

CAF - WORKING PAPER #2021/19

First version: December 16, 2021

This version: August 5, 2026

Human Capital Consequences of Expanding the Supply of Higher Education Institutions in a Developing Country

**Noemí Katzkowicz¹ | Victor Lavy² | Martina Querejeta³
| Tatiana Rosá⁴**

¹Instituto de Economía, FCEA,
Universidad de la República.
noma.katzkowicz@fcea.edu.uy

²Department of Economics, Warwick
University and the Hebrew
University of Jerusalem.
victor.lavy@mail.huji.ac.il

³Instituto de Economía, FCEA,
Universidad de la República.
martina.querejeta@fcea.edu.uy

⁴CAF Development Bank for Latin
America and the Caribbean.
trosa@caf.com

Limited post-secondary infrastructure partly explains low higher education attainment in developing countries. This paper estimates the effects of a program that expanded university access by constructing new campuses across Uruguay. Exploiting temporal and geographic variation in the rollout, we combine a dynamic two-way fixed-effects design with comprehensive administrative records. Reducing the distance to the nearest campus increased university enrollment by 34 percent relative to baseline, with especially large gains among less-privileged, first-generation students. Effects extended up to 30 kilometers and were concentrated among students from less-educated families, reducing spatial inequality. The expansion also increased secondary-school enrollment and completion, an important result given that low secondary completion remains a major barrier to higher education, while leaving tertiary completion rates unchanged. A cost-effectiveness analysis indicates substantial private and social returns, at costs comparable to a scholarship program. Overall, place-based university expansion can cost-effectively reduce spatial inequality and promote intergenerational mobility in developing countries.

KEYWORDS

Spatial inequality, Education, Public policies, Intergenerational mobility, Developing countries

Small sections of text that are less than two paragraphs may be quoted without explicit permission as long as this document is acknowledged. Findings, interpretations and conclusions expressed in this publication are the sole responsibility of its author(s) and cannot be, in any way, attributed to CAF, its Executive Directors or the countries they represent. CAF does not guarantee the accuracy of the data included in this publication and is not, in any way, responsible for any consequences resulting from its use.

CAF - DOCUMENTO DE TRABAJO #2021/19**Primera versión: 16 de diciembre de 2021****Esta versión: 5 de agosto de 2026**

Consecuencias sobre el capital humano de la expansión de la oferta de instituciones de educación superior en un país en desarrollo

**Noemí Katzkowicz¹ | Victor Lavy² | Martina Querejeta³
| Tatiana Rosá⁴**

¹Instituto de Economía, FCEA,
Universidad de la República.
noma.katzkowicz@fcea.edu.uy

²Department of Economics, Warwick
University and the Hebrew
University of Jerusalem.
victor.lavy@mail.huji.ac.il

³Instituto de Economía, FCEA,
Universidad de la República.
martina.querejeta@fcea.edu.uy

⁴CAF Development Bank for Latin
America and the Caribbean.
trosa@caf.com

La limitada infraestructura de educación postsecundaria explica en parte los bajos niveles de educación superior en los países en desarrollo. Este artículo estima el efecto causal de un programa que expandió la educación superior construyendo nuevas sedes universitarias en Uruguay, explotando la variación temporal y geográfica de su implementación con un diseño dinámico de efectos fijos y registros administrativos exhaustivos. El programa aumentó la matrícula universitaria en 34 por ciento, con ganancias especialmente grandes entre estudiantes de primera generación y derrames que se extienden hasta 30 kilómetros de las nuevas sedes. También aumentó la matrícula y la finalización de la secundaria sin reducir la finalización terciaria, con retornos privados y sociales sustanciales y costos comparables a los de un programa de becas. La expansión territorial de la educación superior puede así reducir las desigualdades espaciales y fomentar la movilidad intergeneracional de manera costo-efectiva.

PALABRAS CLAVE

Desigualdad espacial, Educación, Políticas públicas, Movilidad intergeneracional, Países en desarrollo

Pequeñas secciones del texto, menores a dos párrafos, pueden ser citadas sin autorización explícita siempre que se cite el presente documento. Los resultados, interpretaciones y conclusiones expresados en esta publicación son de exclusiva responsabilidad de su(s) autor(es), y de ninguna manera pueden ser atribuidos a CAF, a los miembros de su Directorio Ejecutivo o a los países que ellos representan. CAF no garantiza la exactitud de los datos incluidos en esta publicación y no se hace responsable en ningún aspecto de las consecuencias que resulten de su utilización.

1 | INTRODUCTION

Despite high and rising returns to higher education in developing countries (Montenegro and Patrinos, 2014), tertiary attainment remains limited. In Latin America, only about 30 percent of individuals pursue higher education, compared to more than 60 percent in Europe and North America (UNESCO, 2024). A primary obstacle is the limited supply of institutions. Distance to universities is a key determinant of enrollment (Card, 1993), especially in contexts where institutions are concentrated in capitals and a few large cities. In Latin America, nearly half of the population lives more than 600 kilometers from a capital city, underscoring the magnitude of this spatial barrier. While governments expand access through scholarships for geographic mobility, relocating students to urban centers may exacerbate spatial inequalities.

This paper provides causal evidence on the role of place-based policies in decentralizing higher education. By establishing new campuses outside major metropolitan areas, governments can simultaneously increase enrollment and reduce regional disparities in human capital. Leveraging the staggered construction of campuses outside Uruguay's capital, we show that expanding local access fosters intergenerational mobility by weakening the link between parental background and educational attainment. We further document a "forward-linkage" effect as the mere presence of a local campus generates powerful incentives for students to complete secondary school, even before they reach university age. Although the effects are highly local—concentrated within 30 kilometers of campuses—complementary low-cost policies such as subsidized transportation could further enhance the program's impact, particularly for students from less-educated families. University expansion thus emerges as a powerful tool to address persistent inequalities in access to higher education across the developing world. This suggests that geographic supply is a binding constraint even in relatively small countries.

We exploit a unique policy context in Uruguay, a country where high returns to education and low university attainment exemplify the human capital gap. Two decades ago, fewer than 10 percent of individuals aged 18 or older had a university degree. The situation was especially stark outside the capital: approximately 12 percent of adults in Montevideo had attained tertiary education, compared to fewer than 4 percent in the rest of the country. These low levels of human capital resemble those observed in the United Kingdom in the 1970s and the United States in the early postwar period (Blackburn and Jarman, 1993). To reduce these spatial inequalities, the government established new campuses outside the capital. Public universities in Uruguay are tuition-free and non-selective. This open-access regime provides a rare laboratory for disentangling the effects of university supply from financial constraints and the selective role of secondary school quality, two confounding factors that are difficult to isolate in previous studies, which typically mask the true cost of distance in the literature.

The program began in 2008 and, by 2013, six new campuses had launched. The locations were chosen based on the existence of institutional infrastructure and the minimization of political opposition. This rollout process generated exogenous variation in timing and location, driven by institutional considerations rather than local socioeconomic trends. The university expansion effectively reduced the average distance to a university campus by 300 kilometers, more than half of the longest distance one can traverse in Uruguay (571 km). We estimate impact using a dynamic two-way fixed effects model (Borusyak et al., 2024), leveraging novel administrative data on enrollment at Uruguay's main public university, Universidad de la República, covering 2002–2020, which capture nearly 90 percent of the nation's tertiary students.

Our results proceed in six steps. First, we find that opening a new campus increased the

share of enrolled students among the population aged 18–30 by 0.37 percentage points during the first six years after the campus opening. This effect size is substantial, representing a 34 percent increase relative to the pre-treatment mean.

Second, we document a significant gain in intergenerational mobility. The program increased the share of first-generation university students by 40 percent, suggesting that geographic access is a more significant barrier for students from less-educated families.

Third, we show that the impact extends up to 30 kilometers from a campus, with the effect decaying with distance. Crucially, no significant effects are observed beyond this radius, suggesting that even modest distances remain a prohibitive obstacle for first-generation students.

Fourth, we refute concerns that expanding access dilutes academic quality as degree completion rates remained stable despite the influx of new students.

Fifth, we examine high school outcomes. The program led to an average of 1,082 additional students completing high school per year during the first six years after the campus opening. This suggests that the presence of a local university increases the option value of secondary schooling, effectively mitigating one of the primary barriers to tertiary education.

Finally, we estimate high social returns (18%) and private returns (26%) for the program, confirming that place-based university expansion is an efficient investment. These are similar to returns in other low-income countries (Psacharopoulos and Patrinos, 2018) and suggest that investments in higher education are efficient from both individual and societal perspectives.

This paper contributes to three strands of literature. First, we provide causal evidence on the effects of expanding the supply of higher education in a developing country. Prior work shows that compulsory schooling reforms and investments in school infrastructure increase educational attainment and related outcomes (Angrist and Krueger, 1991; Duflo, 2001; Currie and Moretti, 2003; Oreopoulos, 2006; Frenette, 2009; Alzúa et al., 2015; Akresh et al., 2023), and that distance to a university is an important barrier to enrollment (Card, 1993; Alm and Winters, 2009; Spiess and Wrohlich, 2010). However, most evidence comes from developed countries, and relatively little is known about which policies that reduce distance are cost effective in developing contexts or about their distributional implications. We contribute to the foundational work on distance-to-school (Card, 1993; Duflo, 2001) by isolating geographic supply from financial selection.

Second, we add to the literature on intergenerational mobility in developing countries. While a large body of work documents income mobility (Chetty et al., 2014; Jäntti and Jenkins, 2015; Chetty et al., 2020) and educational mobility (Black et al., 2005; Björklund and Salvanes, 2011; Fleury and Gilles, 2018), evidence from developing countries remains limited (Daude and Robano, 2015; Neidhöfer et al., 2018; Torche, 2019; Bautista et al., 2023). Moreover, the role of higher education policies as mediators of mobility is still underexplored. We show that expanding geographic access to university education increases mobility at the top of the educational distribution.

Finally, the paper provides evidence of broader developmental spillovers. Beyond university enrollment, the expansion increases high school attendance in treated localities, consistent with evidence that educational investments affect earlier educational choices (Jagnani and Khanna, 2020; Rands and Barsanetti, 2024). We demonstrate that supply-side investments can generate powerful "option value" spillovers for secondary education. These effects are identified in a setting with free and non-selective university access, where capacity constraints are not binding, highlighting that distance alone can be a significant barrier to higher education even in relatively small countries. This has implications for other developing economies where spatial constraints are likely to be even more severe.

