Less Is More? How Shorter University Degrees Affect Educational and Early-Career Outcomes

Emma Scandolo*

October 15, 2025

Abstract

This paper examines how students respond when university degrees become shorter, exploiting Portugal's Bologna Process implementation as a natural experiment. Using a regression discontinuity design based on cohort age at implementation, I find that the reform increased university attainment by 1.5pp and average schooling by $\sim$0.3 years. Students shifted away from healthcare programs—where Bologna uniquely lengthened degrees—toward STEM and economics fields, with effects concentrated among females. To capture welfare impacts, I track cumulative earnings by age, which integrates both timing effects (when people enter the labor market) and human capital effects (what they earn). While cumulative earnings are initially negative through the mid-twenties due to foregone wages as more students pursue university, they turn positive by the early thirties. Treated workers experience steeper early-career wage growth and increased sorting into higher-paying firms, particularly among females who shifted into higher-return fields. I show that 24% of the wage increase operates through the workplace channel, whereas 66% is associated exclusively to the individual. Overall, gains from expanding university access and reallocating students toward higher-paying fields outweigh losses from degree compression, validating the "less is more" outcome policymakers intended.

*PhD Candidate, Department of Economics, University College London.

Email: emma.scandolo.21@ucl.ac.uk

1 Introduction

How do students respond when the time required to complete university degrees changes? While an extensive literature examines the effects of extending compulsory schooling—typically through reforms that raise the minimum school-leaving age (Angrist and Krueger, 1991; Meghir and Palme, 1999; Acemoglu and Angrist, 2000; Oreopoulos, 2007; Pischke and Von Wachter, 2008; Portugal et al., 2024)—we know surprisingly little about what happens when higher education degrees become shorter. This gap matters for education policy worldwide, as countries have extensively debated whether the gains from shorter degrees attracting more students to university outweigh the potential human capital losses from compressing the same curricular content into fewer years of study. The fundamental question is whether shorter degrees achieve the "less is more" outcome policymakers hope for—where reducing time requirements actually increases total human capital in the population by bringing in students who would otherwise not attend university. Understanding how students navigate this trade-off when degree lengths change and what the consequences are for their careers is crucial, yet empirical evidence remains scarce.

The Bologna Process—Europe's largest higher education reform—provides an ideal natural experiment to fill this gap, and has been largely understudied (Kroher et al., 2021). This reform pushed over 20 countries to restructure their university systems, transitioning from Continental European models to the Anglo-Saxon one (OECD, 2011). The key difference between these models was the minimum time needed to earn a university degree: Continental European systems typically required 4-5 years of study before students could graduate, while the Anglo-Saxon model offered a shorter 3-year bachelor's degree as a first-cycle credential followed by a 2-year master's degree as a second-cycle credential. This comprehensive restructuring across an entire continent offers three key advantages for identification: the reform's near-universal adoption within countries eliminated selection into treatment, the rapid implementation timeline created clean discontinuities in exposure, and the magnitude of the change—reducing minimum degree requirements by 25-40%—ensures sufficient variation to detect behavioral responses.

The net effect of shorter degrees on educational attainment and earnings is theoretically ambiguous because the reform simultaneously shifts two margins that work in opposite directions. On the extensive margin, shorter degrees lower barriers to university entry: students who previously found a 4- to 5-year commitment too costly might now enroll in the more accessible 3-year program. These marginal students gain access to higher-paying jobs requiring university credentials, but must delay labour-market entry and forego earnings while studying. On the intensive margin, the reform creates an earlier exit option: students who would have completed 4-5 years under the old system might now stop after 3 years with a bachelor's degree. These students enter the labour market sooner and accumulate more work experience by any given age, but potentially sacrifice human capital from the foregone years of education. Which effect dominates—and thus whether shorter degrees increase or decrease total years of schooling and earnings *on net*—is an empirical question that speaks directly to the optimal design of higher education systems.

To assess the consequences of shorter degrees, this paper examines two sets of outcomes. The first part analyzes educational responses: (i) What is the net effect on university graduation rates and years of schooling? (ii) How much of the effect on university graduation is driven by the uptake of shorter degrees? (iii) How do field of study choices shift, and what are the implications for occupational sorting? The second part examines labour market consequences: (iv) What is the effect on cumulative earnings by a specific age? (v) How do changes in wages decompose into worker quality and firm sorting? (vi) Does compressing degrees fundamentally alter the returns to university education?

Portugal offers exceptional conditions to identify these effects. The reform implementation was remarkably sharp, with almost all programs restructured between 2006/07 and 2007/08. Portugal's standard university entry age of 18 implies that cohorts turning 18 in 2007 or later faced the new system when making enrollment decisions, while older cohorts had already decided whether to enter or forego university under the previous system. This age-based discontinuity, combined with Portugal's comprehensive linked employer-employee data (Quadros de Pessoal) that tracks highest educational attainment and labour market outcomes, enables me to trace how the reform affected both educational choices and career trajectories. To address potential selection into private sector employment, I complement the analysis with Labor Force Survey (LFS) data covering the entire population.

I find that the "less is more" hypothesis holds. The Bologna reform increased university graduation rates by 1.5 percentage points ($\sim$ 6.5% relative to baseline), driven primarily by an uptake of shorter degrees. The extensive margin dominates on average, with proxies indicating increased total years of schooling in the population. The reform also induced substantial reallocation across fields. Students shifted away from healthcare programs ($\sim$ 30% decline relative to baseline)—where, uniquely, Bologna actually lengthened degrees to meet European standards—toward STEM ($\sim$ 22% increase relative to baseline) and Economics ($\sim$ 10% increase relative to baseline) that became more accessible with shorter requirements. These effects are particularly pronounced for females, who show both stronger field reallocation and higher overall graduation increases. To estimate the net effects on earnings, I focus on cumulative earnings by a given age. This avoids making wage comparisons at a specific point-in-time or "years-since-entry", as such strategies condition on variables that the reform itself endogenously shifts (how much labour market experience an individual has at a specific point in time, or how old an individual is when they enter the labour market). Cumulative earnings aggregate the reform's complex effects into a single, interpretable measure: it captures both timing effects (delayed entry for new university students versus earlier entry for those who before completed longer degrees and now complete shorter degrees) and human capital effects (gains from more education versus potential losses from compression), providing a welfare-relevant measure that integrates all mechanisms through which the reform operates. I find that cumulative earnings effects are negative in the mid-twenties but turn positive by the early thirties. Initial losses are larger for women; however, women in post-reform cohorts exhibit faster early-career wage growth and are more likely to move to higher-paying firms (higher AKM firm effects) in their first three years. Following Portugal et al. (2024), I decompose wage effects and find that 24% of the wage increase operates through the workplace channel, whereas 66% is associated exclusively to the individual. Lastly, using Lee bounds to address selection in estimating returns to education, I find that by age 30 graduates at the bottom of the earnings distribution gain, while those in the middle and top show no significant decline in returns.

This paper makes three main contributions. First, I provide the first evidence on how the Bologna Process affected graduation rates, field composition, and labour market trajectories—moving beyond existing work that documents only enrollment effects finding a 7-20% increase across European countries (Argentin and Triventi, 2011; Brunori et al., 2013; Cappellari and Lucifora, 2009; Di Pietro and Cutillo, 2008; Bondonio and Berton, 2018; Di Pietro, 2012; Horstschräer and Sprietsma, 2015; Triventi et al., 2017). Second, whereas an extensive body of work examines the effects of extending compulsory schooling (Acemoglu and Angrist, 2000; Angrist and Krueger, 1991; Meghir and Palme, 1999; Oreopoulos, 2007; Pischke and Von Wachter, 2008), I study what happens when higher education requirements are compressed. More recent work by Portugal et al. (2024) has focused on decomposing the returns to schooling into labour market sorting and worker heterogeneity in Portugal, exploiting increases in compulsory schooling as exogenous variation. While I perform a similar decomposition exercise, I rely on a reduction in university length as exogenous variation. Third, my findings on field reallocation speak to the literature on heterogeneous returns to field of study (Bleemer and Mehta, 2022; Hastings et al., 2013; Kirkeboen et al., 2016; Zimmerman, 2014). While those papers identify marginal students through academic ability cutoffs, I identify students at a different margin—those constrained by the time required to complete a given field of study.

The rest of the paper is structured as follows: section 2 describes the reform and institutional background, section 3 the data and section 4 the empirical strategy and identification, section 5 describes the main findings and lastly section 6 concludes.

2 Reform background

The Bologna Process fundamentally represented a choice to converge European higher education toward the Anglo-Saxon model rather than maintaining the diverse Continental European traditions. Three key factors drove this decision. First, Anglo-Saxon countries demonstrated substantially higher tertiary education participation rates—by 2004, 39 percent of Americans aged 25-64 held tertiary degrees compared to just 23 percent of Europeans, suggesting the shorter degree structure successfully expanded access (Aghion, 2006). Second, the Continental European landscape had become increasingly fragmented, with degree lengths varying from four to six years across countries and fields, creating significant barriers to student mobility and degree recognition. In contrast, the Anglo-Saxon three-year bachelor's followed by an optional two-year master's provided a transparent, standardized framework that facilitated cross-border movement of students and workers. Finally, the compressed timeline to a first degree reduced both the financial burden and opportunity cost of higher education, allowing students to enter the labour market sooner or pivot more easily if their initial field choice proved unsuitable—features particularly valuable for students from disadvantaged backgrounds who faced greater constraints in committing to lengthy programs (Jacobs and

Van Der Ploeg, 2006).

