviour of the wage gap.

However, to the best of our knowledge, very few papers directly address the cyclicality of the gaps, and the approach is predominantly macroeconomic (Gupta et al. 2006; Kovalenko and Töpfer 2021; Potts and Roy 2024; Böheim et al. 2021)[4]. This implies that these studies are subject to composition bias problems, as acknowledged by the latter three studies. These problems are similar to those addressed in early macroeconomic studies of wage cyclicality. Indeed, low-skilled workers are typically more represented in employment during expansions than in recessions, introducing a countercyclical bias in wage cyclicality estimates (Solon et al. 1994; Abraham and Haltiwanger 1995; Brandolini 1995). The literature on the cyclicality of wages (Martins et al. 2012; Carneiro et al. 2012; Stüber 2017; Dapi 2020; Pereira et al. 2025), among many others) takes this issue into account using panel data techniques and fixed effects, which control for these and other composition effects over the business cycle.

[4] Some of these papers conclude that the GWG is procyclical (Böheim et al. 2021; Kovalenko and Töpfer 2021) whereas others conclude that it is acyclical (Gupta et al. 2006; Potts and Roy 2024).

In addition, there is evidence that the business cycle is not gender-neutral in terms of workforce composition (Solon et al. 1994). Generally, women are more represented in employment during recessions than in expansions, as they typically work in less cyclical occupations and industries (Hoynes et al. 2012; Razzu and Singleton 2016; Potts and Roy 2024), which is another potential source of bias. Therefore, when unemployment increases, women are more represented in employment and the GWG may mechanically decrease.[5] On the contrary, in upswings, male employment increases faster, which may automatically distort the GWG.[6]

Therefore, to estimate the cyclicality of a gap, such as the GWG, directly, we need to use data from a cross-sectional angle and compute the estimated difference between males' and females' wages at a given point in the wage distribution. However, these estimates are directly affected by composition bias, as previously discussed, and thus may not be

[5] Consider an economy with two sectors, A and B. Sector A is more cyclical and employs predominantly men, while sector B is less cyclical and employs predominantly women. Wages in sector A are higher than in sector B. During periods of economic expansion, wages and employment rise more in sector A than in sector B, leading to an increase in men's average wages and a widening of the gender wage gap. Conversely, during recessions, wages and employment decline more sharply in sector A, which contributes to a narrowing of the gender wage gap. Hence, fluctuations in the business cycle induce variations in the gender wage gap – a procyclical bias – through the differing gender composition of employment across sectors.

[6] There is a branch of the literature that includes microeconomic studies on the cyclicality of wages, which control for the previously mentioned composition issues but analyze the cyclicality of men's and women's wages separately (Park and Shin 2005; Marchand and Olfert 2013). In these studies, the cyclicality of the GWG is inferred indirectly from differences in cyclicality between men's and women's wages. Therefore, it is not the cyclicality of the gap itself (nor its constituent parts from an Oaxaca–Blinder decomposition) that is tested. In any case, the evidence from these studies indicates that both men's and women's wages tend to be procyclical, but men's wages cyclicality is usually higher. Consequently, the GWG as seen from this approach increases in expansions and shrinks in recessions. Park and Shin (2005) and Marchand and Olfert (2013) also try to explain the wage cyclicality differences between genders. The former find that, in the U.S., men are more represented in more cyclical occupations, whereas the latter link the wage cyclicality to differential gender representation in more cyclical industries (and states).

reliable for comparing the gap over time or across countries - especially when those countries are in different phases of the business cycle.

This paper's innovation consists in presenting a first method to estimate the cyclicality of the gender wage gap (GWG) directly from cross-sectional data while controlling for composition changes. The proposed approach can be readily applied to other contexts of wage gap analysis between groups (e.g., races, regions, etc) and can be viewed as an extension of the literature on the cyclicality of wages. The approach we propose to analyse the cyclicality of GWG builds on the reweighting approach of DiNardo et al. (1996) (DFL henceforth), which keeps the sample structure constant over time, regardless of the cycle[7]. Applying this methodology, we do not follow the same individuals over time as longitudinal wage cyclicality studies do, but we draw on a sample of individuals of the same observable characteristics over time. To the best of our knowledge, this is the only way of controlling for workforce cyclical composition changes when having to use cross-sectional data to estimate wage gaps.

We implement the proposed method using data for Portugal, a country with important changes in the composition of the work force since the 1980s. This involved increased participation of women, skill upgrading of both men and women, and closing gender skills gaps (Cardoso et al. 2016; Olivetti and Petrongolo 2016; Pereira 2021). These

[7] The reweighting approach of DFL is a widely used technique in applied Economic studies. Recent examples of application of this approach include: Blau and Kahn (2017), Card and Cardoso (2022), Fortin et al. (2017), among many others.

compositional changes are in line with the international trends regarding the labour market participation of men and women (Blau and Kahn 2017; Albanesi 2025).

We find that the GWG is procyclical, even if it is mostly acyclical at the top of the wage distribution. On top of this, we also study the cyclicality of the GWG constituent parts, as derived from an Oaxaca–Blinder type decomposition. We find that the cyclicality of the GWG has been decreasing over time due to the gender convergence of skills and other observable characteristics. Therefore, the remaining residual cyclicality arises from the unexplained component of the gap and not from compositional differences as suggested from the earlier literature. Moreover, our findings do not support the discrimination hypothesis as the source of the cyclical movements of the GWG, as it would imply a countercyclical behaviour of the unexplained component of the gap, whereas our findings are mostly procyclical or acyclical.

Our findings provide useful insights for DID studies in which the outcome variable is a wage gap. In the case of the GWG that we analyse, special attention should be paid when there are substantial differences in skills and other characteristics between men and women, and when the business cycle shifts from expansion to recession between pre- and post-treatment periods, or vice versa. These insights can be readily extended to other types of gaps. Possible solutions include adding more pre-treatment periods or using robust inference methods.

The remainder of this paper is organized as follows. Section 2 presents the method proposed. Section 3 presents the application - it describes the data, reports the results, and discusses specification issues. Section 4 concludes.

2. METHOD

The method we propose to analyze the cyclicality of a gap, which we apply later to the gender wage gap (GWG), aims to control for (workforce) composition changes without relying on panel data. In studies of wage cyclicality based on panel data (Solon et al. 1994; Abraham and Haltiwanger 1995; Brandolini 1995; Martins et al. 2012; Stüber 2017), workers are followed over time using fixed effects, and wage cyclicality is computed for the same individuals. Usually, additional longitudinal controls are also included, such as firm, occupation, or even the quality of the job match.

Such longitudinal controls are not possible when dealing with inherently cross-sectional concepts such as the GWG or other gaps between groups (e.g., races, regions, etc). These gaps are measured as differences in wages or other outcomes between groups at the same point in the business cycle and therefore depend on who is present in the sample in a given period, which itself may vary with the cyclical position of the economy (Abraham and Haltiwanger 1995). We address these changes in sample structure using the DFL reweighting approach, which ensures a constant sample structure over time by adjusting the weights of the observations. Hence, we do not seek to follow the same individuals over time; instead, observations are reweighted so that the sample remains uniform across periods in terms of observations and their observable characteristics.

The overall procedure for computing the cyclicality of the GWG is carried out in three steps. First, the sample is constructed using the previously described reweighting procedure. Second, the wage gaps are estimated based on this sample. Finally, the estimated wage gaps are regressed on a business cycle indicator to estimate their cyclicality.

We begin by describing the first step in detail. Hence, according to DFL, we have to compute sample weights ($w_{t*,t}$) that rebuild the sample structure of a given time period t* given the sample structure of origin at year t, which depends on the business cycle position of the economy. These sample weights are computed according to the following expression:

$$w_{t*,t} = \frac{\Pr\,(T=t^*|x_t)}{1-\Pr\,(T=t^*|x_t)} \times \frac{1-\Pr\,(T=t^*)}{\Pr\,(T=t^*)} \qquad (1)$$

$\Pr\,(T = t^*|x_t)$ is the probability of a worker with characteristics x at year t being observed in year t*; $\Pr\,(T = t^*)$ and (1-$\Pr\,(T = t^*)$) are the sample proportions of years t* and t, respectively. We estimate $\Pr\,(T = t^*|x_t)$ with a logit model. Finally, as t* in equation 1, we choose 1991, 1993, 1998, 2003, 2012, 2019: three years of growth

(1991, 1998 and 2019) and three years of recession.[8] This is a way of analysing the robustness of the findings to different employment structures[9].

