A Work Project, presented as part of the requirements for the Award of a Master's degree in Economics from the NOVA School of Business and Economics.

UNVEILING THE RELATIONSHIP BETWEEN BACHELORS' FIELDS OF STUDY AND THE GENDER WAGE GAP

Pedro de Vasconcelos Teodósio Nunes dos Reis

Work project carried out under the supervision of:

Professor Pedro Portugal

10-01-2025

Unveiling the relationship between bachelors' fields of study and the gender wage gap

Pedro Reis †

January 10, 2025

Abstract

This paper explores the relationship between bachelors' fields of study and the gender wage gap in Portugal. Results suggest that the STEM fields give relatively large returns and indicate a 14 percent reduction in the estimated gender wage gap once fields of study are considered. Gender differences in returns to fields are statistically significant and generate different gender gaps among bachelors of various fields. Firm sorting explains part of the gender gap among bachelors in some fields. There is some evidence of selection bias in gender wage gaps estimates as individuals self-select into their degrees.

Keywords: Bachelors, Education, Field of study, Firm Fixed Effects, Gelbach Decomposition, Gender Wage Gap, Heckman, Selection, STEM, Worker Fixed Effects

This work used infrastructure and resources funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences DataLab, Project 22209), POR Lisboa (LISBOA-01-0145-FEDER-007722 and Social Sciences DataLab, Project 22209) and POR Norte (Social Sciences DataLab, Project 22209).

†. I want to thank professor Pedro Portugal for his guidance in this research, as well as my family and friends for all their support during this period.

1 Introduction

This paper explores the relationship between the gender wage gap and workers' field of study (college major) in their bachelor's degrees, in Portugal. We want to investigate how the estimated gender wage gap is affected once field of study is accounted for. Further, we want to understand if the gender gap is homogeneous across workers with degrees in different fields of study and how this relationship between the gender gap and the fields of study is related to the workers' own characteristics, discrimination, and firm sorting. Finally, as people self-select into their degrees, this work addresses the issue of selection bias which could hinder the validity of its results.

There is extensive literature on the gender wage gap, and there are also numerous studies on returns to tertiary education that explore how these returns depend on the field of study, but few papers dive into the relationship between both topics. Altonji, Blom, and Meghir (2012), Lemieux (2014), and Card and Payne (2021) address the issue and their empirical evidence suggests that the choice of field of study can explain a significant part of the gender wage gap as men are more likely to join fields that provide larger returns such as the STEM fields (Science, Technology, Engineering, Mathematics). However, these results are often side quests to the central object of these investigations and so they are not evaluated or tested in depth. Nevertheless, Altonji, Blom, and Meghir (2012) explore the subject with more detail and find gender differences in returns to the same fields of study.

The first contribution of this work is therefore to account for the field of study in estimating the gender wage gap and then, by accounting for the gender differences in returns to fields, to estimate the gender wage gap among bachelors of different fields. This contribution is all the more valuable due to the quality of our data. We use *Quadros de Pessoal* which provides longitudinal and matched firm-worker data for all wage earners in Portugal (domestic servants and the public administration are excluded).

Another contribution is the decomposition of the gender gap in the various fields of study into a firm component and a worker component. We follow the method developed by Gelbach (2016) and also applied by Cardoso, Guimarães, and Portugal (2016) to unambiguously (i.e.

without depending on the sequential ordering of covariates) decompose the gender wage gap into these components. However, while Cardoso, Guimarães, and Portugal (2016) decompose the overall wage gap into firm, worker, and occupation components, we decompose, for each field of study, the gender wage gap among bachelors in the same field into a firm component (to control for firm sorting), and a worker component (which may, arguably, measure the role of discrimination). This advances both the literature on returns to tertiary education and the literature concerning the impact of firm sorting and discrimination on the gender wage gap.

A final contribution is that we aggregate the various fields of study into either one of two groups, STEM and Non-STEM fields, and use this aggregation to apply the method developed by Heckman (1979) to deal with a potential selection bias due to individuals self-selection into their bachelor's degrees. This contribution is also methodological since, as reported by Altonji, Blom, and Meghir (2012), most papers do not directly address selection bias and just "use OLS and hope that controls are adequate", while others employ a regression discontinuity approach.

We conclude that some fields of study that are overpopulated by men, in particular the STEM fields, give relatively large returns and that this factor can explain part of the gender wage gap. We show that there is heterogeneity in the gender wage gap among graduates of different fields and provide evidence that this heterogeneity comes from differences across bachelors of different fields in both the worker fixed effects and the firm sorting contributions to the gender wage gap. In some fields, firm sorting explains part of the gender wage gap, while the worker fixed effects contribution - a possible measure of discrimination effects - always explains more than half of the gap. Finally, there is some moderate evidence of selection bias in the estimates of the gender wage gaps as people self-select into their degrees.

The remaining of this paper goes as follows: Section 2 reviews the literature on the gender wage gap and fields of study; Section 3 describes the main data used in this paper; Section 4 presents and discusses the methodology applied in this work; Section 5 presents and discusses the results of the main model, and accesses its robustness to the selection bias; Section 6 concludes.

2 Literature Review

The gender wage gap is among the most well documented facts in economics. In their extensive literature review, Blau and Kahn (2017) recall several traditional explanations for this phenomenon. These include human-capital factors (education and work experience), labour force participation (lower for women), the distribution of roles in the household (which often prioritize the man's career over the woman's), the motherhood child penalty (the tendency for mothers to have lower wages than other women), differences in unionization rates, women shorter careers (on average), and discrimination. Recent studies have looked into the role of sorting into firms and sorting into occupations, which have proven to be significant determinants of the wage gap, both in Portugal and in the United States (Cardoso, Guimarães, and Portugal, 2016; Card, Cardoso, and Kline, 2016).

The role of discrimination is not yet clear. As it is hard to measure, it is often assumed to belong to the unexplained gap (Blau and Kahn, 2017), but the existence of unobservables suggests that the unexplained gap overestimates the discrimination effects - one could make the case that they are underestimated if discrimination affects control variables such as education, occupation sorting, and firm sorting.

