Ca' Foscari University of Venice

Department of Economics

Working Paper

Ylenia Brilli
Elena Cottini
Paolo Ghinetti
Gloria Moroni

University at Your Doorstep: Local Human Capital Effects of Higher Education Expansion

ISSN: 1827-3580
No. 17/WP/2026

Working Papers
Department of Economics
Ca' Foscari University of Venice
No. 17/WP/2026
ISSN 1827-3580

University at Your Doorstep: Local Human Capital Effects of Higher Education Expansion

Ylenia Brilli
Ca' Foscari University of Venice

Elena Cottini
Catholic University of the Sacred Heart

Paolo Ghinetti
University of Eastern Piedmont

Gloria Moroni
Ca' Foscari University of Venice

Abstract
University systems in many countries expanded by establishing new institutions in areas previously lacking higher education. We study Italy's postwar university expansion, which opened the first faculties in provinces that had never hosted one. Exploiting variation in the timing of these openings across provinces and birth cohorts in an event-study design, we find that local access increased graduation rates by 3.2 percentage points on average across treated cohorts (approximately 26 percent relative to the pre-treatment mean), with a sharp discontinuity of 1.2 percentage points at the enrollment-age threshold. The effect is significant across all urbanization levels and increasing in more urbanized provinces, consistent with complementarities between university access and local labour market conditions. Women benefited disproportionately, though gender gaps in labour market outcomes narrowed by less than those in education. Spillover effects to neighbouring provinces exist but are of secondary magnitude, with the local effect approximately twice the size of the neighbour effect. At the province level, these first openings reduced educational disparities between provinces that gained a university and those that remained unserved.

Keywords
Tertiary education, Higher education expansion, Gender gap, spatial inequality

JEL Codes
I23, J16, J21, J24, H55, R12

Address for correspondence:
Ylenia Brilli
Department of Economics
Ca' Foscari University of Venice
Cannaregio 873, Fondamenta S.Giobbe
30121 Venezia - Italy
e-mail: ylenia.brilli@unive.it

University at Your Doorstep: Local Human Capital Effects of Higher Education Expansion*

Y. Brilli[†] E. Cottini[‡] P. Ghinetti[§] G. Moroni[¶]

May 19, 2026

Abstract

University systems in many countries expanded by establishing new institutions in areas previously lacking higher education. We study Italy's postwar university expansion, which opened the first faculties in provinces that had never hosted one. Exploiting variation in the timing of these openings across provinces and birth cohorts in an event-study design, we find that local access increased graduation rates by 3.2 percentage points on average across treated cohorts (approximately 26 percent relative to the pre-treatment mean), with a sharp discontinuity of 1.2 percentage points at the enrollment-age threshold. The effect is significant across all urbanization levels and increasing in more urbanized provinces, consistent with complementarities between university access and local labour market conditions. Women benefited disproportionately, though gender gaps in labour market outcomes narrowed by less than those in education. Spillover effects to neighbouring provinces exist but are of secondary magnitude, with the local effect approximately twice the size of the neighbour effect. At the province level, these first openings reduced educational disparities between provinces that gained a university and those that remained unserved.

Keywords: Tertiary education, Higher education expansion, Gender gap, spatial inequality.

JEL codes: I23, J16, J21, J24, H55, R12.

*We acknowledge funding from Next Generation EU, in the context of the National Recovery and Resilience Plan, Investment PE8 – Project Age-It: "Ageing Well in an Ageing Society". The views and opinions expressed are only those of the authors and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

[†]Università Ca' Foscari Venezia

[‡]Università Cattolica del S. Cuore

[§]Università del Piemonte Orientale

[¶]Università Ca' Foscari Venezia

1 Introduction

Access to tertiary education is a key determinant of individuals' educational trajectories and long run labour market outcomes (Becker, 1964; Mincer, 1974; Card, 1999). By shaping human capital accumulation early in life, investments in higher education can have persistent effects on labour market attachment well beyond the schooling period (Oreopoulos and Petronijevic, 2013; Carneiro and Heckman, 2002; Goldin and Katz, 2008). Understanding how changes in access to universities affect educational attainment and subsequent labour market participation is therefore central to the design of effective education and labour market policies.

A growing literature emphasizes the crucial role of geographic accessibility in tertiary education choices. The financial and opportunity costs of university attendance increase with distance, making local availability of higher education institutions a major determinant of enrollment and completion decisions (Carneiro et al., 2011; Nybom, 2017). Expansions of university supply, particularly through the establishment of universities in previously unserved areas, can substantially lower these barriers and broaden access to higher education (Ichino et al., 2025), though the effects depend on the specific barriers addressed, whether geographic, financial, or ability-related. While evidence on the effects of such expansions on educational attainment is well established, findings on broader labour market consequences remain more limited, and the spatial mechanisms through which these effects operate are not yet fully understood.[1]

More specifically, there is a small but growing strand of literature focusing on the spatial dimension of university access. Evidence from Norway, Sweden, and the Netherlands shows that geographic distance and commuting costs significantly reduce educational attainment, with effects concentrated among students from lower-income families (Falch et al., 2013; Andersson et al., 2004; Kobus et al., 2015). Similarly, Caner et al. (2024) exploit Turkey's large-scale university expansion and find positive enrollment effects that did not, however, close the gender gap — a result that contrasts with the Italian experience documented in our paper. Our paper combines this spatial perspective with a rich institutional setting, an explicit gender dimension, and a comprehensive analysis of the mechanisms through which proximity operates.

This paper studies the expansion of tertiary education in Italy by focusing on the

[1]See Currie and Moretti (2003) for an early contribution on the role of geographic proximity in college enrollment. Fack and Grenet (2015) and Oppedisano (2011) document enrollment effects of university proximity in France and Italy, respectively.

opening of the first university faculty in provinces that previously lacked any higher education institution. These openings—concentrated from the mid-late 1960s to the first half-decade of 2000s—represent a sharp and localized increase in access to tertiary education and provide a source of quasi-exogenous variation across birth cohorts within the same province. Specifically, our main identification strategy exploits the staggered timing of university openings across provinces in an event study design, comparing outcomes of individuals within the same province for whom the university opened before or after they reached enrollment age, using age 19 as the relevant threshold. Such openings predominantly occurred in relatively peripheral, less urbanized areas that had historically limited access to university education. We exploit this variation to examine how local university openings affect university graduation, employment, and inactivity later in life, with an explicit focus on gender differences and on the geographic heterogeneity of the effect.[2]

We find that local university openings increased graduation rates by 1.2 percentage points at the enrollment-age threshold with an average treatment effect of 3.2 percentage points (approximately 26 percent relative to the pre-treatment mean) across all treated cohorts, reflecting the stronger response of cohorts who gained access at younger ages. The effect is increasing in urbanization, consistent with labour market complementarities amplifying returns to local university access. Women benefited disproportionately from university openings, though the translation into labour market outcomes is asymmetric: exposure is associated with higher employment and lower inactivity, but effects are smaller and less precisely estimated than for educational attainment. At the aggregate level, the expansion contributed to a partial convergence in human capital across Italian provinces, with spillovers to neighbouring provinces present but approximately half the size of the direct effect.

Our results are robust across a comprehensive battery of specification checks and sensitivity analyses. The main estimate of 3.2 percentage points survives intact across alternative estimation windows (from [8,28] down to [14,22]), alternative fixed effect structures, and the exclusion of any single treated province, with leave-one-out estimates ranging narrowly between 2.6 and 3.6 percentage points. To address concerns about treatment effect heterogeneity in staggered designs, we complement the baseline with the heterogeneity-robust estimators of Callaway and Sant'Anna (2021) and de Chaisemartin

[2] In this paper, we use the terms "college," "tertiary," and "university" education interchangeably. We focus on tertiary education broadly defined, including undergraduate and postgraduate university education; tertiary non-university education is very limited in Italy. For identification, we primarily leverage local proximity to university premises following the opening of the first university in a province.

and D'Haultfœuille (2020), both of which yield positive and statistically significant estimates. A cohort permutation test — randomly reassigning treatment timing across 500 draws — produces no estimate as large as the observed ATT, yielding a two-sided pp p-value below 0.002, while wild-cluster bootstrap inference corroborates asymptotic standard errors despite the relatively small number of treated provinces. The sharp discontinuity in the treatment effect at the enrollment-relevant age cutoff of 19 — with significant effects at ages 17 and 18 dropping to near zero at age 19 — provides strong support for a causal interpretation. Finally, pre-trend tests based on 1951–1961 Census data show no evidence that provinces opening universities earlier were already on faster educational trajectories.

We make three main contributions to the literature. First, we provide causal evidence on the effects of local university openings on university graduation, and ITT evidence on their association with subsequent labour market outcomes, focusing on both employment and inactivity. By exploiting the staggered timing of the first university opening at the province level, we isolate the extensive margin of tertiary education expansion and trace its long-term consequences across cohorts exposed to improved local access.

Second, we uncover a pronounced geographic gradient in the treatment effect. The effect is positive and significant across all urbanization terciles, and increasing in more urbanized provinces. This pattern is consistent with a hypothesis of labour market complementarities between university access and local labour market conditions: in more urbanized areas, deeper and more diversified labour markets may amplify the returns to obtaining a degree once a local institution becomes available.

Third, we contribute to the literature on gender differences in higher education and labour market outcomes. Women benefited disproportionately from local university openings, consistent with the hypothesis that geographic barriers to higher education weigh more heavily on women due to social norms and family constraints governing female mobility (Boelman, 2021; Braccioli et al., 2023). Yet the translation of educational gains into labour market outcomes is asymmetric: while exposure to local universities is associated with higher employment and lower inactivity at older ages, the effects are smaller and less precisely estimated than for educational attainment, and gender differences in labour market responses persist.

Beyond the individual-level analysis, we document that the university expansion contributed to a partial convergence in human capital across Italian provinces. Late-expansion provinces substantially narrowed their graduation gap relative to historical university towns, while provinces that never received a university fell further behind. Spillover effects to neighbouring provinces exist but are of secondary magnitude: in a horse-race specification that includes both own-province and neighbour opening timing simultaneously, the local effect is approximately twice the size of the spillover, confirming that geographic access operates primarily through direct proximity. Because our data record province of current residence rather than province of birth, we show that residential mobility introduces attenuation bias and calibrate bounds on the true effect using 2001 Census mobility data, under random misclassification, the corrected estimate rises to 4.3 percentage points; graduate-weighted correction yields 4.6 percentage points.

Kyui (2016), Elsayed and Shirshikova (2023), and Si (2022) are the most closely related studies to ours. Kyui (2016) exploit regional variation in student intake capacities during Russia's post-transition expansion, finding strong positive returns in both employment and wages. Elsayed and Shirshikova (2023) examine the construction of public universities in Egypt, finding that local university openings increased degree attainment and improved labour market and marriage outcomes, particularly for women. Si (2022) show that China's 1999 expansion raised attainment for both genders but eroded women's egalitarian views, driven by worsening relative labour market opportunities for women. While these papers provide compelling evidence from transition and developing economies, our paper offers a complementary perspective from a high-income Western European context, where the institutional environment, labour market structure, and social norms differ substantially. We share with Elsayed and Shirshikova (2023) the focus on the extensive margin of university supply, but contribute a richer analysis of spatial mechanisms, a pronounced and positive gender dimension, and evidence on the asymmetric translation of educational gains into labour market outcomes. Moreover, our empirical results contribute to the broader literature on the spatial allocation of public investment in education and its consequences for regional inequality (Moretti, 2004). The findings are particularly relevant for policy debates in countries with geographically concentrated higher education systems and persistent territorial disparities in human capital.

The remainder of the paper is organized as follows. Section 2 provides institutional background on the Italian university expansion. Section 3 describes the data and presents descriptive evidence. Section 4 introduces the empirical strategy. Section 5 presents the main results, including heterogeneity by gender and urbanization. Section 6 discusses robustness checks. Section 7 examines spatial spillovers and convergence in human capital across provinces. Section 8 concludes.

2 Institutional Background

Italy's higher education system has undergone a dramatic geographic transformation over the second half of the twentieth century. Until 1950, university education was concentrated in a small number of historical seats—major cities such as Rome, Milan, Naples, Bologna, Florence, Turin, Padua, and Palermo—leaving vast portions of the country without local access to tertiary education. For individuals residing in provinces without a university, attending higher education required relocating to a distant city, entailing substantial financial, logistic, and personal costs.

Beginning in the 1950s, new university branches and satellite campuses began to open in medium-sized provincial capitals. This process was driven by a sequence of post-war reforms that progressively expanded and restructured the Italian higher education system. Law 910/1969 liberalized university access for students holding a five-year secondary diploma from technical schools, leading to a sharp increase in demand. Law 766/1973 promoted the opening of new universities and the hiring of academic staff. Law 382/1980 reformed university governance and academic careers. Laws 392/1989 and 245–341/1990 modified funding allocation rules, while Law 59/1997 granted universities greater financial and teaching autonomy. As a result, the creation of new faculties accelerated markedly from the 1970s onward, and by the early 2000s the number of provinces hosting at least one university faculty had more than tripled relative to 1950.

Figure 1 reports the timing of first university openings across provinces. The pattern closely mirrors the sequence of higher education reforms, with clusters of new openings following each major legislative change.