In the second step, following Fortin et al. (2011) we estimate Oaxaca–Blinder-type decompositions at several quantiles of the wage distribution (10th, 25th, mean/OLS, 50th, 75th and 90th) on the reweighted samples using the $w_{t*,t}$ weights. With these decompositions we estimate the wage gaps - free of composition effects - and their constituent parts, explained and unexplained. To implement these decompositions, we first estimate wage equations[10] for males and females at those positions of the wage distribution using the recentered influence function (RIF) regression method proposed by Firpo et al. (2009):

$$Y_{g,\Theta,t,t*}=X\beta_{g,\Theta,t,t*}+\epsilon_{g,\Theta,t,t*} \quad (2)$$

where $Y_{g,\Theta,t,t*}$ are the wages for males (g = M) or females (g = F) at point Θ of the wage distribution, time t and employment structure t*; the $\beta_{g,\Theta,t,t*}$ are the corresponding estimates of the returns to observed characteristics for males and for females; $\epsilon_{g,\Theta,t,t*}$ are the corresponding error terms. The RIF dependent variable of the male/female

[8] This procedure involved the estimation of 198 logits. In fact, our sample consists of 34 years of data, ranging from 1985 to 2020 - there is no data available for 1990 and 2001. We consider six of these years as reference structures (t*). For each t*, there are 33 (34-1) years (t) for which a reweighting factor $w_{t*,t}$ must be computed to ensure that the sample structure of the year (t) matches that of the reference year (t*). Since we have six reference years (t*), a total of 198 logits (33 (t) × 6 (t*)) needed to be estimated.

[9] The reweighting approach assumes no general equilibrium effects (Fortin et al., 2011). This implies that while the reweighting procedure alters the sample composition, it maintains workers' returns to observed characteristics constant and assumes no employment effects. This is a relatively strong assumption. Therefore, we conduct these robustness checks to assess whether the findings are sensitive to different employment structures.

[10] The complete procedure involves the estimation of 612 wage equations: 34 (Years of data) ×3 (males, females, non-discriminatory wage structure) ×6 (t*).

equations was estimated separately for each gender. In these estimates we assume that there is no sample selection bias[11]. In section 3.3 we analyse this potential specification issue in more detail.

To implement the wage decompositions, it is necessary to choose a non-discriminatory wage structure. We use a pooled equation of both genders as non-discriminatory wage structure. If, instead, we had chosen the males or the females wage structure as non-discriminatory norm, this would imply necessarily the comparison of the findings with the opposite option, which, as will see, would be impracticable. Hence, according to Neumark (1988) and Fortin et al. (2011), the estimates of the GWG and of the explained and unexplained parts of the wage decomposition at each quantile (Θ), time period (t) and employment structure (t^*) are given according to:

$$wage\ gap_{\Theta,t,t^*} = \{(\bar{X}_{M,t,t^*} - \bar{X}_{F,t,t^*})\beta^*_{\Theta,t,t^*} + \bar{X}_{M,t,t^*}\left(\beta_{M,\Theta,t,t^*} - \beta^*_{\Theta,t,t^*}\right) + \bar{X}_{F,t,t^*}\left(\beta^*_{\Theta,t,t^*} - \beta_{F,\Theta,t,t^*}\right)\} \quad (3)$$

$\bar{X}_{M,t,t^*}$ and $\bar{X}_{F,t,t^*}$ are the vectors of mean levels of observed characteristics at time t and employment structure of t^* for males and females, respectively; β^*_{Θ,t,t^*} is the estimate of the corresponding non-discriminatory wage structure, whereas β_{M,Θ,t,t^*} and β_{F,Θ,t,t^*} are the estimates at Θ, t and t^* of the returns to observed characteristics for males and

[11] Although the established merits of the Quadros de Pessoal data set in terms of international research, which are examples the papers of Martins et al. (2012), Carneiro et al. 2012), Buhai et al. (2014), Card and Cardoso (2022), Snell et al. (2018) , Pereira et al. (2025), among many others, it has no information about family members, the presence of children in the family, or other sources of income than the individual labour income in the private sector, information typically used to control for sample selection bias.

for females, respectively. Hence, according to the wage decomposition, $(\bar{X}_{M,t,t^*} - \bar{X}_{F,t,t^*})\beta^*_{\Theta,t,t^*}$ is the part of the gap attributable to differences of characteristics between men and women at quantile *Θ and time* t, given the employment structure of t^*[12], whereas $\bar{X}_{M,t,t^*}\left(\beta_{M,\Theta,t,t^*} - \beta^*_{\Theta,t,t^*}\right) + \bar{X}_{F,t,t^*}\left(\beta^*_{\Theta,t,t^*} - \beta_{F,\Theta,t,t^*}\right)$ is the unexplained part of the GWG[13] (usually taken as an upper limit of the estimate of discrimination).

The identification of the wage decomposition terms relies on the assumption that men's and women's observed and unobserved characteristics are common to both groups. This means that no characteristic can be used to identify group membership. This is known as the common support assumption (Fortin et al. 2011). The dataset we use is a census of all employees in the private sector in Portugal and, therefore, should not be affected by this potential issue, as both groups are widely represented across all observed characteristics (age, education, tenure, and two-digit-level occupations and industries).

Does the reweighting process influence the common support of the data and the identification of the wage decomposition terms? This is a potential specification issue

[12] If sample selection bias exists in older employment structures (t*), low skilled (low wage) women may be underrepresented in the sample. This may cause a double source of bias on the explained part of the gap. First, on the term $(\bar{X}_{M,t,t^*} - \bar{X}_{F,t,t^*})$, which may be overestimated and therefore contribute to a higher cyclicality of the gap, for a given non-discriminatory wage structure (β^*_{Θ,t,t^*}). Second, the cyclicality of the non-discriminatory wage structure (β^*_{Θ,t,t^*}) may itself suffer from bias as men may be overrepresented in these employment structures and, therefore, their cyclicality contributes more to the cyclicality of the non-discriminatory wage structure.

[13] The possible effects of sample bias on the unexplained part the gap are not so clear as in the explained part as there are possible contradictory effects interacting.

that we may consider. We address these concerns and discuss the respective findings in the section 3.3

With regard to the cyclicality of the GWG, it depends on the cyclicality of the explained and unexplained components of the wage decomposition. The former is influenced by the cyclicality of the non-discriminatory wage structure (β^*_{Θ,t,t^*}), which represents the general wage level, and by the terms $\bar{X}_{M,t,t^*}$ and $\bar{X}_{F,t,t^*}$. These terms should be acyclical in consequence of the reweighting process. However, different employment structures imply a different contribution of these terms to the cyclicality of the gap. In fact, the greater (lower) the difference between the $\bar{X}_{M,t,t^*}$ and the $\bar{X}_{F,t,t^*}$, the greater (lower) is the cyclicality of the gap for a given β^*_{Θ,t,t^*}. The cyclicality of the latter component, in turn, depends on the terms $\left(\beta_{M,\Theta,t,t^*} - \beta^*_{\Theta,t,t^*}\right)$ and $\left(\beta^*_{\Theta,t,t^*} - \beta_{F,\Theta,t,t^*}\right)$. These terms will contribute to the cyclicality of the gap if men's and women's returns to observed characteristics are different and move differently along the cycle.

Having the estimates of the gender wage gap ($gap_{t,t*,\Theta}$) and its explained ($explained_{t,t*,\Theta}$) and unexplained ($unexplained_{t,t*,\Theta}$) components for each employment structure (t^*), we proceed with the third step, regressing each one of these components on a business cycle indicator (BCI) and in a linear time trend (t). We include a time trend in these equations, as previous studies (Blau and Kahn 2017; Goldin, 2014) show that there is a decreasing trend in the gender wage gap over last decades, which reflects, among other factors, the convergence of skills (mainly education, and to a

lesser extent, experience) between men and women[14]. Hence, as the aim of our approach is not to explain the factors behind the long run changes in the gender wage gap, but how it moves during the business cycle, we do not explicit model the long run changes in the gap. For that, there are already established findings and methodologies (Fortin et al. 2011; Goldin 2014; Blau and Kahn 2017). This final step of our approach is quite similar to second step of the "two step-estimator" of the literature on real wage cyclicality (Stüber 2017; Dapi 2020).