Due to lower birth rates and improved education for women, traditional explanations of the gender wage gap explain less than they used to. According to Blau and Kahn (2017), education variables explain about 3 percent of the wage gap in 1980 in the USA, but they increase the unexplained part of the gap in 2010. Portugal exhibits similar trends and the unexplained wage gap in Portugal is increasing again (Cardoso et al., 2016). Therefore, while gender differences in the level of education and labour force participation became smaller (or started favouring women), explanations for the gender wage gap must look into other factors such as the workers' fields of study (college majors) in their bachelor's and master's degrees.

Regression discontinuity approaches using admission cut-offs to tertiary education in Norway (Kirkeboen, Leuven, and Mogstad, 2016) and Chile (Hastings, Neilson, and Zimmerman, 2013) report significant heterogeneity in returns to the field of study. Some research also discusses the presence of comparative advantage and heterogeneity of returns within fields due to matching with individuals and occupations (Paglin and Rufolo, 1990; Hastings, Neilson, and Zimmerman, 2013). In fact, various analyses in different countries suggest that there are larger returns in some fields over-represented by men (typically physical sciences and engineering) over other fields over-represented by women (arts, humanities, and education), although the aggregation of the field of study categories is not consistent across studies (Altonji, Blom, and Meghir, 2012; Zafar, 2013; Kirkeboen, Leuven, and Mogstad, 2016; Card and Payne, 2021). Moreover, Card and Payne (2021) run OLS regressions controlling for age, region, part-time status, and level of education and suggest that the gender gap in STEM fields (Science, Technology, Engineering, Mathematics) may explain up to 20 percent of the wage gap in Canada and the USA. A comparable result to that of Altonji, Blom, and Meghir (2012) whose OLS regressions find a significant decrease in the gender wage gap in the USA once field of study is controlled for and who similarly find different returns to the same fields of study for men and women.

The aforementioned studies establish a correlation between gender and field of study allocation which could explain part of the gender wage gap. However, the (apparent) causal relationship is often not studied in depth nor is it robustly estimated or tested as it is not the central object of any of this research. Thus, the literature does not yet allow us to draw strong conclusions regarding the relationship between the fields of study and the gender wage gap.

3 Data

A strength of our work is the quality of the data. It comes from *Quadros de Pessoal* which comprises annual census data collected in Portugal since 1985. It is mandatory to be filled by every employer with at least one wage earner (apart from public administration and domestic servants), and it provides matched firm and worker data over several years. *Quadros de Pessoal* contains information regarding a worker's age, gender, monthly wage, hours of work, and date of entry into the firm as well as the level of education of the worker and its field of study, the firm's region and other variables. This is therefore a uniquely rich panel dataset as it covers a wide range of variables on a large and representative sample of the whole economy, so it provides a rare opportunity to dive into the relationship between earnings, gender, and fields of study.

The population of interest is people with solely a bachelor's degree, so we exclude people without such a degree and people with a master's degree or a PhD. Our sample period starts in 2006 and ends in 2022 and the data refers always to the month of October of each year. We reduce our sample to mainland Portugal.

Among the remaining 5 048 506 observations, approximately 96 percent have only one job. This percentage is similar for men and women (respectively 94 percent and 97 percent). Hence, moonlighters - i.e. people with more than one job - are removed from the sample as this simplifies the analysis without a relevant cost regarding the object of this study. We are left with 4 837 525 observations. Of these, the largest connected set of worker fixed effects and firm fixed effects is composed of 4 528 830 which we take as the sample of this work. This includes data on 750 453 workers, 87 489 firms, and the years [2006,2022] (not all workers and firms have data on every year).

4 Methodology

4.1 Specifications

We start with a Mincer wage equation including a gender dummy and conventional wage controls. Following Cardoso, Guimarães, and Portugal (2016) and Portugal et al. (2024) we include quadratic terms for age and tenure to control for cohort and experience returns, we also employ year fixed effects (represented by θ_t). Since the population of interest concerns people with a bachelor's degree there is no need to control for the level of education. In line with the literature presented above, we introduce dummy variables to control for the region of the firm (region is defined according to the NUT2 classification). Assuming a total of R regions, we arrive at the following specification for every individual i in year t:

$$\ln w_{it} = c + \gamma gender_i + \beta_1 age_{it} + \beta_2 age_{it}^2 + \beta_3 tenure_{it} + \beta_4 tenure_{it}^2 + \sum_{r=2}^{R} \rho_r region_{it} + \theta_t + u_{it} \quad (1)$$

where u is the error term and w is the real hourly wage.

The absence of controls for fields of study may generate an omitted variable bias regarding the estimator for γ. Thus, specification 2 introduces dummy variables to control for fields of study. Considering a total of J fields, specification 2 would be (s stands for field of study):

$$\begin{aligned}\ln w_{it} = c + \gamma gender_i + \beta_1 age_{it} + \beta_2 age^2_{it} + \beta_3 tenure_{it} + \beta_4 tenure^2_{it} \\ + \sum_{j=2}^{J} \eta_j s_{ji} + \sum_{r=2}^{R} \rho_r region_{it} + \theta_t + u_{it}\end{aligned} \quad (2)$$

Recall that the choice of field of study is mutually exclusive. In the example above, the omitted field of study is field s_1 and γ is the gender wage gap.

Specification 3 adds interactions between gender and field of study to capture gender heterogeneity in returns to the same field:

$$\begin{aligned}\ln w_{it} = c + \gamma gender_i + \beta_1 age_{it} + \beta_2 age^2_{it} + \beta_3 tenure_{it} + \beta_4 tenure^2_{it} \\ + \sum_{j=2}^{J} \eta_j s_{ji} + \sum_{j=2}^{J} \delta_j s_{ji} gender_i + \sum_{r=2}^{R} \rho_r region_{it} + \theta_t + u_{it}\end{aligned} \quad (3)$$

In the specification above, γ is the gender wage gap among bachelors with a degree in the omitted field of study.