The decision to open a new university branch was not part of a centralized national plan. Rather, it was typically driven by a combination of local political pressure, regional development objectives, and initiatives from existing universities seeking to expand their catchment area. The timing of openings varied considerably across provinces for reasons largely unrelated to local educational outcomes. This variation in timing, after conditioning on province and cohort fixed effects, is plausibly exogenous to individual-level educational outcomes, and forms the basis of our empirical strategy.

Figure 2 displays the geographic distribution of university openings, classifying provinces by the period of first opening. The map reveals that the expansion process does not follow a simple North-to-South pattern. Historical seats are distributed across the entire peninsula, while subsequent waves of expansion occurred both in northern provinces (e.g.,

Figure 1: Number of provinces where the first university opened, by year of opening

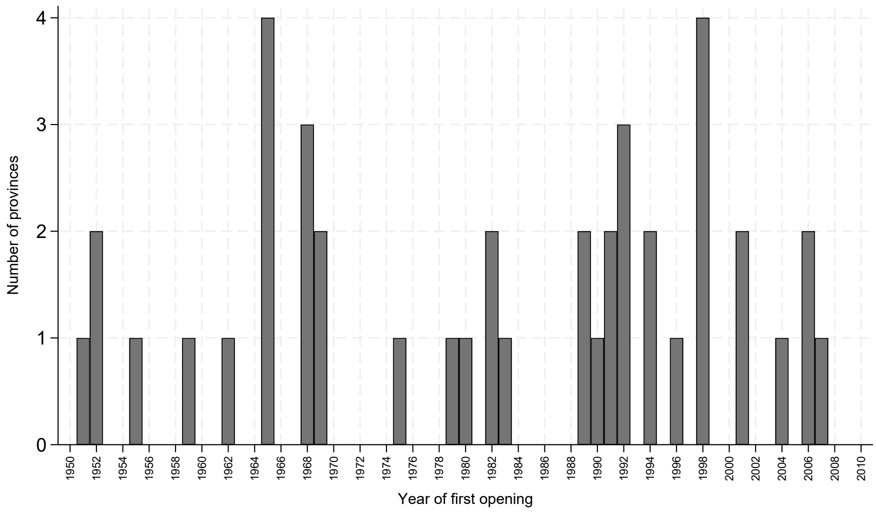

Note: Own calculations using the HIU dataset.

Brescia, Varese, Bolzano, Vercelli) and in southern ones (e.g., Cosenza, Potenza, Campobasso, Foggia). A substantial number of provinces, predominantly small and inland, never received a university branch during our sample period. This geographic dispersion of treatment timing ensures that the variation we exploit is not confounded by single regional trends.

The expansion of university access occurred against the backdrop of a profound transformation in women's educational attainment in Italy. In the 1950s and 1960s, female university enrollment was extremely low, particularly in southern and rural provinces where social norms discouraged women from leaving home to study in distant cities (Braccioli et al., 2023). The cost of geographic mobility for university education was thus disproportionately borne by women, for whom distance represented not only a financial barrier but also a social and cultural barrier. The opening of local university branches potentially reduced these gender-specific barriers to access, a hypothesis we test directly in Section 5.

Over the same period, tertiary education enrollment and attainment increased steadily in Italy. Census data document that the share of individuals aged six and above holding a university degree rose from 1.8 percent in 1971 to 11.2 percent in 2011, representing more than a sixfold increase over four decades. Figure A-1, in the Appendix, provides descriptive evidence on how graduation rates evolved across cohorts for different types of provinces. All province groups exhibit rising graduation rates, reflecting the nationwide

Figure 2: Year of first university opening by province

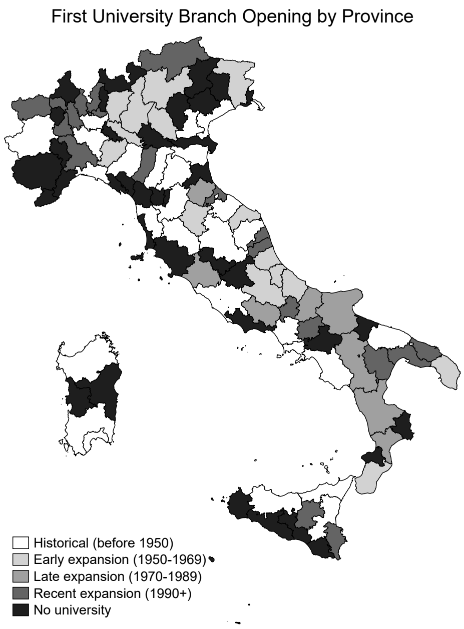

Note: Own elaborations based on the HIU dataset (Cottini et al., 2026). Provinces are classified by the year of first faculty opening. White: historical university seats (before 1950); light grey: early expansion (1950–1969); medium grey: late expansion (1970–1989); dark grey: recent expansion (1990 and later); black: no university by 2010.

trend. However, provinces that received a new university show a steeper increase for younger cohorts, converging toward historical seats, while provinces that never received a university lag behind. This descriptive pattern is consistent with a positive effect of local university presence on graduation rates, though our causal identification strategy—described in Section 4—relies on within-province variation across cohorts rather than on these between-province comparisons.

3 Data and descriptive evidence

University expansion data. We combine two main data sources. The first is the History of Italian Universities (HIU), assembled by Cottini et al. (2026). This database provides a comprehensive registry of university openings in Italy, recording the dates of creation of each faculty within each university from Italian unification in 1861 up to 2010.[3]

[3]The terms university/college and faculty/school may have different meanings across countries and education systems. Throughout the paper, we use college and university interchangeably to indicate a single administrative unit offering degree programs in multiple fields through different faculties or schools (e.g. economics, humanities, engineering). We also use faculty and school interchangeably to

Using the HIU, we construct a province-level panel dataset which, for each year since 1950, records the number of faculties operating in each province and tracks the timing of every new faculty opening.

LFS data on education and labour outcomes. The second data source is the Italian Labour Force Survey (LFS, *Rilevazione trimestrale sulle forze di lavoro*) administered by ISTAT for the period 2009–2019. The LFS provides detailed information on individuals' labour market status, educational attainment, and geographic location (province and region of residence at the time of the survey).

We select individuals aged 35 to 65 observed between 2009 and 2019.[4] Given this age restriction, the youngest cohort observed in the data is that born in 1975. To ensure that individuals are observed for a sufficient number of years during the later stages of their labour market career, we also restrict the sample to cohorts born from 1950 onwards. Under these selection criteria, the baseline dataset is a repeated cross-section comprising approximately 2,151,000 observations, of which 13 percent are university graduates and about 52 percent are women. The average age in the sample is 50.5 years. In terms of labour market status, 64 percent of individuals are employed, while 31 percent are inactive.

A limitation of the LFS is that we observe individuals' province of *current residence* rather than province of birth. If individuals who benefited from a local university subsequently moved to a different province, treatment assignment is measured with error. Using microdata from the 2001 Population Census, we document that 76 percent of individuals aged 25–54 reside in their province of birth, with this share rising to 87 percent in the South. Since this type of measurement error likely attenuates the estimated treatment effect toward zero, our estimates should be interpreted as a lower bound of the true effect. We provide a detailed analysis of this issue, including formal bounds following Manski (1990), in Section 6. Moreover, the most plausible form of selective migration, graduates from peripheral treated provinces relocating to larger cities in search of better employment

denote individual teaching units within a university. Typically, each faculty offers one or more degree programs in a specific field. We focus on institutions offering undergraduate education. Institutions specialized exclusively in postgraduate education or targeting foreign students are excluded, as proximity to a university is likely to matter primarily for undergraduate students, who are less geographically mobile.

[4]We exclude individuals aged 25–34 to focus on the long-term implications of tertiary education for labour market outcomes. In Italy, the average age at university graduation was close to 28 until recently, implying that early labour market experiences are often unstable and less representative of the core working career, especially for graduates.

opportunities, would itself generate attenuation bias, since these individuals are reclassified as controls in the destination province (if it is among the ones used in the analysis) or they exit the estimation sample. This reinforces the lower-bound interpretation of our estimates rather than threatening their validity. The reverse concern—selective in-migration of educated workers into treated provinces generating upward bias—is structurally limited in our setting. Treated provinces are by definition those that previously lacked a university and are predominantly peripheral and less urbanized. In the Italian context, net migration flows run from peripheral areas toward metropolitan centres, which are the historical university seats already excluded from our sample. Selective in-migration toward treated provinces is therefore unlikely to be a quantitatively important source of bias, and the dominant measurement error remains the out-migration of graduates, which attenuates our estimates toward zero.[5]

The combined LFS–HIU dataset. We merge LFS microdata with the HIU panel using the individual's province of residence. This allows us to reconstruct the local history of university faculty openings, including the number of faculties operating in a province in a given year, the year of the first university opening, and the individual's age at first exposure to tertiary education.

Figure A-2 in the Appendix shows the average number of faculties available in the province at age 18 by birth cohort. Individuals born in 1950 had on average 3.5 faculties available; this rises to 4 for the 1960 cohort and 5 for the 1970 cohort.

For each individual, we observe whether a faculty was already present by their 19th birthday, the typical age of tertiary education choice in Italy. Cohorts within the same province differ in their age at the time of the first university opening, generating variation in exposure to local tertiary education. How this variation is exploited for causal identification is discussed in Section 4.

Our analysis focuses on provinces that initially lacked a university and uses the timing of the first faculty opening to define the treatment. We restrict attention to individuals aged 8–28 at the time of the first opening. Individuals younger than 19 at the opening are classified as treated, those 19 or older as controls, resulting in a balanced 8–28 age window across cohorts.[6] The relatively wide window captures the gradual effects of establishing

[5]Differential mobility by gender reinforces this interpretation: women are less geographically mobile than men, implying less misclassification and estimates closer to the true effect. The larger treatment effect for women therefore reflects a genuine difference in educational responses, not a measurement artefact.

[6]For example, the 1970 cohort contributes to the event-study from age 8 (1978) to 28 (1998), with

a university in previously unserved provinces.

We focus on the extensive margin of university expansion: first faculty openings. Additional faculties in provinces that already hosted a university are excluded, as are provinces that never experience a first opening. This ensures clean identification driven by the timing of initial exposure across cohorts within the same province. Overall, 33 provinces (about 30 percent of the total) contribute to the analysis, yielding a final sample of about 283,000 observations.[7]

Descriptive evidence. Table 1 shows summary statistics for the main variables, overall and by gender. The sample is broadly similar to the LFS population. The balanced 8–28 age window ensures equal representation of treated (51 percent) and control individuals.

Table 1: Summary statistics by gender

	All		Men		Women	
	mean	sd	mean	sd	mean	sd
Age	49.81	7.84	49.72	7.83	49.90	7.84
Female	0.52	0.50	0.00	0.00	1.00	0.00
Graduate	0.12	0.33	0.11	0.32	0.13	0.34
Employed	0.62	0.48	0.75	0.43	0.50	0.50
Inactive	0.33	0.47	0.20	0.40	0.46	0.50
Treated	0.51	0.50	0.51	0.50	0.51	0.50
Obs.	282,729		136,350		146,379	

Note: own elaborations on the LFS-HIU dataset. The variable Treated is missing for 12,102 observations (individuals aged 19 at the opening, used as the normalization group in the event-study, see Section 4).

Figure 3 illustrates the graduation probability by cohort and gender. In the peripheral provinces considered, college attainment was initially low (about 10 percent) and similar for men and women. From cohorts born in the mid-1960s, graduation rates rose rapidly, especially among women, who surpassed men by about 2 percentage points.

Table 2 reports means by treatment status and gender. Graduation differences are small. Employment declines with age for all groups and is systematically lower for treated individuals, especially women, while inactivity shows the opposite pattern.

treated ages corresponding to <19. Cohorts 1950–1975 are fully observed over this window, ensuring comparable exposure probabilities.

[7] They are: Alessandria, Ancona, Aosta, Arezzo, Ascoli Piceno, Benevento, Bergamo, Bolzano, Brescia, Campobasso, Catanzaro, Chieti, Como, Cosenza, Foggia, Forlì-Cesena, Frosinone, Isernia, Novara, Pescara, Potenza, Reggio Di Calabria, Reggio Nell'Emilia, Rimini, Siracusa, Taranto, Teramo, Trento, Udine, Varese, Vercelli, Verona, Viterbo.

Figure 3: Graduation probability, by cohort and gender (all cohorts).