2.1 Institutional Background in Portugal

Portugal's implementation of this Anglo-Saxon convergence provides an exceptional natural experiment for studying the effects of compressing higher education degrees. Portugal maintained a binary system comprising universities, which focused on general academic education and research, and polytechnic schools, which were traditionally more vocationally oriented. Both sectors offered the licenciatura as the first degree, followed by optional postgraduate programs including the mestrado (master's degree). The Bologna framework, initiated by the 1999 Bologna Declaration and reinforced through subsequent ministerial conferences in Prague (2001), Berlin (2003), and Bergen (2005), aimed to create a European Higher Education Area with comparable degree structures across countries. The core requirement was adopting a three-cycle system: a bachelor's degree (typically three years), master's (normally two years), and doctorate. While the standard model prescribed a three-year first cycle, Portugal retained flexibility through the "integrated master" option—a combined five to six-year program that students could pursue in fields where professional requirements or academic traditions warranted longer training.

The Portuguese Ministry of Science, Technology and Higher Education (MSTHE) implemented Bologna as part of a broader reorganization of the higher education system, occurring against a backdrop of significant market pressures. The sector faced growing imbalances between supply and demand: while enrollment in public universities had increased by 62% during the 1990s, demographic changes subsequently reduced the candidate pool, admission standards tightened with the reintroduction of national exams and minimum grades, and capacity expanded through public sector investments. This created an increasingly competitive environment where institutions needed to differentiate themselves to attract students. The reform's implementation schedule offered institutions strategic flexibility that proved revealing. Universities could begin implementing Bologna-compliant curricula in the 2006/07 academic year, with mandatory completion by 2008/09, yet virtually all programs had converted by the 2007/08 academic year. As documented by Cardoso et al. (2008) in the reform's first year, approximately 43% of study programs adapted to Bologna principles by 2006/07 and only one institution delayed the implementation of curricula changes to the 2008/09 academic year.

For my empirical analysis, this implementation structure offers several advantages. First, the compressed timeline—with virtually all programs transitioning by 2007/08—creates a sharp discontinuity rather than the gradual phase-in that complicates identification in other countries. Second, Portugal's standard university entry age of 18 aligns perfectly with the reform timing: cohorts turning 18 in 2007 or later faced the new system when making enrollment decisions, while older cohorts had already entered or foregone university under the previous system. This age-based discontinuity, combined with Portugal's comprehensive administrative data on education and employment, enables clean identification of the reform's causal effects on both educational attainment and labour market outcomes.

3 Data

3.1 Quadros de Pessoal (QP)

Quadros de Pessoal (QP) is a rich and comprehensive linked employer-employee administrative data collected annually by Portugal's Ministry of Employment. It is a panel dataset that encompasses all private sector establishments with at least one employee, with wage data recorded each October. While it excludes civil servants, the self-employed, and domestic workers, it captures the complete universe of workers and firms in manufacturing and private services.

The following variables are reported on each worker: gender, birth date, educational attainment, occupation, firm entry date, monthly compensation, working hours, and hierarchical position (such as manager or highly qualified employee). Educational variables capture both the highest degree completed and the field of study. Employer characteristics include firm size and geographic location. To construct hourly wages, I divide total monthly compensation—comprising base salary, seniority premiums, and other regular payments—by standard monthly hours worked. All wage measures are deflated to 2010 euros using the consumer price index. I use information stretching from 2000 to 2023. However, no worker data are available for 2001. Table 2 in the Appendix presents the main descriptive statistics.

I construct two important measures related to earnings and wages. As the reform comes in place in 2007, post-reform individuals will be observed latest at age 34, whereas pre-reform individuals are observed up to age 41. Similarly, post-reform individuals will be observed, on average, at later calendar years compared to pre-reform cohorts (as shown in Table 2 in the Appendix). To control for any time-specific trends that might confound wages, I index wages using older cohorts observed in the same calendar year and of the same gender. A more detailed description of the construction of relative wages can be found in Section A.1 in the Appendix. Secondly, I construct cumulative annual earnings by a given age. For each individual observed in the dataset, I set annual earnings to zero if the person has no observed earnings at a specific age, and I compute cumulative earnings at specific age A by aggregating all the cash flows earned up to age A. A more detailed description of the construction of cumulative earnings can be found in Section A.2 in the Appendix.

3.2 Labor Force Surveys (LFS)

The Labor Force Survey (LFS) provides a nationally representative sample of Portugal's entire population, collected quarterly by the National Statistics Institute. Unlike the QP data, the LFS is a repeated cross-section rather than a panel—I observe different individuals in each survey wave and cannot track the same person over time. However, this dataset offers crucial advantages for benchmarking: it covers all employment statuses including public sector workers, the unemployed, and those out of the labour force, providing an unselected view of the population's educational and employment outcomes. The survey records detailed information on educational attainment, field of study, employment status, occupation, and wages, using similar educational classifications as the QP data. This compatibility allows me to validate that my QP-based results on educational attainment are not driven by selection into private sector employment. I use LFS data spanning the same period from 2000 to 2023, enabling consistent comparisons across both datasets throughout the pre- and post-reform periods. While the repeated cross-sectional nature prevents me from analyzing individual career trajectories, the LFS serves as an essential robustness check to ensure my findings reflect population-wide effects rather than compositional changes in private sector employment.

4 Empirical Strategy

Using the data described in the previous section, I construct a running variable $Z_i = c_i - 2007$, where $c_i = \text{Year of birth} + 18$, which captures the distance, in years, between the year individual i was 18 and the year in which the reform was fully implemented, i.e. 2007. To identify treated individuals, i.e. those more exposed to the reform, I use the treatment dummy $D_i = \mathbf{1}[Z_i \geq 0]$.

To estimate the impact of the reform, I estimate the following Fuzzy Regression Discontinuity equation:

$$Y_i = \alpha_0 + \alpha_1 D_i + \alpha_2 Z_i + \alpha_3 Z_i \times D_i + \alpha_4 \mathbf{X}_i + \eta_i \tag{1}$$

where α_1 is the coefficient of interest, α_2 is the slope of the relationship between the running variable and the outcome before the reform, for $D_i = 0$, and α_3 is the slope after the reform. $\mathbf{X}_i$ is a vector of controls including age polynomials. Y_i captures education and labour market outcomes. Education outcomes are a series of dummies taking value one if an individual's highest education is: (i) a university degree; (ii) a first-cycle university degree; (iii) a university degree in a given field. To proxy years of schooling, I use the age at which an individual reports to have started to work in the QP dataset while I use the age at which an individual reports leaving their highest education from the LFS. Labour market outcomes are: (i) log of cumulative annual earnings and relative hourly wages; (ii) dummies for whether the occupation matches the field of study; (iii) a dummy for whether the individual has a highly qualified position (manager, quadro, or highly qualified employee); (iv) wage growth in first 3 years; (v) firm premium growth in first 3 years.

Identification assumptions The main identification assumption on which my empirical strategy is based is that individuals cannot manipulate the running variable in order to benefit from the reform. As individuals do not have the power to change the year in which they were born, it is reasonable to assume that the density is smooth around the cutoff. To formally test this, I estimate separate regressions of the log density on each side of the cutoff and test for a jump at the cutoff, $Z_i = 0$. Specifically, I compute the density in each cohort bin as $\hat{f}(z_j) = N_j/(N \cdot h)$ where $h = 1$ cohort year is the bin width, and run

$$\log \hat{f}(z_j) = \alpha_L + \beta_L z_j + u_j \quad (z_j < 0), \qquad \log \hat{f}(z_j) = \alpha_R + \beta_R z_j + v_j \quad (z_j \geq 0), \tag{2}$$

using weights N_j (bin counts) and a triangular kernel within a symmetric window around the cutoff. The parameter of interest is the discontinuity in log density, $\Delta = \alpha_R - \alpha_L$, with null $H_0 : \Delta = 0$. I report the point estimate $\hat{\Delta}$, its standard error, and p-value. Results for this exercise are reported in Figure 1, where there is no visual break in the density at the cutoff. A formal test finds that the change in the log density around the cutoff is not statistically significant.

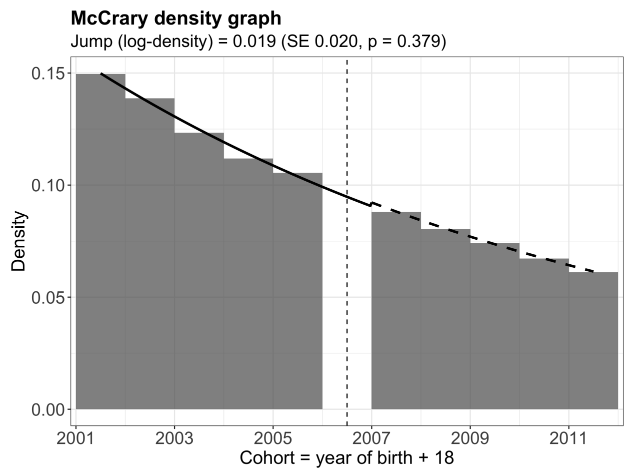

Figure 1: McCrary density test for manipulation of the running variable

Importantly, however, individuals have the power to decide at what age to enroll into university, and cohorts who were 18 years old before the reform was implemented could actively choose to delay their university enrollment in order to benefit from shorter degrees. This would be particularly relevant if individuals could anticipate the reform implementation. I test this assumption by estimating equation (1) using as an outcome the share of individuals that are observed to be working between graduating from high-school and graduating from university. Intuitively, if we were to see an increase in the share of people who take up a job before going to university, this would suggest that the reform likely delayed the timing of university enrollment. Results for this test are reported in Table 3 in the Appendix, where we see that the reform has no effect on the share of people who work between graduating high-school and graduating university, and that this is true for both the female and male sample.