The full specification contains the same controls (henceforth represented as vector $\boldsymbol{x_{it}}$) and worker fixed effects α_i and firm fixed effects ϕ_f. The worker fixed effects control for time-invariant characteristics such as innate ability, race, and gender. Notice that gender, field of study, and the interaction variables are absorbed by the worker fixed effects and so are dropped out of the vector $\boldsymbol{x_{it}}$. Thus, the full specification is the following (f stands for firm):

$$\ln w_{it} = \boldsymbol{x_{it}\beta} + \sum_{r=2}^{R} \rho_r region_{it} + \theta_t + \alpha_i + \phi_f + u_{it} \quad (4)$$

4.2 The Gelbach decomposition

The inclusion of relevant variables from a limited specification to a full specification changes the estimated parameters of the initial specification due to the reduction of the omitted variable bias. Gelbach (2016) proposed a decomposition to unambiguously quantify the change of a parameter estimate which should be attributed to each of the new variables.

We will employ Gelbach's decomposition to disentangle the contribution of the covariates which were added in specification 4 (firm and worker fixed effects) to the gender wage gaps estimated in specification 3. Cardoso, Guimarães, and Portugal (2016) show that this decomposition is still possible when, as in our case, the main variable of interest (gender) is absorbed by the fixed effects.

Let $\hat{\gamma}_{base}$ be the estimator of γ in the baseline specification (specification 3). γ is the wage gap in the field of study s_1. It can be shown that we can regress each of the fixed effects on the explanatory variables of the baseline decomposition to obtain their contribution to the previously found gender gaps, and we have:

$$\hat{\alpha}_i = c^{\alpha} + \gamma^{\alpha} gender_i + \beta_1^{\alpha} age_{it} + \beta_2^{\alpha} age_{it}^2 + \beta_3^{\alpha} tenure_{it} + \beta_4^{\alpha} tenure_{it}^2 + \sum_{r=2}^{R} \rho_r^{\alpha} region_{it} + \sum_{j=2}^{J} \eta_j^{\alpha} s_{ji} + \sum_{j=2}^{J} \delta_j^{\alpha} s_{ji} gender_i + \theta_t^{\alpha} + u_{it}^{\alpha} \quad (5)$$

$$\hat{\phi}_i = c^{\phi} + \gamma^{\phi} gender_i + \beta_1^{\phi} age_{it} + \beta_2^{\phi} age_{it}^2 + \beta_3^{\phi} tenure_{it} + \beta_4^{\phi} tenure_{it}^2 + \sum_{r=2}^{R} \rho_r^{\phi} region_{it} + \sum_{j=2}^{J} \eta_j^{\phi} s_{ji} + \sum_{j=2}^{J} \delta_j^{\phi} s_{ji} gender_i + \theta_t^{\phi} + u_{it}^{\phi} \quad (6)$$

From these auxiliary regressions we can take the estimates $\tilde{\gamma}^{\alpha}$ and $\tilde{\gamma}^{\phi}$. Gelbach's theorem implies that:

$$\hat{\gamma}^{base} - \hat{\gamma}^{full} = \tilde{\gamma}^{\alpha} + \tilde{\gamma}^{\phi} \quad (7)$$

In this decomposition, $\tilde{\gamma}_{\alpha}$ and $\tilde{\gamma}_{\phi}$ are the estimated contributions of, respectively, the the worker fixed effects and the firm fixed effects to the gender wage gap in the omitted field of study. We can work out the decomposition of the parameters δ_j to decompose the gender gap in any field.

In this case we cannot estimate γ in the full specification as gender is absorbed by the worker fixed effects. As in (Cardoso, Guimarães, and Portugal, 2016), $\hat{\gamma}^{full}$ is absorbed by $\tilde{\gamma}^{\alpha}$, but the decomposition is possible and we have:

$$\hat{\gamma}^{base} = \tilde{\gamma}^{\alpha} + \tilde{\gamma}^{\phi} \tag{8}$$

Although gender is absorbed by the worker fixed effects, firm fixed effects may still be important to eliminate the bias in any estimate of the effects of gender. This is intuitive if one considers that firm sorting can be correlated with gender and wages. Hence, removing firm fixed effects would produce a bias, even when controlling for worker fixed effects (because, in such a case, firm sorting could be correlated with worker fixed effects).

4.3 Robustness

As individuals self-select into their field of study, we want to evaluate the robustness of our main results to selection bias. For that, we will apply the Heckman (1979) method to the individuals' selection into fields of study.

First, we simplify the individual's choice from a multivariate choice where the individual can rank and choose among multiple fields of study (as defined in *Quadros de Pessoal*) to a binary choice where the individual can choose between either a STEM or a Non-STEM field. Thus, we split the twenty-two fields of study defined in *Quadros de Pessoal* into eleven STEM fields and eleven Non-STEM fields. [1]

Second, we adapt our main model, its four specifications, and the Gelbach (2016) decomposition to this binary choice framework.

We start by working only with the wages of STEM bachelors. For simplicity, let $\boldsymbol{x_{it}}$ be the vector of regressors in specification 1 and β be the respective vector of regression coefficients. We would like to estimate the parameters of the following wage equation among STEM bachelors:

$$\ln w_{it}{}^{STEM} = \boldsymbol{x_{it}\beta^{STEM}} + u_{it}^{STEM} \tag{9}$$

1. STEM fields: Agriculture, Forestry, and Fisheries; Architecture and Construction; Computer Science; Engineering and Related Technologies; Environmental Protection; Health; Life Sciences; Manufacturing Industries; Mathematics and Statistics; Physical Sciences; Veterinary Sciences.

Non-STEM fields: Arts; Business Sciences; Teaching; Humanities; Information and Journalism; Law; Personal Services; Security Services; Social and Behavioural Sciences; Social Services; Transportation Services

Notice this is specification 1 with field aggregation and only concerning STEM bachelors. If there is selection bias, u^{STEM} will not be orthogonal to the explanatory variables, and the estimates of the parameters will be biased. We therefore include a selection variable λ^{STEM} to control for selection into STEM. Our goal is now to estimate the parameters of the following specification which we will call "STEM Heckman Specification":

$$\ln w_{it}{}^{STEM} = \boldsymbol{x_{it}}\boldsymbol{\beta^{STEM}} + \zeta^{STEM}\lambda_{it}^{STEM} + \epsilon_{it}^{STEM} \tag{10}$$

Where, after the inclusion of λ^{STEM}, ϵ_{it}^{STEM} is the error term and holds to the standard OLS assumptions.