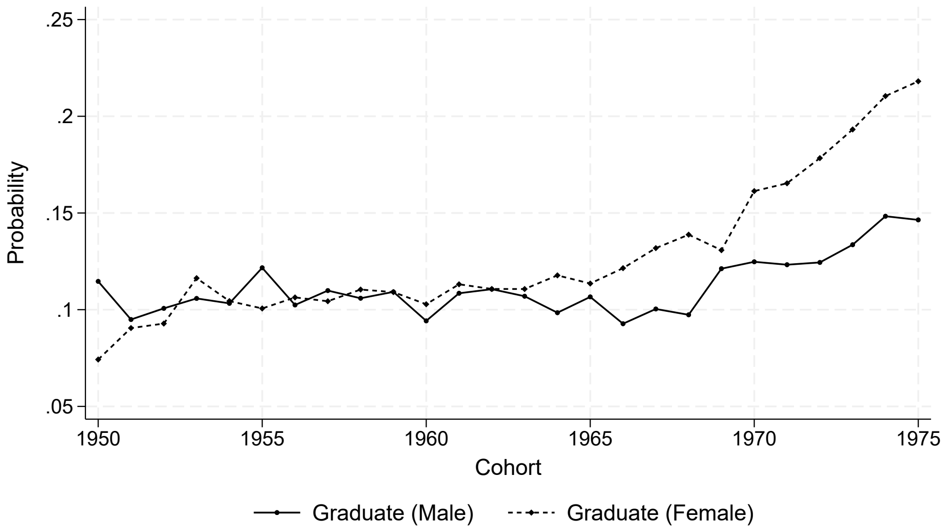

Note: own elaborations on the LFS-HIU dataset.

Table 2: Descriptive statistics: Treated vs Control by age group and gender

	All		Men		Women	
	Treated	Control	Treated	Control	Treated	Control
Graduate (all ages)	0.12	0.13	0.11	0.12	0.13	0.14
Employed (all ages)	0.56	0.68	0.70	0.80	0.43	0.57
Age 35–44	0.67	0.78	0.82	0.88	0.53	0.67
Age 45–54	0.65	0.70	0.81	0.83	0.49	0.59
Age 55–59	0.55	0.57	0.67	0.70	0.44	0.45
Age 60–65	0.31	0.40	0.40	0.50	0.24	0.31
Inactive (all ages)	0.39	0.27	0.25	0.14	0.53	0.39
Age 35–44	0.26	0.17	0.11	0.06	0.41	0.28
Age 45–54	0.30	0.24	0.13	0.11	0.46	0.37
Age 55–59	0.42	0.40	0.29	0.25	0.55	0.53
Age 60–65	0.68	0.58	0.59	0.47	0.75	0.68

Note: own elaborations on the LFS-HIU dataset. Means are reported for all variables.

Figure A-1 in the Appendix shows complementary evidence on employment and inactivity rates by educational attainment and gender. In particular, the employment gender gap behaves differently depending on the level of education as individuals age: it tends to decrease with age for individuals with a high-school degree or lower and reaches its lowest point for the 60–65 age group. Conversely, the gender gap increases with age for college-educated individuals, reaching its maximum in the 60–65 cohort. These patterns reflect heterogeneous labour market behaviors at the end of working life across education groups and are consistent with findings that women—even those with college degrees—may exit the labour market earlier than men due to care responsibilities or other constraints (Brilli and Bassoli, 2025; Ciani, 2016; Maura and Profeta, 2025).

A deeper understanding of these descriptive patterns motivates the econometric analysis discussed in the next section.

4 Empirical strategy

To analyse the effect of higher education expansion on individuals' educational and labour market outcomes, we estimate an event-study model of the form:

$$Y_{icat} = \alpha_a + \eta_c + \mu_{ac} + \beta_t + \sum_{\substack{s=8 \\ s\neq 19}}^{28} \gamma_s \cdot \mathbf{1}(s = \text{Year opening}_a - \text{Year birth}_i) + \boldsymbol{X}_{it}\delta + \epsilon_{icat} \quad (1)$$

Here, Y_{icat} denotes the outcome of individual i at time t, from cohort c, residing in province a, and can be either a binary indicator for university graduation or for employment/inactivity. The model includes province fixed effects α_a, while year fixed effects β_t capture shocks common to all provinces in a given year.[8] Cohort fixed effects η_c, defined over five-year birth intervals, control for cohort-specific shocks. μ_{ac} is a full set of cohort $\times$ province indicator variables (cohort $\times$ province fixed effects), absorbing all cohort-specific differences within each province. $\boldsymbol{X}_{it}$ is a vector of individual characteristics, including gender, age, and age squared, and ϵ_{icat} is the error term; standard errors are clustered at the province level.[9]

The coefficients of interest, γ_s, measure the effect of exposure to the first university

[8] Since the LFS is collected quarterly, we also include quarter fixed effects to absorb within-year seasonality.

[9] age and age squared enter only in the employment and inactivity specifications.

opening at age s, normalized to $\gamma_{19} = 0$, corresponding to the typical age of tertiary education choice. The time window spans ages 8 to 28 relative to the opening, including 11 treated cohorts (ages 8–18 at opening), 9 control cohorts (ages 20–28 at opening), plus the normalization cohort at age 19 (excluded from the sum, $\gamma_{19} \equiv 0$). This structure allows us to examine pre-trends and trace dynamic effects of treatment across the life course. Figure A-3 in the Appendix shows that observations are reasonably balanced across this age window.

In other words, the empirical strategy exploits the fact that university openings occurred at different points in time, generating variation in age at exposure across cohorts within the same province, while we identify effects exclusively from differences in age at exposure to the first university opening within province and birth cohort.[10] The inclusion of cohort $\times$ province fixed effects absorbs all cross-cohort and cross-province heterogeneity, so that the coefficients γ_s are identified by comparing birth cohorts within the same province that differed in their age at the time of the first university opening. As a result, the identifying assumption is that, absent the first faculty opening, outcomes would have evolved smoothly with age within a given cohort $\times$ province cell. Under the parallel trends assumption, any divergence in outcomes after exposure is attributable to the causal effect of local university expansion, net of cohort-specific provincial characteristics. To assess the plausibility of this assumption, we examine the evolution of outcomes in the pre-treatment period for treated and control cohorts, allowing for a visual check of parallel pre-trends. This complementary province-level check is in Table A-2, where we regress the 1951–1961 change in four provincial outcomes on the year of first faculty opening within the 33 treated provinces.[11] A significant coefficient would indicate that provinces already on a faster educational trajectory opened their universities earlier, threatening identification. We find no such evidence for the main outcomes: educational attainment ($p = 0.252$), illiteracy ($p = 0.662$), and female employment ($p = 0.658$) show no pre-trend.[12]

[10]While university openings are staggered over time across provinces, the event-study does not rely on comparing treated provinces to historical university seats. Instead, it exploits within-province variation across birth cohorts in age at exposure to the first local university opening. Older cohorts in the same province were already past the typical enrollment age when the opening occurred, whereas younger cohorts were exposed before making tertiary education decisions.

[11]These balance tests use the ISTAT *8000 Census* historical Census database, which records municipality-level data from the Italian Population Censuses of 1951 and 1961, including educational attainment, illiteracy, employment by sector and gender, population, and surface area. Municipality-level observations are aggregated to the province level using population-weighted averages.

[12]Agricultural employment is marginally significant ($p = 0.065$), though small in magnitude and unrelated to the educational outcomes of interest.

When Y_{icat} is the employment or inactivity indicator, equation (1) estimates intention-to-treat (ITT) effects, capturing the impact of university expansion on labour market outcomes. In this respect, tertiary education can be interpreted as a natural mediator.

To complement the event-study analysis, we also estimate a difference-in-differences (DiD) specification that aggregates the age-specific treatment effects into a single binary treatment indicator, equal to one for individuals younger than 19 at the time of the first faculty opening. While more restrictive, this specification provides a compact summary of the average treatment effect and facilitates comparison across subgroups.

We assess the robustness of our results to alternative estimation windows, different control specifications and estimation strategies. These checks are presented in Section 6.

5 Results

In this section, we present causal estimates of the effects of tertiary education expansion on university graduation, and ITT estimates of its association with labour market outcomes. We first assess the impact on tertiary education completion, then examine heterogeneity by gender and urbanization, and finally turn to labour market outcomes.

5.1 University graduation

Figure 4 presents the estimated event-study coefficients for the probability of obtaining a university degree, by age at which individuals were exposed to the first faculty opening in their province of residence. The pre-treatment coefficients provide a visual check of the parallel trends assumption, while the post-treatment coefficients trace the dynamic effects of exposure across ages. Coefficients prior to exposure are small and statistically indistinguishable from zero, providing no evidence of differential pre-trends and supporting the validity of the identification strategy.

Following exposure, results for the full sample in Panel (a) suggest that the probability of completing tertiary education increases substantially, with statistically significant effects emerging around the typical age of university enrollment. The largest effects—an increase of about 5–6 percentage points in graduation probability—are observed when the first faculty opened while individuals were in the lower part of the age window (8 to 10 years old). This pattern is consistent with the idea that newly opened faculties take time to become fully operational and salient to prospective students, gradually reducing informational and access barriers to higher education. Overall, these results are consistent with the idea that the geographic expansion of universities substantially lowered barriers to tertiary education and translated into higher rates of university graduation for exposed cohorts.

5.2 Heterogeneity by gender

Figure 4, Panel (b) presents the estimated event-study coefficients separately by gender. The increase in university graduation following the opening of new faculties is markedly stronger for women than for men. For women, the effect rises steeply with exposure, with the largest age-specific coefficients approaching 10 percentage points for the youngest female cohorts. For men, the effect is positive but substantially smaller and less precisely estimated throughout the post-treatment window.

This asymmetry is consistent with the hypothesis that geographic barriers to higher education disproportionately affected women. In the Italian context of the 1960s and 1970s, attending a university in a distant city required not only financial resources but also a degree of geographic independence that social norms and family expectations made less accessible to young women than to young men. The opening of a local university branch potentially reduced this gender-specific barrier, consistent with the hypothesis that women who otherwise would not have relocated were able to enroll. Boelman (2021) documents a similar pattern in the context of Germany's university expansion, finding that women's lower geographic mobility is a key driver of gender differences in the response to local university availability. The observed gender gradient is consistent with both geographic mobility constraints rooted in social norms and with higher returns to education for marginal female students in this context; our data do not allow us to distinguish between these two channels.

The DiD estimates in Appendix Table A-3 confirm the gender gap quantitatively: the treatment effect on graduation is 0.037 for women ($p < 0.01$) and 0.027 for men ($p < 0.05$), implying that women's gains were roughly 40 percent larger than men's. Overall, these results indicate that the geographic expansion of university supply contributed substantially to closing the gender gap in higher education observed across cohorts born between 1950 and 1975.

Figure 4: Event-study: Probability of graduating by age of first faculty opened in the province (full sample and separately by gender)

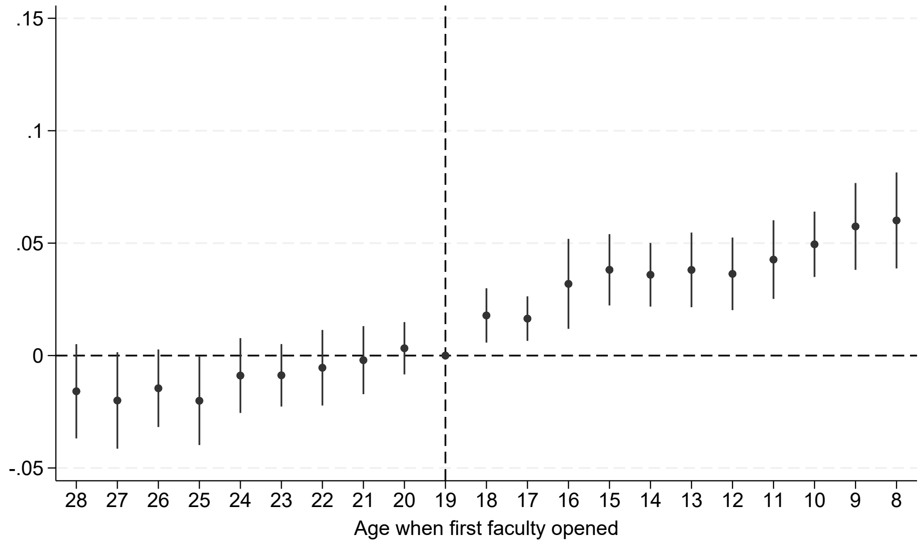

(a) Full sample.

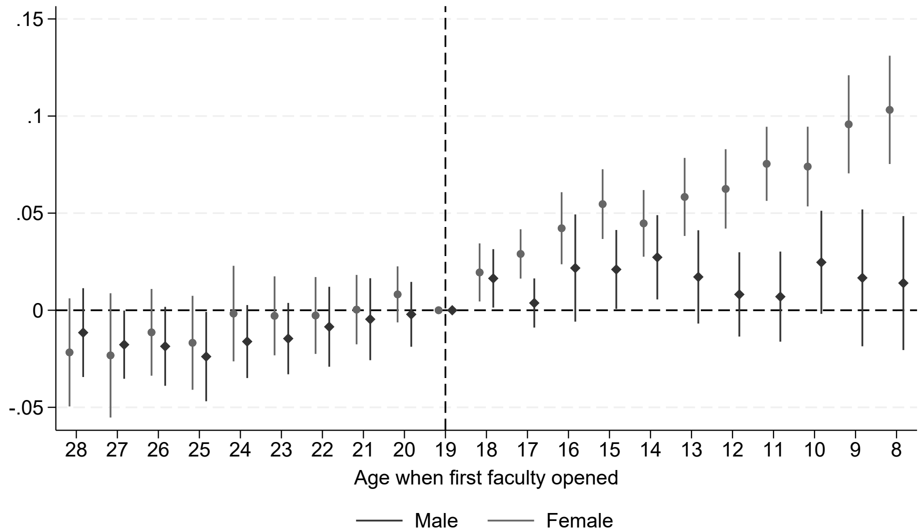

(b) Women and men separately.