The rest of the paper proceeds as follows. Section 2 describes the relevant institutional context, including details on the University's geographic expansion program. Section 3 provides details on the data and sample selection, and Section 4 presents the identification strategy. Section 5 presents the main results regarding the impact of the university campuses' expansion program, including heterogeneity analysis, and Section 6 discusses the spillover effects on high school attendance and completion. Section 7 provides evidence on the robustness of the main results, and Section 8 provides evidence on the private and social returns of investing in higher education. Finally, Section 9 concludes.

2 | INSTITUTIONAL SETTING

2.1 | An Overview of the Uruguayan Educational System

In Uruguay, compulsory basic education spans fourteen years, comprising two years of preschool, six years of primary education, and six years of secondary education. Upon completion, students may continue their formal studies at the tertiary or university level. Typically, students enter higher education at around 18 years of age. The country stands out in the Latin American context for its strong welfare state, characterized by universal access to free public education at all levels, including higher education, and further distinguished by the absence of admission exams.

Educational coverage at the compulsory levels has expanded substantially over the past decade. However, while coverage is nearly universal up to age 14, educational attainment declines markedly in upper secondary school (ages 15-17). Moreover, socioeconomic disparities remain pronounced: students from low socioeconomic backgrounds exhibit consistently lower attainment across nearly all educational levels (INEEd, 2018). As a result, despite relatively high attainment through lower secondary education compared to other Latin American countries, Uruguay's overall secondary school completion rate remains low—around 40 percent—compared to a regional average of 62 percent (INEEd, 2020). This low completion rate reflects persistent structural characteristics of the education system, including high rates of grade repetition, early labor market entry among students from low-income households, and strong intergenerational persistence in educational attainment. These patterns are consistent with broader evidence from Latin America, where intergenerational educational mobility has historically been limited and secondary school completion has remained a key bottleneck for upward mobility, strongly associated with parental education and socioeconomic status (Berniell et al., 2021). In this context, completing high school represents a non-trivial margin for a large share of the youth population.

Access to higher education is also limited: fewer than 20 percent of adults have completed any university-level education. This is particularly striking given the substantial labor market returns to tertiary education, with university graduates earning, on average, a 30 percent wage premium over individuals who enrolled but did not complete a degree (Amarante et al., 2025).

Uruguay's main university, Universidad de la República (Udelar), has historically been concentrated in the capital city, Montevideo. This geographic centralization, together with the low secondary school completion rates, contributed to persistently low levels of university attainment. This outcome is particularly striking given Uruguay's universal education system and overall socioeconomic development.

This study examines the effects of a university expansion program designed to reduce barriers to higher education access by addressing geographic inequalities.

2.2 | The University Expansion Program

In 2008, authorities at Udelar launched a major university expansion program to improve access to tertiary education outside Montevideo. Before the reform, public university education was highly centralized in the capital, creating large geographic barriers for students from other regions, especially those from disadvantaged backgrounds. The main objectives of the expansion were to reduce spatial inequality in access to higher education, lower student migration costs, and increase enrollment among students living outside the capital city. The program focused on establishing new regional campuses offering undergraduate degrees previously available only in Montevideo.

The expansion was not designed to target specific sectors or local labor market needs. Instead, campuses offered a broad range of degree programs similar to those in Montevideo, reflecting the broad inclusion goal of the decentralization policy.

By 2020, university campuses had been established in six of Uruguay's nineteen departments (Figure 1). Campus locations were chosen mainly based on existing infrastructure and the minimization of political opposition, rather than local socioeconomic trends or changes in educational demand.¹ No other major national policies during this period directly affected the geographic availability of university education or varied differentially across regions in ways that coincide with the timing and location of campus openings. This staggered rollout, therefore, provides plausibly exogenous variation in the timing and location of campus openings. Spatial variation was driven by institutional and political considerations rather than local socioeconomic characteristics, whereas temporal variation reflected administrative sequencing independent of local demand shocks. Assuming that contested and uncontested localities are comparable in unobserved outcome determinants—conditional on pre-expansion covariates—the staggered implementation approximates a natural experiment. Subsection 4.1 presents further evidence supporting the quasi-experimental nature of the reform.

The university's geographic expansion was followed by a marked increase in enrollment among students residing outside Montevideo.² Panel (a) of Figure 2 shows the evolution of first-year enrollment by students' region of residence prior to entering university. In 2002, students from outside the capital accounted for 44.5 percent of total entrants, and by 2020 this share had risen to 58 percent. The program also significantly reduced student migration to Montevideo. As shown in Panel (b) of Figure 2, the share of first-year students enrolling at campuses outside the capital increased from just 0.2 percent in 2002 to 15 percent in 2020. As a result, the university expansion operated primarily through the removal of distance-based frictions.

3 | DATA

3.1 | Individual-Level Administrative Records and Coverage

To implement our empirical design, we exploit rich individual-level administrative records from Uruguay's primary public higher education institution, the Universidad de la República

¹Before this structural reform, the interior departments of Paysandú, Salto, and Rivera possessed very limited, disconnected tertiary outposts. However, these regional locations were only formalized and consolidated as fully comprehensive university campuses in 2010, 2011, and 2013 through targeted national budgetary expansions and broader academic program offerings.

²Udelar enrolls approximately 86% of all university students, with the remaining 14% attending private institutions (Udelar, 2020).

(Udelar), covering the full census of student enrollment cohorts from 2002 to 2020.³ Our core dataset comprises exhaustive registry information on entering students integrated with microdata from the mandatory student census administered systematically during their initial year of university study. Additionally, we track academic trajectories using institutional completion registries available for cohorts graduating between 2006 and 2020. This merged database contains comprehensive demographic and socioeconomic profiles, including granular measures of parental educational attainment and pre-university geographic locations. We determine each student's baseline geographic locality of origin from the specific secondary school where they completed their final year of upper secondary education; this predetermined location proxy is non-missing and precisely validated for approximately 93 percent of our initial administrative sample, providing an exceptionally clean spatial identifier that is free from self-reporting or recall bias.

3.2 | Primary Variables and Operational Definitions

Our primary independent treatment variable is a time-varying binary indicator denoting the operational launch of a new regional campus within a given locality, constructed according to the institutional rollout timeline mapped in Figure 1. The key outcome variables driving our causal analysis are university enrollment, degree completion, and first-generation student (FGS) status. At the micro-level, the enrollment indicator equals 1 for all individuals in our administrative database (by definition, all Udelar matriculants). The completion indicator equals 1 if an enrolled student successfully earns a university degree within 5 years of initial enrollment, which corresponds to the standard theoretical program duration for most academic majors in Uruguay. Finally, we define the FGS indicator as a dummy variable equal to 1 for students who are the first in their immediate families to attend higher education, specifically those whose parents have no history of university enrollment.

3.3 | Locality-Level Aggregations

Because individual administrative registries naturally cover only those individuals who select into higher education, extensive-margin population effects cannot be estimated directly from raw individual files. We overcome this limitation by collapsing these individual-level indicators into cell means at the locality level, indexing units by the precise geographic locality of the secondary school from which the student graduated. A locality is defined as a census-delineated, sub-departmental geographic unit representing a population center utilized by Uruguay's National Institute of Statistics (Instituto Nacional de Estadística, INE) for enumeration and policy reporting. While Uruguay is administratively divided into 19 macro-departments, the 2011 Population Census identifies 615 distinct localities across the national territory.

Using these spatial units, we estimate the impact of the infrastructure expansion program on three core locality-level outcomes: the share of enrolled university students relative to the total youth population aged 18 to 30, the share of enrolled FGS relative to the total population aged 18 to 30, and the share of FGS within the aggregate local university enrollment pool.⁴

³Enrollment registries strictly exclude students who entered after the onset of the COVID-19 pandemic in March 2020 to prevent health shocks and remote learning mandates from confounding baseline structural enrollment trends.

⁴We focus on the population aged 18–30 because it aligns with the United Nations definition of youth and captures 91 percent of total first-year university enrollees in Uruguay. While students typically matriculate at ages 18 or 19 immediately following high school graduation, older individuals residing in peripheral departments face severe historical geographic frictions and frequently seize the opportunity to return to formal schooling when a new campus opens locally, as detailed in Online Appendix Figure A.1.

We interpret the latter two metrics as structural indices of absolute intergenerational mobility, as they isolate the shifting volume of individuals ascending the educational ladder relative to their parental baselines. By construction, these directional measures do not capture downward educational mobility, instead, they serve as explicit signals of upward mobility localized at the right tail of the human capital distribution.

To address potential endogeneity stemming from contemporaneous migration or population shocks within specific localities during the treatment period, we freeze our outcome denominators using historical population data fixed before policy implementation. Specifically, the outcome variables are normalized as a share of the total population aged 18 to 30 recorded in the 1996 Population Census, the last comprehensive national census conducted before the decentralization program began.

3.4 | Sample Restrictions and Sample Composition

Our primary estimation sample is constructed using several transparent filtering criteria designed to ensure data consistency and clean identification. First, we restrict the sample to individuals aged 18 to 30 at the exact time of university entry. Second, we include only those individuals who appear in the administrative enrollment registries and have fully completed the mandatory first-year student census. Third, we exclude any students whose baseline pre-university geographic origin cannot be definitively verified. Finally, we systematically omit students originating from the departments of Montevideo and Canelones. Montevideo is excluded because it contains the historic, centralized flagship campus of Udelar, while Canelones is dropped due to its immediate geographic proximity to the capital, which gives its residents radically different baseline transit access, labor market characteristics, and socioeconomic profiles. The populations in these two metropolitan areas exhibit markedly different baseline secondary schooling completion rates and transport options compared to the rest of the country. Applying these standard filters yields a final estimation sample of 64,820 students from 75 distinct localities, resulting in 1,225 locality-year observations over our panel window.