Hence, the test specifications of the cyclicality of the GWG and its constituent parts are given by the following equations:

$$gap_{t,t*,\Theta} = \hat{a}_{gap,\Theta,t*} + \hat{b}_{gap,\Theta,t*} \times BCI_t + \hat{\gamma}_{gap,\,\theta,t*} \times\ t + \epsilon_{t,t*,\Theta} \quad (4)$$

$$explained_{t,t*,\Theta} = \hat{a}_{expl,\Theta,t*} + \hat{b}_{expl,\Theta,t*} \times BCI_t + \hat{\gamma}_{expl,\,\theta,t*} \times t + \psi_{t,t*,\Theta} \quad (5)$$

$$unexplained_{t,t*,\Theta} = \hat{a}_{unexpl,\Theta,t*} + \hat{b}_{unexpl,\Theta,t*} \times BCI_t + \hat{\gamma}_{unexpl,\,\theta,t*} \times t + \zeta_{t,t*,\Theta} \quad (6)$$

The estimates $\hat{b}_{gap,\Theta,t*}$, $\hat{b}_{expl,\Theta,t*}$ and $\hat{b}_{unexpl,\Theta,t*}$ are the parameters of interest. They measure the sensitivity of the GAP, as well as its explained and unexplained components, to the business cycle, depending on the position of the wage distribution (Θ) and the employment structure (t*) These parameters can be interpreted as the change (in log points) in each of these components when the unemployment rate increases by one percentage point. The coefficients $\hat{\gamma}_{gap,\,\theta,t*}$, $\hat{\gamma}_{expl,\,\theta,t*}$ and $\hat{\gamma}_{unexpl,\,\theta,t*}$

[14] The fit of the estimates improves considerably when a linear time trend is included in equations 4 to 6. For example, the R^2 of the regression of the non-reweighted GAP on the lagged unemployment rate increases from 40,1% without time trend to 85,67% when the time trend is added to the regression.

capture the effect of the time trend in each one of the components. Finally, $\epsilon_{t,t*,\theta}$, $\psi_{t,t*,\theta}$ and $\zeta_{t,t*,\theta}$ are the error terms.

3. APPLICATION

3.1 DATA

We use *Quadros de Pessoal* (QP) data covering the period 1985 to 2020. QP is a matched employer-employee data set collected by the Ministry of Employment of Portugal, covering all enterprises in the private sector with at least one employee. It includes information on the characteristics of workers and firms, including their unique identifiers. Our sample is restricted to individuals who worked at least 120 hours in the reference month (October) and are 18 to 64 years old.[15]

As independent variables x in the logit model to estimate the sample weights (equation 1)) we include age, age squared, seven education dummies, all interactions between age and age squared and the education dummies, tenure and tenure squared, a gender dummy, and two-digit industry and occupation dummies. The gender dummy captures, among other effects, the big changes that took place in the Portuguese labour market regarding women presence in employment that increased from about 30% in 1985 to about 46% in 2020.

[15] There is no workers' information in 1990 and in 2001. Workers in agriculture, unpaid family workers, apprentices, and workers whose base wage is less than 80% of the national minimum wage were excluded from the analysis. This is the minimum admissible wage for apprentices and trainees. See Ferreira et al. (2026) for a recent study of the GWG and (Pereira et al. 2025) for a recent study of real wage cyclicality both using this data set.

With respect to the wage equations (equation 2), we used as measure of wages, log hourly earnings, which were regressed on age, age squared, education dummies, tenure, tenure squared and two-digit industry and occupation dummies. As *BCI* in equations (4) to (6) we chose the lagged unemployment rate, the usual business cycle indicator adopted in studies of the Portuguese labour market (Carneiro et al. 2012). Finally, Table A.1 in appendix A provides the summary statistics of the main variables used.

3.2 RESULTS

3.2.1 THE LONG RUN CHANGES IN THE GAP AND ITS COMPONENTS

First, we describe the long-run changes in the gap and its constituent parts. This analysis helps to clarify how the cyclicality of the gap evolves under different employment structures, by presenting the evolution of men's and women's observable characteristics together with the gap and its explained and unexplained components.

Figure 1 displays the estimates of the GWG based on equation 3. The estimated gap shows a clearly decreasing trend, particularly at the lower half of the wage distribution. For instance, the estimated mean GWG ranges from 23,3 log points in 1985 to 11,8 log points in 2020[16]. At higher positions of the wage distribution the speed of closing the

[16] The mean raw gap ranges from about 32 log points in 1985 to about 16 log points in 2020. The mean raw gap in Cardoso et al. (2016) for the period 1986-2008 is of about 24,2 log points.

gap is lower, but the trend is still clearly decreasing. Blau and Kahn (2017) present estimates for the U.S. consistent with those for Portugal.

[Insert figure 1, about here]

Next, we examine how each component of the gap (according to equation 3) evolved in figures 2 and 3. They show that the explained component of the gap is the main factor explaining the convergence of the GWG. The gap in characteristics between men and women has significantly decreased since the mid-1990s and has remained very small since the mid-2000s[17]. The convergence of skills and other labour market characteristics between men and women is an international trend which was also reported by other authors (Blau and Kahn 2017; Goldin 2014).

On the contrary, the unexplained component of the gap only evidences the same behaviour at lower and middle positions of the wage distribution (figure 3). At the 75th and 90th quantiles there exists a resilient residual gap between men's and women's wages.

[Insert figures 2 and 3, about here]

Detailing the evolution of the explained part of the gap by groups of covariates (age, education, tenure, occupations and industries), we can understand what groups of covariates are responsible for the convergence of this component. This analysis is

[17] These findings are consentaneous with those of Centeno and Novo (2014) and Pereira (2021) also for Portugal.

displayed in figure 4 for the mean of the wage distribution[18], which indicates that men and women have very similar levels of age, tenure and distribution by occupations at the end of the period. However, women are better endowed in terms of average level of education, whereas men are more represented in industries paying higher wages. Nonetheless, as the aggregated composition effect has been decreasing and is very small at the end of the period, the contribution of the explained component for the cyclicality of the gap should be small and decreasing, particularly when assessed with more recent employment structures (see first term of equation 3).

[Insert figure 4, about here]

3.2.2 THE CYCLICALITY OF THE GAP AND ITS COMPONENTS

Next, we present the cyclicality estimates of the GWG and its components (explained and unexplained), both adjusted and unadjusted for composition changes (Table 1). In doing so, we analyze how the gender wage gap varies over the business cycle, specifically in relation to the evolution of the lagged unemployment rate, our business cycle indicator.

We find that, independently of the employment structure considered, the GWG is mostly procyclical (Table 1-A)[19]. For example, a percentage point increase in the unemployment rate changes the gap in between 0.0150 and -0.578 percentage points

[18] For other positions of the wage distribution the findings are quite similar.
[19] Figure C.1 in appendix C shows the effect of the reweighting process on the relationship between the wage gap and the business cycle.

at mean of the wage distribution[20] depending on the employment structure adopted. Hence, women seem to be more protected from recessions than men but also benefit less from wage increases in expansions. Moreover, the cyclicality of the gap is also lower, or even non-significant, for more recent employment structures, which, as we saw in figures 2 and 4, are more similar for both genders than the older.[21] The estimates unadjusted for composition changes are somewhat different from those that have a constant employment structure. Despite this, they seem to be closer to estimates based on more recent employment structures, likely because that type of structure weights more on all the sample (the number of individuals has been increasing in QP over time).

[insert table 1, about here]

Table 1-B displays the cyclicality of the explained component of the gap. For this component be cyclical men and women should have an unequal distribution of skills and other labour market observable characteristics (see equation 3) and the non-discriminatory wage structure (β^*_{θ,t,t^*}) should be cyclical. The estimates in table 1B) show that the cyclicality of the explained component of the gap decreases when we

[20] These estimates of the cyclicality of the gap are somewhat lower than those of Böheim et al. (2021) for Austria. However, beyond the potential differences between the two countries (Portugal and Austria), there are other distinctions between the studies that make the estimates not directly comparable. First, the analysis by Böheim et al. (2021) does not control for composition changes over the business cycle, whereas ours does—and even our unadjusted estimates (i.e., not accounting for composition changes) are lower than theirs. Second, we include a time trend to capture long-run changes in the gap, while the Austrian study does not. Finally, the Portuguese sample spans substantially more years (34) than the Austrian one (13).