Note: Estimated γ_s coefficients (dots) and their 95% confidence intervals (bars) from equation (1), when Y_{icat} is a dummy for graduating. OLS regression, standard errors clustered at the province level.

5.3 Heterogeneity by urbanization

A natural question is whether the effect of university expansion varies with local characteristics. We investigate this by estimating the baseline specification separately by tercile of provincial urbanization.[13] Figure 5 shows the event-study coefficients by urbanization level.

Figure 5: Event-study by urbanization level

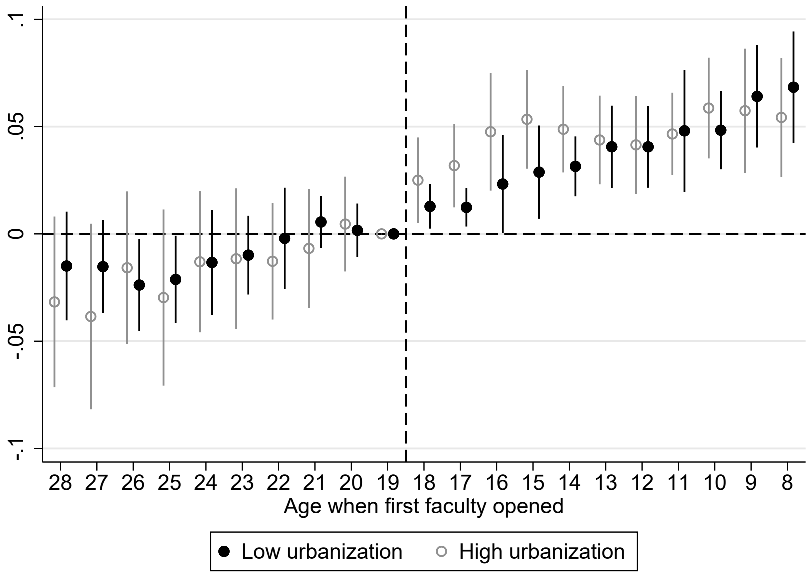

Note: Estimated γ_s coefficients from equation (1), cohorts born 1950–1975, estimated separately by urbanization tercile (low, medium, high). 95% confidence intervals shown. Standard errors clustered at the province level.

The DiD estimates (Appendix Table A-4) show positive and significant effects across all urbanization terciles, with point estimates rising from 0.018 ($p < 0.05$) in low-urbanization provinces to 0.041 ($p < 0.10$) and 0.045 ($p < 0.05$) in more urbanized areas. The gradient is consistent with the idea that broader and more diversified local labour markets amplify the returns to higher education once a local university becomes available, though the upper-tercile estimates carry wider confidence intervals and the underlying mechanism warrants further investigation.

Appendix Table A-5 examines whether the urbanization gradient differs by gender.

[13] Province-level population density is measured as inhabitants per km^2, averaged over the LFS sample period 2009–2019 (source: ISTAT). Low-urbanization provinces: Aosta, Isernia, Potenza, Bolzano, Campobasso, Vercelli, Trento, Foggia, Viterbo, Arezzo, Cosenza. Medium-urbanization provinces: Udine, Alessandria, Benevento, Chieti, Catanzaro, Frosinone, Teramo, Forlì-Cesena, Ascoli Piceno, Reggio di Calabria, Siracusa. High-urbanization provinces: Reggio nell'Emilia, Taranto, Ancona, Pescara, Brescia, Novara, Verona, Rimini, Bergamo, Como, Varese.

Two patterns emerge. First, in low-urbanization provinces, only women show a significant effect (0.023, $p < 0.01$), while the male coefficient is small and insignificant (0.013). This is consistent with a cost-of-access mechanism operating primarily through female mobility constraints: in peripheral areas, where geographic barriers were most binding, local university openings disproportionately relaxed the constraints faced by women. Second, in high-urbanization provinces, women again drive the effect (0.062, $p < 0.05$), while men show a smaller and less precisely estimated gain (0.029, $p < 0.10$). This suggests that labour market complementarities also amplify female gains more than male gains, possibly because urban labour markets offer a broader range of graduate-level occupations that are accessible to women once the educational barrier is removed. Taken together, these results indicate that the two mechanisms underlying the urbanization gradient—geographic access and labour market complementarities—operate primarily through the female margin, though the data do not allow us to distinguish between them.

5.4 Labour market outcomes

We then turn to labour outcomes. Full sample event-study results (ITT effects) for employment and inactivity probability are in Figure A-4 in the Appendix. Overall, coefficient estimates are relatively imprecise, but suggest that exposure to a university opening is associated with a higher probability of employment later in life: increased access to tertiary education is associated with stronger labour market attachment later in life, although these ITT estimates are less precise than the education effects and may also reflect broader local general-equilibrium effects. This is summarized by the DiD estimate of Table A-6 that indicates an average employment gain of 2.1 percentage points, and a reduction of inactivity by about 3 percentage points. These associations may reflect improved labour market outcomes for graduates, but could also capture general equilibrium effects of the university opening.

Figure 6 presents event-study estimates by gender: Panel (a) reports ITT coefficients for the employment equation, while Panel (b) shows the corresponding results for inactivity. Unlike educational attainment, gender differences in labour market responses to university expansion are more nuanced, also reflecting less precisely estimated coefficients.

For employment, the estimated effects are generally positive for women across exposure ages, whereas they are smaller and not statistically significant for men, with effects concentrated at younger ages and fading as cohorts approach the pivotal age of 19. By contrast, the results for inactivity are more precisely estimated and indicate a significant reduction in inactivity for both genders, with larger effects for women.

Overall, the stronger gains in educational attainment observed for women do not translate into uniformly stronger labour market outcomes relative to men, suggesting the presence of gender-specific frictions in converting human capital into sustained employment.

Summary DiD estimates reported in Table A-6 broadly confirm the event-study results and support the main message: university expansion raised educational attainment substantially, especially for women, but the translation into labour market outcomes is incomplete and gender-asymmetric.[14]

This result is consistent with the descriptive evidence presented in Section 3, showing that highly educated women are more likely than men to exit employment in the years approaching retirement. Institutional features of the labour market, care responsibilities, and retirement incentives may therefore limit the extent to which educational gains transfer into sustained employment for women later in life.

6 Robustness

We subject our main results to an extensive battery of robustness checks, addressing concerns related to the timing of university openings, aggregation of treatment effects, pre-trends, specification sensitivity, and the influence of individual provinces.

Residence misclassification. A potential concern is that our data record province of current residence rather than province of birth, introducing measurement error in the treatment assignment. If individuals who benefited from a local university opening subsequently moved to a different province, the observed effect would understate the true causal impact. We address this using 2001 Census data on inter-provincial mobility for

[14]For employment effects, the larger and statistically significant estimate for men (0.029, $p < 0.01$), compared to the smaller and insignificant estimate for women (0.017, $p > 0.10$), differs from the event-study pattern. This reflects the aggregation of heterogeneous age-specific effects into a single average in the DiD specification. In the event-study, employment effects are imprecisely estimated and vary across exposure ages for both genders; as a result, the corresponding average effect is sensitive to the weighting structure and may differ across groups.

Figure 6: Event-study: Probability of being employed and inactive by age of first faculty opened in the province (separately for women and men).

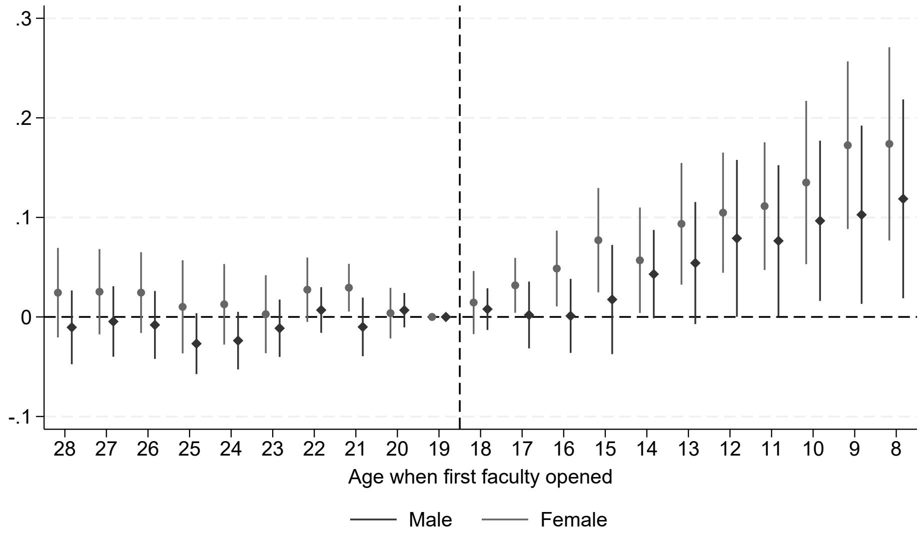

(a) Employment.

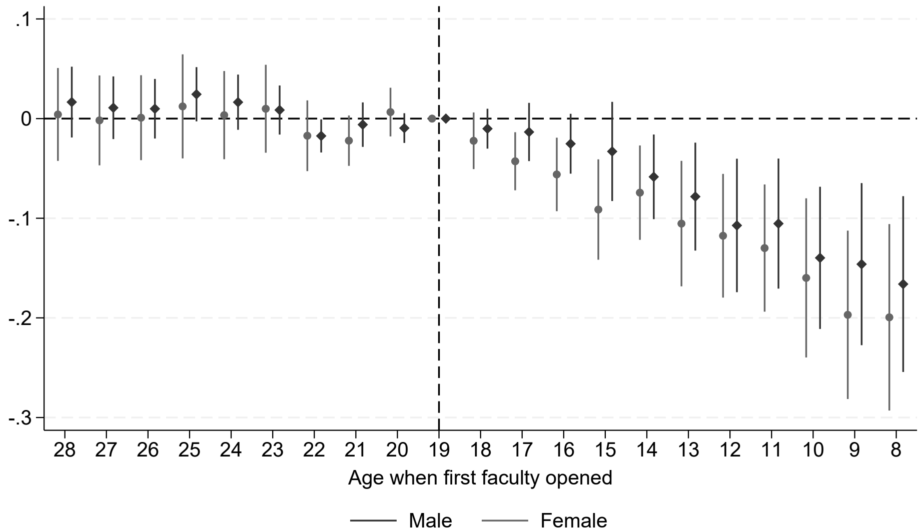

(b) Inactivity.

Note: Estimated γ_s coefficients (dots) and their 95% confidence intervals (bars) from equation (1), when Y_{icat} is a dummy for being employed or inactive. OLS regression, standard errors clustered at the province level.

Italian citizens born between 1947 and 1976.[15] Table A-7 reports the results. Panel A shows that 76% of the full sample and 87% of Southern residents still live in their birth province. Graduates are more mobile than non-graduates (71% vs. 77% residing in their birth province), which is the relevant margin for attenuation bias. Panel B presents three sets of bounds. Under random misclassification, the corrected estimate is $\hat{\beta}/\pi = 0.043$; using the graduate-specific mobility rate yields $\hat{\beta}/\pi_g = 0.046$. These are modestly larger than the observed effect of 0.032, suggesting attenuation of approximately 30–40%.[16] The worst-case Manski bounds without additional assumptions are wide (Panel B), but imposing the monotonicity restriction that university openings weakly increase graduation—supported by our event-study evidence—tightens the lower bound to zero (Panel C). Taken together, these results indicate that our baseline estimates are, if anything, conservative, and the true effect likely lies in the range $[0.032, 0.046]$.

Exogeneity of opening timing. Table A-8 regresses the year of first faculty opening on 1961 Census characteristics within the 33 treated provinces. Among all covariates, only log population is statistically significant ($\beta = -12.0$, $p < 0.01$): larger provinces opened universities earlier, as expected. Educational attainment, illiteracy, female employment, agricultural share, and population density are all individually insignificant. The joint F-test ($p = 0.004$) is driven entirely by log population, which is subsumed by province fixed effects in the main specification. We find no evidence that opening timing is systematically related to pre-existing educational trajectories in the treated provinces, beyond province size.

Aggregation of heterogeneous treatment effects. A general concern with two-way fixed effects estimators in settings with staggered treatment timing is that, under treatment effect heterogeneity, comparisons across groups exposed at different times may induce biased estimates (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Callaway and Sant'Anna, 2021). In our setting, identification relies exclusively on

[15] The Census records both municipality of residence and municipality of birth, allowing us to compute the share of individuals residing in their birth province by cohort and education level. However, the Census records province of birth as a relative indicator with respect to current residence, and birth year is available only in five-year age classes rather than exact years. These limitations preclude computing the age-at-opening variable that underlies our identification strategy. We therefore use the Census to calibrate the attenuation correction rather than to estimate the treatment effect directly.

[16] This correction follows Manski (1990). The key insight is that measurement error in a binary treatment variable attenuates the OLS coefficient toward zero, so the observed effect is a lower bound under the assumption that misclassification is non-differential.

within-province variation across cohorts in age at exposure to the first university opening, while province-by-cohort fixed effects absorb all cross-cohort differences within provinces. This substantially limits concerns related to cross-group comparisons that typically arise in staggered DiD designs with two-way fixed effects estimation, where already-treated units may serve as controls for later-treated ones, leading to biased estimates when treatment effects are heterogeneous over time.