3.5 | Descriptive Statistics

Table 1 provides comprehensive descriptive statistics for our primary variables at both the individual and aggregated locality levels. In the sample, 63 percent of enrolled students are female, and the average age at matriculation is 19.02 years. Regarding parental backgrounds, 12 percent of entering students have parents with primary education as their highest completed level, 51 percent come from households with secondary education, 18 percent have parents with non-university tertiary vocational training, and only 18 percent have parents with university degrees.⁵ Consequently, a striking 82 percent of the student body is the first in their family to enroll in a university track. At the locality-level, the average annual enrollment per geographic unit is 112.93, representing an average enrollment share of 2% of the local youth population. On average, 89.38 of these local enrollees are first-generation students, establishing a mean FGS enrollment share of 80 percent (up from a pre-reform cross-sectional baseline of 73 percent). Finally, the raw locality-average probability of completing a university degree within five years of enrollment is 9 percent across the entire sample, a completion rate that remains identical when restricted to the FGS sub-sample.

⁵ Approximately 4 percent of students have missing registry data on paternal education, and 0.5 percent have missing data on maternal education. For these observations, the maximum completed level of the non-missing parent is utilized to define the household's educational background. No individual in our final estimation sample has missing educational data for both parents.

When extending the tracking window longitudinally to evaluate long-run graduation, the cumulative proportion of students who ever complete their university degree reaches 36 percent among cohorts entering between 2002 and 2015.⁶

This 36 percent baseline graduation rate must be conceptualized strictly as an on-time completion measure. Because Udelar operates under a completely tuition-free, open-access institutional design with no admission sorting exams, it successfully lowers entry barriers but mechanically increases baseline academic preparation variance. This open structure correlates with a high prevalence of stop-out behaviors, delayed re-entry intervals, and concurrent work-while-studying choices, all of which mechanically extend the observed time-to-degree. Consequently, while Uruguay's on-time completion rate appears low in absolute terms, it aligns closely with structural graduation dynamics documented in analogous open-access public tertiary systems across the developing world.⁷

3.6 | Secondary Data Inputs

To complement our core administrative files, we integrate independent microdata from Uruguay's National Household Survey (Encuesta Continua de Hogares, NHS), collected annually by the INE from 2002 to 2019.⁸ Incorporating these household-level survey panels allows us to evaluate the program's downstream consequences for broader educational margins, such as intermediate secondary school attainment, within a nationally representative framework, while providing an independent database for rigorous identification, household baseline balancing, and structural robustness tests.

4 | EMPIRICAL STRATEGY

4.1 | Identification Strategy

To identify the causal effects of expanding the supply of higher education institutions on university enrollment, degree completion, and the socioeconomic composition of the student body, we exploit the staggered temporal and geographic variation arising from the program's rollout across localities. The primary threat to our empirical design is the potential endogeneity of campus placement and timing. For instance, if authorities preferentially targeted regions with rising secular trends in educational demand or accelerating local economic development, standard two-way fixed effects (TWFE) estimates would be upwardly biased. Conversely, if campuses were strategically deployed as regional equalization tools in lagging or structurally disadvantaged areas, our estimates would be systematically biased downward.

As established by the institutional narrative in Section 2, the selection of treatment sites and the administrative sequencing of campus openings were governed by institutional constraints, political negotiations, and pre-existing local infrastructure, rather than local trajectories in higher-education demand or underlying socioeconomic shocks. Furthermore, the absence of concurrent nationwide shocks targeting the selected localities (e.g.,

⁶Entering student cohorts from 2016 onward are strictly censored from these long-run completion parameters because our administrative graduation data tracks outcomes only up to calendar year 2020.

⁷Low on-time completion rates are endemic across Latin America, where conditional graduation rates are modest and delays are common (Ferreyra et al., 2017). Similar patterns characterize comparable institutional regimes, such as in Chile, where academic persistence is low and graduation delays are substantial (von Hippel and Hofflinger, 2020; Moraga-Pumarino, 2023), with only 15 to 20 percent of students graduating within the formal program duration (OECD, cited in Moraga-Pumarino, 2023).

⁸The NHS is a nationally representative, cross-sectional survey. Full microdata, sampling weights, and institutional documentation are made public and accessible via the official INE portal through this [link](#).

no major infrastructure investments or labor market reforms coincided with the expansion) strengthens the causal interpretation. Assuming that localities are comparable in unobserved outcome-relevant characteristics, conditional on pre-expansion covariates, the staggered implementation of the program approximates a natural experiment.

To formally validate the plausibility of this exogeneity assumption, we evaluate whether baseline locality-level characteristics systematically predict the timing of campus inaugurations. Specifically, we regress baseline locality-level characteristics measured before the program's inception through Equation 2. Figure 3 provides evidence that treated and non-treated localities did not differ in pre-determined characteristics before the program implementation. These characteristics include local employment rates, average household income, and population composition measures. Table 2 provides a descriptive cross-sectional comparison by treatment status, never treated vs treated localities, and early vs late adopters among the treated localities. Results show no systematic differences across groups.⁹ This absence of systematic correlation strongly supports our core identifying assumption that the rollout approximates a quasi-natural experiment.

4.2 | Two-Way Fixed Effects (TWFE) Specification

We formalize this identification framework by estimating a generalized difference in differences specification that accommodates staggered treatment timing across a balanced panel of localities.¹⁰ The baseline econometric model is structured as follows:

$$Y_{l,t} = \alpha_0 + \gamma \text{Post}_{l,t} + \mu_l + \mu_t + \varepsilon_{l,t} \quad (1)$$

where $Y_{l,t}$ is the outcome of interest measured in locality l in calendar year t , such as the total university enrollment share, the first-generation enrollment share, or the proportion of first-generation students within the enrolled cohort. The independent variable of interest, $\text{Post}_{l,t}$, is a time-varying indicator that takes a value of one if a university campus has been established and is actively operating in locality l in or before year t , and zero otherwise. This is computed according to the timeline presented in Figure 1. The parameter μ_l represents a comprehensive vector of locality-fixed effects that absorb time-invariant regional unobservables, such as geographic distance to metropolitan markets, regional wealth distributions, and historical educational infrastructure baselines. The μ_t denotes a vector of year fixed effects that capture macro-level shocks, national business cycles, or universal educational policy adjustments that affect all localities symmetrically. Finally, $\varepsilon_{l,t}$ is the idiosyncratic error term.

We estimate the static TWFE model following the imputation approach of [Borusyak et al. \(2024\)](#). The coefficient of interest is γ , which identifies the average intent-to-treat (ITT) effect of the regional university expansion on our localized human capital outcomes. This can be viewed as a lower bound of the treatment effect. To account for potential serial correlation in unobserved economic or educational shocks within specific geographic areas over time, standard errors are conservatively clustered at the locality level across all specifications.

⁹While some differences are statistically significant, their magnitudes are small in economic terms. Moreover, many of these characteristics are largely time invariant. As a result, they are absorbed by locality fixed effects in our empirical specification and are unlikely to bias the estimated treatment effects. We find no significant differences in observable characteristics between early and late adopters, further supporting the comparability of treated units.

¹⁰See [Goodman-Bacon \(2021\)](#) for a review of difference in differences with variation in treatment timing.

4.3 | Dynamic Event-Study Design

To map the dynamic trajectory of the treatment effects over time, we extend our baseline TWFE specification into a flexible event-study framework following [Borusyak et al. \(2024\)](#) imputation method. The dynamic empirical model is specified as follows:

$$Y_{l,t} = \alpha_0 + \sum_{h=-a}^{h=b-1} \gamma_h 1[K_{l,t} = h] + \gamma_b 1[K_{l,t} = b] + \mu_l + \mu_t + \epsilon_{l,t} \quad (2)$$

where E_l is the year of the campus inauguration in treated locality l . $K_{l,t} = t - E_l$ is the number of periods since the event date E_l , namely the relative time variable in an event study design. Therefore, the indicator variables $1[K_{l,t} = h]$ represent event-time dummies tracking the number of years relative to the campus opening date, where $h = 0$ identifies the exact calendar year of inauguration. In this dynamic configuration, the coefficients γ_{-6} through γ_{-2} provide a direct, pre-treatment test for pre-trends. If our quasi-experimental identification strategy is valid, these coefficients should be statistically indistinguishable from zero. Conversely, the coefficients γ_0 through γ_6 trace out the dynamic post-treatment path of the campus openings. This allows us to observe whether the regional supply shock triggers an immediate, step-function adjustment in higher-education demand, or whether human capital gains accumulate gradually as successive cohorts spend a larger share of their secondary schooling careers in close proximity to a campus. Our main specification considers the period between six years before and six years after the program implementation, using a balanced sample of localities observed throughout the period (the last university campus opened in 2013).

Under this approach untreated potential outcomes are first estimated using not-yet-treated and never-treated observations, and treatment effects at each relative time are then computed as the difference between observed outcomes and the corresponding imputed untreated outcomes. Unlike in a conventional TWFE event-study regression, this approach does not require omitting one event-time indicator for normalization, so all event-time coefficients are reported directly relative to the imputed untreated baseline.

5 | THE EFFECT OF UNIVERSITY CAMPUSES' EXPANSION

5.1 | Distance and Enrollment

Before estimating the causal effects of campus openings, we first provide evidence that distance affects enrollment. To that end, we regress our main outcomes on locality and year fixed effects and a measure of distance between each locality and the nearest campus which varies by year.

Table 3 shows that an increase of 100 km to the campus reduces the share of enrolled students by 0.09 percentage points (pp). Results also indicate that the distance reduces the share of FGS in total enrollment by 1.62 pp. Consistent with [Card \(1993\)](#), our findings show that distance to the campus significantly explains variations in university enrollment and intergenerational mobility. This enhances the importance of estimating the effects of the Uruguayan program in reducing the distance to university campuses.

5.2 | University Enrollment and Intergenerational Mobility

Figure 4 and Table A.1 in the Appendix show the estimates of Equation 2 for the share of enrolled students relative to the population aged 18-30 during the first six years of the program (Panel a). The share of enrolled students in treated localities during the six years following the campus opening increased by 0.37 pp relative to the share in non-treated localities. However, the effect decreases in years 5 and 6 after campus opening. To check whether this decrease persists over time we run the same specification but including up to ten years after the program (Panel b). Results show that the decline reverses from year 7, implying that the program had a positive and significant effect on enrollment in the long run. It is worth recalling that our panel is balanced up to six years after the program implementation (the last campus opened in 2013), and so the effects using a ten-year window are less precisely estimated. For that reason, our main results are based on a six-year window around the opening of new campuses. At longer horizons, the estimates are based on an increasingly unbalanced sample as late-adopter localities are not included in the analysis. Therefore, the available data allow us to explore dynamic effects in the first years after campus openings, when identification is strongest. The average effect during the six years after the program (0.37 pp) is sizable, representing an increase in university enrollment by nearly 34% of the pre-treatment mean.