[21] Albanesi (2025), in a study on the United States, also finds that the cyclicality of men's and women's working hours has become more synchronized since the mid-1990s. The convergence of observable characteristics between genders plays an important role in these findings. It also appears that the cyclicality of the gender wage gap is influenced by these factors. The decreased cyclicality of the gap we observe can be viewed as part of the broader phenomenon known as the Great Moderation, which is associated with the reduced volatility of macroeconomic indicators since the early 1980s (Stock and Watson 2003).

consider more recent employment structures and is even non-significant (across all the wage distribution) when the most recent are considered (2012 and 2019). This is fundamentally explained by the convergence of skills and other labour market characteristics between genders across the time (figure 2 and 4), which makes this source of cyclicality almost irrelevant (see equation 3). Therefore, the established explanation for the higher cyclicality of men's wages based on their greater representation on more cyclical occupations (Park and Shin 2005) or industries (Marchand and Olfert 2013) does not seem to be more valid for Portugal, and also likely for other countries with a similar evolution of observable characteristics of both genders. Therefore, the cyclicality that still exists in the GWG should arise from the unexplained (residual) part of the gap.

Table 1-C displays the cyclical behaviour of the residual component of the wage decomposition, which is also typically procyclical, relatively stable for different employment structures, but whose effect is decreasing across the wage distribution, being mostly acyclical/countercyclical at the top of the wage distribution. These findings do not provide relevant support for the hypothesis that wage discrimination increases with unemployment (Biddle and Hamermesh 2013). If that were the case, the wage gap should increase (decrease) during recessions (expansions), being countercyclical. However, in the upper part of the wage distribution (Q90), there is some partial and weak evidence that the unexplained portion of the gap may move countercyclically. In any case, this does not necessarily imply that discrimination is the source of this cyclical behaviour. In fact, it is difficult to rationalize that, within the same labour market, highly

skilled and highly paid women would face discrimination while low-skilled and low-paid women would not—especially given that the former typically have greater bargaining power than the latter.

Therefore, it is likely that other mechanisms may be operating, such as the bargaining effect on rent sharing found in a study for Portugal by Card et al. (2016). These authors found that women in Portugal receive lower wage premiums than men at the same firms - about 90% of those that men receive -, but this ratio is roughly constant over time. As these variable components of pay are more important for high-skilled and high-paid workers, this effect may contribute to closer wage cyclicality of men's and women's wages at higher positions of the wage distribution and therefore may be one of the likely explanations for the mostly acyclical unexplained component of the wage gap at higher positions of the wage distribution. From this perspective, the lower wage premiums observed among women are not attributable to gender-based discrimination, but rather to a lower propensity to engage in wage negotiations with employers and/or reduced effectiveness in such negotiations (Stuhlmacher and Walters 1999; Card et al. 2016; Blau and Kahn 2017).

Table 2 presents the tests of equal cyclicality of the gap (2A) and of its unexplained part (2B) across the wage distribution[22]. These tests confirm that there are differences of cyclicality (gap and unexplained component) across the wage distribution. The major

[22] We do not present these tests for the explained component because it is almost acyclical across the wage distribution (table 1-B).

and significant arise between positions in the lower halve of the wage distribution and the 90th percentile. So, these findings suggest that what explains the cyclicality at the top of the wage distribution may be not the same what explains that cyclicality in the lower half of the wage distribution. At the top, as we mentioned above, it is likely that the acyclical mechanisms related to rent sharing between men and women had a relevant role. In other parts of the wage distribution, the small cyclicality that still exists arises from other mechanisms included in the unexplained part as the explained part is acyclical for more recent employment structures.

[insert table 2, about here]

Another consequence of the cyclical behaviour of the GWG is that it may introduce confounding effects when evaluating policies aimed at reducing the GWG—such as transparency policies (e.g., Bennedsen et al. 2023; Cullen 2024) — especially when using difference-in-differences (DID) methodologies. This issue can be particularly severe when the business cycle shifts between the pre- and post-treatment periods, such as a transition from expansion to recession, or vice versa. In fact, if the economy moves from expansion to recession between these periods, studies may mistakenly attribute observed reductions in the GWG to the policy, when in reality the narrowing of the gap may be driven by the economic downturn itself. In such cases, the estimated policy effects are likely to be positively biased. Conversely, if the cycle shifts from recession to expansion, the GWG tends to widen due to the mechanical effects of the cycle, potentially resulting in a negative bias in the estimated policy impact—that is, the cycle may obscure the policy's true effect. A similar bias may also arise if the growth rate of

the economy changes significantly between the pre- and post-treatment periods. More generally, some degree of bias is likely to exist irrespective of whether the economy is expanding or contracting.

The pre-trend tests may fail to detect such issues, as they typically have low statistical power (Roth 2022). Given that we find the cyclicality of the GWG is more pronounced in contexts with substantial differences in skills and other job characteristics between men and women, a possible remedy is to include additional pre-treatment periods (Roth 2022). An extended pre-treatment period can facilitate the detection of violations of the parallel trends assumption, particularly if it coincides with a period of convergence in observable characteristics between men and women. In such contexts, robust inference methods (Rambachan and Roth 2023) may also be employed.

Finally, similar problems may arise in other contexts, for example, if wages of other two different groups (or regions) react differently to the business cycle. So, the findings of this paper may also be useful in other contexts.

3.2.3 POTENTIAL SPECIFICATION ISSUES – SAMPLE SELECTIVITY BIAS AND COMMON SUPPORT

In this sub-section we analyse two possible specification issues that potentially could influence our findings, sample selection bias and the request of the common support for the wage decompositions being carried out. Regarding the first issue, this should not be

matter of major concern as women labour market participation and employment rates are high and have been increasing in Portugal (Olivetti and Petrongolo 2016). Furthermore, as (Cardoso et al. 2016) show in a study on Portugal using the same dataset for the period 1986–2008, the pattern of labour force participation over the life cycle is quite similar for men and women, and the incidence of part-time employment is low in Portugal.

Nonetheless, to be conservative, we can acknowledge the existence of some degree of sample selection bias, particularly in the early years of the sample, when men's and women's participation in employment was still somewhat unbalanced. This potential specification issue is likely to affect low-skilled (low-wage) women, who may be underrepresented in employment. Therefore, older employment structures (t*), especially those from the 1990s, may reflect different selection mechanisms for men and women in the labour market, and the reweighting process may propagate these differences throughout the entire sample. This could lead to the underrepresentation of women in the labour market, particularly in low-skilled jobs, which in turn may contribute to an overestimation of the explained component of wage decomposition and, consequently, to an overestimation of the cyclicality of the wage gap in older employment structures (see first term of equation 3)). This bias may therefore be propagated through the reweighting process to the entire sample. That said, in older employment structures, the average skill levels of men and women was very different. Thus, while the finding that the gender wage gap in these structures is cyclical should not be questioned, its magnitude may be somewhat overestimated

Finally, with respect to the common support analysis, we computed propensity scores for group membership based on the observable characteristics included in the wage equations. The findings show that the vast majority of observations - on average, about 97% - fall within the region of common support; that is, they share similar propensity score values (Glewwe and Todd 2022). This conclusion remains unchanged after the reweighting process. Details of this analysis can be found in appendix B.

4. CONCLUSIONS

In this paper, we propose a simple method to analyse the cyclicality of the gender wage gap free of composition effects. The method can be understood as an extension of the literature on the cyclicality of wages to contexts where panel data is not available or cannot be used, such as the study of gaps between groups (e.g., races, regions). Therefore, it is readily applicable to other contexts.

We apply this method to the study of the gender wage gap (GWG) in Portugal. As in many other countries, the GWG has been decreasing over time in Portugal. Besides this long run trend, we show that there is also a cyclical component in the GWG, even if small. This finding was obtained considering several employment structures and therefore is robust to business cycle composition effects (which can affect men and women differently). Therefore, comparisons of the gap across countries or industries or

(within countries) over different years may be strongly influenced by their macroeconomic context.

As, in general, the GWG is procyclical, women are more protected from wage losses in recessions than men, but also benefit less from wage growth in expansions. Previous evidence about such GWG cyclicality was indirect, but also suggested that the GWG was procyclical. An explanation provided before was that the cyclicality was driven by the stronger presence of men in more cyclical industries and/or occupations. Our findings show that these explanations may not be longer valid due to the gender convergence of skills and other observed labour characteristics in Portugal. As this is an international trend, it is likely that the same conclusion can extend to other labour markets.