4.1 AKM model for firm sorting analysis

To examine whether the reform changes worker quality and worker sorting across firms, I implement the Abowd, Kramarz, and Margolis (AKM) model (Abowd et al., 1999). This approach decomposes log wages into worker characteristics, firm-specific pay premiums, and residual variation. The AKM methodology requires restricting the analysis to the largest connected set of firms—that is, firms that are linked through worker mobility. Two firms are connected if at least one worker moves between them during the sample period, and by extension, all firms form a connected set if they can be linked through chains of such moves. This restriction ensures that worker and firm effects are separately identified through the mobility patterns in the data. In Portugal's relatively integrated labour market, the largest connected set encompasses the vast majority of firms and workers.

By including firm and worker fixed effects I am able to control for unobservables at the firm and worker level that capture a substantial amount of wage variation, while mitigating potential endogeneity problems. More specifically, I estimate the following equation:

$$y_{it} = X_{it}\beta + \alpha_i + \theta_{F(i,t)} + \mu_t + \epsilon_{it} \tag{3}$$

where y_{it} is the log hourly wage for worker i in year t, X_{it} includes a quadratic on the experience and tenure of the worker, α_i are worker fixed effects, $\theta_{F(i,t)}$ represents a fixed effect for the firm employing worker i at time t, and μ_t are year fixed effects. By tracking changes in the average firm and worker fixed effect for workers before and after the reform, I can determine whether treated cohorts earn more because they sort into higher-paying firms or because they are, on average, better workers.

I also consider an alternative AKM specification, following Portugal et al. (2024), which includes occupation and firm-worker match effects:

$$y_{it} = X_{it}\beta + \lambda_{O(i,t)} + \phi_{i\times F} + \mu_t + \epsilon_{it} \tag{4}$$

where $\lambda_{O(i,t)}$ is an occupation fixed effect, and $\phi_{i\times F}$ is a fixed effect that accounts for all firm-worker unique combinations. This is the same as adding a worker fixed effect, a firm fixed effect, and a firm-worker fixed effect. They refer to the latter as the "pure matching effect". The model presented in equation (4) nests what is presented in equation (3) in that it includes an additional occupation and match quality fixed effect.

4.1.1 Decomposing wage effects

To interpret the reform's impact on wages through the lenses of the AKM model, I decompose the discontinuity in log wages at the cutoff into the contributions of worker composition, firm sorting, and a remainder term. After estimating equation (3) once on the whole sample to obtain a single, common set of worker effects $\hat{\alpha}_i$ and firm effects $\hat{\theta}_{f(i,t)}$ under a connected set, I construct, for each observation, the identity $y_{it} = \hat{\alpha}_i + \hat{\theta}_{f(i,t)} + (X_{it}\hat{\beta} + \hat{\varepsilon}_{it})$. I then run the same RD regression as specified in equation (1) separately on y_{it}, on $\hat{\alpha}_i$, on $\hat{\theta}_{f(i,t)}$, and on the remainder $Other := y_{it} - \hat{\alpha}_i - \hat{\theta}_{f(i,t)}$ (which aggregates $X_{it}\hat{\beta}$ and the residual). By linearity of OLS and because all components are analyzed on the same sample with the same right-hand side, the estimated jump in wages satisfies

$$\Delta y \ = \ \Delta\overline{\alpha} \ + \ \Delta\overline{\theta} \ + \ \Delta\overline{(X\hat{\beta})} \ + \ \Delta\overline{\varepsilon} \tag{5}$$

up to rounding. In this decomposition, $\Delta\overline{\alpha}$ captures changes in the average worker effect, $\Delta\overline{\theta}$ captures changes in the average firm pay premium (sorting to higher- or lower-paying firms), and the "Other" term captures changes in observables and unexplained residual variation. I report results for the whole sample and by gender (using the same $\hat{\alpha}_i$ and $\hat{\theta}_{f(i,t)}$ to preserve additivity) and present each component's share of the total wage jump (Δcomponent$/\Delta y$).

Similarly, decomposing AKM effects estimated in the specification outlined in equation (4) satisfies:

$$\Delta y \ = \ \Delta\overline{\lambda} \ + \ \Delta\overline{\phi} \ + \ \Delta\overline{(X\hat{\beta})} \ + \ \Delta\overline{\varepsilon} \tag{6}$$

and so the remainder becomes $Other := y_{it} - \hat{\lambda}_{O(i,t)} - \hat{\phi}_{i\times F}$.

5 Results

5.1 Education

Table 1 presents the main results on educational attainment using LFS and QP data. The Bologna reform increased overall university graduation rates by 1.8 percentage points in the LFS sample and 1.5 percentage points in the QP sample. As expected, the baseline share of people completing a university degree is lower in the LFS sample (18% vs 23.6%). In both datasets, the increase is stronger for the female sample. The second outcome shows the effects of the reform on the probability of having completed a first-cycle university degree (bachelor's) as highest education. In the LFS sample, this increase outweighs the overall increase in university attainment (2.3pp vs 1.8pp), indicating that not only the reform operated almost entirely through the extensive margin, but also that the share of people completing a master's degree declined and thus that the reform also operated on the intensive margin. In the QP sample, this pattern can be observed for the male sample only - a 1.4pp increase in bachelor's attainment compared to a 1pp increase in overall university attainment. Females in QP data instead show a 1.2pp increase in bachelor's attainment compared to a 2pp increase in overall university attainment. The last outcome of Table 1 shows the effects on the age in which individuals report finishing their highest education (LFS data) or the age in which individuals report starting to work (QP) as a proxy for years of schooling. In both cases, we see a positive net effect on this years of schooling proxy, indicating that the reform likely increased overall educational attainment in the population. Overall, both datasets highlight a strong extensive margin effect of the reform, as well as heterogeneity by gender, suggesting that overall schooling in the population increased with the reform. The QP data, however, seems to show stronger effects on master's completion for females, highlighting a possible over-representation of female master graduates compared to the whole population.

Field of study reallocation Figure 2 document substantial reallocation across fields of study. The most striking pattern is the shift away from healthcare programs (excluding medicine), which experienced a 1.5 percentage point decline, representing a 31% decrease from the baseline of 4.8%. Notably, healthcare (excluding medicine) was the only field where Bologna actually lengthened the degree requirement—nursing and other health programs that previously required only 3 years were extended to 4 years to meet European standards for professional qualifications. This reversal of the general reform pattern likely explains the sharp decline: while other fields became more accessible through shorter degrees, healthcare became less attractive due to the increased time commitment. This decline in healthcare is matched by an increase of the same magnitude in STEM and Economics fields. STEM fields saw a 1.2 percentage point increase (22% relative to baseline). Economics saw a 0.3 percentage point increase (10% relative to baseline). Figure 8 in the Appendix shows that both the uptake of STEM and Economics degrees and the decline in Healthcare degrees is again driven by the female sample.

Table 1: Education outcomes: LFS vs QP

	Probability that highest education is:				Age when:	
	Any university		First-cycle uni		Finish Educ	Start Work
	LFS	QP	LFS	QP	LFS	QP
			Whole sample			
RD Jump	0.018***	0.015***	0.023***	0.013***	0.254***	0.327***
	(0.003)	(0.001)	(0.003)	(0.001)	(0.032)	(0.005)
Baseline	0.180	0.236	0.123	0.175	18.066	21.997
			Female			
RD Jump	0.027***	0.020***	0.032***	0.012***	0.244***	0.286***
	(0.005)	(0.001)	(0.004)	(0.001)	(0.044)	(0.007)
Baseline	0.226	0.295	0.157	0.230	18.395	22.067
			Male			
RD Jump	0.008**	0.010***	0.015***	0.014***	0.264***	0.363***
	(0.004)	(0.001)	(0.003)	(0.001)	(0.045)	(0.008)
Baseline	0.135	0.183	0.089	0.126	17.741	21.933
p-value Female = Male	0.002	0.000	0.003	0.057	0.747	0.000

* p<0.10, ** p<0.05, *** p<0.01. Standard errors in parentheses.

Notes: Entries show donut RD jumps at the 2007 cohort (excluding 2006), with linear trends in the running variable fitted separately pre/post. Controls include age polynomials. "Baseline" is the mean in the cohort immediately before the reform. LFS = Labour Force Survey; QP = Quadros de Pessoal. The first dependent variable is the probability of an individual's highest education being a university degree. The second one is the probability of an individual's highest education being a first-cycle degree. The third outcome is **age at completion of highest education** for LFS and **age started working** in the private sector for QP.