Then, we build a selection equation which we will use to estimate λ^{STEM}. Along with $\boldsymbol{x_{it}}$, we also include an instrumental variable denoted by z_i in the selection equation: the ratio of the number of new STEM bachelors over new Non-STEM bachelors in the civil year the worker i becomes 21 years old - henceforth called STEM-Non-STEM Ratio.[2] This variable is likely to be exogenous to wages (once fields of study are accounted for). Yet, the choice of STEM fields should be correlated with the STEM-Non-STEM ratio. Hence, the selection equation is the following:

$$STEM_{it} = \mathbf{1}[\boldsymbol{x}_{it}\boldsymbol{\beta}^{selection} + \psi z_i + v_{it} > 0] \tag{11}$$

Where v follows a standard normal distribution and is correlated with u^{STEM}.

We proceed to estimate the parameters in the selection equation using probit. The inverse Mills ratio for each observation gives us the estimates of λ^{STEM} which we will use to control for selection. We then employ these results to estimate the parameters using OLS in the STEM Heckman Specification.

Finally, we repeat exactly the same procedure and keep the same variables (including the STEM-Non-STEM Ratio), this time building the "Non-STEM Heckman Specification" where we work only with the wages of Non-STEM bachelors.

2. In Portugal, if a student does not fail or skip any year and does not take any gap year, he will finish high school during the civil year he becomes 18 and he will finish obtain bachelor's degree in the civil year he becomes 21 (the vast majority of bachelor's degrees in Portugal require 3 years of study). As in Portugal, the academic year starts in the Fall, a student born in 1993, for instance, will, under these conditions, finish high school in the academic year of 2010/11 and graduate by the end of the academic year of 2013/14.

One should note that we only collected data to calculate the STEM-Non-STEM Ration for individuals born during or after 1993, for we dropped all the remaining observations in all these estimates. [3]

4.4 A discussion about biases

All specifications control for age, tenure, region, and year fixed effects. There is no need to control for the level of education as there is no variation of it in our sample. There is a strong case in the literature that the choice of field of study affects wages and is correlated with gender, so specification 1 will probably produce a biased estimate of the gender gap. Specification 2 accounts for the field of study allocation, but its estimate of the gender gap can be biased if returns to any given field are not equal for men and women. Specifications 1 and 2 may suffer too from the omitted variable bias we can find in specification 3 which we will discuss hereafter.

Specification 3 ignores worker heterogeneity and worker allocation into firms. If men and women work, on average, in firms with different pay policies or different productivity, their wages could differ. Such is the finding of Cardoso, Guimarães, and Portugal (2016) who conclude that firm sorting explains a significant share of the gender gap in Portugal.

Specification 3 also fails to account for unobservable features of individuals such as mental ability, psychological traits, and physical traits. Of these, the main feature that can affect productivity and wages is ability. If ability is uncorrelated with gender, it should not produce any bias regarding the gender gap. One can conjecture that ability is correlated with the field of study allocation, and its omission will bias the estimated returns to fields.

Lemieux (2014) states that "studies that use IV strategies to estimate the causal effect of education on earnings have generally concluded that the ability bias was small". This is unsurprising if, for instance, ability is a determinant of the level of education but has otherwise no strong effects on earnings. It is plausible that a similar pattern occurs regarding the causal effect of the field of study. Moreover, we can expect any ability bias in our population of interest to be small as it is only made of bachelors and therefore the variation in ability should be lower

3. This particular data was not retrieved from *Quadros de Pessoal* but from the Brighter Future database which collected, worked, and presented data regarding the bachelors who graduated in the 2013/14 academic year and up to the 2022/23 academic year as reported by the Direção-Geral de Estatísticas da Educação e Ciência (Directorate-General for Education and Science Statistics).

than in studies that include several levels of education.

All in all, it can be argued that the ability bias on the estimated effects of field of study returns is small, but we cannot exclude its presence. Moreover, other features such as physical and psychological traits are not accounted for, and these may affect wages and be correlated with fields of study and gender. Nevertheless, within a given field, there should be no ability bias on the estimated gender gap, and the remaining biases (from unobservable traits) should be small as the variation in unobservable variables should be lower than in the general population.

Considering all these, the full specification (specification 4) includes controls for both individual and firm fixed effects. Although gender and field of study are individual-specific and time-invariant and are, therefore, absorbed by the fixed effects, the application of Gelbach's decomposition will allow us to recover the estimated gender gap in each field and unambiguously estimate how it can be explained by the new controls introduced in the full specification (worker and firm fixed effects).

Firm fixed effects capture the wage premium (penalty) enjoyed by all workers who work at a better (worse) paying firm, regardless of these workers' features. Thus, firm fixed effects measure the impact of firm sorting on wages. If correlated with gender, firm sorting explains part of the gender gaps found in the baseline specification (specification 3). The remaining part is explained by worker fixed effects.

If there is no ability or other biases on the gender wage gap, as we expect ability to be uncorrelated with gender, and if there is no discrimination in firm sorting, the worker fixed effects contribution can be interpreted as a measure of the effects of discrimination, as interpreted by Cardoso, Guimarães, and Portugal (2016).

This measure of discrimination effects may still have flaws. On the one hand, it could overestimate discrimination effects as it assumes that there are no relevant differences between men and women regarding features such as preferences and athleticism, which are captured by the worker fixed effects and which may arguably affect earnings. This is, in general, a reasonable assumption as most jobs in today's Portuguese economy do not demand athleticism, and as the differences in preferences and other psychological traits within graduates of the same field should be less pronounced than within the population as a whole, especially as we also control for age and region. In some specific fields, notably Security Services, a relevant proportion of graduates may indeed perceive careers that demand high levels of athleticism or strength - in such cases, the worker fixed effects contribution probably overestimates the impact of discrimination as worker fixed effects are affected by athleticism and strength too. On the other hand, this measure of discrimination may underestimate its effects if there is discrimination in firm sorting (i.e. better-paying firms being more discriminatory towards women than the other firms). Nevertheless, the biases tend to have opposite signs and should be, in most cases, small. We will therefore follow Cardoso, Guimarães, and Portugal's (2016) approach.