The small degree of cyclicality that we still observe in the GWG arises from the unexplained component of the wage decomposition. This component serves as an upper bound estimate of labour market discrimination between men and women. However, our findings do not support the discrimination hypothesis, as the unexplained part of the GWG tends to be procyclical, whereas discrimination would be expected to follow a countercyclical pattern. Only at the top of the wage distribution (90*th* quantile) there is some evidence that the unexplained part of the gap is acyclical / countercyclical and, therefore, that discrimination may have some role in explaining GWG cyclicality. However, it is more likely that a constant rent sharing between men and women, a hypothesis validated by Card et al. (2016) also for Portugal, could explain the mostly acyclical behaviour of the unexplained component of the gap at the top of the wage

distribution. Below the top of the wage distribution, the small procyclical component of the unexplained component can have several explanations, including premiums for long and continuous hours of work (Goldin 2014), typically more likely to be attributed to men. These premiums may have a cyclical component, which is left for further research.

Finally, our findings also contribute to the discussion of the parallel trends hypothesis in difference-in-differences (DID) studies involving the GWG (Bennedsen et al. 2023) or other gaps between groups. The decreased cyclicality may make it more difficult to detect violations of the parallel trends assumption, especially given the frequently low power of pre-trends tests (Roth 2022). Therefore, particular attention should be paid in cases where significant differences in skills (or other characteristics) exist between men and women, or when between the pre - and post-treatment periods the business cycle shifts from expansion to recession (or vice versa), or the rate of economic growth changes significantly. In such contexts, the cyclical nature of the gap has greater potential to introduce bias into the evaluation of the policy under analysis. Possible solutions include adding more pre-treatment periods (Roth 2022) or adopting more robust inference methods (Rambachan and Roth 2023).

REFERENCES

Abraham KG, Haltiwanger JC (1995) Real Wages and the Business Cycle. J Econ Lit 33:1215–1264

Albanesi S (2025) Changing Business Cycles: The Role of Women's Employment. American Economic Journal: Macroeconomics 17:345–390. https://doi.org/10.1257/mac.20210138

Baker M, Halberstam Y, Kroft K, et al (2023) Pay Transparency and the Gender Gap. Am Econ J Appl Econ 15:157–183. https://doi.org/10.1257/app.20210141

Bennedsen M, Larsen B, Wei J (2023) Gender wage transparency and the gender pay gap: A survey. J Econ Surv 37:1743–1777. https://doi.org/10.1111/joes.12545

Biasi B, Sarsons H (2022) Flexible Wages, Bargaining, and the Gender Gap. Quarterly Journal of Economics 137:215–266. https://doi.org/10.1093/qje/qjab026

Biddle JE, Hamermesh DS (2013) Wage discrimination over the business cycle. IZA J Labor Policy 2:. https://doi.org/10.1186/2193-9004-2-7

Blau FD, Kahn LM (2017) The gender wage gap: Extent, trends, & explanations. J Econ Lit 55:789–865. https://doi.org/10.1257/jel.20160995

Böheim R, Fink M, Zulehner C (2021) About time: the narrowing gender wage gap in Austria. Empirica 48:803–843. https://doi.org/10.1007/s10663-020-09492-4

Brandolini A (1995) In Search of a Stylised Fact: Do Real Wages Exhibit A Consistent Pattern Of Cyclical Variability? J Econ Surv 9:103–163

Buhai I, Portela M, Teulings C, Vuuren A (2014) Returns to Tenure or Seniority? Econometrica 82:705–730. https://doi.org/10.3982/ecta8688

Cameron A, Trivedi P (2022) Microeconometrics Using Sata, 2nd Edition. Stata Press, College Station, Texas

Card D, Cardoso AR (2022) Wage Flexibility under Sectoral Bargaining. J Eur Econ Assoc 20:2013–2061

Card D, Cardoso AR, Heining J, Kline P (2018) Firms and Labor Market Inequality: Evidence and Some Theory. Journal ofLabor Economics 36:S13–S70. https://doi.org/https://doi.org/10.1086/694153

Card D, Cardoso AR, Kline P (2016) Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. Quarterly Journal of Economics 131:633–686. https://doi.org/10.1093/qje/qjv038

Cardoso AR, Guimarães P, Portugal P (2016) What drives the gender wage gap? A look at the role of firm and job-title heterogeneity. Oxf Econ Pap 68:506–524. https://doi.org/10.1093/oep/gpv069

Carneiro A, Guimarães P, Portugal P (2012) Real Wages and the Business Cycle: Accounting for Worker, Firm, and Job Title Heterogeneity. American Economic Journal: Macroeconomics 4:133–152

Centeno M, Novo Á (2014) When supply meets demand: wage inequality in Portugal. IZA Journal of European Labor Studies 3:23:

Cullen Z (2024) Is Pay Transparency Good. Journal of Economic Perspectives 38:153–180. https://doi.org/10.1257/jep.38.1.153

Dapi B (2020) Wage Cyclicality and Composition Bias in the Norwegian Economy. Scandinavian Journal of Economics 122:1403–1430

Dinardo J, Fortin NM, Lemieux T (1996) Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach. Econometrica 64:1001–1044

Ferreira A, Martins P, Pereira J, et al (2026) Can transparency mandates reduce the gender wage gap? Firm size discontinuity evidence. mimeo

Firpo S, Fortin N, Lemieux T (2009) Unconditional Quantile Regressions. Econometrica 77:953–973. https://doi.org/10.3982/ecta6822

Fortin N, Lemieux T, Firpo S (2011) Decomposition Methods in Economics. In: Handbook of Labor Economics. pp 1–102

Fortin NM, Bell B, Böhm M (2017) Top earnings inequality and the gender pay gap: Canada, Sweden, and the United Kingdom. Labour Econ 47:107–123. https://doi.org/10.1016/j.labeco.2017.05.010

Freyaldenhoven S, Hansen C, Shapiro JM (2019) Pre-event trends in the panel event-study design. American Economic Review 109:3307–3338. https://doi.org/10.1257/aer.20180609

Glewwe P, Todd P (2022) Impact Evaluation in International Development: Theory, Methods and Practice. The World Bank, Washington

Goldin C (2014) A grand gender convergence: Its last chapter. American Economic Review 104:1091–1119. https://doi.org/10.1257/aer.104.4.1091

Gupta ND, Oaxaca RL, Smith N (2006) Analysing trends in US and Danish gender wage gaps in the 1980s and 1990s. Appl Econ Lett 13:643–647. https://doi.org/10.1080/13504850600684825

Hoynes H, Miller DL, Schaller J (2012) Who suffers during recessions? In: Journal of Economic Perspectives. pp 27–48

Kovalenko T, Töpfer M (2021) Cyclical dynamics and the gender pay gap: A structural VAR approach. Econ Model 99:. https://doi.org/10.1016/j.econmod.2021.03.007

Marchand J, Olfert S (2013) The us gender gap through the great recession using an alternative approach to cyclicality. Appl Econ Lett 20:276–281. https://doi.org/10.1080/13504851.2012.692868

Martins PS, Solon G, Thomas JP (2012) Measuring What Employers do about Entry Wages over the Business Cycle: A new Approach. American Economic Journal: Macroeconomics 4:36–55. https://doi.org/10.1257/mac.4.4.36

Neumark D (1988) Employers' Discriminatory Behavior and the Estimation of Wage Discrimination. Source: The Journal of Human Resources 23:279–295

Olivetti C, Petrongolo B (2016) The Evolution of Gender Gaps in Industrialized Countries. Annu Rev Econom 8:405–434. https://doi.org/10.1146/annurev-economics-080614-115329

Park S, Shin D (2005) Explaining procyclical male-female wage gaps. Econ Lett 88:231–235. https://doi.org/10.1016/j.econlet.2005.01.021

Pereira JM (2021) Did wage inequality increase in Portugal? Yes, and for good reasons. Appl Econ Lett 28:973–977. https://doi.org/10.1080/13504851.2020.1789057

Pereira J, Ramos R, Martins PS (2025) Wage Cyclicality and Labor Market Institutions. Industrial Relations: A Journal of Economy and Society. https://doi.org/10.1111/irel.12387

Potts T, Roy J (2024) The heterogeneous impact of uncertainty shocks on labour market outcomes for men and women. Appl Econ. https://doi.org/10.1080/00036846.2024.2399811

Rambachan A, Roth J (2023) A More Credible Approach to Parallel Trends. Review of Economic Studies 90:2555–2591. https://doi.org/10.1093/restud/rdad018

Razzu G, Singleton C (2016) Gender and the business cycle: An analysis of labour markets in the US and UK. J Macroecon 47:131–146. https://doi.org/10.1016/j.jmacro.2015.12.006