Table 4 in the Appendix shows the comparison between LFS and QP samples. To maintain comparability across the samples due to different classifications of fields of study, I aggregated in the QP data: (i) economics with social sciences and law, and (ii) healthcare with medicine. We see very similar patterns in the composition of the fields of study using LFS and QP data, i.e. an increase in STEM and economics and

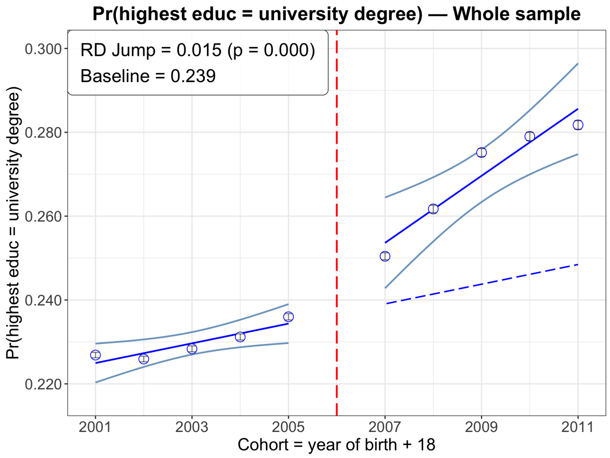

(a) Pr(highest education = university)

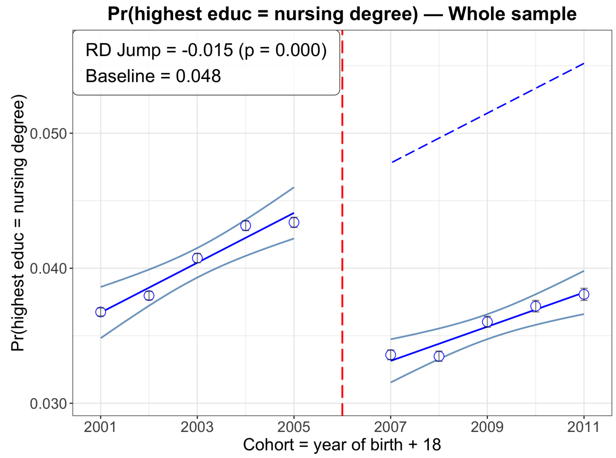

(b) Pr(highest education = university - Healthcare)

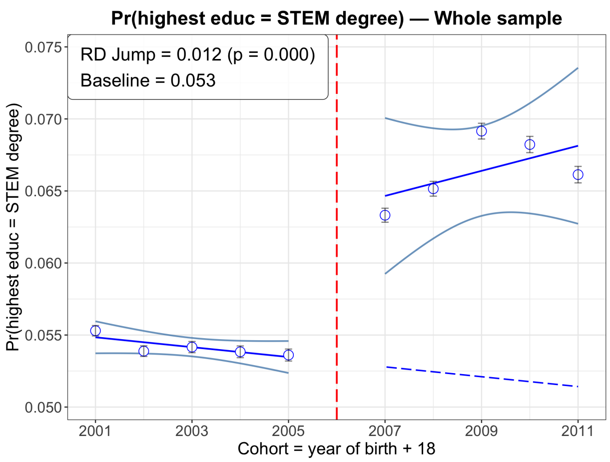

(c) Pr(highest education = university - STEM)

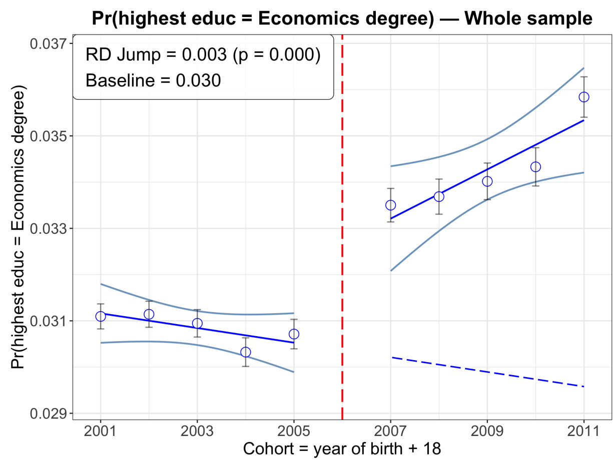

(d) Pr(highest education = university - Economics)

Figure 2: RD graphs on highest education and field of study

social sciences as well as a decline in healthcare; in both datasets, the effects are stronger for the female sample.

Absorption of new graduates A critical question is whether the labour market successfully absorbed the excess supply of university graduates. For this analysis, I deliberately focus only on university graduates rather than the full population. This restriction is important because I want to capture the effect of the reform, including the consequences of increased competition among graduates. By examining only those with university degrees, I can assess whether the expansion in graduate supply—driven by the reform's extensive margin—led to deterioration in job match quality or occupational downgrading. This speaks directly to concerns that making university more accessible might devalue degrees through oversupply. Figure 3 presents

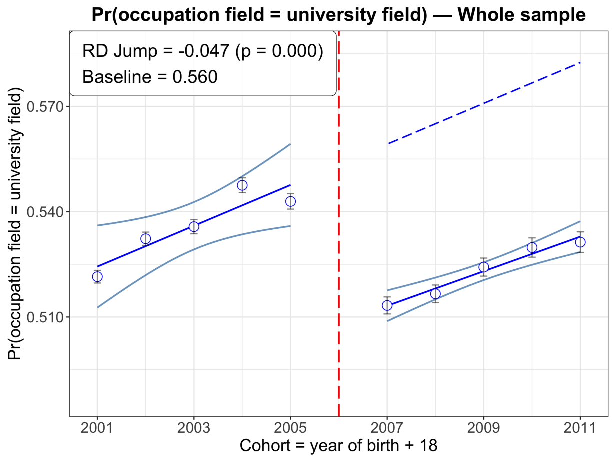

(a) Pr(occupation=field | university)

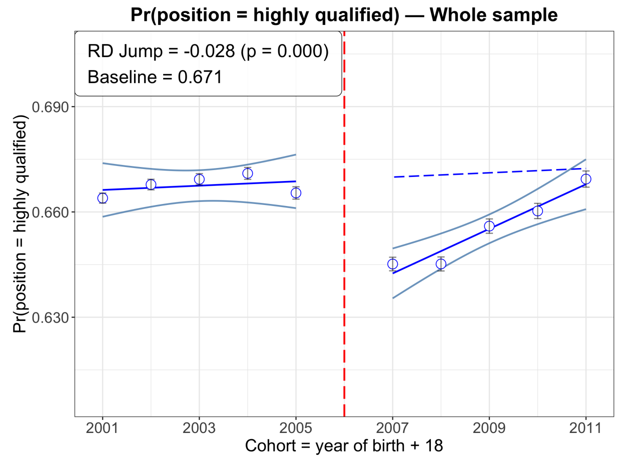

(b) Pr(position = highly qualified | university)

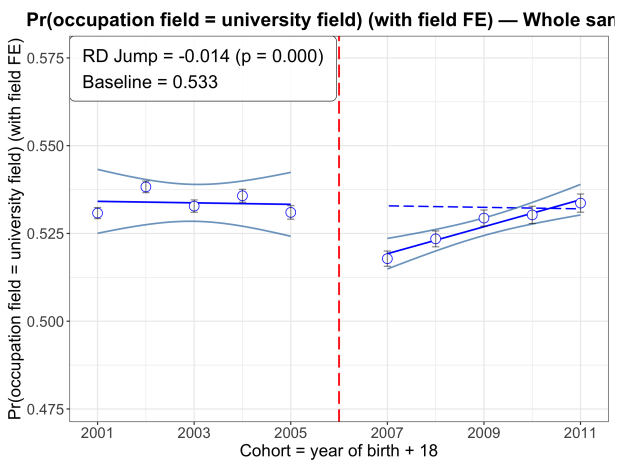

(c) Pr(occupation=field | university, **field**)

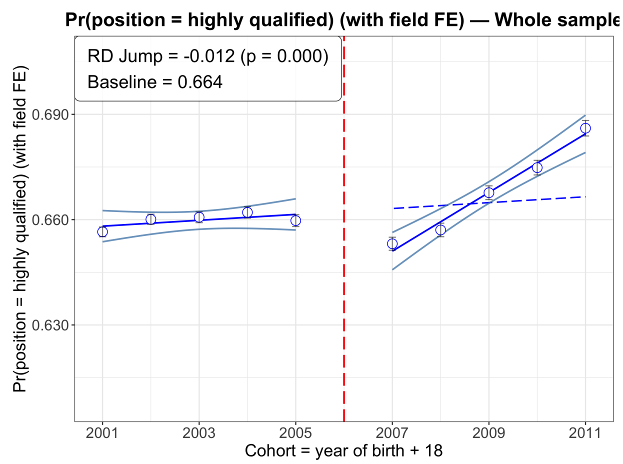

(d) Pr(position = highly qualified | university, **field**)

Figure 3: RD graphs on match rates and positions within the firm

two measures of job match quality for university graduates: the probability of working in an occupation that reflects the field studied at university (occupation-field match rate) and the probability of securing a highly qualified position. Without controlling for field composition, both measures show concerning declines. The occupation-field match rate fell by 4.7 percentage points from a baseline of 56% (Figure 4a). Similarly, the share of graduates in highly qualified positions declined by 2.8 percentage points from a baseline of 67.1% (Figure 4b). These raw effects suggest the labour market struggled to maintain job quality as the supply of graduates increased. However, these aggregate patterns are largely explained by the compositional shift in fields of study. The match rate decline shrinks to just 1.4 percentage points (Figure 4c), while the highly qualified position effect becomes 1.2 percentage points (Figure 4d). The negative effects appear concentrated in the first two post-reform cohorts (2007-2008) only, suggesting a temporary adjustment period as the labour market absorbed the initial wave of restructured degrees. The reduction in these negative effects after controlling for field composition reflects two mechanisms. First, healthcare fields exhibit exceptionally high occupation-field match rates by nature—nurses become nurses, not economists. The shift away from healthcare mechanically reduces average match rates even if graduates in other fields maintain their typical matching patterns. Second, the expansion in fields like social sciences and business, which have inherently more diffuse occupational pathways, naturally produces lower match rates without necessarily indicating worse job quality or misallocation. Indeed, the temporary nature of the highly qualified position effect—evident in the convergence back toward baseline by 2009-2010—suggests that employers quickly adapted to the new credentials. The small initial skepticism about three-year degrees appears to have dissipated as the first cohorts proved their productivity in the labour market, consistent with a transitional rather than permanent devaluation of university credentials.