Figure 5 shows estimates of Equation 2. The share of enrolled FGS is also statistically significant and positive (Panel a). Opening a campus in a given locality increased the share of FGS in the population aged 18-30 by 0.32 pp. This result suggests that in localities where a campus opened, the share of individuals surpassing their parents' educational level increased significantly. The results also show that the program increased the share of FGS over total enrollment by 6.1 pp (Panel b). This suggests that the program also affected the enrollment composition, increasing the share of first-generation students among the total number of students. Overall, we provide evidence that the supply expansion program enhanced educational intergenerational mobility.

Finally, Figure 6 shows estimates of Equation 2 for the total number of enrolled students and total FGS. Results indicate an average increase of over 100 enrolled students in treated localities, of whom around 90 are the first generation in their families to enter University. Taken together, these results demonstrate how the program succeeded in increasing enrollment both in relative terms (shares) and in absolute numbers. Examining shares is important because it reveals changes in composition and access equity, while examining total numbers is crucial for assessing the overall expansion of opportunities created by the program. The fact that the increase is especially concentrated among first-generation students highlights the program's effectiveness in broadening access to higher education for historically underrepresented groups.

A potential concern is whether the estimated effects reflect net new enrollment or crowding-out from other post-secondary options. To assess this, we rely on household survey data and estimate two additional regressions that mirror our baseline specification. Results show that campus openings increase total university enrollment and do not decrease tertiary non-university enrollment, indicating net expansion rather than substitution (Figure A.2 in the Online Appendix).

Recent advances in the econometric literature demonstrate that in panel settings with staggered treatment timing and heterogeneous treatment effects across cohorts or over time, standard TWFE estimates can suffer from severe bias due to negative weights assigned to early-treated units. While we implement the imputation approach proposed by [Borusyak et al. \(2024\)](#), Table A.2 in the Online Appendix shows that our results are robust to the choice of difference in differences estimator and with different weighting schemes. The

doubly-robust estimator of Callaway and Sant'Anna (2021) and a standard TWFE/OLS regression with a single post-treatment dummy deliver very similar magnitudes and the same pattern of significance for the three outcomes.

The preceding results underscore the program's impact on disadvantaged students. Our findings provide compelling evidence that reducing the distance to the University can help overcome various barriers to university enrollment, such as financial constraints. This increases enrollment and positively enhances intergenerational mobility in educational attainment at the top of the educational distribution.

5.3 | Heterogeneous Effects

We next examine the heterogeneous effects on FGS enrollment by parental education. Figure 7 and Table A.3 in the Appendix show the estimates of Equation 2 for subgroups of students according to their parent's educational level: those whose parents have some primary education (Panel a), those whose parents have some secondary education (Panel b), and those whose parents have some tertiary education (Panel c). Overall, expanding university campuses had a larger effect on students from less educated families.

Panel (a) of Figure 7 shows that, in the first six years following the university expansion, the share of enrolled students whose parents had only some primary education increased by 0.067 pp. This represents a substantial effect, accounting for 67% of the pre-treatment mean enrollment share of 0.10% (Table A.3). Panel (b) shows the same pattern for students whose parents reached secondary education. After the program's implementation, the proportion of these students increased by 0.22 pp which represents 43% of the pre-treatment mean (0.51%). Finally, Panel (c) shows a smaller effect of 0.034 pp on the enrollment share of students with parents with tertiary non-university education. However, the effect is not statistically different from zero at a 10% significance level for most years following the program, which accounts for nearly 20% of the pre-treatment mean (0.18%).

Our results suggest that the program has heterogeneous effects depending on students' parental backgrounds. While the larger estimate is among students whose parents have secondary education, the higher effect relative to the pre-treatment mean is observed among students whose parents only achieved primary education. These results are consistent with the program's greater impact on students from more disadvantaged backgrounds. While we do not have information about the channels through which the program operated, this heterogeneity could be attributed to the fact that less educated parents may be less able to financially support their children while in college (an income effect mechanism) or may not have their children's higher education as a high priority (change in preferences mechanism).

5.4 | Regional Spatial Spillovers and Distance Dynamics

Although university campuses were physically established within specific host localities, the localized supply shock may generate geographic externalities across surrounding non-host regions. To systematically evaluate these potential spatial spillovers, we estimate an alternative empirical specification where treatment exposure is defined by a locality's geographic proximity to the nearest infrastructure site. Using geo-referenced data, we calculate the geodetic distance between each locality and the nearest new campus in each year.¹¹ We construct two alternative treatment radii to capture the geographic boundaries of this shock: an inner catchment zone within 30 km of a campus and an outer catchment zone

¹¹Geodetic distance is defined as the length of the shortest curve between two points along the surface of a mathematical model of the Earth. Distances are computed using the *Stata* command *geodist*, exploiting precise latitude and longitude coordinates for each unique locality and campus destination.

between 30 and 50 km, strictly excluding the campus-hosting locality itself to isolate pure spillovers. Figure 1 maps the treated and neighboring localities by their relative distance bands.

Table 4 reports the static TWFE estimates derived from this spatial measurement framework. The indicator variable $\text{Radius}_{0-30\text{km}}$ equals one for non-host localities situated within a 30 km radius of an active campus, while $\text{Radius}_{30-50\text{km}}$ equals one for localities located between 30 and 50 km away. The baseline omitted reference category comprises all remaining untreated localities nationwide. We evaluate these spatial treatments across our three primary locality-level outcomes: the total university enrollment share relative to the baseline population aged 18–30 (Column 1), the share of first-generation students (FGS) relative to the population aged 18–30 (Column 2), and the relative share of FGS within the total enrolled student cohort (Column 3).

The empirical results reveal statistically significant, positive spatial spillovers within the inner 30 km radius, driven by simultaneous expansions in both total and first-generation university enrollment relative to the population aged 18–30 (Columns 1 and 2). For these two outcomes, the point estimates within the 30 km buffer are very similar in magnitude to the baseline effects estimated for the actual campus-hosting localities, and the differences are not statistically significant. A clear distance-decay pattern emerges, however, for the share of FGS within total enrollment (Column 3): while the host-locality effect is large and highly significant (6.2 pp), the corresponding spillover coefficient for non-host localities within 30 km is roughly an order of magnitude smaller and statistically indistinguishable from zero. Finally, we detect no statistically significant effects within the outer 30–50 km geographic zone for any of the three outcomes, indicating that the human capital spillovers of the regional campuses fade out beyond a 30 km commuting threshold.

To confirm the structural validity of this spatial design, we estimate corresponding dynamic event-study specifications under these alternative proximity-based treatment definitions. Reassuringly, the pre-treatment coefficients across all event leads are small and statistically indistinguishable from zero, validating the parallel-trends assumption in the spatial counterfactual pool. The post-treatment event trajectories mirror the gradual expansion documented in our baseline results, with the corresponding estimates and plots detailed in Figures A.3–A.5 and Tables A.4–A.6 of the Online Appendix.

Taken together, the attenuation of the treatment effect among first-generation students and the statistical disappearance of the program's impact on the relative share of FGS within total enrollment carry important socioeconomic equity implications. These patterns demonstrate that students whose parents lack university experience face higher distance-elasticities of demand than their peers from more advantaged backgrounds, whose families are already integrated into the higher education system.

While several unobserved structural factors may differentiate campus-hosting municipalities from surrounding rural catchments, such as population density or local economic opportunities, physical commuting time emerges as the primary friction. Although a 30 km radius represents a short geographic distance, deficient public transportation infrastructure and limited inter-city transit schedules across Uruguay's interior departments amplify travel costs. Consequently, even modest geographic separation poses a steep barrier for marginal first-generation students, highlighting the critical role that complementary, low-cost policy interventions, such as targeted transit subsidies or localized student housing, could play in reducing residual friction for nearby peripheral communities.

A remaining potential threat to our spatial identification is that the infrastructure expansion might trigger general-equilibrium spillovers, such as selective demographic sorting or migratory responses across localities, which could contaminate both our spillover parameters and our baseline host estimates. To formally test this migratory channel, we leverage

nationally representative data from the National Household Survey (NHS) to estimate the program's causal effect on the share of the youth population aged 16–30 residing in each locality. We find no statistically significant changes in local population shares or demographic composition following campus openings (see Appendix Figure A.6). This flat demographic response confirms that our human capital findings are not driven by the physical sorting or reallocation of young individuals across geographic lines, but rather reflect genuine behavioral adjustments in higher education demand among existing local residents. Reassuringly, these findings hold identically when restricting the analysis to the primary 18–30 cohort window. Taken together, these robustness checks rule out general equilibrium migration distortions across localities.

5.5 | University Completion Rate

Because Udelar operates under a completely tuition-free, open-access institutional model with no entry-sorting exams, a major potential concern is that universalizing access might come at the cost of lowering aggregate academic standards, leading to a decline in university completion rates. To rigorously examine whether the influx of marginal students compromised graduation efficiency, we analyze the program's impact on the probability that an enrolled student completes their degree within a five-year window. We restrict this analysis sample to student cohorts who enrolled between 2002 and 2015, ensuring a minimum five-year tracking window prior to our final administrative data cutoff in 2020.

Figure 8 and Appendix Table A.7 present the estimates derived from Equation 2 across two completion outcomes: the graduation rate within five years among the universe of enrolled students (Panel a), and the corresponding completion rate restricted to first-generation students (Panel b). In this framework, the event-time variable is anchored to the individual's enrollment year, such that event year $\tau = 0$ captures the five-year completion rate of cohorts who entered the university during the exact calendar year the local campus opened. These results prove that the infrastructure expansion had zero negative effect on university completion rates.¹² Across both panels of Figure 8, the dynamic event-study coefficients are tightly clustered around zero and statistically insignificant, indicating that the universalization of access did not trigger the unintended adverse consequence of diluting graduation efficiency or inflating dropout rates.

Finally, we exploit the underlying individual-level microdata to estimate micro-level regressions using the exact specification detailed in our aggregate models. We run these individual regressions across two primary student milestones: the conditional probability of being a first-generation student and the conditional probability of degree completion within five years. The micro-level results, presented in Appendix Figure A.8 and Table A.8, confirm that the program significantly increased the probability that an enrolled student originates from a first-generation background. Crucially, while the conditional graduation rate remains statistically unchanged, this pattern does not imply that the policy has no effect on human capital completion for marginal students. Because these marginal individuals would never have entered the university system in the absence of a regional campus opening, the policy increases their absolute probability of obtaining a tertiary degree from zero to the stable baseline average for completion, representing a substantial net increase in aggregate human capital accumulation.