Roth J (2022) Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends. Am Econ Rev Insights 4:305–322. https://doi.org/10.1257/aeri.20210236

Roth J, Sant'Anna PHC, Bilinski A, Poe J (2023) What's trending in difference-in-differences? A synthesis of the recent econometrics literature. J Econom 235:2218–2244. https://doi.org/10.1016/j.jeconom.2023.03.008

Snell A, Martins P, Stüber H, Thomas JP (2018) Bias in Returns to Tenure When Firm Wages and Employment Comove: A Quantitative Assessment and Solution. Source: Journal of Labor Economics 36:47–74. https://doi.org/10.2307/26553302

Solon G, Barsky R, Parker JA (1994) Measuring the Cyclicality of Real Wages: How Important is Composition Bias?*. Q J Econ 109:1–25

Stock JH, Watson MW (2003) Has the Business Cycle Changed and Why? In: NBER Macroeconomics Annual 2002. MIT, pp 159–218

Stüber H (2017) The Real Wage Cyclicality of Newly Hired and Incumbent Workers in Germany. The Economic Journal 127:522–546

Stuhlmacher A, Walters AE (1999) Gender differences in negotiation outcome: A meta-analysis. Pers Psychol 3:653–677. https://doi.org/https://doi.org/10.1111/j.1744-6570.1999.tb00175.x

TABLES AND FIGURES

Table 1: Sensitivity of the wage gap and the unexplained component to unemployment, Portugal, 1985-2020

	Employment structure (t*)						
	unadjusted	1991	1993	1998	2003	2012	2019
Points of the wage distribution (Q)	**A) Wage gap**						
Q10	-0.398**	-0.657***	-0.636***	-0.547***	-0.514***	-0.392**	-0.336**
Q25	-0.430**	-0.655***	-0.642***	-0.588***	-0.545***	-0.424***	-0.370***
mean	-0.275*	-0.578***	-0.571**	-0.496***	-0.415***	-0.121	0.0150
Q50	-0.307*	-0.583**	-0.587**	-0.555**	-0.529**	-0.326	-0.134
Q75	-0.468*	-0.669*	-0.696*	-0.613**	-0.587**	-0.421*	-0.497
Q90	-0.107	-0.471**	-0.436**	-0.361**	-0.158	0.329	0.356
	B) Explained component						
Q10	-0.160	-0.235*	-0.230*	-0.199*	-0.198*	-0.0407	0.0477
Q25	-0.0407	-0.227*	-0.216*	-0.144*	-0.113	0.0112	0.105
mean	-0.117	-0.388**	-0.385**	-0.328**	-0.222*	-0.0158	0.103
Q50	-0.0097	-0.263*	-0.263*	-0.206**	-0.156*	0.0568	0.281
Q75	-0.135	-0.443*	-0.446*	-0.404*	-0.264*	-0.0870	-0.142
Q90	-0.588	-0.969*	-0.954*	-0.900*	-0.622*	-0.422	-0.567
	C) Unexplained component						
Q10	-0.238**	-0.422***	-0.407***	-0.348***	-0.316***	-0.352***	-0.383***
Q25	-0.390***	-0.428***	-0.425***	-0.444***	-0.432***	-0.435***	-0.475***
mean	-0.158**	-0.190**	-0.185**	-0.168*	-0.193**	-0.105	-0.0883
Q50	-0.297**	-0.320**	-0.324**	-0.349**	-0.374**	-0.383**	-0.415**
Q75	-0.333**	-0.226*	-0.250*	-0.210	-0.323**	-0.334**	-0.355**
Q90	0.482	0.498*	0.518*	0.539	0.463	0.751	0.923*
N	34	34	34	34	34	34	34

Notes: In each cell, we present the coefficient estimates from equations (4), (5), and (6), which regress the components of the wage decomposition — wage gap (Block A), explained part (Block B), and unexplained part (Block C) — on the lagged unemployment rate (the business cycle indicator) and in a linear time trend. These estimates are obtained for different employment structures (t* = 1991, 1993, 1998, 2003, 2012, 2019) and at various points of the wage distribution (Q10, Q25, mean [representing the mean OLS results], Q50, Q75, and Q90). In Block A, we report the coefficient estimates from equation (4) related to the business cycle indicator. For example, the coefficient at the intersection of row Q50 and the 1998 column measures the effect of a one-percentage-point increase in the lagged unemployment rate on the wage gap at the median of the wage distribution (Q50) considering the employment structure of 1998 (t*). Specifically, an increase of one percentage point in the lagged unemployment rate reduces the gender wage gap by 0.555 percentage points at that point in the distribution, admitting that the employment structure is that of 1998. We present results using the mean OLS estimation as a benchmark. Blocks B and C display the estimates corresponding to the explained and unexplained components of the wage decomposition, respectively. Significance levels: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table 2A: tests of equal cyclicality across the wage distribution, wage gap, Portugal, 1985-2020.

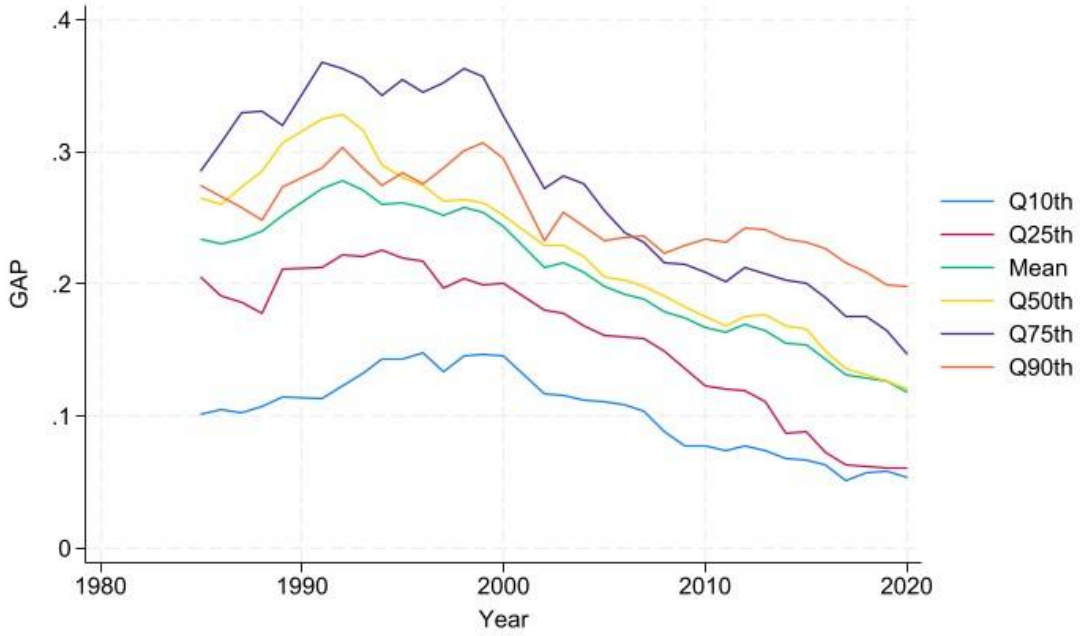
Employment structure	Points of the wage distribution (Q)							
	Q10=…Q50=..=Q90	Q10=Q50	Q10=Q90	Q90=Q50	Q25=Q50	Q25=Q75	Q25=Q90	Q50=Q75
1991	10.584**	.318	4.504**	1.063	.759	.012	5.324**	.821
1993	13.141**	.142	5.039**	1.866	.413	.169	6.618**	1.529
1998	15.615***	.005	8.288***	2.699	.122	.064	8.05***	.984
2003	11.983**	.019	6.967***	4.234**	.025	.173	5.908**	.916
2012	13.45***	.427	10.173***	6.751***	.741	.001	9.354***	.756
2019	11.323**	1.769	3.32*	1.104	2.623	.168	3.095*	.85

Table 2B: tests of equal cyclicality across the wage distribution, unexplained component

1991	37.419***	.777	16.648***	8.128***	2.13	8.456***	12.186***	3.61*
1993	32.681***	.493	16.97***	8.391***	1.765	6.531***	12.876***	2.213
1998	19.228***	0	13.608***	8.961***	1.564	11.554***	13.33***	3.704*
2003	17.859***	.306	13.216***	9.814***	.523	3.272*	14.424***	.848
2012	15.37***	.154	11.274***	9.709***	.424	2.022	12.493***	.304
2019	18.658***	.167	11.559***	10.108***	.608	2.256	12.82***	.32