5.2 Cumulative Earnings

The challenge of capturing full welfare effects is that the reform operates on both the extensive and the intensive margin of university attendance. On the extensive margin, more students pursuing university delays labour-market entry (reducing early cumulative earnings through forgone wages) but potentially increases later earnings through higher human capital. On the intensive margin, students exiting after shorter degrees gain additional years of work experience by any given age (boosting cumulative earnings through more years in the labour force) but may sacrifice human capital from foregone education (potentially reducing their wage levels). Cumulative earnings by a given age captures both these timing effects (when people enter the labour market) and human capital effects (what they earn once there), providing a net welfare measure that integrates all mechanisms through which the reform operates. Rather than attempting to decompose and separately identify each channel through which the reform operates, cumulative earnings measures the net effect: the total income generated by treated cohorts from labour market entry through a given age, compared to control cohorts over the same lifecycle period. This approach naturally incorporates both the costs and benefits of the reform. Years spent in additional education appear as foregone earnings (zeros or low values during schooling), while subsequent returns to human capital are captured through the full earnings trajectory post-graduation. Figure 4 presents the RD estimates for log cumulative earnings at ages 25 and 31,

(a) Cumulative earnings by age 25

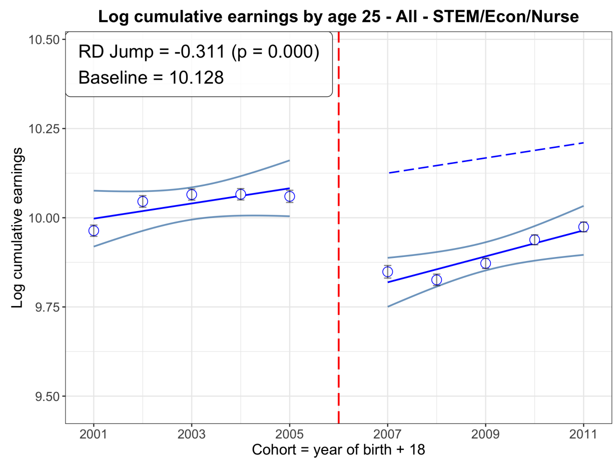

(b) Cumulative earnings by age 25 - STEM/Econ/Nurse

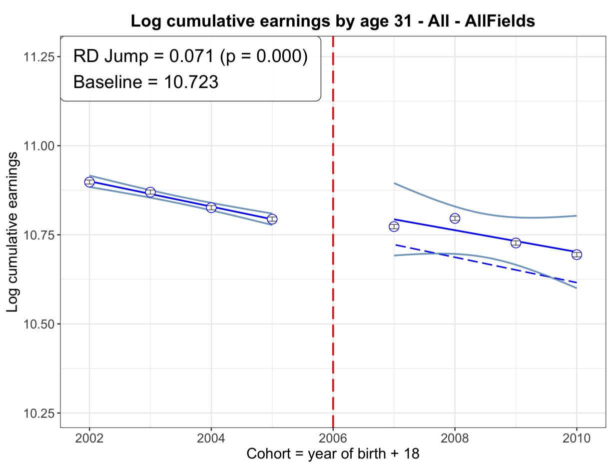

(c) Cumulative earnings by age 31

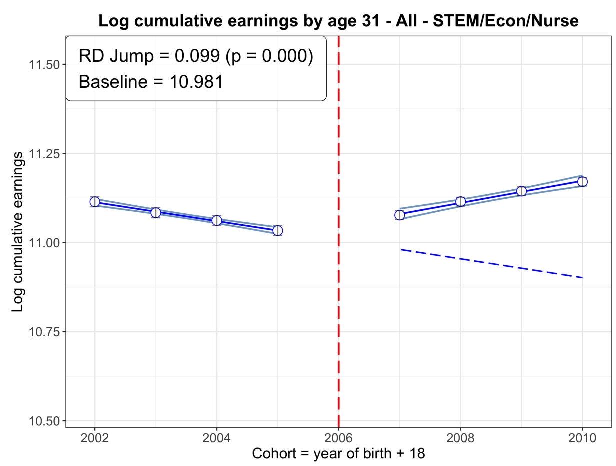

(d) Cumulative earnings by age 31 - STEM/Econ/Nurse

Figure 4: RD graphs on cumulative earnings by age 25 and 31, constructed as outlined in Section A.2.

while Table 5 in the Appendix includes the RD estimates for all ages between 25 and 31, by field and gender. The results reveal a striking intertemporal pattern consistent with human capital investment dynamics. At age 25, treated cohorts show substantially lower cumulative earnings, with a 30.1% reduction for all fields and an even larger 31.1% reduction for individuals who completed STEM/Economics/Nursing fields, i.e.

where we saw the starkest switch in field of study. These negative effects reflect the immediate costs of the reform: more individuals pursuing higher education rather than entering the labour market, resulting in foregone earnings during their early twenties. By age 31, however, this pattern completely reverses. Treated cohorts now show 7.1% higher cumulative earnings across all fields, with STEM/Economics/Nursing fields exhibiting even stronger gains at 9.9%. This transformation from negative to positive effects demonstrates that the reform's human capital gains eventually dominate the initial opportunity costs. Table 5 in the Appendix includes the RD estimates for all ages between 25 and 31, by field and gender. For females, STEM/Economics/Nursing fields show substantially larger initial penalties than the overall female sample (-37.4% vs -28.6% at age 25), followed by stronger recovery (6.4% vs 4.8% at age 31). This pattern aligns with the previous findings on females having a strong switch from healthcare towards STEM and economics programs, which require greater upfront educational investment but yielding higher returns later on.

5.3 Earnings, firm and worker quality

Decomposing wage effects The aim of the current step of the analysis is to quantify how much of the effect on hourly wages operates through workers and firms. Table 6 decomposes overall effects on hourly wages into worker FE and firm FE estimated through equation (3), as explained in the decomposition equation (5). First, however, note that hourly wages increase by 3.8% for the whole sample. A simple back-of-the-envelope calculation merging education and wage results provides useful context: dividing the 3.8% wage increase by the 0.327-year increase in schooling (proxied by delayed labour market entry) in Table 1 yields an implied return of 11.6% per year of education; alternatively, dividing by the 0.254 coefficient from the LFS proxy for years of schooling implies a return of 15%. Both estimates are broadly consistent with existing literature on the returns to schooling. For instance, Portugal et al. (2024) find a 9.8% return per year of education using changes in compulsory schooling. However, this calculation assumes the entire wage effect operates through increased schooling, which may not hold if returns to education changed or if there are spillover effects.

Turning now to the decomposition exercise, I find that 84% of the effect of the reform on wages is an individual component, indicating that much of the effect of the reform operates through an improvement in worker quality, whilst 8% of this effect operates through the allocation of workers to firms. Heterogeneity analysis by gender reveals, firstly, greater wage gains for females, as well as differences in the decomposition exercise: for females, 81% of the effect operates through worker quality and 12% through the firm; for males, 91% is through the worker and only 3% operates through the firm channel.

Following Portugal et al. (2024)'s full specification, adding occupation and match effects to the AKM model (as outlined in equation (4)), I find that 66% (compared to 60% in Portugal et al. (2024)) of the wage effect is an individual component purged of sorting into firms and occupations, and 24% (compared to 30% in Portugal et al. (2024)) of the effect operates via the allocation of workers to firms. Occupation and the interaction effect between worker and firm do not play any role in explaining the effect of the reform on wages.