¹²As this analysis relies on a different sample, we provide evidence that the university enrollment results estimated from this restricted sample are free of sample selection bias. Figure A.7 in the Appendix shows that the enrollment and intergenerational mobility results remained unchanged when considering this reduced sample.

6 | SUPPLY-SIDE SPILLOVERS ON SECONDARY SCHOOL ATTAINMENT

Thus far, the empirical evidence demonstrates that the university expansion program has exerted a powerful positive impact on higher education enrollment, particularly for first-generation students, without triggering a corresponding decline in graduation efficiency. Beyond these direct tertiary effects, however, mitigating spatial barriers can generate profound backward-linkage spillovers that alter human capital investments earlier in the educational pipeline. According to classic human capital theory, an individual's sequential decision to invest in intermediate schooling levels depends intrinsically on the expected future returns and option value of subsequent academic stages. By establishing local campuses and expanding the geographic availability of higher education, the policy reduces future migration and opportunity costs, thereby elevating the perceived benefits of clearing lower-level milestones. Consequently, we may expect the program to influence lower-level educational attainment by enhancing the net option value of secondary schooling, given the newly available opportunities for continuing formal studies. Stated alternatively, this spatial shock can alter the forward-looking incentives of younger cohorts, inducing behavioral adjustments that encourage them to remain in school and complete secondary education (see, e.g., [Lavy \(1996\)](#), [Lincove \(2009\)](#), and [Mukhopadhyay and Sahoo \(2016\)](#)).

To formally evaluate these supply-side spillovers, we leverage nationally representative microdata from the NHS covering the period 2002–2019. This database allows us to evaluate educational attainment and enrollment indicators for cohorts that were young enough to be exposed to a local campus during their high school years. We restrict our primary analysis sample to the young population and construct two key locality-level outcome metrics: the secondary school attendance for individuals aged 15 to 18 and completion for individuals aged 18 to 30.¹³

Figure 9 and Table A.9 in the Appendix show the dynamic event-study estimates derived from Equation 2. Panel (a) plots the event trajectory for secondary enrollment, while Panel (b) traces the trajectory for upper secondary completion. Reassuringly, for both outcomes, most pre-treatment coefficients across the event leads are tightly bound around zero and statistically insignificant.¹⁴ This flat pre-trend profile confirms that the expansion program did not selectively target localities already experiencing secular upward trajectories in high school graduation rates, thereby validating the core identification assumption. Following the campus opening, the post-treatment event coefficients reveal a gradual, monotonic increase in secondary attainment that mirrors the long-run development patterns documented for our primary university enrollment outcomes.

The estimated effects indicate an increase in both secondary school margins in localities where new campuses opened. Almost one thousand (992) additional high school students enrolled in high school per year during the first six years after the program's implementation. This implies an increase of around 25% relative to the pre-treatment mean. Furthermore, the opening of a local campus increases high school graduation, and the structural effect size indicates that the regional infrastructure program produced, on average, approximately 1,082 additional high school graduates per year during the first six years following reform

¹³The selection of these age groups is based on the fact that ages 15 to 18 represent the majority of upper high school students. In contrast individuals aged 18 to 30 are expected to have already completed high school.

¹⁴Although only the coefficient in the period immediately before treatment is statistically significant for secondary attendance, while none are significant for completion, a pattern consistent with anticipation effects, we acknowledge that standard pre-trend tests may have limited power. We therefore complement the event-study evidence with the [Rambachan and Roth \(2023\)](#) honestdid sensitivity analysis, which shows that the positive effect at event time 0 is robust to moderate departures from parallel trends, specifically to deviations up to one-half of the largest pre-treatment deviation.

implementation.

The effects on high school attendance and completion should not be interpreted as implying that campus openings are the sole driver of changes in secondary school outcomes. Rather, the expansion of local university supply likely operates as one mechanism among others, increasing the perceived returns to completing secondary education and reducing geographic and informational barriers, within a context characterized by low baseline completion rates and strong socioeconomic constraints. These backward spillovers carry profound implications for public policy design as low secondary completion has historically served as the primary structural bottleneck restricting higher education access and upward mobility in Uruguay (Berniell et al., 2021). Because Udelar operates under a completely tuition-free, open-access regime, high school graduation represents the single most critical institutional filter remaining in the educational architecture. By inflating the forward-looking option value of schooling and encouraging marginal students to graduate high school, the place-based program structurally expands the underlying pool of university-eligible youth.¹⁵

7 | ROBUSTNESS CHECKS: ALTERNATIVE IDENTIFICATION STRATEGIES

7.1 | Difference in Differences Approach

To ensure that our primary empirical findings are structurally robust and not driven by idiosyncratic modeling assumptions or baseline specification choices, we implement a series of alternative identification strategies and sensitivity checks. First, we depart from our baseline aggregate locality-level dataset to leverage individual-level microdata drawn from the National Household Survey (NHS) spanning the cross-sectional years 2018 to 2019. This alternative micro-level sample allows us to construct a distinct treatment framework based on precise, birth-cohort-level policy exposure. Specifically, we define treated or exposed cohorts as individuals aged 14 to 18 in the calendar year a university campus opened in their respective locality. We contrast these individuals against a control group of non-exposed cohorts, defined as older individuals aged 30 to 35 residing in the same geographic boundaries who had already completed their typical human capital accumulation life-cycles before the reform took place.¹⁶

We provide evidence on the robustness of the main results based on an alternative identification strategy: a difference in differences strategy widely implemented in similar studies (Duflo, 2001). Specifically, we estimate the program's causal effect by comparing differences between exposed and non-exposed cohorts across treated and non-treated localities, formalized in the following equation:

$$Y_{i,l,t} = \beta_0 + \gamma(\text{Exposed}_i * \text{Campus}_{l,t}) + \beta_1 X_i + \mu_l + \mu_{a(i)} + \mu_t + \epsilon_{i,l,t} \quad (3)$$

where $Y_{i,l,t}$ are the educational outcomes of interest (the number of years of education and the probabilities of having high, middle, or low level education), $\text{Exposed} * \text{Campus}$ is the

¹⁵While analyzing spillovers on labor market outcomes would also be informative, data limitations restrict such analysis. The available data only allow us to study short- to medium-run enrollment responses, which are insufficient to observe individuals from enrollment to graduation and transition to the labor market. However, as discussed in Section 8, while crucial, the labor market returns represent only one dimension of the broader welfare implications of expanded access to higher education.

¹⁶Table A.10 in the Appendix compares key outcomes and characteristics across groups. Preliminary comparisons suggest that the program improved educational outcomes for exposed cohorts in treated localities, while having no significant impact on demographics.

interaction of the two dummy variables, Exposed that takes value 1 if the individual i is aged between 14 and 18 the year the campus opened or in 2008 if a campus never opened in locality l , and 0 if aged between 30 and 35 at that time, and Campus that takes value one if there is a campus in locality l in year t and 0 otherwise. X_i is a set of variables that account for individual characteristics, including gender, marital status, number of children under 12 years old, and years of experience. Then μ_l , $\mu_{a(i)}$, and μ_t stand for a locality, individual age at the time the campus opened, and calendar year fixed effects. Finally, $\epsilon_{i,l,t}$ represents the model's idiosyncratic shock.

Table 5 reports the difference in differences parameters derived from this alternative cohort specification. As shown in Column 1, geographic exposure during formative high school years significantly increases cumulative educational attainment by 0.443 years of schooling. To explore where along the distribution this human capital accumulation is concentrated, Columns 2 through 4 decompose educational attainment into discrete binary indicators. Column 2 documents a statistically significant 3.2 percentage-point increase in the probability of achieving some higher-education milestones (defined as completing more than 12 years of formal schooling). Concurrently, Column 4 reveals a significant 4.5 percentage-point reduction in the probability of remaining in the low-education bracket (defined as 6 or fewer years of schooling). The margin for middle education (7 to 12 years of schooling) reported in Column 3 is small and statistically insignificant, a pattern consistent with marginal individuals shifting sequentially out of the lowest educational categories and directly into advanced tertiary tracks.¹⁷

7.2 | Synthetic Difference in Differences Approach

Second, we address potential methodological concerns regarding the underlying validity of our staggered TWFE imputation model by replicating our main enrollment results using the Synthetic Difference in Differences (SDID) framework (Arkhangelsky et al., 2021). While the imputation approach of Borusyak et al. (2024) effectively immunizes our baseline coefficients against negative weighting traps and heterogeneous treatment dynamics, the SDID framework provides a critical layer of verification by computing cohort-specific DiD estimates and then calculating a weighted average of the treatment effects (as described by Clarke et al. (2024)), thereby combining the dual advantages of synthetic control matching and traditional difference in differences.

Table 6 presents a direct side-by-side comparison between our baseline static TWFE estimates and the Staggered SDID parameters. The empirical parameters demonstrate remarkable stability across both econometric approaches. Under our baseline TWFE setup (Borusyak et al., 2024), the average program effect on the university enrollment share is 0.37 percentage points and on the the first-generation student enrollment share is 0.32 percentage points. When executing the SDID approach, the estimated impact on university enrollment remains highly stable and precisely estimated at 0.31 percentage points and the corresponding impact for the FGS is 0.27 percentage points. The alignment in both economic magnitude and statistical significance between these two distinct estimation frameworks strongly confirms that our primary causal conclusions are not artifacts of specific weighting configurations and reinforces the validity of our main results.

¹⁷Expanding the exposed cohorts as those aged 14-22 yields smaller effects only on some higher education consistent with older cohorts having already made their post-secondary education decisions before the campus expansion took place; and a placebo exercise comparing cohorts aged 30-35 to those aged 36-40 shows no significant effects. These results support the validity of the treatment definition and the identification strategy.

8 | ECONOMIC RETURNS TO HIGHER EDUCATION

Since the emergence of the human capital theory, the empirical literature in the economics of education has focused on estimating the returns to education as a means to understand educational investment decisions (see [Psacharopoulos and Patrinos \(2018\)](#) for a review). These returns can be assessed from at least two distinct perspectives, depending on who is evaluating the policy and making the decisions ([Psacharopoulos, 2024](#)). In a cost-benefit analysis framework, the private perspective is useful for showing why students might or might not invest on their own. Still, the social perspective is the one governments use to decide if subsidizing the policy makes sense.