However, when we collapse the event-study into a single treatment indicator, the resulting DiD estimate aggregates potentially heterogeneous effects across ages at exposure. To assess the robustness of this aggregation, we complement the analysis with the estimators proposed by Callaway and Sant'Anna (2021) (C&S) and de Chaisemartin and D'Haultfœuille (2020) (dC&H), which rely only on not-yet-treated observations as controls and are robust to heterogeneous and dynamic treatment effects.[17]

When we apply C&S and dC&H at the province-cohort aggregate level, the C&S simple ATT is 0.013 (SE = 0.006) and the dC&H average cumulative effect is 0.012 (SE = 0.004), both positive and statistically significant. Their lower estimates reflect a different estimand – equal-weighted ATT across province-cohort groups using only not-yet-treated provinces as controls – rather than a bias correction to our main specification.[18] Both remain positive and statistically significant, confirming the direction and significance of the effect (Table A-9). Pre-treatment placebo tests for both estimators show no evidence of differential pre-trends: the C&S average pre-treatment effect is 0.002 ($p = 0.36$) and the joint test of the dC&H placebo coefficients yields $p = 0.096$, supporting the parallel trends assumption.

[17]Our DiD estimator (as any aggregate DiD) can be expressed as a weighted average of the event-study coefficients γ_s. In this sense, the DiD strategy does not identify an independent parameter, but collapses the dynamic treatment effects captured by the event-study into a single scalar summary. Likewise, for example, the C&S estimator can be interpreted in this setting as a re-weighted aggregation of the same underlying treatment effects, differing in the weighting scheme used to combine cohort-specific effects. While our DiD estimator aggregates the γ_s coefficients using weights proportional to the empirical distribution of exposure ages among treated cohorts, the C&S estimator assigns weights based on the relative size of each exposure group and the amount of variation available in the corresponding not-yet-treated comparison groups at each point in the life cycle, where groups with stronger identification power (i.e. those with a richer set of valid not-yet-treated comparisons) receive greater weight relative to purely frequency-based aggregation in the DiD specification.

[18]Unsurprisingly, these results are similar to the event-study coefficients for the treated cohorts closer to age 19. The reason is that the age groups more 'central' and close to the pivotal age have the larger comparison group (within-province controls, not-yet-treated), while groups with less support of controls contribute less.

Randomization inference. With 33 provinces included in the analysis, asymptotic inference may be unreliable even with clustered standard errors. We complement the wild-cluster bootstrap (see below) with a cohort permutation test that directly examines whether the age-19 cutoff is special. Under the null hypothesis of no effect, the specific cohort boundary that defines treatment should be arbitrary: any set of 11 age values drawn from the range 8–28 should produce an ATT indistinguishable from the actual estimate. We therefore randomly draw 11 age values (out of 21 available) as the "treated" cohorts, exclude the adjacent age as the reference category, and re-estimate the baseline specification on each of 500 permutations. Across all permutations, not a single draw produces an absolute ATT as large as the actual estimate, yielding a two-sided p-value < 0.002. Appendix Figure A-5 displays the null distribution alongside the actual ATT (dashed line). The permutation distribution is tightly centered around zero with a standard deviation of 0.003, while the actual estimate lies far in the right tail, confirming that the estimated effect is driven by the specific age-19 cutoff rather than by chance variation in cohort composition.

Leave-one-out. To verify that our results are not driven by any single province, we re-estimate the baseline specification dropping one province at a time. Appendix Figure A-6 reports the 33 estimated coefficients, ranked from smallest to largest, together with their 95% confidence intervals. The estimates range from 0.026 to 0.036, with the full-sample estimate of 0.032 lying near the centre of the distribution. All leave-one-out estimates remain positive and statistically significant. Appendix Table A-12 provides summary statistics for the leave-one-out distribution.

Specification sensitivity. We assess the sensitivity of our results to several features of the empirical specification. First, we vary the estimation window from the baseline [8, 28] to narrower windows ([10, 26], [12, 24], [14, 22]). As shown in Appendix Table A-13, the estimated treatment effect is virtually unchanged across all windows (0.0323–0.0324), indicating that the result is not driven by observations at the tails of the age distribution. Second, we consider alternative specifications of province-by-cohort controls. Appendix Table A-14 shows that replacing province-by-cohort fixed effects with region-by-cohort fixed effects yields a coefficient of 0.019 ($p < 0.05$), while removing province-by-cohort controls entirely gives 0.019. These results confirm that our findings are robust to the particular granularity of the fixed effects structure, though the most demanding specification (province-by-cohort) provides the most precise estimates.

Inference with few clusters. With 33 provinces contributing to identification, asymptotic cluster-robust standard errors may be less reliable (Cameron et al., 2008). We address this concern by computing wild-cluster bootstrap p-values with Rademacher weights and 999 replications (Roodman et al., 2019). The bootstrap p-value for the full-sample baseline estimate is 0.009, well below the 1% level (95% bootstrap CI: [0.012, 0.056]). For the gender subgroups, bootstrap p-values are 0.011 for women and 0.025 for men. The wild-cluster bootstrap therefore corroborates asymptotic inference and allays concerns about few-cluster bias in both the full-sample and subgroup estimates.

Age cutoff sensitivity. Our baseline specification defines treatment as having a university available by age 18. Appendix Figure A-7 examine the sensitivity of the DiD coefficient to this choice by varying the age cutoff from 17 to 21. The coefficient is positive and statistically significant for cutoffs at age 17 and 18, equal to 0.023 ($p < 0.001$) and 0.032 ($p = 0.001$) respectively. At age 19, the coefficient drops sharply to 0.013 and becomes statistically insignificant ($p = 0.108$), remaining close to this value for higher cutoffs. This discontinuity at exactly the enrollment-relevant age provides strong support for the causal interpretation: the effect operates through university enrollment at the appropriate age rather than through confounding trends or selective migration patterns that would not exhibit such a sharp age-specific threshold.

Treatment intensity. If the effect operates through local access, provinces where, when the first university settled in the territory, it did so opening more faculties should exhibit stronger effects. We test this by replacing the binary treatment indicator with the number of faculties available in the province when the individual turned 19, set to zero for control groups and untreated cohorts. Appendix Table A-11 reports the results by gender. Panel B shows that each additional faculty is associated with a 1.1 percentage point increase in graduation probability for the full sample, with the effect nearly twice as large for women (1.5 pp) as for men (0.7 pp). Panel A (dose bins) reveals a monotonically increasing pattern: across all bins, women's coefficients are consistently larger than men's, ranging from 2.7 pp (one faculty) to 5.7 pp (three or more faculties) for women, compared to 2.7 pp and 3.3 pp for men. While pairwise differences across bins are not individually significant—reflecting the limited number of treated provinces—the gender gradient in the dose-response is consistent with the interpretation that geographic access barriers weigh more heavily on women: as more programs become available locally, the gains from removing distance constraints are amplified for those who face stronger mobility constraints.

Cohort window sensitivity. Appendix Table A-10 examines how the main results change as the cohort window used for the analysis is progressively extended from considering individuals born until 1965 to 1975. Panel A shows that the ATT is remarkably stable at approximately 0.018 for windows 1950–1965 and 1950–1970, then jumps to 0.032 when cohorts born 1970–1975 are included. These cohorts are the youngest in the treated sample: on the one hand, they were exposed for a longer period as treated to the openings that also affected older cohorts (likely in primary or middle school when their province's university opened, meaning the institution was available throughout their entire secondary schooling trajectory); on the other hand, they experienced more openings as compared to individuals born, say, in 1965 or 1970. This maximum-exposure profile is consistent with stronger responsiveness among those who could plan their educational investments from an early age with full knowledge of local university availability and where this availability is higher.

7 Spatial Analysis

The individual-level analysis presented in Section 5 establishes that university openings increased graduation rates in the provinces where new institutions were established. Two natural extensions follow: whether these effects spilled over to neighboring provinces, and whether the expansion contributed to a convergence in human capital across Italian provinces. We address the first question using the same individual-level LFS data as the baseline, and the second using province-level data collapsing individual graduation outcomes into province-by-cohort-group cells.[19]

Spillover effects. If proximity to a university lowers the cost of attendance, individuals living in provinces adjacent to those that received a new institution may also benefit. We investigate this possibility by replicating the baseline individual-level specification but replacing—or supplementing—the own-province treatment timing with the timing of the earliest expansion opening among contiguous provinces.[20] Table A-16 reports four

[19]For the convergence analysis we define five cohort groups of approximately equal size: 1950–1954, 1955–1959, 1960–1964, 1965–1969, and 1970–1975.

[20]The adjacency matrix is based on shared provincial borders as defined by ISTAT administrative boundaries. For each province we record the year of the earliest expansion opening (i.e. the first post-

specifications using the same estimation window $[8, 28]$ and the same fixed-effect structure as the baseline.

Column (1) replicates the baseline for the 33 treated provinces using own-province opening timing ($\widehat{\text{ATT}} = 0.032^{***}$). Column (2) applies the same specification to the 15 *never* provinces —those that never received a university of their own, but benefited by an opening in a contiguous province— using the timing of the earliest expansion opening in a contiguous province as the treatment date. The resulting coefficient (0.033^{**}) is virtually identical in magnitude to the baseline. Taken in isolation, this finding is ambiguous: it is consistent with genuine spillovers, but may also reflect the fact that never provinces tend to be geographically clustered near treated provinces and thus share common cohort trends.

Columns (3) and (4) sharpen the identification. Column (3) restricts attention to the 27 of 33 treated provinces that also have a contiguous expansion province whose opening falls in the window $[8, 28]$, using *only* the neighbour opening as the treatment date. The estimated effect is 0.021^{**}—positive and significant, but meaningfully smaller than the baseline. Column (4) runs the horse race on the same 27 provinces, including both own-province and neighbour treatment indicators simultaneously. The own-province coefficient remains significant (0.024^{*}), while the neighbour coefficient falls to 0.015^{*} and is significant only at the 10 per cent level. The two effects are thus separable: local access is the dominant channel, but a genuine, smaller spillover is also present.

Taken together, these results indicate that the effect of university openings is primarily local. In the horse-race specification, the own-province effect is approximately twice as large as the neighbour effect, while the positive and marginally significant neighbour coefficient points to moderate spillovers to adjacent provinces.[21] The pattern is consistent with a proximity channel of moderate strength: some individuals in bordering provinces can access a university in a neighbouring province, but the reduction in travel costs from having an institution in one's own province is substantially larger.

Convergence in human capital. Italy provides a particularly informative setting for studying educational convergence, given its well-documented geographical divide in

1950 university opening) among all contiguous provinces, and define the neighbour treatment age as this date minus year of birth.

[21]This finding qualifies the SUTVA assumption. Interference across units exists but is small relative to the direct effect; our baseline estimates should therefore be interpreted as a lower bound of the total effect, combining the direct effect with the partial spillover to contiguous provinces.

human capital accumulation and local labour market outcomes (Paci and Pigliaru, 1997; Dalmazzo and De Blasio, 2007). We compare graduation trends in provinces receiving the first university after 1968 (late- and recent-expansion provinces, in the terminology of Figure 2) with provinces never receiving the first university in the observation period (never provinces) using a standard difference-in-differences framework. Table A-17 reports the results. The interaction between late-expansion status and post-1968 cohorts yields a coefficient of 0.012, significant at the 5% level with province and cohort fixed effects (column 3). This estimate is remarkably close to the individual-level baseline, providing province-level corroboration of the micro-econometric results.[22] Figure A-8 visualizes these dynamics by plotting the graduation gap of each province type relative to historical university provinces (those with a university before 1945). Early-expansion and late-expansion provinces maintain a roughly stable gap relative to historical universities, while never provinces fall substantially further behind — diverging by nearly 2.5 percentage points over the sample period.[23] These findings are suggestive but indicate that the university expansion contributed to a partial convergence in educational attainment, with the largest gains in previously underserved areas.

8 Conclusion

We investigate the impact of Italy's university expansion, specifically the opening of a first faculty in provinces that had never hosted a university, on educational attainment and provide complementary ITT evidence on later-life labour market outcomes. Exploiting variation in the timing of first university openings across provinces in an event-study framework, we provide causal evidence that a first local university substantially raised the probability of obtaining a university degree. The event-study estimate at the enrollment-age threshold is 1.2 percentage points; aggregating across all treated cohorts, the difference-in-differences ATT is 3.2 percentage points (approximately 26% relative to the pre-treatment mean). Both are robust to a comprehensive battery of specification checks, alternative estimators, and randomization inference. Randomly permuting the treatment timing among the 33 treated provinces yields an ATT as large as the observed one in zero out of 500 draws (two-sided $p = 0.002$), providing strong evidence that the es-

[22] The concordance is reassuring, as the two approaches differ in unit of analysis, fixed effect structure, and identifying variation. The province-level approach compares only late-expansion versus never provinces, while the individual-level event-study exploits the full staggered timing.

[23] Population weights (column 4) yield a virtually identical coefficient (0.012), ruling out compositional effects from differential population growth across province types.

timated effect is not attributable to chance variation in the timing of university openings.