Early career To understand the mechanisms driving these wage differences, Figure 9a examines wage growth during the critical first three years after labour market entry for all workers. Women in treated cohorts experienced 0.7 percentage points higher annual wage growth, while men showed no significant change in wage dynamics. Figure 9b tracks changes in firm pay premiums—estimated through equation (3)

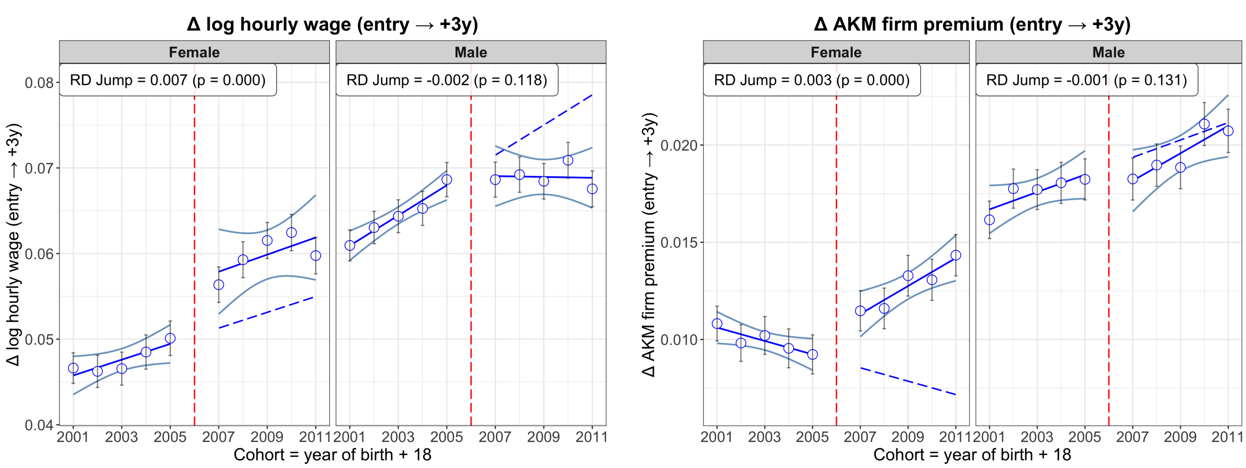

(a) Δ Log(relative hourly wages)

(b) Δ AKM firm FE

Figure 5: Early career changes in wages and firm premium. Firm FE as estimates as outlined in equation (3).

—from labour market entry to three years later. Women in treated cohorts experience a 0.3 percentage point larger increase in their firm premium when changing jobs, indicating they successfully sort into higher-paying firms. Men show no such improvement in firm sorting. This gender difference in strategic job mobility could reflect several factors. Figure 9 in the Appendix reveals that the increase in wage growth and firm premium that females experience post-reform is stronger for females who studied STEM, economics and nursing, suggesting that it is precisely the field of study switch that is driving the faster growth results.

5.4 Lee Bounds on returns to university

The population-level wage gains documented above cannot directly identify how the reform affected returns to university education. The 7% increase in graduation rates means that post-reform university graduates are a different, larger group than pre-reform graduates, creating a fundamental selection problem. The marginal students—those induced to attend university only because of the shorter, more accessible degrees—may differ systematically from always-takers in both observable and unobservable characteristics that affect wages. To bound the effect on returns to education, I employ the trimming procedure developed by Lee (2009), which provides sharp bounds on treatment effects in the presence of sample selection. The key insight is that while we cannot identify which specific individuals are marginal students, we know their exact proportion: the reform increased graduation rates by approximately 6% of post-reform graduates (1.5/25.4) would not have attended university absent the reform. The Lee bounds approach makes no assumptions about the selection process beyond monotonicity—that the reform only induced students into university, never out of it. Instead, it considers two extreme scenarios that bound the true effect:

- **Lower Bound:** Assumes marginal students are positively selected—they would have earned the highest wages among university graduates had they attended pre-reform. To construct this bound, I trim the top 6% of the post-reform graduate wage distribution and compare the remaining distribution to pre-reform graduates. This provides the most conservative estimate, as it assumes the reform attracted exceptionally high-ability students who happened to be constrained by the longer degrees.
- **Upper Bound:** Assumes marginal students are negatively selected—they would have earned the lowest wages among graduates. I trim the bottom 6% of the post-reform distribution, providing the most optimistic estimate under the assumption that the reform primarily attracted lower-ability students previously excluded by the time and financial costs.

The true effect on incumbent students—those who would have attended university regardless—must lie within this range. If even the lower bound remains close to zero or positive, it would indicate that returns to university education were maintained despite the expansion in supply and compression in study time.

Figure 6 presents the Lee bounds on returns to university education across the wage distribution, tracking graduates from age 24 through 30. The results reveal a dynamic pattern that varies across the distribution. Across all quartiles, the actual wage effect for university graduates is substantially negative in the immediate post-graduation period, ranging from -8% at the lower quartile to -15% at the upper quartile at age 24. Even the upper bounds remain negative or near zero at these early ages, indicating that this initial wage penalty is not driven by selection. This suggests that employers were initially skeptical of the new three-year degrees, possibly viewing them as providing insufficient training compared to the traditional longer programs. The wage trajectories show marked improvement with labour market experience. By age 30, the

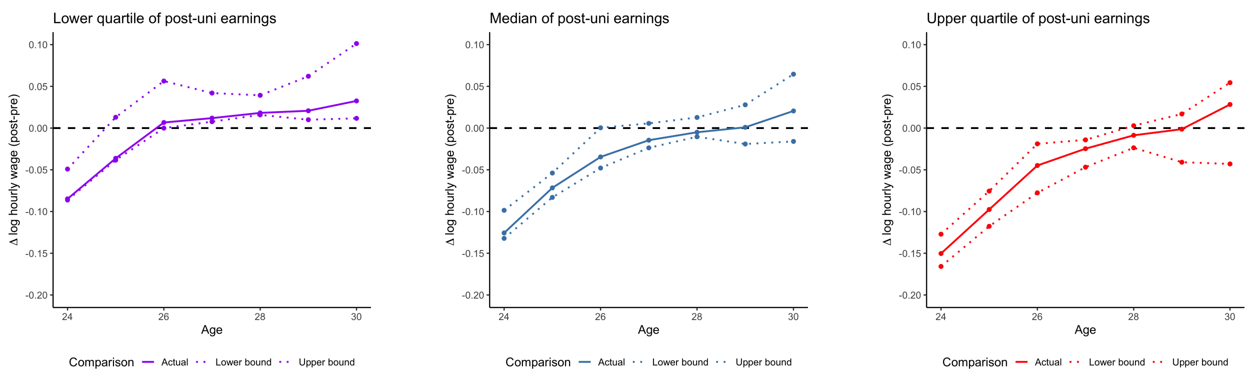

Figure 6: Lee Bounds on returns to university

actual effects converge to approximately 2-3% across all quartiles of the distribution. For graduates in the lower quartile, returns become unambiguously positive by age 28-29, with even the lower bound exceeding zero. By age 30, these workers experience 2-3% higher wages, with tight bounds suggesting minimal selection concerns. Notably, the lower bounds at age 30 remain close to zero or slightly positive across the distribution, while upper bounds reach 5-10%. This indicates that even under the most conservative assumptions about selection, the reform did not permanently reduce returns to university education.

The recovery is remarkably uniform across the wage distribution. While higher earners (upper quartile) experience the deepest initial penalty, they also show the steepest recovery trajectory. By age 30, the bounds are relatively tight across all quartiles, suggesting that the true effect is close to the actual estimate of a small positive return. These patterns are consistent with a market learning story: employers initially discounted the new credentials but revised their beliefs upward as Bologna graduates demonstrated their productivity through actual work performance. The fact that returns fully recover and even turn positive by age 30 suggests that the shorter degrees ultimately maintained their value despite the compression in study time and the expansion in graduate supply. This finding is in line with previous results showing that the ratio of wages to years of schooling remains unchanged with the reform.

6 Conclusion

This paper provides the first comprehensive evidence on how students respond when university degrees become shorter, exploiting Portugal's implementation of the Bologna Process as a natural experiment. The reform's effects were theoretically ambiguous ex-ante, with competing forces potentially pushing educational attainment and earnings in opposite directions. My findings resolve this ambiguity and reveal several important patterns. First, the extensive margin dominates: the reform increased university graduation rates by 1.5 percentage points and raised average years of schooling in the population. This suggests that degree length represents a more important barrier to university access than previously recognized. Second, shortening degrees induced substantial reallocation across fields. Students shifted away from healthcare—where Bologna actually lengthened degrees to meet European standards—toward STEM and Economics fields where the new three-year structure created previously unavailable options. This reallocation proves crucial for understanding labor market impacts, as earnings gains and wage growth concentrate among those who completed fields most affected by the reform. Third, while the reform initially reduced cumulative earnings through age 25 due to delayed labor market entry, human capital gains eventually dominate: by age 31, treated cohorts have higher cumulative earnings. This reversal is driven by steeper early-career wage growth and increased sorting to higher-paying firms. Decomposing the wage effects reveals that 66% of the gains stem from improved worker quality—the reform attracted and developed higher-quality workers—while 24% comes from better firm sorting.

These effects are particularly pronounced for women, who leveraged the reform's flexibility to both shift fields and pursue more master's degrees. This gender heterogeneity suggests that degree compression may help reduce educational and labor market gaps when traditional longer programs pose differential constraints. The findings have direct policy relevance as countries worldwide debate optimal degree structures. Shorter first-cycle degrees can expand access with positive effects on earnings, on average—achieving the "less is more" outcome that reformers intended. However, success appears to depend on maintaining flexibility for continued study and ensuring labor market recognition of the new credentials. Future research using administrative education data will examine the socioeconomic backgrounds of marginal graduates and decompose enrollment versus completion effects. Understanding who benefits most from shorter degree options will further inform the design of higher education systems.