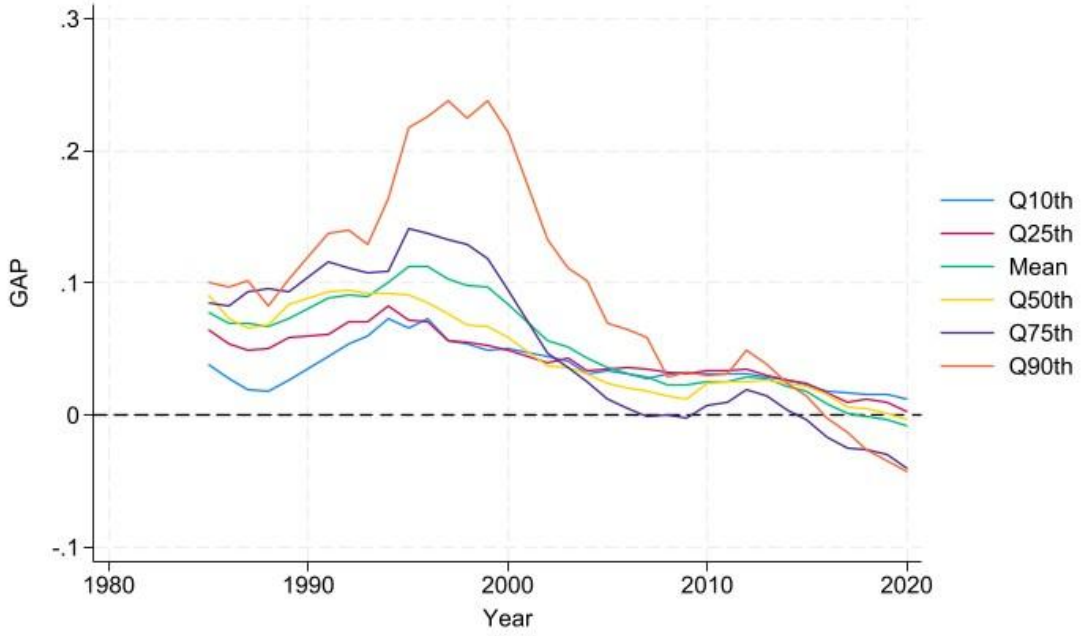
Notes: in each cell we display the chi squared statistic for the null hypothesis that the cyclicality of each component of the wage distribution (estimated gap, block 2A; unexplained component, block 2B) at given point of the wage distribution is equal to the cyclicality in another point (or points) of the wage distribution. Hence, we test the equality of coefficients displayed on table 1, across different points of the wage distribution (in line), but for the same employment structure (t* - in column). For example, the statistic of test in block A, cell (2003; Q90=Q50) = 4.234**, tests the hypothesis that the cyclicality of the gap at the median of the wage distribution (Q50) is equal to the cyclicality of the gap at 90*th* percentile (Q90) of the wage distribution when the employment structure considered is that of 2003 (from Table 1, cell (Q90, t*=2003)=cell(Q50, t*=2003)= -0.158=-0.529**). As these tests involve parameters estimated in different models (different points of the wage distribution), but on the same data, we carry out these tests with the Stata commands *test* and *suest.* The latter of these builds a simultaneous matrix of variance-covariance for the estimates in different equations based on the sandwich estimator (Cameron and Trivedi 2022). Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Figure 1: estimated wage gap, Portugal, 1985-2020



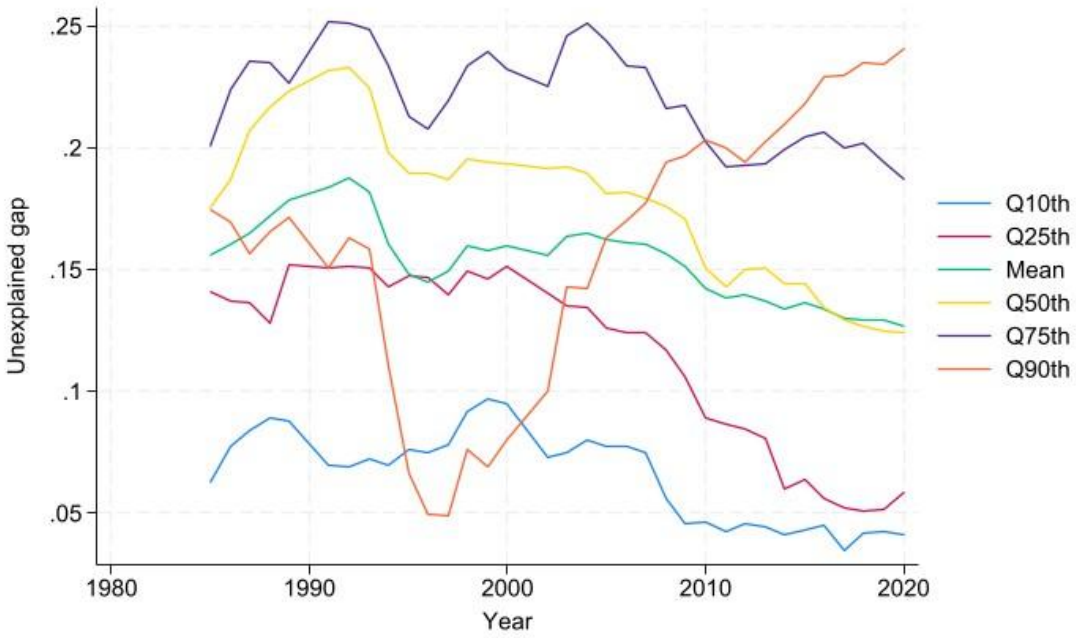
Note: This figure presents the estimated wage gap from equation 3, without reweighting.

Figure 2: Estimated wage gap, explained component, Portugal, 1985-2020



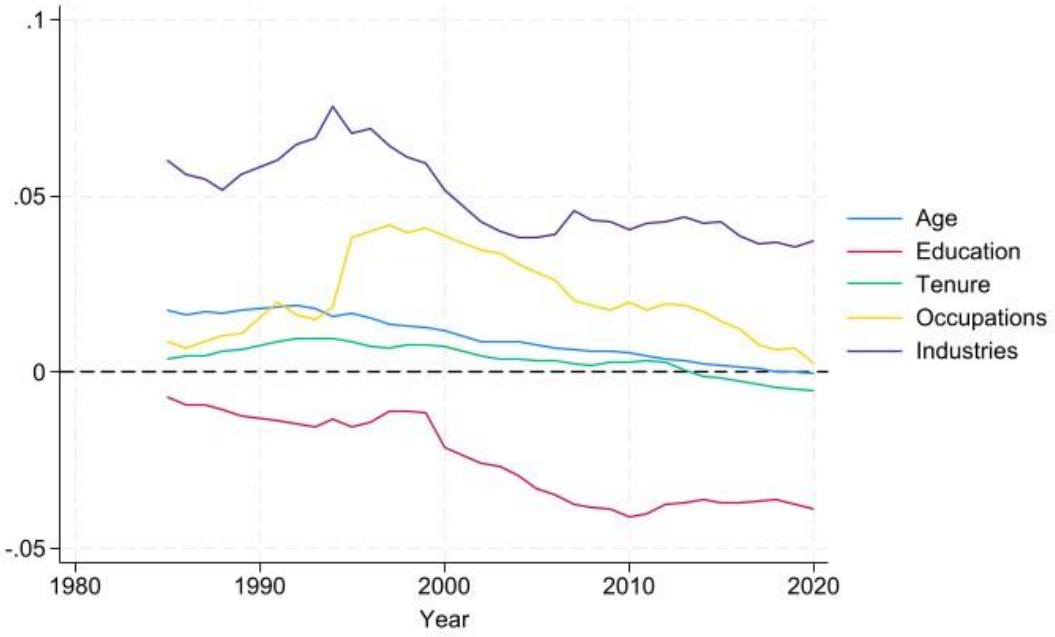
Note: This figure presents the estimate of the explained component of the wage gap from equation 3, without reweighting.

Figure 3: Estimated wage gap, unexplained component, Portugal, 1985-2020



Note: This figure presents the estimate of the unexplained component of the wage gap from equation 3, without reweighting.

Figure 4: Estimated wage gap, explained component, detailed, mean of the wage distribution Portugal, 1985-2020



Note: This figure presents the estimates of the explained component of the wage gap from equation 3 (without reweighting), detailed by component, that is, the contribution in absolute value of each group of covariates to the explained component of the gap.