The private returns to education compare the costs and benefits of education from the individual's perspective. The latter includes the increase in wages induced by obtaining an additional level of education, while the costs include tuition fees, transportation costs, other direct private expenditures, and forgone wages while studying.

The social returns to education adopt the point of view of the society as a whole and thus compare the costs and benefits of education, including also public resources invested in education within the cost, such as infrastructure, maintenance, and personnel expenses; and externalities and indirect effects, such as improved health, reduced crime, among other positive micro- and macroeconomic effects as social benefits.

We estimate both the private and social rates of return to investment in education following the so-called short-cut method first proposed by [Psacharopoulos and Mattson \(1998\)](#). We rely on information on wages of 18-65 year-old adults by level of education from the National Household Survey and budget information from the Ministry of Economy as reported in Table A.11 in the Appendix, and estimate the following equations:

$$r_{\text{private}} = \frac{\bar{W}_u - \bar{W}_s}{5 * \bar{W}_s} \quad (4)$$

$$r_{\text{social}} = \frac{\bar{W}_u - \bar{W}_s}{5 * (\bar{W}_s + C_u)} \quad (5)$$

where r is the rate of return to higher education, $\bar{W}_u$ and $\bar{W}_s$ are the annual average net wages of individuals who completed university or secondary schooling, respectively, and C_u are the annual average public resources invested in higher education per student (including infrastructure, maintenance, and personnel expenses, among other direct costs). Note that the term $5 * \bar{W}_s$ included in both equations refers to the forgone wages while studying, assuming a 5-year degree duration. In addition, students typically incur private out-of-pocket expenses (such as books, materials, and transport). We abstract from these since they are relatively modest compared to forgone earnings and public expenditures. Note also that we are not considering the social benefits thus likely underestimating the social returns to higher education.

Table A.11 in the Appendix shows estimates of the private and social returns to education in Uruguay. The economic returns of higher education are positive both from the private (26.0%) and social (18.6%) perspectives. Results from a Mincer regression of the return to university schooling provide similar estimates. These rates of return are high by international comparison and similar to the pattern observed for low-income countries ([Psacharopoulos and Patrinos, 2018](#)).¹⁸ Taken together, these positive cost-benefit ratios

¹⁸When focusing on geographical regions affected by the program, the private and social returns to higher education estimates remain of similar magnitude (25.5% and 21.5% respectively). Moreover, if taking a more conservative approach considering the annual average public resources invested in higher education per graduate—instead of per student—, then the social return to higher education is 11%.

suggest that investment in higher education is an efficient decision from both the individual and social point of view. Moreover, this considering we are likely overestimating the social costs as not considering the fiscal returns of increased wages, and underestimating the social benefits as overlooking the indirect effects of the program in improved health (Oreopoulos, 2007; Grossman, 2006; Currie and Moretti, 2003), reduced crime (Wang et al., 2022; Brugård and Falch, 2013), improved local development (Lehnert et al., 2024), and citizenship (Milligan et al., 2004), as well as the dynamic effects of increasing human capital (Oreopoulos, 2007).

Finally, we benchmark the cost of university expansion and decentralization against alternative policies aimed at increasing enrollment and promoting social mobility. When compared to a scholarship program covering the full duration of a university degree, the monetary costs are broadly similar. Providing financial support through the scholarship program amounts to roughly USD 4,000 per student per year—comparable to the average cost of pursuing a degree outside the main capital. However, studying in the capital entails additional social and economic burdens, including relocation, housing, transportation, and the personal costs of moving away from one's family and community. Moreover, relocating students to congested and expensive urban centers through scholarships may exacerbate rather than reduce spatial inequalities. In this context, investing in local higher education infrastructure emerges as a cost-effective strategy relative to other policies designed to expand university access.

This analysis is based on interpreting a correlation as returns to education. Avenues for future research include examining the labor market effects of expanded access to higher education, which we are unable to assess due to data limitations and the relatively short post-expansion window in our data. Understanding these impacts would be important for evaluating the full welfare implications of increased enrollment. However, potential benefits may also arise through other channels such as health, crime, civic participation, and local development.

9 | CONCLUSIONS

This paper evaluates the human capital consequences of a large-scale, place-based policy that expanded public tertiary education infrastructure beyond the capital city of Uruguay, Montevideo. By exploiting the staggered temporal and geographic rollout of new regional university campuses, we isolate the causal impacts of mitigating spatial barriers to higher education access. Our empirical framework relies on rich administrative student registries and household survey data, providing an exceptionally comprehensive look at how supply-side expansions reshape educational investments in a developing country context.

Our main empirical results deliver three core insights. First, reducing geographic distance frictions generates a profound expansion in higher education demand and attainment. Exposure to a regional campus increases university enrollment and educational mobility. Crucially, this enrollment spike did not translate into declining completion rates in tertiary education, demonstrating that localized expansions can attract students capable of completing tertiary milestones without expanding aggregate dropout pools or compromising degree match quality.

Second, the structural consequences of expanding tertiary infrastructure extend backward along the educational pipeline, generating substantial positive human capital spillovers. Proximity to a local university increases secondary school enrollment and upper secondary school completion among younger cohorts. This behavioral response confirms that expanding localized access alters the perceived option value of schooling, creating forward-looking

incentives that encourage youth to clear lower-level academic thresholds. This is crucial as low high school completion rates have been identified as a bottleneck towards greater intergenerational educational mobility.

Third, these human capital gains yield substantial private and social returns, suggesting that university expansion is an efficient investment. The consistently positive cost-benefit analysis suggests that investing in higher education expansion yields substantial net gains, for individuals and society.

From a policy perspective, these findings offer critical lessons for design frameworks in developing countries, where low higher education attainment and high spatial inequality frequently coexist. Our results indicate that supply-side constraints and geographic distance frictions operate as severe barriers to opportunity, particularly for disadvantaged, first-generation students who face high migration and opportunity costs.

Although Uruguay's context of free, non-selective higher education with unconstrained capacity differs from settings with binding admission quotas or tuition fees, the results speak to broader challenges in the developing world, where distance and spatial inequality remain major obstacles to educational attainment. While the analysis focuses on a specific policy in Uruguay, the findings address broader questions regarding the role of geographic access in higher education. That said, we do not claim that the magnitude of the effects estimated here necessarily generalizes to other institutional settings. Differences in baseline enrollment rates, labor market opportunities, and higher education systems are likely to shape the impact of similar policies elsewhere.

Crucially, our analysis reveals that place-based infrastructure investments can complement standard demand-side interventions, such as student scholarships or cash transfers. By simultaneously raising secondary school completion rates and university attainment, localized institutional expansions can serve as a powerful tool to break the intergenerational persistence of poverty and foster balanced regional economic growth. Future research should continue to track these cohorts to evaluate long-run intergenerational human capital spillovers and broader structural transformations in peripheral labor markets.

ACKNOWLEDGEMENTS

This research project was carried out with financial and scientific support from CAF. We thank Lucila Berniell and all participants in the Academic Workshop "RED 2022: Intergenerational mobility in Latin America" for excellent technical suggestions and comments. We gratefully acknowledge the central offices at Universidad de la República for providing us with the data and valuable information on the program, comments, and suggestions. We also thank participants at seminars in Instituto de Economía at Universidad de la República in Uruguay, Pontificia Universidad Católica de Chile, Universidad Alberto Hurtado and Universidad de Chile, EAFIT University in Colombia, Center for Human Development Studies of the Universidad de San Andrés in Argentina, WIDER development annual conference in Colombia, the Annual Meeting of the Chilean Economic Society (SECHI) in Chile, the Bristol 2024 Conference on Higher Education, the LACEA 2024 conference in Uruguay, and the 11th Conference of the Society for the Study of Economic Inequality (ECINEQ) for useful comments and suggestions.

REFERENCES

- Akresh, R., Halim, D. and Kleemans, M. (2023) Long-term and intergenerational effects of education: Evidence from school construction in Indonesia. *The Economic Journal*, **133**, 582–612.
- Alm, J. and Winters, J. V. (2009) Distance and Intrastate College Student Migration. *Economics of Education Review*, **28**, 728–738. URL: <https://www.sciencedirect.com/science/article/pii/S0272775709000922>.
- Alzúa, M. L., Gasparini, L. and Haimovich, F. (2015) Education Reform and Labor Market Outcomes: The Case of Argentina's Ley Federal De Educación. *Journal of Applied Economics*, **18**, 21–43. URL: <https://www.sciencedirect.com/science/article/pii/S1514032615300027>.
- Amarante, V., Carrasco, P., Echeberria, M. E., Katzkowicz, N. and Querejeta, M. (2025) Retornos económicos a la formación universitaria: un análisis de ingresos laborales en Uruguay. *Tech. rep.*, Instituto de Economía.
- Angrist, J. D. and Krueger, A. (1991) Does Compulsory School Attendance Affect Schooling and Earnings? *The Quarterly Journal of Economics*, **106**, 979–1014.
- Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W. and Wager, S. (2021) Synthetic difference-in-differences. *American Economic Review*, **111**, 4088–4118.
- Bautista, M. A., González, F., Martínez, L. R., Muñoz, P. and Prem, M. (2023) The Intergenerational Transmission of Higher Education: Evidence from The 1973 Coup in Chile. *Explorations in Economic History*, **90**, 101540.
- Berniell, L., Bonavida, C., de la Mata, D. and Schargrodsy, E. (2021) La Movilidad Educativa Intergeneracional en el Siglo XX en América Latina y el Caribe. *CAF Working Paper Series*.
- Björklund, A. and Salvanes, K. G. (2011) Education and Family Background: Mechanisms and Policies. vol. 3 of *Handbook of the Economics of Education*, chap. 3, 201–247. Elsevier. URL: <https://www.sciencedirect.com/science/article/pii/B978044453429300003X>.
- Black, S. E., Devereux, P. J. and Salvanes, K. G. (2005) Why the Apple Doesn't Fall Far: Understanding Intergenerational Transmission of Human Capital. *American economic review*, **95**, 437–449.
- Blackburn, R. M. and Jarman, J. (1993) Changing inequalities in access to British universities. *Oxford review of education*, **19**, 197–215.
- Borusyak, K., Jaravel, X. and Spiess, J. (2024) Revisiting event-study designs: Robust and efficient estimation. *The Review of Economic Studies*, **91**, 3253–3285.
- Brugård, K. H. and Falch, T. (2013) Post-compulsory education and imprisonment. *Labour Economics*, **23**, 97–106. URL: <https://www.sciencedirect.com/science/article/pii/S0927537113000572>.
- Callaway, B. and Sant'Anna, P. H. (2021) Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, **225**, 200–230.
- Card, D. (1993) Using Geographic Variation in College Proximity to Estimate the Return to Schooling. *NBER Working Papers 4483*, National Bureau of Economic Research.
- Chetty, R., Friedman, J. N., Saez, E., Turner, N. and Yagan, D. (2020) Income Segregation and Intergenerational Mobility Across Colleges in the United States. *The Quarterly Journal of Economics*, **135**, 1567–1633.
- Chetty, R., Hendren, N., Kline, P. and Saez, E. (2014) Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States. *The Quarterly Journal of Economics*, **129**, 1553–1623.