Appendix

A Constructing Relative Wages and Cumulative Earnings

A.1 Relative Wages

To address the concerns on age and calendar time changing across birth cohorts, I construct relative wages that net out year-specific effects. I use the 1987-1999 cohorts as a comparison group, as these individuals were already 26-38 years old when Bologna was implemented and thus unaffected by the reform. For each individual in my cohorts of interest (2000-2012), I compute their wage relative to the average wage of older individuals observed in the same calendar year and of the same gender, and hence subject to the same macro-level shocks. For example, consider measuring wages at age 30:

- An individual from the 2000 cohort reaches age 30 in 2012
- Individuals from the 1987 cohort are age 43 in that same year (2012)
- I express the 30-year-old's wage as a ratio to the average wage of 43-year-old workers observed in the same calendar year and of the same gender

This is graphically shown in Figure 7.

This approach effectively controls for any gender and year-specific shocks that affect all workers regardless of their Bologna exposure, including business cycle effects, general wage trends, or other contemporaneous policy changes. The identifying assumption is that year-specific shocks affect different age groups proportionally within the same calendar year and gender, a reasonable assumption given that I compare prime-age workers (30-year-olds with 43-year-olds in the example above).

A.2 Cumulative earnings

For each individual i and age a, let y_{ia} denote real annual labour earnings (deflated to a common base year); we set $y_{ia} = 0$ when the person has no observed earnings (e.g., studying). We first pad the panel to one row per (i, a) over a pre-specified age grid (e.g., from school-leaving age up to $A \in \{25, 27, 30\}$) so that years

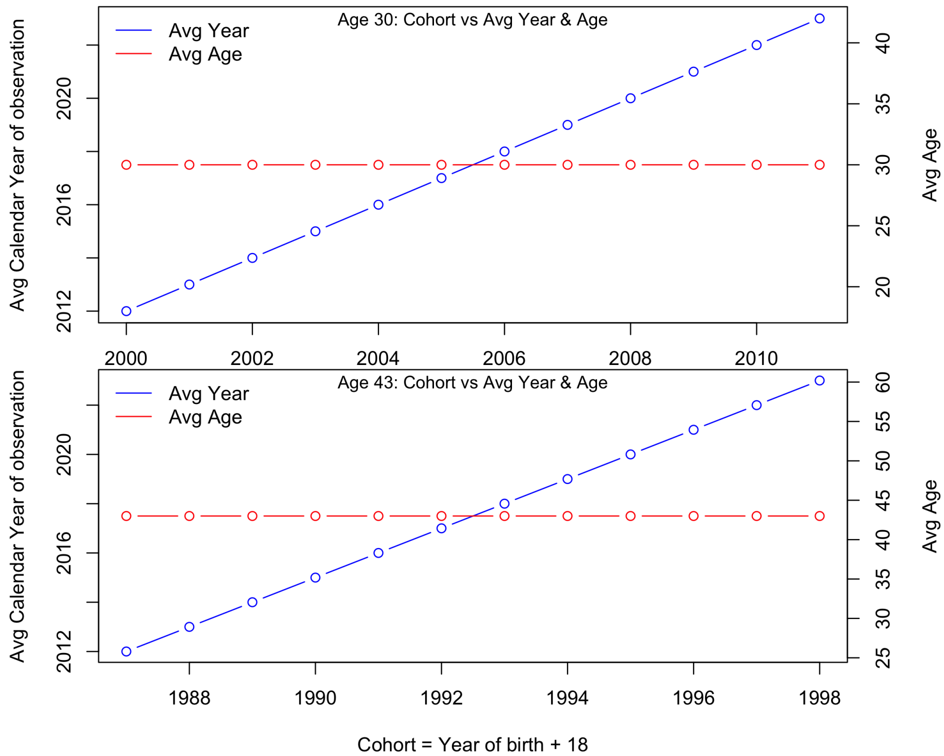

Figure 7: Explaining the construction of relative wages for individuals aged 30

outside employment contribute zeros rather than missing values. We then compute cumulative earnings in levels without discounting as

$$CE_i(A) = \sum_{a=a_0}^{A} y_{ia}$$

which aggregates all cash flows earned up to age A. To obtain a logarithmic outcome that is defined at zero and less sensitive to extreme values, we use the transformation

$$LCE_i(A) = \log\bigl(1 + CE_i(A)\bigr).$$

This construction treats forgone earnings while out of the labour market as true zeros (capturing delayed entry on the extensive margin) and additional earnings from earlier entry as positive contributions (capturing intensive-margin shifts), while avoiding post-treatment conditioning on employment or experience.

Table 2: Descriptive Statistics on QP data

	Control			Treatment		
	N = 8,074,453			N = 4,574,038		
VARIABLE	mean	min	max	mean	min	max
Age	30	17	41	27	17	34
Year	2014	2000	2023	2018	2006	2023
Female	0.48	0.00	1.00	0.47	0.00	1.00
University (any)	0.25	0.00	1.00	0.26	0.00	1.00
University (first-cycle)	0.20	0.00	1.00	0.18	0.00	1.00
Hourly Wage	5.13	0.07	1036	5.93	0.13	1790
Age at first Observation	22.66	17.00	41.00	22.19	17.00	41.00

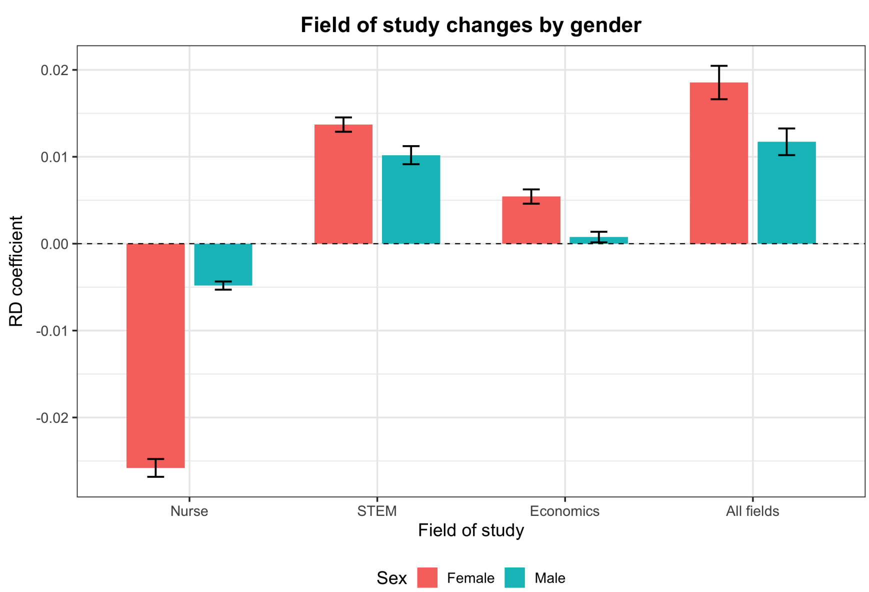

Figure 8: Field of study results using QP data

Table 3: Testing the identification assumption on the reform affecting the decision of when to enroll in university

	Pr(employed between graduating high-school and graduating university)
	Whole sample
RD Jump	0.001
	(0.001)
Baseline	0.105
	Female
RD Jump	0.002
	(0.002)
Baseline	0.128
	Male
RD Jump	-0.000
	(0.002)
Baseline	0.085
p-value Female = Male	0.283

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. Standard errors in parentheses.

Notes: Entries show donut RD jumps at the 2007 cohort (excluding 2006), with linear trends in the running variable fitted separately pre/post. Controls include age polynomials. “Baseline” is the mean in the cohort immediately before the reform.

Table 4: Field of study results: LFS vs QP

	Probability that highest education is a uni degree in:							
	Health		STEM		LawSocSci		Other	
	LFS	QP	LFS	QP	LFS	QP	LFS	QP
	Whole sample							
RD Jump	-0.008***	-0.015***	0.004**	0.012***	0.011***	0.006***	0.011***	0.011***
	(0.001)	(0.000)	(0.002)	(0.000)	(0.002)	(0.000)	(0.002)	(0.000)
Baseline	0.037	0.043	0.038	0.054	0.042	0.056	0.038	0.083
	Female							
RD Jump	-0.014***	-0.025***	0.006***	0.014***	0.016***	0.012***	0.019***	0.019***
	(0.003)	(0.001)	(0.002)	(0.000)	(0.003)	(0.001)	(0.003)	(0.001)
Baseline	0.062	0.073	0.027	0.035	0.060	0.079	0.049	0.109
	Male							
RD Jump	-0.003**	-0.005***	0.001	0.010***	0.007***	0.002***	0.005***	0.003***
	(0.001)	(0.000)	(0.002)	(0.001)	(0.002)	(0.000)	(0.002)	(0.000)
Baseline	0.013	0.016	0.049	0.071	0.025	0.036	0.027	0.059
p-value Female = Male	0.002	0.000	0.077	0.000	0.013	0.000	0.000	0.000

* p<0.10, ** p<0.05, *** p<0.01. Standard errors in parentheses.

Notes: Entries show donut RD jumps at the 2007 cohort (excluding 2006), with linear trends in the running variable fitted separately pre/post. Controls include age polynomials. "Baseline" is the mean in the cohort immediately before the reform. LFS = Labour Force Survey; QP = Quadros de Pessoal. The first dependent variable is the probability of an individual's highest education being a university degree in Healthcare. The second one is the probability of an individual's highest education being a university degree in STEM. The third one is the probability of an individual's highest education being a university degree in law and social sciences. The fourth one is the probability of an individual's highest education being a university degree in other fields.