Four sets of findings merit emphasis. First, the effect exhibits a significant geographic gradient: it is positive and significant across all urbanization terciles and increasing in more urbanized provinces, consistent with complementarities between university access and local labour market conditions. Dose-response analysis is consistent with both the access and labor market complementarity channels: replacing the binary treatment indicator with the number of faculties available by age 18 yields a monotonically increasing pattern—the effect rises from 2.7 percentage points in provinces with one faculty to 4.5 percentage points in provinces with three or more. The sharp drop in the effect for cohorts above typical enrollment age (19 and older) provides further support for the causal interpretation. Second, heterogeneity by gender reveals that women benefited disproportionately from local university openings, consistent with the hypothesis that geographic barriers weigh more heavily on women due to social norms and family constraints governing female mobility. The dose-response gradient is also stronger for women: each additional faculty is associated with a 1.5 percentage point increase in female graduation probability, compared to 0.7 for men, consistent with the interpretation that mobility constraints are more binding for women not only at the extensive margin of access but also along the intensive margin of breadth of supply.

Yet, despite narrowing the gender gap in education, increased access did not fully translate into comparable labour market gains. Third, the treatment effect varies systematically across cohort windows. For cohorts born before 1971 the effect is stable at approximately 1.8 percentage points, with the impact concentrated in less urbanized provinces, consistent with a cost-of-access mechanism. When the youngest cohorts (born 1973–1975) are included, the overall ATT rises to 3.2 percentage points and the urbanization gradient reverses, reflecting the stronger response of cohorts who had full exposure to local universities throughout their secondary schooling. Fourth, at the province level, the expansion contributed to a partial convergence in human capital across Italian regions. Late-expansion provinces substantially narrowed their graduation gap relative to historical university towns, while provinces that never received a university fell further behind. Spillover effects to neighbouring provinces exist but are of secondary magnitude, confirming that local access is the dominant channel.

These findings contribute to an active policy debate on the role of geographic access in shaping educational attainment and its long-run labour market consequences. The causal evidence that establishing a first university in previously unserved provinces raised graduation rates meaningfully — with effects larger for women, consistent with geographic barriers being more binding where mobility constraints were strongest — suggests that place-based investment in higher education can be an effective tool for expanding access in underserved areas. The weaker translation of women's educational gains into labour market outcomes signals that access expansion alone is insufficient: complementary policies addressing occupational segregation and care responsibilities remain important.

References

Andersson, R., Quigley, J. M., and Wilhelmsson, M. (2004). University decentralization as regional policy: The Swedish experiment. Journal of Economic Geography, 4(4):371–388.

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

Boelman, B. (2021). Women's missing mobility and higher education: Evidence from Germany's university expansion. Working Paper, University College London.

Braccioli, F., Ghinetti, P., Moriconi, S., Naguib, C., and Pellizzari, M. (2023). Education expansion, college choice and labour market success. Working Paper 10842, CESifo.

Brilli, Y. and Bassoli, E. (2025). How raising the full retirement age affects women's early retirement choices: insights from the interaction of two policies. WorkINPS Paper 107, INPS, Roma, Italia.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. Journal of Econometrics, 225(2):200–230.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. Review of Economics and Statistics, 90(3):414–427.

Caner, A., Derebasoglu, M., and Okten, C. (2024). Attainment and gender equality in higher education: Evidence from a large-scale expansion. Journal of Human Capital, 18(3):469–530.

Card, D. (1999). The causal effect of education on earnings. In Ashenfelter, O. and Card, D., editors, Handbook of Labor Economics, volume 3, chapter 30, pages 1801–1863. Elsevier.

Carneiro, P. and Heckman, J. J. (2002). The evidence on credit constraints in post-secondary schooling. The Economic Journal, 112(482):705–734.

Carneiro, P., Heckman, J. J., and Vytlacil, E. J. (2011). Estimating marginal returns to education. American Economic Review, 101(6):2754–2781.

Ciani, E. (2016). Retirement, pension eligibility and home production. Labour Economics, 38:106–120.

Cottini, E., Ghinetti, P., and Moriconi, S. (2026). Keeping up with the joneses? the rise of modern universities and local economic development in italy. Journal of Economic Growth.

Currie, J. and Moretti, E. (2003). Mother's education and the intergenerational transmission of human capital: Evidence from college openings. Quarterly Journal of Economics, 118(4):1495–1532.

Dalmazzo, A. and De Blasio, G. (2007). Social returns to education in italian local labour markets. Annals of Regional Science, 41(1):51–69.

de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. American Economic Review, 110(9):2964–2996.

Elsayed, A. and Shirshikova, A. (2023). The women-empowering effect of higher education. Journal of Development Economics, 163:103101.

Fack, G. and Grenet, J. (2015). Improving college access and success for low-income students: Evidence from a large need-based grant program. American Economic Journal: Applied Economics, 7(2):1–34.

Falch, T., Lujala, P., and Strøm, B. (2013). Geographical constraints and educational attainment. Regional Science and Urban Economics, 43(1):164–176.

Goldin, C. and Katz, L. F. (2008). The Race Between Education and Technology. Belknap Press of Harvard University Press, Cambridge, MA.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. Econometrica, 89(5):2261–2290.

Ichino, A., Rustichini, A., and Zanella, G. (2025). College, cognitive ability, and socioeconomic disadvantage: policy lessons from the UK in 1960–2004. The Review of Economic Studies, pages 1–45. Forthcoming.

Kobus, M., Van Ommeren, J. N., and Rietveld, P. (2015). Student commute time, university presence and academic achievement. Regional Science and Urban Economics, 52:129–140.

Kyui, N. (2016). Expansion of higher education, employment and wages: Evidence from the russian transition. Labour Economics, 39:68–87.

Manski, C. F. (1990). Nonparametric bounds on treatment effects. American Economic Review, 80(2):319–323.

Maura, F. and Profeta, P. (2025). Women's caring penalty at retirement in europe. SHARE Working Paper Series 95.

Mincer, J. (1974). Schooling, Experience, and Earnings. National Bureau of Economic Research / Columbia University Press, New York.

Moretti, E. (2004). Estimating the social return to higher education: evidence from longitudinal and repeated cross-sectional data. Journal of Econometrics, 121(1-2):175–212.

Nybom, M. (2017). The distribution of lifetime earnings returns to college. Journal of Labor Economics, 35(4):903–952.

Oppedisano, V. (2011). The (adverse) effects of expanding higher education: Evidence from italy. Economics of Education Review, 30(5):997–1009.

Oreopoulos, P. and Petronijevic, U. (2013). Making college worth it: A review of the returns to higher education. The Future of Children, 23(1):41–65.

Paci, R. and Pigliaru, F. (1997). Structural change and convergence: an italian regional perspective. Structural Change and Economic Dynamics, 8(3):297–318.

Roodman, D., Nielsen, M. Ø., MacKinnon, J. G., and Webb, M. D. (2019). Fast and wild: Bootstrap inference in Stata using `boottest`. Stata Journal, 19(1):4–60.

Si, W. (2022). Higher education expansion and gender norms: Evidence from china. Journal of Population Economics, 35(4):1821–1858.

Appendix

Figure A-1: Graduation rates by birth cohort and province type

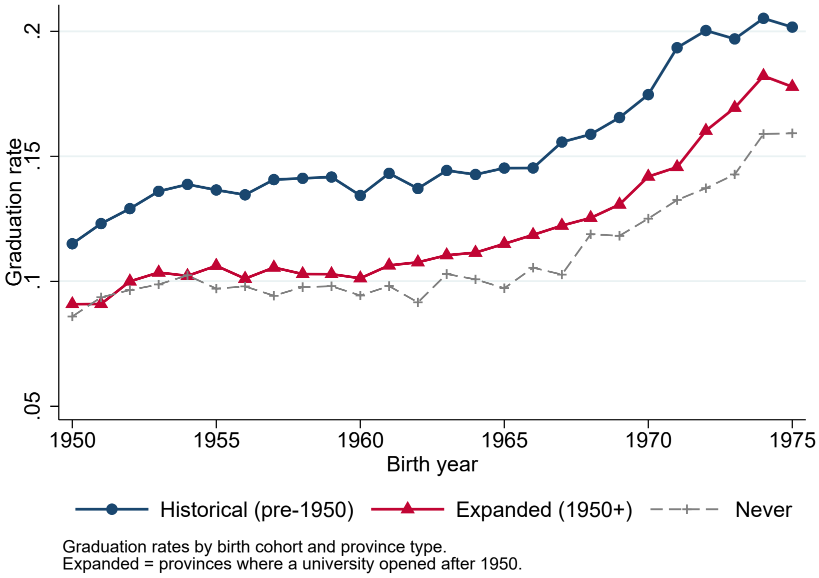

Note: Own elaborations on the LFS-HIU dataset. Historical: provinces with a university before 1950. Expanded: provinces where the first university opened after 1950. Never: provinces that never received a university by 2010. Graduation rates are computed for five-year birth cohorts.

Figure A-2: Number of faculties by province and cohort

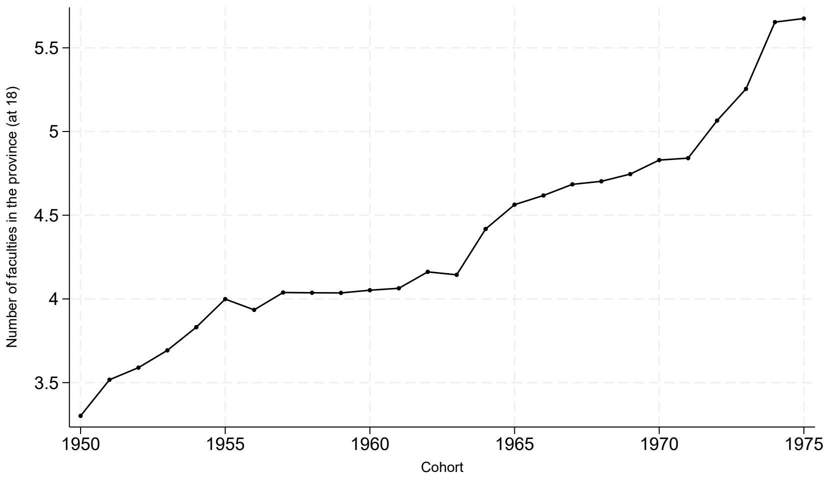

Note: Own elaborations based on the combined LFS-HIU dataset.

Table A-1: Descriptive statistics: Graduated vs Not graduated by age group and gender

	All		Men		Women	
	Graduate	No grad.	Graduate	No grad.	Graduate	No grad.
Employed (all ages)	0.85	0.59	0.91	0.73	0.81	0.45
Age 35–44	0.87	0.70	0.93	0.85	0.83	0.56
Age 45–54	0.90	0.65	0.95	0.81	0.85	0.50
Age 55–59	0.88	0.52	0.93	0.66	0.84	0.40
Age 60–65	0.66	0.30	0.75	0.38	0.56	0.22
Inactive (all ages)	0.12	0.36	0.07	0.21	0.17	0.50
Age 35–44	0.09	0.23	0.04	0.09	0.13	0.38
Age 45–54	0.08	0.29	0.03	0.13	0.12	0.45
Age 55–59	0.11	0.45	0.06	0.30	0.15	0.59
Age 60–65	0.34	0.69	0.24	0.59	0.43	0.77

Note: own elaborations on the LFS-HIU dataset. The values are the means of the corresponding variables.

Figure A-3: Distribution of observations by age at which the first faculty opened in the province.

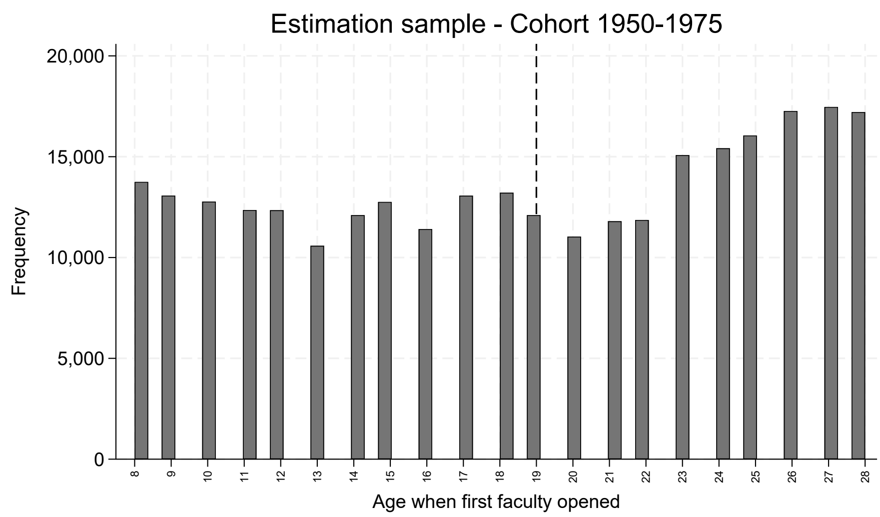

Note: Own calculations using the LFS-HIU dataset. The observations aged 8-18 are treated; from 19 to 28 are the controls. Cohorts born between 1950 and 1970 (included).