- Clarke, D., Pailaño, D., Athey, S. and Imbens, G. (2024) On Synthetic Difference-in-Differences and Related Estimation Methods in Stata. *Stata Journal*, **24**, 557–598. URL: <https://www.sciencedirect.com/science/article/pii/S0272775709001198>.
- Currie, J. and Moretti, E. (2003) Mother's Education and the Intergenerational Transmission of Human Capital: Evidence from College Openings. *The Quarterly journal of economics*, **118**, 1495–1532.
- Daude, C. and Robano, V. (2015) On Intergenerational (im)mobility in Latin America. *Latin American Economic Review*, **24**, 115–135. URL: <https://doi.org/10.1007/s40503-015-0030-x>.
- Duflo, E. (2001) Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment. *The American Economic Review*, **91**, 795–813.
- Ferreira, M. M., Avitabile, C., Botero Alvarez, J., Haimovich, F. and Urzua, S. (2017) *At a Crossroads: Higher Education in Latin America and the Caribbean*. Washington, DC: World Bank.
- Fleury, N. and Gilles, F. (2018) The Intergenerational Transmission of Education. A Meta-Regression Analysis. *Education Economics*, **26**, 557–573. URL: <https://doi.org/10.1080/09645292.2018.1517863>.
- Frenette, M. (2009) Do Universities Benefit Local Youth? Evidence from the Creation of New Universities. *Economics of Education Review*, **28**, 318–328. URL: <https://www.sciencedirect.com/science/article/pii/S0272775708000927>.
- Goodman-Bacon, A. (2021) Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics*, **225**, 254–277. URL: <https://www.sciencedirect.com/science/article/pii/S0304407621001445>.
- Grossman, M. (2006) Chapter 10 Education and Nonmarket Outcomes. vol. 1 of *Handbook of the Economics of Education*, 577–633. Elsevier. URL: <https://www.sciencedirect.com/science/article/pii/S1574069206010105>.
- von Hippel, P. T. and Hofflinger, A. (2020) The Data Generating Process of College Enrollment and Graduation: Evidence from Chile. *AERA Open*, **6**, 1–15.
- INEEd (2018) Reporte del Mirador Educativo 1. Desigualdades en el Acceso a la Educación Obligatoria. *Tech. rep.*, INEEd, Montevideo.
- (2020) Reporte del Mirador Educativo 6. 40 años de Egreso de la Educación Media en Uruguay. *Tech. rep.*, INEEd, Montevideo.
- Jagnani, M. and Khanna, G. (2020) The effects of elite public colleges on primary and secondary schooling markets in India. *Journal of Development Economics*, **146**, 102512. URL: <https://www.sciencedirect.com/science/article/pii/S0304387820300870>.
- Jäntti, M. and Jenkins, S. P. (2015) Income Mobility. In *Handbook of income distribution*, vol. 2, 807–935. Elsevier.
- Lavy, V. (1996) School Supply Constraints and Children's Educational Outcomes in Rural Ghana. *Journal of Development Economics*, **51**, 291–314. URL: <https://www.sciencedirect.com/science/article/pii/S0304387896004166>.
- Lehnert, P., Dell, M., Backes-Gellner, U. and Bettinger, E. (2024) The Effect of Postsecondary Educational Institutions on Local Economies: A Bird's-Eye View. *NBER Working Papers* 32679, National Bureau of Economic Research.
- Lincove, J. A. (2009) Determinants of Schooling for Boys and Girls in Nigeria Under a Policy of Free Primary Education. *Economics of Education Review*, **28**, 474–484.

- Milligan, K., Moretti, E. and Oreopoulos, P. (2004) Does Education Improve Citizenship? Evidence from the United States and the United Kingdom. *Journal of Public Economics*, **88**, 1667–1695. URL: <https://www.sciencedirect.com/science/article/pii/S0047272703002056>.
- Montenegro, C. E. and Patrinos, H. A. (2014) Comparable Estimates of Returns to Schooling Around the World. *World Bank policy research working paper*.
- Moraga-Pumarino, F. (2023) Profiles of university students who graduate on time. *Behavioral Sciences*, **13**, 582.
- Mukhopadhyay, A. and Sahoo, S. (2016) Does Access to Secondary Education Affect Primary Schooling? Evidence from India. *Economics of Education Review*, **54**, 124–142.
- Neidhöfer, G., Serrano, J. and Gasparini, L. (2018) Educational Inequality and Intergenerational Mobility in Latin America: A new Database. *Journal of Development Economics*, **134**, 329–349. URL: <http://www.sciencedirect.com/science/article/pii/S0304387818306618>.
- Oreopoulos, P. (2006) Estimating Average and Local Average Treatment Effects of Education when Compulsory Schooling Laws Really Matter. *American Economic Review*, **96**, 152–175.
- (2007) Do Dropouts Drop Out Too Soon? Wealth, Health and Happiness from Compulsory Schooling. *Journal of Public Economics*, **91**, 2213–2229. URL: <https://www.sciencedirect.com/science/article/pii/S0047272707000138>.
- Psacharopoulos, G. (2024) Returns to education: a brief history and an assessment. *Education Economics*, **32**, 561–565. URL: <https://doi.org/10.1080/09645292.2024.2370119>.
- Psacharopoulos, G. and Mattson, R. (1998) Estimating the returns to education: A sensitivity analysis of concepts and methods. *Journal of Educational Planning and Administration*, **12**, 271–287.
- Psacharopoulos, G. and Patrinos, H. A. (2018) Returns to investment in education: a decennial review of the global literature. *Education Economics*, **26**, 445–458. URL: <https://doi.org/10.1080/09645292.2018.1484426>.
- Rambachan, A. and Roth, J. (2023) A more credible approach to parallel trends. *Review of Economic Studies*, **90**, 2555–2591.
- Rands, T. and Barsanetti, B. (2024) College access spillover on high school performance: Evidence from college openings. *Tech. rep.*
- Spiess, C. K. and Wrohlich, K. (2010) Does Distance Determine Who Attends a University in Germany? *Economics of Education Review*, **29**, 470–479. URL: <https://www.sciencedirect.com/science/article/pii/S0272775709001198>.
- Torche, F. (2019) Educational Mobility in Developing Countries. *Tech. rep.*, WIDER Working Paper 2019/88 Helsinki: UNU-WIDER.
- Udelar (2020) *Propuesta al País 2020-2024. Plan Estratégico de Desarrollo de la Universidad de la República*. URL: <https://udelar.edu.uy/portal/2020/09/propuesta-al-pais-2020-2024-plan-estrategico-de-desarrollo/>.
- UNESCO (2024) Higher education participation data. UNESCO Institute for Statistics. <https://uis.unesco.org>. Accessed: 2024-02-09.
- Wang, C., Liu, X., Yan, Z. and Zhao, Y. (2022) Higher education expansion and crime: New evidence from China. *China Economic Review*, **74**, 101812. URL: <https://www.sciencedirect.com/science/article/pii/S1043951X22000700>.


10 | FIGURES AND TABLES

FIGURE 1 Timeline of University Expansion Program, Including Localities 30 and 50 km Away From New Campuses.


2008	2010	2011	2012	2013 onward
Maldonado	Maldonado	Maldonado	Maldonado	Maldonado
Rocha	Rocha	Rocha	Rocha	Rocha
La Paloma*	Paysandú	Paysandú	Paysandú	Paysandú
Pan de Azucar*	La Paloma*	Salto	Salto	Salto
Piriapolis*	Pan de Azucar*	La Paloma*	Tacuarembó	Tacuarembó
Punta del Este*	Piriapolis*	Pan de Azucar*	La Paloma*	Rivera
San Carlos*	Punta del Este*	Piriapolis*	Pan de Azucar*	La Paloma*
Aigua**	San Carlos*	Punta del Este*	Piriapolis*	Pan de Azucar*
	Aigua**	San Carlos*	Punta del Este*	Piriapolis*
	Quebracho**	San Antonio*	San Carlos*	Punta del Este*
	San Javier**	Aigua**	San Antonio*	San Carlos*
		Quebracho**	Aigua**	San Antonio*
		San Javier**	Quebracho**	Aigua**
			San Javier**	Quebracho**
			Tranqueras**	San Javier**
				Tranqueras**

Notes: Information taken from Udelar's official documents. *Localities included in the 30 km Radius, **Localities included in the 50 km Radius. Unless specified, campuses are located in the department's capital city. Back to Institutional Context 2.2, Back to Data 3, Back to Identification Strategy 4, Back to Spatial Spillovers 5.4.


FIGURE 2 First-Year Students by Geographical Region of Origin and Destination, Years 2002-2020



Notes: The Figures show (Panel a) the number of entry students by region where they lived before entering university and (Panel b) the share of entry students enrolled in a university campus outside the capital city. All new entry students with information in administrative records and the Census in the enrollment year were considered. The sample was drawn from the university administrative records and census for the enrollment year. Back to Institutional Context [2.2](#).


FIGURE 3 Balancing Tests of Pre-Treatment Characteristics

Notes: The Figures show event study estimates of pre-determined locality characteristics before the program implementation. 'Dependency ratio' relates the number of children under 18 to the population aged 30-60. Shaded areas represent 95% confidence intervals. Standard errors are clustered at the locality level. All specifications include calendar year and locality fixed effects. All regressions at the locality level weighted using the population between 18 and 30 years old, based on pre-reform data (2002 to 2007). The sample was drawn from the NHS years 2002 to 2019 using survey weights computed by the National Institute of Statistics. Back to Subsection 4.1.