Table 5: RD estimates by age and subgroup

	log(cumulative annual earnings) by age:						
	25	26	27	28	29	30	31
				Whole sample			
All	-0.301***	-0.276***	-0.270***	-0.178***	-0.076***	-0.012**	0.033***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
STEM/Econ/Nurse	-0.311***	-0.262***	-0.213***	-0.063***	0.024**	0.075***	0.086***
	(0.013)	(0.012)	(0.011)	(0.011)	(0.011)	(0.010)	(0.011)
				Female			
All	-0.286***	-0.263***	-0.261***	-0.167***	-0.062***	0.006	0.048***
	(0.007)	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)	(0.007)
STEM/Econ/Nurse	-0.374***	-0.325***	-0.269***	-0.109***	-0.015	0.043***	0.064***
	(0.016)	(0.016)	(0.015)	(0.014)	(0.014)	(0.014)	(0.014)
				Male			
All	-0.318***	-0.294***	-0.282***	-0.191***	-0.091***	-0.030***	0.019***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
STEM/Econ/Nurse	-0.224***	-0.176***	-0.141***	-0.008	0.069***	0.111***	0.110***
	(0.020)	(0.018)	(0.017)	(0.016)	(0.016)	(0.016)	(0.016)

* p<0.10, ** p<0.05, *** p<0.01. Standard errors in parentheses.

Notes: Entries are RD estimates by age (columns) for two field sets (rows) within each subgroup panel. The first field set includes everyone. The second field set includes individuals who graduated from STEM, economics or nursing degrees at university.

Table 6: Decomposing wage effects into worker and firm AKM fixed effects

	RD Jump				Share of wage effect (%)		
	Log(hwage rel)	Worker FE	Firm FE	Other	Worker FE	Firm FE	Other
				Whole sample			
RD Jump	0.038***	0.032***	0.003***	0.003***	84.2	7.9	7.9
	(0.001)	(0.000)	(0.000)	(0.000)			
				Female sample			
RD Jump	0.042***	0.034***	0.005***	0.003***	81.0	11.9	7.1
	(0.001)	(0.001)	(0.000)	(0.000)			
				Male sample			
RD Jump	0.033***	0.030***	0.001***	0.002***	90.9	3.0	6.1
	(0.001)	(0.001)	(0.000)	(0.000)			
p-value Female = Male	0.000	0.000	0.000	0.319			

* p<0.10, ** p<0.05, *** p<0.01. Standard errors in parentheses.

Notes: Donut RD at the 2007 cohort cutoff (excluding 2006) with linear trends on either side. Worker and firm FE are estimates using the AKM model reported in equation (3), and the decomposition is done as outlined in equation (5). RD jumps are estimates using equation (1). Shares are computed by dividing Worker FE, Firm FE, or the remainder component by the overall increase in log hourly relative wages. Relative hourly wages are constructed as outlined in Section A.1.

Table 7: Decomposing wage effects into worker and firm AKM fixed effects

	RD Jump						Share of wage effect (%)				
	Log(hwage rel)	Worker FE	Firm FE	Occupation FE	Match FE	Other	Worker FE	Firm FE	Occupation FE	Match FE	Other
						Whole sample					
RD Jump	0.038***	0.025***	0.009***	-0.000	-0.000	0.004***	65.8	23.7	0.0	0.0	10.5
	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)					
						Female sample					
RD Jump	0.042***	0.027***	0.011***	-0.000***	-0.000	0.004***	64.3	26.2	0.0	0.0	9.5
	(0.001)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)					
						Male sample					
RD Jump	0.034***	0.021***	0.008***	0.000***	0.000	0.004***	61.8	23.5	0.0	0.0	11.8
	(0.001)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)					
p-value(F=M)	0.000	0.000	0.002	0.251		0.000					

* p<0.10, ** p<0.05, *** p<0.01. Standard errors in parentheses.

Notes: Donut RD at the 2007 cohort cutoff (excluding 2006) with linear trends on either side. Worker and firm FE are estimates using the AKM model reported in equation (4), and the decomposition is done as outlined in equation (6). RD jumps are estimates using equation (1). Shares are computed by dividing Worker FE, Firm FE, or the remainder component by the overall increase in log hourly relative wages. Relative hourly wages are constructed as outlined in Section A.1.

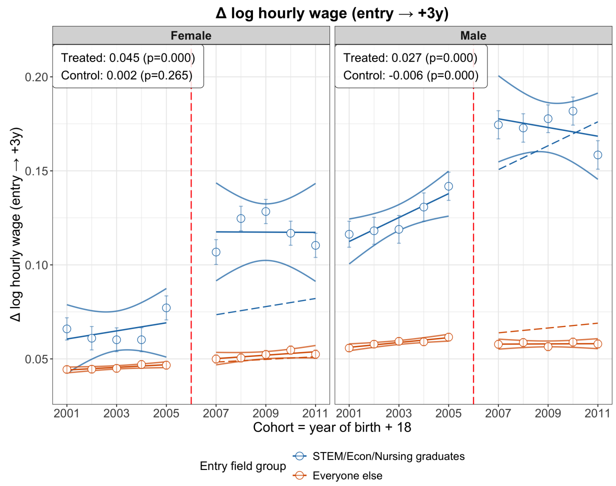

(a) Δ Log(relative hourly wages) by group

(b) Δ AKM firm FE (from equation (3)) by group.

Figure 9: Early career changes, by group

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Acemoglu, D. and Angrist, J. (2000). How large are human-capital externalities? evidence from compulsory schooling laws. *NBER macroeconomics annual*, 15:9–59.

Aghion, P. (2006). A primer on innovation and growth. bruegel policy brief 2006/06, october 2006.

Angrist, J. D. and Krueger, A. B. (1991). Does compulsory school attendance affect schooling and earnings? *The Quarterly Journal of Economics*, 106(4):979–1014.

Argentin, G. and Triventi, M. (2011). Social inequality in higher education and labour market in a period of institutional reforms: Italy, 1992–2007. *Higher education*, 61:309–323.

Bleemer, Z. and Mehta, A. (2022). Will studying economics make you rich? a regression discontinuity analysis of the returns to college major. *American Economic Journal: Applied Economics*, 14(2):1–22.

Bondonio, D. and Berton, F. (2018). The impact of degree duration on higher education participation: Evidence from a large-scale natural experiment. *Oxford Bulletin of Economics and Statistics*, 80(5):905–930.

Brunori, P., Peragine, V., and Serlenga, L. (2013). The bologna process and fairness in university education: Evidence from italy. *CESifo DICE Report*, 11(2):19–22.

Cappellari, L. and Lucifora, C. (2009). The "bologna process" and college enrollment decisions. *Labour economics*, 16(6):638–647.

Cardoso, A. R., Portela, M., Sá, C., and Alexandre, F. (2008). Demand for higher education programs: the impact of the bologna process. *CESifo Economic Studies*, 54(2):229–247.

Di Pietro, G. (2012). The bologna process and widening participation in university education: new evidence from italy. *Empirica*, 39(3):357–374.

Di Pietro, G. and Cutillo, A. (2008). Degree flexibility and university drop-out: The italian experience. *Economics of Education Review*, 27(5):546–555.

Hastings, J. S., Neilson, C. A., and Zimmerman, S. D. (2013). Are some degrees worth more than others? evidence from college admission cutoffs in chile. Technical report, National Bureau of Economic Research.

Horstschräer, J. and Sprietsma, M. (2015). The effects of the introduction of bachelor degrees on college enrollment and dropout rates. *Education Economics*, 23(3):296–317.

Jacobs, B. and Van Der Ploeg, F. (2006). Guide to reform of higher education: a european perspective. *Economic Policy*, 21(47):536–592.

Kirkeboen, L. J., Leuven, E., and Mogstad, M. (2016). Field of study, earnings, and self-selection. *The Quarterly Journal of Economics*, 131(3):1057–1111.

Kroher, M., Leuze, K., Thomsen, S. L., and Trunzer, J. (2021). Did the" bologna process" achieve its goals? 20 years of empirical evidence on student enrolment, study success and labour market outcomes. Technical report, IZA Discussion Papers.

Meghir, C. and Palme, M. (1999). Assessing the effect of schooling on earnings using a social experiment. Technical report, IFS Working Papers.

OECD (2011). Education at a glance 2011. *Editions OECD*, 90.

Oreopoulos, P. (2007). Do dropouts drop out too soon? wealth, health and happiness from compulsory schooling. *Journal of public Economics*, 91(11-12):2213–2229.

Pischke, J.-S. and Von Wachter, T. (2008). Zero returns to compulsory schooling in germany: Evidence and interpretation. *The Review of Economics and Statistics*, 90(3):592–598.

Portugal, P., Reis, H., Guimarães, P., and Cardoso, A. R. (2024). What lies behind the returns to schooling: The role of labor market sorting and worker heterogeneity. *Review of Economics and Statistics*, pages 1–45.

Triventi, M., Vergolini, L., and Zanini, N. (2017). Do individuals with high social background graduate from more rewarding fields of study? changing patterns before and after the 'bologna process'. *Research in Social Stratification and Mobility*, 51:28–40.

Zimmerman, S. D. (2014). The returns to college admission for academically marginal students. *Journal of Labor Economics*, 32(4):711–754.