Figure A-4: Event-study: Probability of being employed and inactive by age of first faculty opened in the province (full sample)

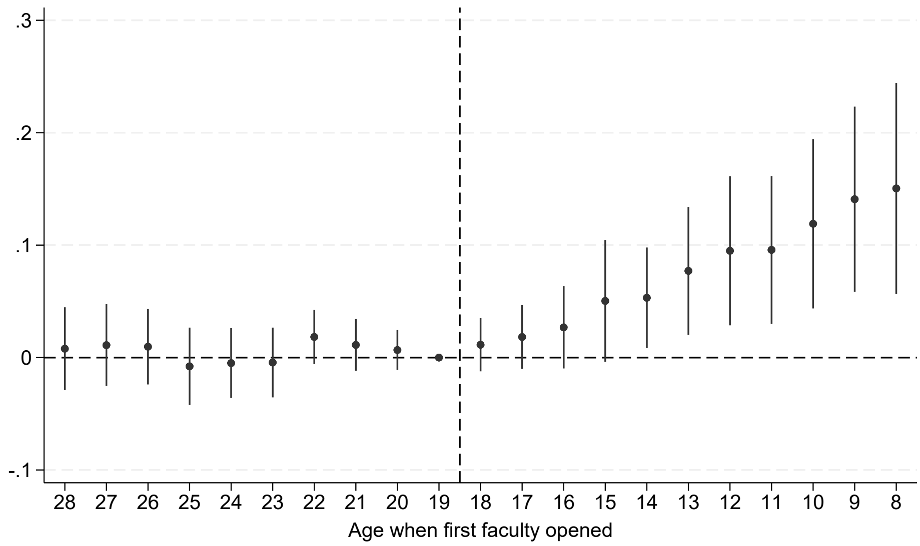

(a) Employment.

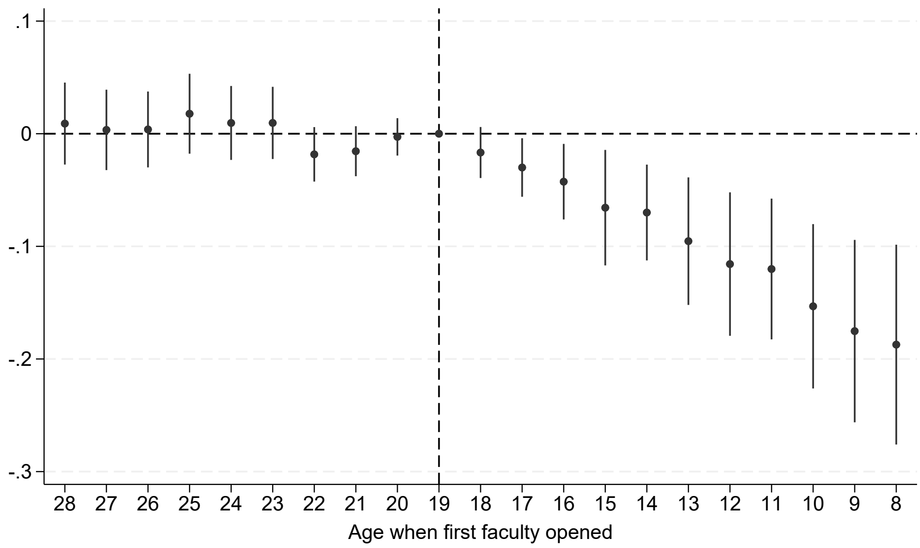

(b) Inactivity.

Note: Estimated γ_s coefficients (dots) and their 95% confidence intervals (bars) from equation (1), when Y_{icat} is a dummy for being employed or inactive. OLS regression, standard errors clustered at the province level.

Table A-2: Pre-trend test: 1951–1961 educational growth and opening timing

	(1) Δ Graduates	(2) Δ Illiteracy	(3) Δ Fem. emp.	(4) Δ Agri. emp.
Year of first opening	-0.005	0.020	0.018	0.107*
	(0.005)	(0.045)	(0.041)	(0.056)
Observations	33	33	33	33
R-squared	0.049	0.006	0.005	0.090

Note: OLS. Dependent variable: change in population-weighted provincial outcome between Census 1951 and 1961, before the university expansion. Sample: the 33 treated provinces in the main estimation. Δ Graduates = change in share of graduates and diploma holders (% pop. aged 6+); Δ Illiteracy = change in illiteracy rate; Δ Fem. emp. = change in female employment rate; Δ Agri. emp. = change in agricultural employment share. Source: ISTAT 8000 Comuni Census database. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A-5: Randomization inference: permutation distribution of ATT

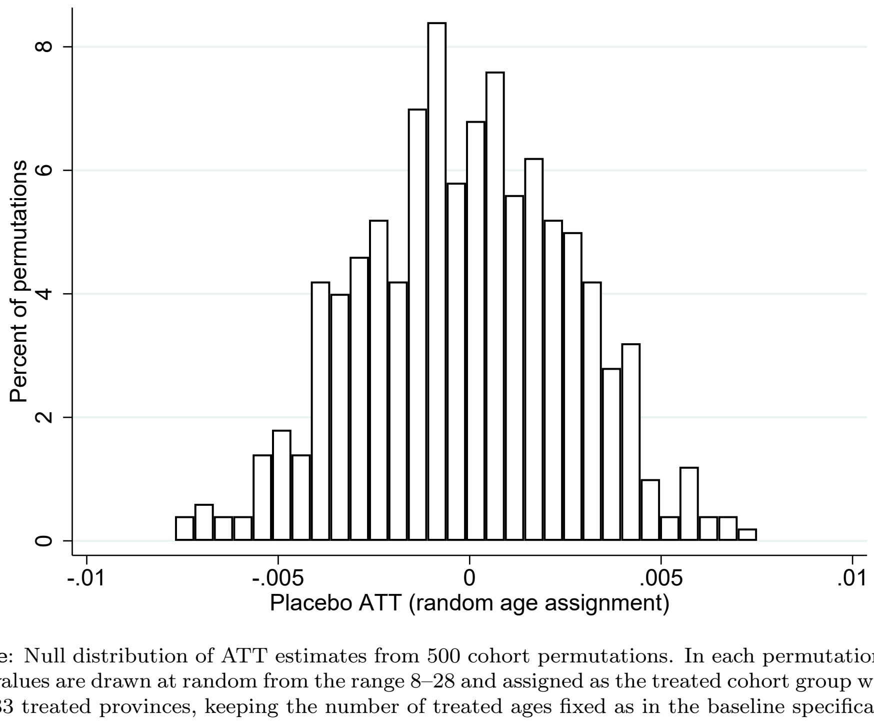

Note: Null distribution of ATT estimates from 500 cohort permutations. In each permutation, 11 age values are drawn at random from the range 8–28 and assigned as the treated cohort group within the 33 treated provinces, keeping the number of treated ages fixed as in the baseline specification. The distribution is centered at zero (mean = -0.0001, SD = 0.003). The actual ATT of 0.032 lies far to the right of the displayed range; no permutation produces an absolute effect as large, yielding a two-sided p-value < 0.002.

Figure A-6: Leave-one-out: treated coefficient

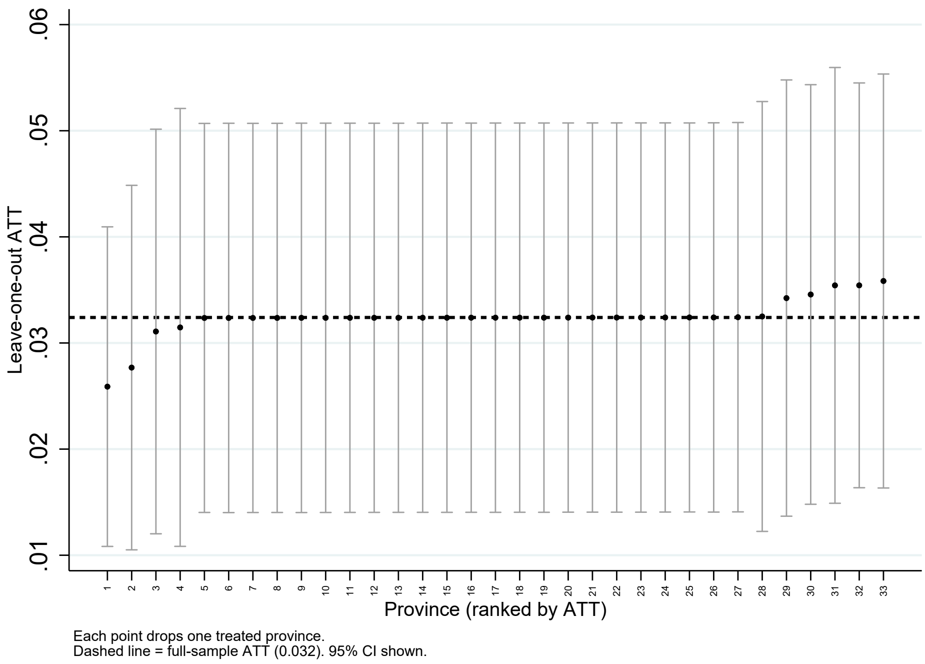

Note: Each point represents the coefficient on Treated from the baseline DiD specification estimated excluding one province. Dashed red line: full sample estimate. Vertical bars: 95% confidence intervals. Standard errors clustered at the province level.

Figure A-7: DiD coefficient by age cutoff defining treatment

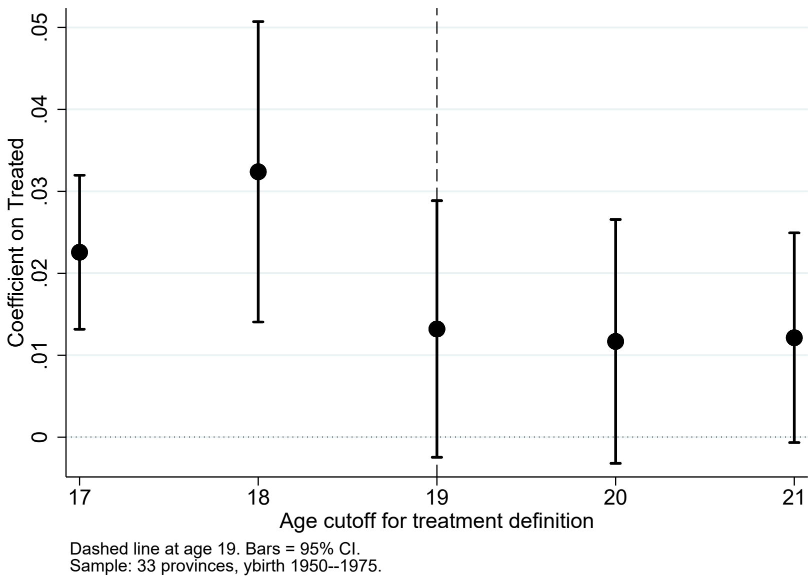

Note: Each point is the coefficient on Treated from the baseline DiD, varying the age cutoff from 17 to 21. Vertical bars: 95% confidence intervals. SE clustered at the province level.

Table A-3: DiD model, treatment effects for graduate equation

	All	Men	Women
Treated (Graduate)	0.0324***	0.0272**	0.0372***
	(0.0094)	(0.0113)	(0.0113)
Observations	270,627	130,454	140,173
Cohort $\times$ Prov FE	Yes	Yes	Yes
Mean dep. var. (Graduate)	0.125	0.114	0.135
Mean dep. var. (Employed)	0.622	0.751	0.501
Mean dep. var. (Inactive)	0.333	0.198	0.458

Note: Treated = 1 if individual was aged ≤ 18 when the first faculty opened in their province; age 19 at opening (normalization cohort) excluded. Birth cohorts 1950–1975, age at opening in [8, 28]. All specifications include province, cohort-bin, year, and quarter fixed effects, cohort-bin $\times$ province fixed effects, and gender (All column). Standard errors clustered at the province level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A-4: DiD estimates by urbanization tercile

	(1) Low	(2) Medium	(3) High
Treated	0.018**	0.041*	0.045**
	(0.006)	(0.020)	(0.019)
Observations	134,132	77,375	59,120
Cohort $\times$ Prov FE	Yes	Yes	Yes
Mean dep. var.	0.126	0.121	0.125

Note: Urbanization terciles defined at province level. All specifications include province, cohort-bin, year, and quarter FE, cohort-bin $\times$ province FE, and gender. SE clustered at the province level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A-5: DiD estimates by gender and urbanization tercile

	Urbanization tercile		
	Low	Medium	High
Panel A: Men			
Treated	0.0127	0.0491	0.0288*
	(0.0127)	(0.0299)	(0.0156)
Observations	65,108	36,789	28,557
Mean dep. var.	0.117	0.107	0.116
Panel B: Women			
Treated	0.0234***	0.0340	0.0615**
	(0.0025)	(0.0212)	(0.0272)
Observations	69,024	40,586	30,563
Mean dep. var.	0.136	0.133	0.134
Cohort $\times$ Prov FE	Yes	Yes	Yes

Note: Treated = 1 if aged ≤ 18 at university opening in the province; age 19 excluded. Birth cohorts 1950–1975, age at opening [8, 28]. All specifications include province, cohort-bin, year, and quarter FE, cohort-bin $\times$ province FE. Urbanization terciles defined at the province level. SE clustered at the province level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A-6: DiD model, treatment effects for employment, and inactivity equations