Finally, all these specifications may suffer from selection bias as the allocation to fields of study is not random, but it is chosen by individuals, and it also depends on other factors like the number of vacancies available. Although some of the controls that are present in all specifications should already be removing some selection bias (e.g. gender, region).[4] As the full specification is controlling for worker fixed effects, the worker's time invariable characteristics which affect their personality, skills, and preferences and which may therefore affect the field of study selection are being accounted for - as our discussion regarding the ability bias illustrates. However, preferences and other selection determinants are not time-invariant and this may generate some bias.

To assess the robustness of our main model, we will apply Heckman's selection model and analyse if we are drawn to comparable conclusions.

5 Results

5.1 Part 1 - Main Model

In all specifications, age, and tenure exhibit positive and diminishing returns, which is intuitive and consistent with other literature for Portugal (Portugal et al., 2024). The gender gap in specification 1 (Table 1) is approximately 23 log points, which equals the estimates of Cardoso, Guimarães, and Portugal (2016) who also look at Portugal and control for age and tenure

4. Although the region variable in the specifications refers to the region of the firm, and not the region of the worker, these variables should be correlated; region of the worker may affect selection due to features such as preferences and culture.

Table 1: Specifications 1 and 2

	(1)	(2)
Age/100	4.272***	4.721***
	(0.529)	(0.053)
$(\text{Age}/100)^2$	-2.584***	-3.087***
	(0.071)	(0.070)
Tenure/100	0.264***	0.258***
	(0.002)	(0.002)
$(\text{Tenure}/100)^2$	-0.0386***	-0.036***
	(0.001)	(0.001)
Female	-0.228***	-0.195***
	(0.001)	(0.001)
Teaching		omitted
Arts		0.041***
		(0.004)
Humanities		0.026***
		(0.003)
Social and Behavioural Sciences		0.092***
		(0.003)
Information and Journalism		0.009*
		(0.004)
Business Sciences		0.145***
		(0.002)
Law		0.159***
		(0.005)
Life Sciences		0.132***
		(0.005)
Physical Sciences		0.158***
		(0.006)
Mathematics and Statistics		0.249***
		(0.005)
Computer Science		0.222***
		(0.003)
Engineering and Related Technologies		0.188***
		(0.003)
Manufacturing Industries		0.128***
		(0.010)
Architecture and Construction		0.105***
		(0.005)
Agriculture, Forestry, and Fisheries		0.019**
		(0.006)
Veterinary Sciences		0.017*
		(0.008)
Health		0.237***
		(0.002)
Social Services		-0.066***
		(0.003)
Personal Services		-0.047***
		(0.006)
Transportation Services		0.312***
		(0.162)
Environmental Protection		0.024**
		(0.008)
Security Services		0.212***
		(0.022)
N	4 528 829	4 528 829
R^2	0.354	0.371
Year fixed effects	✓	✓
Region fixed effects	✓	✓

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Age is measured in years, tenure is measured in months.
Robust standard errors to clusters at the worker level are depicted in parenthesis.

(although their population of interest is not restricted to bachelors). All these results are coherent with the literature, suggesting that the results and conclusions of this research may perhaps be extended to broader groups and compared with other studies on the gender wage gap and on returns to education, even though this work has a relatively narrow population of interest as it only concerns college graduates.

Specification 2 (Table 1) controls for fields of study and assumes that men and women enjoy equal returns to each field. As expected, the STEM fields tend to give larger wages, whilst Teaching and the Humanities receive smaller returns. For instance, returns to Engineering and Related Technologies are 18.8 log points higher than returns to Teaching. It is worth noting that Altonji, Blom, and Meghir (2012), Hastings, Neilson, and Zimmerman (2013), and Lemieux (2014) obtain similar results to ours regarding returns to fields and the advantage of STEM graduates.

Importantly, the distribution of graduates across fields is different for men and women . About 32 percent of men graduated in Engineering and Related Technologies, while only 8 percent of women did so.[5] In fact, the average gender gap in specification 2 drops to approximately 20 log points. This is a 14 percent reduction of the original gap, which is slightly lower than from the results of Card and Payne (2021) where field of study may explain up to one fifth of the gender gap in Canada and the USA. We thus infer that it is important to control for fields of study when estimating the gender wage gap. We should notice however that, despite this, we cannot conclude that field of study explains 14 percent of the gender gap among bachelors in Portugal because, as shown by Gelbach (2016), this method of decomposition of the gender gap is ambiguous since it depends on the sequential ordering of covariates.

These results suggest that field of study heterogeneity may as well be a useful factor to understand the glass ceiling reported in the literature (Blau and Kahn, 2017): if graduates in STEM are over-represented among the highest earnings quantiles, we should expect a glass ceiling to occur.

Specification 3 (first column of Table 2) controls for gender differences in fields premia.[6]

5. For more descriptive statistics see Table 6 and Table 7 in the Appendix

6. The joint-significance test F test on the new controls (interactions between gender and each field) rejects the null hypothesis at a 1 percent significance level, from which we conclude that gender differences in returns to the same fields are statistically significant.

Specification 3 does not tell us how the gender gap is reduced by accounting for the field of study. Instead, it shows that there is not one common gender gap among graduates, but different gender gaps among the graduates of different fields because male and female returns vary within the same field, a result which Altonji, Blom, and Meghir (2012) also obtained. The gender gaps always favour men and are statistically significant. These vary substantially as they range from 7 log points among graduates in Social Services to 57 log points in Security Services graduates.

As discussed before, these estimates regarding the gender wage gap within fields can be biased due to the lack of controls. We should recall that this bias should come mostly from firm sorting, as ability is unlikely to be correlated with gender and as the variation in psychological traits between workers in the same field should be smaller than in the general population.

The full specification (specification 4) includes firm fixed effects, and worker fixed effects to control for firm sorting, worker ability and other unobservable firm and individual time-invariant effects on wages (physical appearance, psychological features, etc.). We are now faced with the issue that some variables in specification 3 - gender, field of study, and the interaction terms - are absorbed by the fixed effects. The Gelbach decomposition (Table 2) estimates the contributions of the new controls - individual and firm fixed effects - to the gender wage gaps estimated in the baseline specification (specification 3).