FIGURE 4 Event Study Estimated Effects on the Share of Enrolled Students

(a) Share of enrollment over population 18-30 (b) Share of enrollment over population 18-30


Notes: The figures report estimates of Equation 2 for the share of total enrollment over the population aged 18–30, for up to six years after treatment (Panel A) and up to ten years after treatment (Panel B). Shaded areas represent 95% confidence intervals. The static TWFE estimates reported in the figures correspond to the average causal effect obtained from Equation 1. Standard errors are clustered at the locality level. All specifications include calendar year and locality fixed effects. Localities are weighted by the pre-reform population aged 18–30 (Census 1996). The sample combines information from university administrative records and the population census in the corresponding enrollment year. Back to Results 5.2.

FIGURE 5 Event Study Estimated Effects on the Educational Intergenerational Mobility


(a) FGS in Population 18-30

(b) FGS in total Enrollment

Notes: The figures report estimates of Equation 2 for (Panel A) the share of FGS over the population aged 18–30, and (Panel B) the share of FGS over total enrollment. Shaded areas represent 95% confidence intervals. Standard errors are clustered at the locality level. All specifications include calendar year and locality fixed effects. Localities are weighted by the pre-reform population aged 18–30 (Census 1996). The sample combines university administrative records with census data from the corresponding enrollment year. Back to Results 5.2.


FIGURE 6 Event Study Estimated Effects on Total and FGS Enrollment

Notes: The figures report estimates of Equation 2 for (Panel A) the total number of students enrolled, and (Panel B) the total number of FGS enrolled. Shaded areas represent 95% confidence intervals. The static TWFE estimates reported in the figures correspond to the average causal effect obtained from Equation 1. Standard errors are clustered at the locality level. All specifications include calendar year and locality fixed effects. Localities are weighted by the pre-reform population aged 18–30 (Census 1996). The sample combines university administrative records with census data from the corresponding enrollment year. Back to Results 5.2.

FIGURE 7 Event Study Estimated Effects on the Share of FGS by Parental Background


Notes: The figures show estimates of Equation 2 for (Panel A) the share of enrolled students with parents who completed only primary education over the population aged 18–30, (Panel B) the share of enrolled students with parents who completed secondary education over the population aged 18–30, and (Panel C) the share of enrolled students with parents who completed tertiary (non-university) education over the population aged 18–30. Shaded areas represent 95% confidence intervals. The static TWFE estimates reported in the figures correspond to the average causal effect obtained from Equation 1. Standard errors are clustered at the locality level. All specifications include calendar year and locality fixed effects. Localities were weighted using population between 18 and 30 years old, based on pre-reform data (Census 1996). The sample was drawn from the university administrative records and census at the enrollment year. Back to Subsection 5.3.

FIGURE 8 Event Study Estimated Effects on Completion of University Degree: All and First-Generation Students



Notes: The figures show estimates of Equation 2 for (Panel A) the probability of a student obtaining a degree within five years of enrolling, and (Panel B) the probability of a FGS obtaining a degree within five years of enrolling. Shaded areas represent 95% confidence intervals. The static TWFE estimates reported in the figures correspond to the average causal effect obtained from Equation 1. Standard errors are clustered at the locality level. All specifications include calendar year and locality fixed effects. Localities are weighted by the pre-reform population aged 18–30 (Census 1996). The sample combines university administrative records with census data at the enrollment year, and is restricted to students who enrolled in university between 2002 and 2015. Back to Subsection 5.5.

FIGURE 9 Event Study Estimated Effects on High School Outcomes



Notes: The figures show estimates of Equation 2 for (Panel A) the number of students attending high school, and (Panel B) the number of students completing high school education. Shaded areas represent 95% confidence intervals. The static TWFE estimates reported in the figures correspond to the average causal effect obtained from Equation 1. Standard errors are clustered at the locality level. All specifications include calendar year and locality fixed effects. Localities were weighted using population between 18 and 30 years old, based on pre-reform data (2002 to 2007). The sample was drawn from the NHS years 2002 to 2019 using survey weights computed by the National Institute of Statistics. Back to Subsection 6.

TABLE 1 Descriptive Statistics

Panel A: Individual Level			
	Mean	SD	Obs.
<i>Student's Characteristics</i>			
Female	0.63	0.48	64,820
Age	19.02	2.31	64,820
<i>Parents' Education</i>			
Illiterate	0.00	0.05	64,820
Primary	0.12	0.33	64,820
Secondary	0.51	0.50	64,820
Tertiary	0.18	0.39	64,820
University	0.18	0.39	64,820
<i>Outcomes</i>			
First generation university	0.82	0.39	64,820
Completion	0.36	0.48	34,390
Panel B: Locality Level			
	Mean	SD	Obs.
<i>Outcomes</i>			
Enrollment	112.93	97.32	1,225
Share Enrollment	0.02	0.01	1,225
FG enrollment	89.38	77.25	1,225
Share FG (enroll)	0.80	0.10	1,225
Share FG before treatment	0.73	0.09	50
Completion	0.09	0.07	876
FG completion	0.09	0.08	870

Notes: The Table presents the means, standard deviations, and valid observations for the main characteristics of the estimation sample. The sample was drawn from the university administrative records and the enrollment year census. For completion, the sample was restricted to students who enrolled in university between 2002 and 2015. Data at the locality level is weighted using population between 18 and 30 years old, based on pre-reform data (Census 1996). Back to Section 3.

TABLE 2 Balancing Tests of Pre-Treatment Characteristics (2002-2007)

Variable	Never treated	Treated	(NT-T)	Early adopters	Late adopters	(E-L)
% women	0.513	0.528	0.016***	0.528	0.529	-0.001
Average age	35.680	34.826	-0.854*	34.525	35.127	-0.602
% aged 18-30	0.165	0.176	0.011***	0.172	0.179	-0.007
Employment rate (18-50)	0.685	0.676	-0.009	0.666	0.686	-0.020
Formality rate (18-50)	0.605	0.579	-0.026	0.558	0.600	-0.042*
Average household income	15,135.918	14,937.268	-198.651	14,415.204	15,459.331	-1,044.128
% no health coverage	0.025	0.037	0.012***	0.045	0.028	0.018**
Years of education (18-30)	8.852	9.357	0.505***	9.373	9.340	0.033

Notes: The table presents the means and mean-difference tests for selected demographic and socioeconomic variables before the first campus opened, by treatment status and locality. Early adopters are Paysandú, Salto, and Rivera. *** Significant at the 1% level, ** at the 5% level, and * at the 10% level. The sample was drawn from the 2002–2007 NHS waves using survey weights computed by the National Institute of Statistics. It is restricted to individuals living in municipal capitals, excluding Montevideo and Canelones. Back to Subsection 4.1.

TABLE 3 Regression Estimates of Distance and Enrollment

	Share Enrollment	Share FG (pop)	Share FG (enroll)
Distance to closest campus	-0.0009*** (0.00)	-0.0009*** (0.00)	-0.0162*** (0.00)
Pre-treatment Mean	0.017	0.014	0.808
Obs.	1225	1225	1225

Notes: The Table presents the marginal effect of distance (in hundreds km) on enrollment and intergenerational mobility. All specifications include calendar year and locality fixed effects. Standard errors clustered at the locality level in parentheses. ***significant at the 1% level, **5% level, *10% level. Localities were weighted using population between 18 and 30 years old based on pre-reform data (Census 1996). The sample was drawn from the university administrative records and the enrollment year census. Back to Subsection 5.1.

TABLE 4 Main Effects Under an Alternative Treatment Definition

	Share Enrollment	Share FG (pop)	Share FG (enroll)
Treated locality	0.0038*** (0.00)	0.0033*** (0.00)	0.0621*** (0.02)
Radius _{0–30km}	0.0048* (0.00)	0.0031** (0.00)	0.0083 (0.04)
Radius _{30–50km}	0.0003 (0.00)	0.0020 (0.00)	0.0435 (0.04)
Obs.	1225	1225	1225
R-squared	0.800	0.790	0.527

Notes: The Table shows the static TWFE estimates of Equation 1. All specifications include calendar year and locality fixed effects. Standard errors clustered at the locality level in parentheses. ***significant at the 1% level, **5% level, *10% level. Localities were weighted using population between 18 and 30 years old, based on pre-reform data (Census 1996). The sample was drawn from the university administrative records and census at the enrollment year. Back to Results 5.4.

TABLE 5 Diff-in Diff Estimates of the Program Effect on Educational Outcomes

	Years of Education	High-Education	Middle-Education	Low-Education
Exposed x Campus	0.443** (0.21)	0.032* (0.02)	0.013 (0.03)	-0.045* (0.03)
Obs.	8184	8184	8184	8184
R-squared	0.049	0.038	0.018	0.038

Notes: The Table shows the average causal effects of the program. 'High Education' is a dummy variable taking value 1 if the individual has more than 12 years of education, 'Middle Education' takes value 1 if the individual has between 7 and 12 years of education, and 'Low Education' takes value 1 if the individual has 6 or less years of education. All specifications include as controls calendar year, locality, cohort fixed effects, gender, marital status, and number of children under 12 in the household. Standard errors clustered at the locality level in parentheses. ***significant at the 1% level, **5% level, *10% level. The sample was drawn from the NHS years 2018 to 2019 using survey weights computed by the National Institute of Statistics. Sample is restricted to individuals in municipal capitals, except Montevideo and Canelones. Back to Section 7.1.

TABLE 6 Event-study estimates: TWFE vs. Staggered SDID

	TWFE				Staggered SDID			
	Share Enrollment		Share FG (pop)		Share Enrollment		Share FG (pop)	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
Average Effect	0.0037	0.0005	0.0032	0.0004	0.0031	0.0010	0.0027	0.0008

Notes: The Table shows estimates of Equation 1 using [Borusyak et al. \(2024\)](#) TWFE estimator (Columns 1–4) and [Arkhangelsky et al. \(2021\)](#) SDID estimator (Columns 5–8). All specifications include calendar year and locality fixed effects. Standard errors clustered at the locality level. Localities were weighted using population between 18 and 30 years old, based on pre-reform data (Census 1996). The sample was drawn from the university administrative records and census at the enrollment year. Back to Results 7.2.