	All	Men	Women
Treated (Employed)	0.0211**	0.0288***	0.0169
	(0.0086)	(0.0072)	(0.0139)
Treated (Inactive)	−0.0294***	−0.0250***	−0.0362***
	(0.0080)	(0.0064)	(0.0113)
Observations	270,627	130,454	140,173
Cohort × Prov FE	Yes	Yes	Yes
Mean dep. var. (Graduate)	0.125	0.114	0.135
Mean dep. var. (Employed)	0.622	0.751	0.501
Mean dep. var. (Inactive)	0.333	0.198	0.458

Note: Treated = 1 if individual was aged ≤ 18 when the first faculty opened in their province; age 19 at opening (normalization cohort) excluded. Birth cohorts 1950–1975, age at opening in [8, 28]. All specifications include province, cohort-bin, year, and quarter fixed effects, cohort-bin × province fixed effects, and gender (All column). Standard errors clustered at the province level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A-7: Bounds on the Treatment Effect Under Province Misclassification

	Observed effect (lower bound)	Random misclass. $\hat{\beta}/\pi$	Graduate-weighted $\hat{\beta}/\pi_g$	Manski (worst case)
Panel A: Mobility rates (Census 2001, % in birth province)				
Full sample	76.0%	77.4% (non-grad)	71.1% (grad)	
South	86.8%	87.5% (non-grad)	80.3% (grad)	
Panel B: Bounds on treatment effect				
Full sample	0.0324	0.0426	0.0456	$[-0.273, 0.358]$
South	0.0287	0.0331	0.0357	$[-0.119, 0.185]$
Panel C: With monotonicity ($\beta_{true} \geq 0$)				
Full sample	0.0324	0.0426	0.0456	$[0.000, 0.358]$
South	0.0287	0.0331	0.0357	$[0.000, 0.185]$

Note: Bounds on the true treatment effect under province misclassification (province of residence observed; province of birth required). Mobility rates from the 2001 Census (Italian citizens, cohorts 1947–1976). *Random misclassification*: assumes classical (non-differential) measurement error in the binary treatment; the observed coefficient is attenuated by π, so $\hat{\beta}_{\text{true}} = \hat{\beta}_{\text{obs}}/\pi$, where π is the share of individuals residing in their birth province. *Graduate-weighted*: uses π_g (the graduate-specific retention rate) as the relevant attenuation factor, since misclassification is more likely among graduates. *Manski worst-case bounds*: following Manski (1990), since $Y \in \{0, 1\}$ (binary outcome), the worst-case bounds are $\left[(\hat{\beta}_{\text{obs}} - (1-\pi))/\pi,\ (\hat{\beta}_{\text{obs}} + (1-\pi))/\pi\right]$; for the full sample ($\hat{\beta}_{\text{obs}} = 0.0324$, $\pi = 0.760$) this yields $[-0.273, 0.358]$. *Monotonicity*: imposes $\beta_{\text{true}} \geq 0$ (university openings weakly increase graduation), truncating the lower bound at zero.

Table A-8: Timing regression: year of first opening on Census 1961 characteristics

	(1) Raw coefficients	(2) Standardised regressors
Graduates + diplomas (%)	1.006	0.691
	(3.365)	(2.310)
Illiteracy rate (%)	0.094	0.679
	(0.481)	(3.490)
Female employment rate (%)	0.605*	4.400*
	(0.312)	(2.270)
Agricultural employment (%)	−0.175	−2.590
	(0.333)	(4.917)
Population density	−0.002	−0.591
	(0.017)	(4.198)
Log population	−11.994***	−6.022***
	(4.078)	(2.047)
Observations	33	33
R-squared	0.303	0.303
Joint F-test (p-value)	0.004	

Note: OLS. Dependent variable: year of first faculty opening. Sample: the 33 treated provinces in the main estimation. Regressors: population-weighted province averages, ISTAT Census 1961. Column (2): regressors standardised to mean 0, SD 1. The joint F-test ($p = 0.004$) is driven entirely by log population; all other covariates are individually insignificant. Log population is absorbed by province fixed effects in the main specification. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A-9: Robustness to Staggered Treatment: Heterogeneity-Robust Estimators

	Callaway & Sant'Anna (2021)	de Chaisemartin & D'H. (2020)
ATT (graduate)	0.013**	0.012***
	(0.006)	(0.004)
Pre-trend test (p-value)	0.360	0.096
Provinces	33	33
Cells	377	377

Note: Both estimators use not-yet-treated provinces as the control group. Callaway & Sant'Anna (2021): simple ATT aggregated across groups and periods (`csdid`, `notyet` option). De Chaisemartin &D'Haultfoeuille (2020): average cumulative effect per treatment unit(`did_multiplegt_dyn`), weighted by cell size, 8 post-treatment periods. Pre-trend test: C&S average pre-treatment ATT (p-value from t-test); dCH joint test of 4 placebo coefficients. Outcome: graduate degree. Standard errors clustered at province level. TWFE baseline estimate: 0.032 (SE = 0.007).

Table A-10: Sensitivity to cohort window: ATT, distance gradient, and urbanization gradient

	Cohort window		
	(1)	(2)	(3)
	yb≤1965	yb≤1970	yb≤1975
Treated (ATT)	0.0179**	0.0178**	0.0324***
	(0.0066)	(0.0066)	(0.0094)
Observations	143,421	203,096	270,627
Cohort × Prov FE	Yes	Yes	Yes
Mean dep. var.	0.106	0.111	0.125

Note: Each column restricts to the indicated cohort window. Panels B and C restrict to treated provinces.
All specifications include province FE, cohort×province FE, year FE, quarter FE, and gender dummy.
SE clustered by province. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A-11: Dose-response: treatment intensity by number of faculties, by gender

Panel A: Dose bins (reference = 0 faculties by age 18)

	(1) All	(2) Women	(3) Men
1 faculty	0.0270**	0.0271**	0.0270
	(0.0121)	(0.0114)	(0.0182)
2 faculties	0.0313***	0.0406**	0.0214**
	(0.0114)	(0.0190)	(0.0100)
3+ faculties	0.0453***	0.0565***	0.0327**
	(0.0103)	(0.0161)	(0.0140)
Observations	270,627	140,173	130,454
Cohort × Prov FE	Yes	Yes	Yes
Mean dep. var.	0.125	0.135	0.114

Panel B: Continuous treatment intensity (number of faculties)

	(1) All	(2) Women	(3) Men
N faculties (count)	0.0114***	0.0154***	0.0072*
	(0.0038)	(0.0046)	(0.0038)
Observations	282,729	146,379	136,350
Cohort × Prov FE	Yes	Yes	Yes
Mean dep. var.	0.124	0.134	0.114

Note: Outcome: probability of university graduation. Treatment intensity = number of faculties in province when individual turned 18, set to zero for control groups and untreated cohorts (age_fac_open > 18). All specifications include province, cohort-bin, year, quarter FE, cohort-bin × province FE. SE clustered by province. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A-12: Leave-one-out robustness: summary statistics

	Coefficient on Treated
Full sample	0.0324
	(0.0094)
Leave-one-out statistics	
Mean	0.0324
Minimum	0.0259
Maximum	0.0358
N provinces	33

Note: Each LOO estimate drops one treated province and re-estimates the baseline. SE clustered by province.

Table A-13: Robustness to estimation window

	(1) [8,28]	(2) [10,26]	(3) [12,24]	(4) [14,22]
Treated (age_fac_open≤18)	0.0324***	0.0324***	0.0323***	0.0324***
	(0.0094)	(0.0094)	(0.0093)	(0.0094)
Observations	270,627	209,123	150,672	97,250
Cohort x Province FE	Yes	Yes	Yes	Yes
Window	[8,28]	[10,26]	[12,24]	[14,22]
Mean dep. var.	0.125	0.123	0.123	0.124

Note: Treat = university open by age 18. Different age windows. SE clustered by province.

Table A-14: Robustness to alternative control specifications

	(1) Baseline	(2) Region x Coh	(3) Minimal
Treated (age_fac_open≤18)	0.0324***	0.0191**	0.0193**
	(0.0094)	(0.0077)	(0.0072)
Observations	270,627	270,627	270,627
Prov-cohort controls	Cohort x Prov	Region x Cohort	None
Mean dep. var.	0.125	0.125	0.125

Note: SE clustered by province. All include province, cohort, year, quarter FE and gender.

Table A-15: DiD estimates by age cutoff defining treatment

	(1) Age 16	(2) Age 17	(3) Age 18	(4) Age 19	(5) Age 20	(6) Age 21
treat_cut	0.0251***	0.0309***	0.0324***	0.0123	0.0105	0.0163*
	(0.0069)	(0.0073)	(0.0094)	(0.0106)	(0.0082)	(0.0087)
Observations	269,661	269,517	270,627	271,692	270,931	270,867
Cohort x Province FE	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	0.124	0.124	0.125	0.124	0.124	0.125

Note: Treat = university open by age X (omits age X+1). SE clustered by province.

Table A-16: Spillover effects: own vs. neighbour university opening

	(1) Baseline	(2) Never provinces	(3) Treated (neigh)	(4) Horse race
Treated (own opening)	0.0324***			0.0241*
	(0.0094)			(0.0120)
Treated (neigh opening)		0.0328**	0.0213**	0.0145*
		(0.0111)	(0.0078)	(0.0084)
Observations	270,627	98,151	289,313	282,144
Provinces	33	15	27	27
Cohort × Prov FE	Yes	Yes	Yes	Yes
Sample	Treated prov.	Never prov.	Treated prov.	Treated prov.
Mean dep. var.	0.125	0.119	0.122	0.122

Notes. Col. (1): 33 treated provinces (own opening timing). Col. (2): 15 never provinces with adjacent expansion province; treatment timing = earliest adjacent expansion opening − birth year. Col. (3): treated provinces that also have a neighbouring expansion opening in window $[8, 28]$; treatment timing = neighbour opening only. Col. (4): same sample as col. (3); horse race with both own and neighbour opening timing included simultaneously. All specifications include province FE, cohort-bin FE, year FE, quarter FE, cohort-bin × province FE, and a female indicator. Estimation window $[8, 28]$, birth cohorts 1950–1975. Standard errors clustered by province in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A-17: University Openings and Human Capital Convergence

	(1) OLS	(3) FE	(4) W-FE
Late/Recent expansion	0.0054*		
	(0.0029)		
Post (cohorts 1965+)	0.0271***		
	(0.0041)		
Late/Recent expansion × Post	0.0123*	0.0123**	0.0116**
	(0.0064)	(0.0049)	(0.0049)
South			
Province FE	No	Yes	Yes
Cohort FE	No	Yes	Yes
Observations	350	350	350
R-squared	0.289	0.808	0.828

Note: Unit of observation: province × cohort group. Sample restricted to late- and recent-expansion and never provinces. Late/Recent expansion: first university opened after 1968. Post: birth cohorts 1965 and later. Even columns weighted by cell population. Column 2 includes acontrol for Southern provinces. Columns 5–6 restricted to Southern provinces (ISTAT code > 59). Standard errors in parentheses, clustered by province in columns 3–6. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A-8: Graduation rate gap relative to historical university provinces, by province type and birth cohort

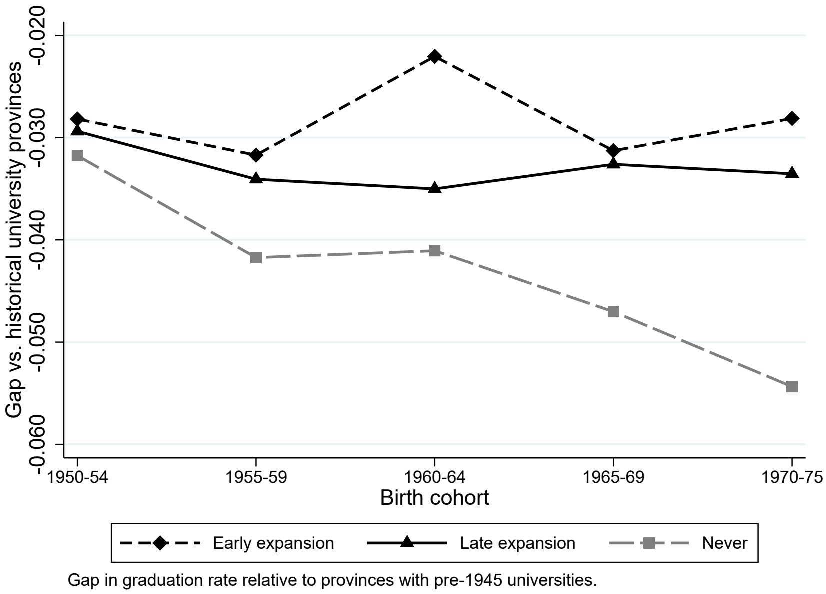

Note: Graduation rate gap relative to provinces with a university before 1945, by province type and birth cohort. Early-expansion provinces nearly close the gap by the final cohort group; late-expansion provinces narrow it substantially; never provinces fall further behind. The divergence between late-expansion and never provinces for post-1965 cohorts is consistent with the causal effect documented in Table A-17.