The gender gap among Teaching bachelors which was estimated in the baseline specification is 20.6 log points. Once we account for firm sorting, this gap is reduced by 9.8 log points, almost half of the original gap; the remaining part of the gap is explained by worker fixed effects. In other words, if men and women were randomly distributed across firms, the gender gap among Teaching bachelors would be 10.9 log points. An analogous interpretation can be made for the results concerning graduates in any other field, although, the contribution of firm sorting to the gender gap is usually much lower in absolute terms as well as in relation to the contribution of the worker-fixed effects component.

The estimated contribution of discrimination as measured by the worker component explains, always, more than half of a field's gender gap. Firm sorting generally favours men, but there are some fields, such as Computer Science, where its contribution to the gender gap is irrelevant. In fact, among graduates in the Humanities and in Mathematics and Statistics, firm sorting favours

Table 2: Decomposition of the Gender Wage Gaps

	Gender wage gap	Worker contribution	Firm contribution
Teaching	-0.206***	-0.109***	-0.098***
	(0.07)	(0.006)	(0.004)
Arts	-0.119***	-0.112***	-0.012**
	(0.008)	(0.007)	(0.005)
Humanities	-0.111***	-0.125***	0.015***
	(0.007)	(0.006)	(0.004)
Social and Behavioural Sciences	-0.260***	-0.234***	-0.028***
	(0.005)	(0.005)	(0.002)
Information and Journalism	-0.108***	-0.097***	-0.010**
	(0.009)	(0.008)	(0.004)
Business Sciences	-0.229***	-0.210***	-0.018***
	(0.003)	(0.003)	(0.001)
Law	-0.152***	-0.144***	-0.007*
	(0.009)	(0.009)	(0.004)
Life Sciences	-0.147***	-0.117***	-0.030***
	(0.010)	(0.009)	(0.005)
Physical Sciences	-0.186***	-0.172***	-0.013**
	(0.011)	(0.010)	(0.006)
Mathematics and Statistics	-0.133***	-0.168***	0.034***
	(0.010)	(0.009)	(0.005)
Computer Science	-0.120***	-0.116***	-0.005
	(0.054)	(0.005)	(0.003)
Engineering and Related Technologies	-0.226***	-0.200***	-0.026***
	(0.003)	(0.003)	(0.002)
Manufacturing Industries	-0.288***	-0.272***	-0.016
	(0.020)	(0.020)	(0.010)
Architecture and Construction	-0.255***	-0.202***	-0.052***
	(0.008)	(0.007)	(0.005)
Agriculture, Forestry, Fisheries	-0.204***	-0.200***	-0.007
	(0.012)	(0.011)	(0.007)
Veterinary Sciences	-0.192***	-0.133***	-0.059***
	(0.020)	(0.019)	(0.013)
Health	-0.122***	-0.116***	-0.006***
	(0.004)	(0.004)	(0.002)
Social Services	-0.070***	-0.046***	-0.025***
	(0.011)	(0.010)	(0.007)
Personal Services	-0.201***	-0.159***	-0.045***
	(0.013)	(0.012)	(0.007)
Transportation Services	-0.416***	-0.361***	-0.049**
	(0.034)	(0.030)	(0.019)
Environmental Protection	-0.071***	-0.067***	-0.010
	(0.017)	(0.016)	(0.010)
Security Services	-0.570***	-0.440***	-0.123***
	(0.039)	(0.033)	(0.015)
N	4 528 829	4 528 829	4 528 829
Control variables	✓	✓	✓

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Control variables: quadratic in age, quadratic in tenure, year fixed effects, region fixed effects, field of study fixed effects.

Robust standard errors to clusters at the worker level are depicted in parenthesis. Singletons were kept in the estimation of specification 4, thus likely overstating the statistically significance of the decomposition (Correia, 2014).

The coefficients regarding the fields of study represent the estimated gender gap among graduates of each field according to specification 3. The gender gaps and the returns to fields are not estimated in specification 4 as they are absorbed by the worker fixed effects. The results of specification 4 are used to apply Gelbach's (2016) decomposition method to estimate the worker fixed effects and firm fixed effects contribution to each gender gap.

The R^2 of specification 3 is 0.373, the R^2 of specification 4 is 0.897, and the within-R^2 of specification 4 is 0.139.

women, implying that the estimated discrimination effect is larger than the gender wage gap. Thus, the gender gap among graduates in these two fields would favour women in the absence of discrimination.

Cardoso, Guimarães, and Portugal (2016) also control for age, tenure, and time fixed effects. They find that, in Portugal, about one-fifth of the remaining gender gap is explained through firm sorting (another fifth through job segregation, which we did not control for, and the remaining gap is explained by worker effects). Our results do not necessarily contradict these, but they show that the contribution of firm sorting varies strongly across the field of study of graduates. Firm sorting explains more than one-fifth of the gender gap among graduates in seven of the twenty-two fields (e.g. Teaching, Security Services), suggesting that men with those degrees tend to go to better-paying firms than women. At the other end of the spectrum, firm sorting explains less than ten percent of the gender gap among graduates of most other fields. Among graduates in Mathematics and Statistics and in the Hummanities fields, firm sorting actually favours women.

One could argue that one of the reasons for the variability in the effects of firm sorting on the gender gap is that discrimination is not homogeneous across sectors. Indeed, if technological markets (who should tend to look for STEM graduates) are competitive, firms who discriminate in hiring are penalized by the market and have to drop out or start hiring and offer wages according to worker productivity. This mechanism would not be as relevant in less competitive markets. We could possibly extend this argument to more industries that look for STEM graduates as, judging by the returns of these fields, these graduates are - or are perceived to be - more productive. Our approach does not provide sufficient evidence to evaluate this claim, but heterogeneity in hiring discrimination across industries could be an area for further research.

5.2 Part 2 - Robustness

Table 3 shows the results of specifications 3 and 4 with the aggregation of fields into either STEM or Non-STEM fields, concerning workers born during or after 1993 who hold a bachelor's degree. The gender gap is 8.6 log points among Non-STEM bachelors and 4.5 log points among STEM bachelors, this goes in line with the previous results that the gender gap tends to be lower among

Table 3: Specifications 3 and 4 with Field Aggregation

	(3)	(4)
Age/100	8.497***	17.728***
	(0.986)	(1.145)
$(\text{Age}/100)^2$	-13.392***	-30.013***
	(1.991)	(2.161)
Tenure/100	0.529***	2.533
	(0.015)	(0.016)
$(\text{Tenure}/100)^2$	-0.384***	-0.142***
	(0.026)	(0.021)
Female	-0.086***	
	(0.004)	
STEM	0.135***	
	(0.004)	
Female*STEM	0.041***	
	(0.005)	
N	200 041	200 041
R^2	0.192	0.874
Year fixed effects	✓	✓
Region fixed effects	✓	✓
Worker fixed effects		✓
Firm fixed effects		✓

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Robust standard errors to clusters at the worker level are depicted in parenthesis.