APPENDIX A: DESCRIPTIVE STATISTICS

Table A: summary statistics, selected variables, Portugal, 1985-2020

	Men	**Women**
Log real hourly earnings	0.355 (0.565)	0.181 (0.509)
Age (years)	39.043 (11.038)	37.617 (10.499)
Tenure (years)	8.680 (10.540)	8.075 (9.589)
Education (non-defined)	0.011	0.009
Education (less than basic school)	0.023	0.017
Education (basic school)	0.287	0.220
Education (preparatory)	0.186	0.168
Education (lower secondary)	0.214	0.199
Education (upper secondary)	0.177	0.230
Education (university degree)	0.097	0.151
N	37,585,348 (58.3%)	26,851,803 (41.7%)

Note: Standard deviations are in parentheses. Source: authors' computations based on *Quadros de Pessoal* data set.

APPENDIX B- COMMON SUPPORT ANALYSYS

In this appendix we analyze the common support assumption for men and women, the groups considered in our decomposition analysis (Step 2 of the proposed methodological approach). This analysis is carried out by computing propensity scores for men and women, and identifying the set of observations that share similar values of these scores. This set of observations represents the region of common support (Glewwe and Todd 2022)

Hence, in the first step, we estimate the probability of being male or female in each year of our sample using a logit model, conditional on the set of variables included in our wage equations: age, age squared, tenure, tenure squared and, at two digits level, occupations and industries.

In the second step, we compute the minimum and maximum values of the estimated propensity scores and identify the observations for both genders that fall within this range. This defines the common support region.

Table B.1 presents the results of this analysis, showing that the vast majority of observations fall within the common support region. The minimum share of observations within the common support is 89.75% (in 1993), the maximum is 99.99%

(in 2007), and the average is about 97%. These results clearly indicate that the decomposition approach satisfies the common support assumption

Table B.1: Common support - sample without reweighting

Year	share of observations	N
1985	0,934385998	1288856
1986	0,92078819	1234778
1987	0,92714844	1258669
1988	0,911658689	1282820
1989	0,913734983	1383875
1991	0,913061475	1415713
1992	0,904541567	1453177
1993	0,897465753	1457435
1994	0,898437649	1463308
1995	0,960646681	1527368
1996	0,963235405	1526115
1997	0,963676827	1649085
1998	0,969398761	1667939
1999	0,969109731	1750616
2000	0,972019555	1838534
2002	0,981045902	1943432
2003	0,9908339	2001942
2004	0,99768785	2072097
2005	0,998051564	2234099
2006	0,999567274	2271643
2007	0,99985636	2332225
2008	0,999474196	2386822
2009	0,9997973	2279227
2010	0,999769717	2275462
2011	0,999541297	2236743
2012	0,999648256	2089586
2013	0,998933268	2067999
2014	0,999070168	2118663
2015	0,999008048	2172484
2016	0,998802529	2247236
2017	0,998850467	2337471
2018	0,999146304	2423578
2019	0,999179049	2404528
2020	0,998818497	2343626

The values in this table should be read as follows: In 2000, for example, approximately 97,2% of the observations (N= 1838534) are within the common support region.

Finally, we also analyze whether the observations included in the common support region are affected by the reweighting process. To assess this, we compute different common support regions for each year based on the various reweighting factors $w_{t*,t}$. Specifically, for each year, we estimate logit models and the corresponding propensity scores associated with each reweighting factor $w_{t*,t}$.

We present only selected examples[23] - specifically, cases in which the reweighting process is more likely to affect the common support region due to significant differences in employment structures between distant years. These include, for example, the years 1985 (the initial year of the sample) and 2015, when the data are reweighted to represent employment structures as distant as those of 2019 or 1991, respectively.

This analysis is presented in Tables B.2 and B.3. The main conclusion is that the common support regions obtained after reweighting are very similar to those presented in Table B.1 without reweighting.

[23] Other cases are available on request.

Table B.2: common support, 1985
Reweighted sample

share of observations	wt*, t
.9344504	w1991, 1985
.93446591	w1993, 1985
.93459937	w1998, 1985
.93468006	w2003, 1985
.93466764	w2012, 1985
.93458074	w2019, 1985

The values in this table should be read as follows: In 1985, considering, for example, the employment structure of 2012, 93,466764 % of the observations are within the common support region.

Table B.3: common support, 2015
Reweighted sample

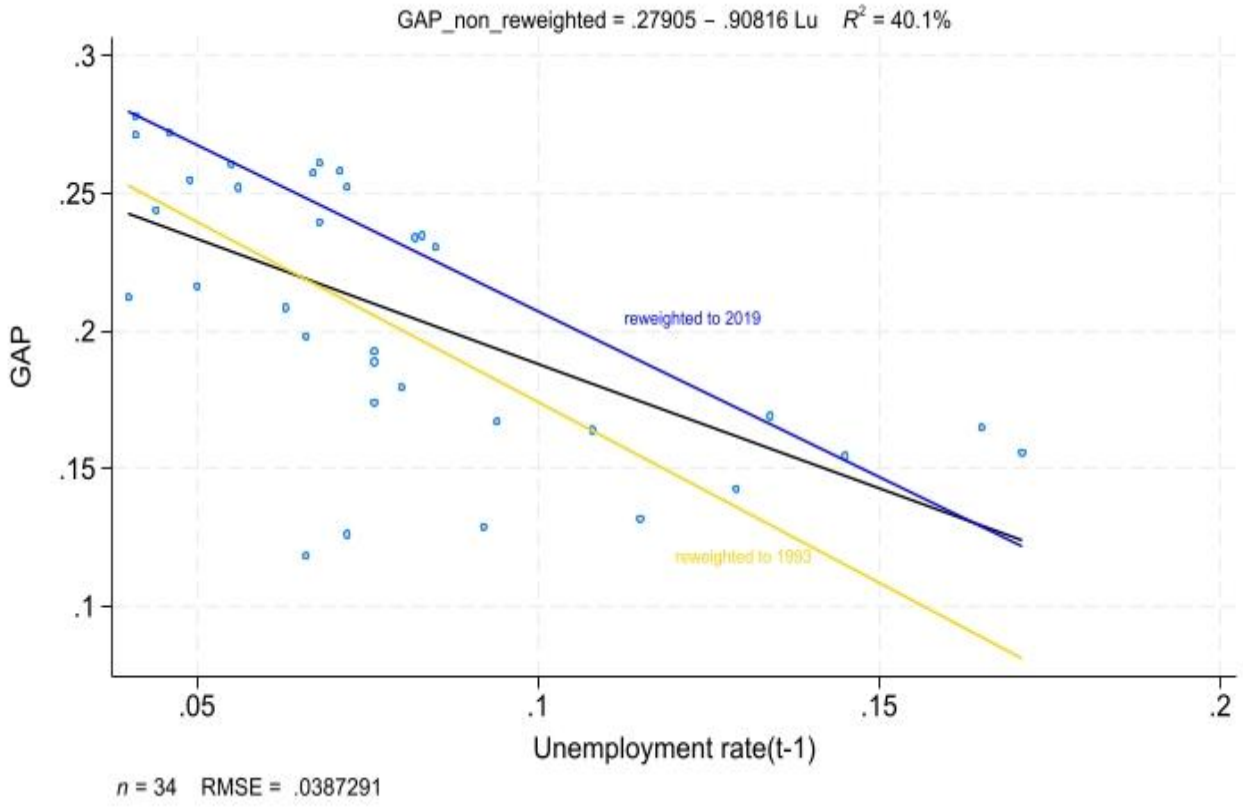
share of observations	wt*, t
.99874199	w1991, 2015
.99878066	w1993, 2015
.99886259	w1998, 2015
.99909827	w2003, 2015
.99903152	w2012, 2015
.99899102	w2019, 2015

The values in this table should be read as follows: In 2015, considering, for example, the employment structure of 1991, 99,874199 % of the observations are within the common support region.

APPENDIX C: EFFECT OF THE REWEIGHTING ON THE RELATION BETWEEN THE ESTIMATED WAGE GAP AND THE BUSINESS CYCLE

Figure C.1 shows the effect of the lagged unemployment rate on the wage gap, considering both reweighted and non-reweighted estimates.

Figure C.1: Wage gap and the business cycle, reweighted and non-reweighted estimates



These estimated linear regressions do not include a linear time trend, although visual inspection of the data clearly suggests one exists. This is one of the reasons why we include a linear time trend in our base estimates. The figure also shows that the sensitivity of the wage gap to the business cycle depends on the employment structure adopted.