Singletons were kept in the estimation of specification 4, thus underestimating its standard errors.

The within-R^2 of specification 4 with field aggregation is 0.201.

STEM bachelors. Both gender gaps are much smaller than most gender gaps we found before, but this is probably due to sample selection and reflects how the gender gap is smaller for younger generations.

Table 4 shows the results of the Gelbach Decomposition of the gender wage gaps estimated under this field aggregation and concerning these same workers. Coherently with the pattern found in the main results, the gender gap is larger among Non-STEM bachelors compared to STEM bachelors, and firm sorting is a relevant component of the gender gap, but only among Non-STEM bachelors.

We can see that, despite the field aggregation and the decreased number of observations, these results are coherent with the results found in the main model. The estimated gender gap is different, but the impact of the fields of study on wages and the relationship between the gender wage gap and the choice of field of study is very similar. The advantage of the aggregation of the fields is that we can apply the Heckman (1979) model to control for the selection into fields. These results are depicted in Table 5.

Table 4: Decomposition of the Gender Wage Gaps with Field Aggregation

	Gender wage gap	Worker contribution	Firm contribution
STEM	-0.045***	-0.048***	-0.002
	(0.003)	(0.003)	(0.003)
Non-STEM	-0.086***	-0.055***	-0.031***
	(0.004)	(0.003)	(0.003)
N	200 041	200 041	200 041
Control variables	✓	✓	✓

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Control variables: quadratic in age, quadratic in tenure, year fixed effects, region fixed effects, field of study fixed effects.

Robust standard errors to clusters at the worker level are depicted in parenthesis. Singletons were kept in the estimation of specification 4, thus likely overstating the statistically significance of the decomposition (Correia, 2014).

Table 5: Selection Model

	STEM Bachelors	STEM Selection	Non-STEM Bachelors
Selection Equation			
STEM-Non-STEM Ratio		0.095***	
		(0.024)	
Female		-0.326***	
		(0.006)	
Sample Size			
N		200,041	
Wage Equation			
Female	-0.079***		-0.140***
	(0.019)		(0.027)
Mills Ratio			
Lambda	0.167*		-0.262**
	(0.099)		(0.122)
Sample Size			
Selected	111,320		88,721
Control variables	✓	✓	✓

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Control variables: quadratic in age, quadratic in tenure, year fixed effects, region fixed effects.

The selection equation was estimated using a probit model, and the wage equation was estimated with OLS. Standard errors are depicted in parenthesis.

The correlation between the error terms in the selection equation (i.e. $\hat{v}$) and the wage equation without the selection control (i.e. $\hat{u}$) is +0.481 in the STEM Bachelors specification and -0.687 in the Non-STEM Bachelors specification.

The selection equation is in line with our expectations. First, gender is a predictor of the field of study (men are more likely to join a STEM field). Second, the STEM-Non-STEM ratio is another predictor of the field of study allocation and thus a proper instrument. Notice that age, tenure, year, and region are being controlled for.

Regarding the wage equation, we find that, once selection is controlled for, the gender gaps increase from 4.5 log points to 7.9 log points among STEM bachelors and from 8.6 log points to 14 log points among Non-STEM bachelors. On the one hand, this supports the findings of the previous models where the gender gap among STEM bachelors was found to be lower than among Non-STEM bachelors. On the other hand, it suggests that the gender wage gaps estimated in the previous models may be underestimated due to selection bias. In fact, the estimated selection effect of λ (i.e. the estimate of coefficient ζ in equation 10) is statistically significant at a 10 percent level in both STEM and Non-STEM specifications, but only in the Non-STEM specification is it significant at a 5 percent level.

The estimated selection effect is positive in the STEM specification, meaning that the average STEM bachelor would earn more than the average bachelor in the (artificial) case were both go to STEM. In other words, on average, the characteristics that make someone more likely to graduate in STEM (one may speculate if a high mental ability is one such characteristic) also make this person more likely to receive receive higher wages than the average. Furthermore, the estimated selection effect is negative in the Non-STEM specification, meaning that the average Non-STEM bachelor would earn less than the average in the (artificial) case were both go to a Non-STEM field. Thus, people who graduate in STEM hold, on average, an absolute advantage in both fields than people who graduate in Non-STEM fields.

We can therefore conclude that there is some evidence of selection bias regarding the estimated returns to fields of study and the gender wage gaps. However, we should be cautious in interpreting the results of these selection models due to the limited sample upon which we applied them and due to their lack of controls, for instance, for firm sorting and worker fixed effects, besides the loss of information and other problems which can arise from the aggregation of twenty-two fields into only two categories. It could be interesting to see what would happen to the selection bias and to the firm contribution to the gender wage gaps once both selection and firm fixed effects were considered (a work which we will not develop given the multitude of firm fixed effects and the already very truncated sample would not allow for any strong conclusion).

6 Conclusion

We can summarize our results in four main findings.

First, once we account for the workers' fields of study, there is a modest reduction of 14 percent of the estimated gender wage gap (controlling for age, tenure, year, and region). This is consistent with the literature on returns to fields of study and reflects the fact that, compared to women, men tend to go to areas that exhibit higher returns, in particular the STEM fields.

Second, there is a strong heterogeneity in the gender wage gap among bachelors of different fields of study due to gender differences in returns within fields. Therefore, when investigating the relationship between wage differences between men and women and their bachelor's field of study it is more appropriate to refer to various gender wage gaps rather than a homogeneous wage gap for bachelors in all fields.

Third, as estimated by the Gelbach decomposition, the contribution of the worker fixed effects to the gender wage gaps (a possible measure of discrimination), explains more than half of the estimated gender gap in all fields. Firm sorting has a smaller, although still significant, explanatory power among bachelors in some fields. For instance, firms sorting explains 9.8 log points of the 20.6 log point gender gap among Teaching bachelors but it only explains 0.5 log points of the gender gap among Computer Science bachelors. There is some evidence that the contribution of firm sorting to the gender wage gap among STEM bachelors is null.

Fourth, there is some evidence that accounting for the field of study but not controlling for selection into fields of study produces biased estimates of the gender wage gap as individuals self-select into their fields of study.

From these conclusions, we can present promising issues for further research. First, one could extend our analysis to a broader population of interest and include different education levels rather than just bachelors. Second, one could look at the impacts of firm and worker fixed effects on the gender gap in different fields of study while also controlling for selection in the main model with the full sample. Finally, one could aim to disentangle the contribution of worker fixed effects to the various gender gaps into features such as mental ability, physical traits, psychological traits, and preferences, to produce a better measurement of the effects, and types of discrimination.

References

Altonji, Joseph G., Erica Blom, and Costas Meghir. 2012. "Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers." *Annual Review of Economics* 4 (September): 185–223.

Blau, Francine D., and Lawrence M. Kahn. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations." *Journal of Economic Literature* 55, no. 3 (September): 789–865.

Brighter Future. *Explorar Cursos e Áreas.* Accessed: 2024-11-30. https://www.brighterfuture.pt/explorar/cursos.

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *The Quarterly Journal of Economics* 131, no. 2 (May): 633–686.

Card, David, and A. Abigail Payne. 2021. "High School Choices and the Gender Gap in STEM." *Economic Inquiry* 59, no. 1 (January): 9–28.

Cardoso, Ana Rute, Paulo Guimaraes, Pedro Portugal, and Pedro S Raposo. 2016. "The sources of the gender wage gap." *Banco de Portugal Economic Studies* 2, no. 2 (April): 47–66.

Cardoso, Ana Rute, Paulo Guimarães, and Pedro Portugal. 2016. "What drives the gender wage gap? A look at the role of firm and job-title heterogeneity." *Oxford Economic Papers* 68, no. 2 (April): 506–524.

Correia, Sérgio. 2014. *Singletons, Cluster-Robust Standard Errors and Fixed Effects: A Bad Mix.* Technical Note. Duke University, November.

Gelbach, Jonah B. 2016. "When Do Covariates Matter? And Which Ones, and How Much?" *Journal of Labor Economics* 34, no. 2 (April): 509–543.

Hastings, Justine S, Christopher A Neilson, and Seth D Zimmerman. 2013. *Are Some Degrees Worth More than Others? Evidence from college admission cutoffs in Chile.* Working Paper, Working Paper Series 19241. National Bureau of Economic Research, July.

Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47, no. 1 (January): 153–161.

Kirkeboen, Lars J., Edwin Leuven, and Magne Mogstad. 2016. "Field of Study, Earnings, and Self-Selection." *The Quarterly Journal of Economics* 131, no. 3 (August): 1057–1111.

Lemieux, Thomas. 2014. "Occupations, Fields of Study and Returns to Education." *Canadian Journal of Economics/Revue canadienne d'économique* 47, no. 4 (November): 1047–1077.

Paglin, Morton, and Anthony M. Rufolo. 1990. "Heterogeneous Human Capital, Occupational Choice, and Male-Female Earnings Differences." *Journal of Labor Economics* 8, no. 1 (January): 123–144.

Portugal, Pedro, Hugo Reis, Paulo Guimarães, and Ana Rute Cardoso. 2024. "What Lies behind the Returns to Schooling: The Role of Labor Market Sorting and Worker Heterogeneity." *The Review of Economics and Statistics* (August): 1–45.

Zafar, Basit. 2013. "College Major Choice and the Gender Gap." *The Journal of Human Resources* 48, no. 3 (July): 545–595.

Appendix - Descriptive Statistics

Table 6: Men Field of Study Distribution

	Absolute Frequency	Relative Frequency (%)
Teaching	34 987	1.79
Arts	42 050	2.15
Humanities	49 012	2.50
Social and Behavioural Sciences	128 582	6.57
Information and Journalism	29 949	1.53
Business Sciences	440 266	22.50
Law	46 589	2.38
Life Sciences	23 118	1.18
Physical Sciences	24 628	1.26
Mathematics and Statistics	27 591	1.41
Computer Science	207 102	10.58
Engineering and Related Technologies	631 702	32.28
Manufacturing Industries	8 352	0.43
Architecture and Construction	60 983	3.12
Agriculture, Forestry, and Fisheries	20 657	1.06
Veterinary Sciences	5 862	0.30
Health	140 536	7.18
Social Services	8 121	0.42
Personal Services	10 957	0.56
Transportation Services	5 432	0.28
Environmental Protection	6 391	0.33
Security Services	3 883	0.20
Total	1 956 750	100.00

These numbers concern the number of "male observations", meaning that a man who graduated in Teaching and is recorded in the dataset during 3 years contributes with 3 units to the absolute frequency of male Teaching bachelors.

Table 7: Women Field of Study Distribution

	Absolute Frequency	Relative Frequency (%)
Teaching	238 104	9.26
Arts	51 363	2.00
Humanities	143 878	5.59
Social and Behavioural Sciences	295 281	11.48
Information and Journalism	65 781	2.56
Business Sciences	591 805	23.01
Law	86 633	3.37
Life Sciences	56 894	2.21
Physical Sciences	31 949	1.24
Mathematics and Statistics	49 474	1.92
Computer Science	56 647	2.20
Engineering and Related Technologies	198 528	7.72
Manufacturing Industries	8 791	0.34
Architecture and Construction	36 239	1.41
Agriculture, Forestry, and Fisheries	20 691	0.80
Veterinary Sciences	14 003	0.54
Health	487 993	18.97
Social Services	96 921	3.77
Personal Services	24 157	0.94
Transportation Services	1 519	0.06
Environmental Protection	12 112	0.47
Security Services	3 317	0.13
Total	2 572 080	100.00

These numbers concern the number of ”female observations”, meaning that a woman who graduated in Teaching and is recorded in the dataset during 3 years contributes with 3 units to the absolute frequency of female Teaching bachelors.