Funding Instruments and Effort Choices in Higher Education *

Guillem Foucault† Juan Pal‡

Job Market Paper

October 31, 2026
Click HERE for latest version

Abstract

This paper examines the effects of tuition-free higher education (TF) policies on student enrollment and academic performance, focusing on the 2016 Chilean reform that eliminated higher education tuition fees for students from the lowest five income deciles. Using a two-way fixed effects (TWFE) approach, we find that TF increased enrollment among eligible students by around 4 percentage points—concentrated among lower-ability students—and both reduced dropout rates and increased graduation rates among students enrolled before the reform. For post-reform cohorts, these effects diminish: graduation and on-time graduation gains persist but are smaller, and the dropout reduction vanishes. To disentangle the role of student effort from composition effects, we develop a dynamic structural model in which students choose effort levels in response to financial incentives. The model reproduces the enrollment response estimated by TWFE and its concentration among lower-ability students. Counterfactual simulations show that reintroducing annual progress requirements reduces the cost per additional graduate from $108,000 to $70,000, but widens socioeconomic gaps in graduation. Minimum entry scores would raise the cost per additional graduate nearly threefold and leave the socioeconomic gap almost unchanged. Phasing out the benefit for middle-income students and providing a stipend to the poorest raises graduation as much as TF does, at lower cost and with a smaller socioeconomic gap.

Keywords: Tuition-free higher education, Financial aid, Student effort
JEL Codes: I22, I23, I24, I28

*We are very grateful for the guidance from Matteo Bobba, Olivier De Groote, and Ana Gazmuri. We also thank Paul Diegert, Margaux Luflade, Thierry Magnac, Maximilian Müller, and François Poinas for their helpful comments. We are also grateful to workshop and seminar participants at the Toulouse School of Economics and at the University of Pennsylvania. Finally, we are indebted to Rodrigo Rolando Meneses and Iván Atencio Abarca at the Higher Education Information Service (SIES) for granting us access to the data, and more broadly to Chile's Ministry of Education (MINEDUC) and the Department for Educational Evaluation, Measurement and Registry (DEMRE). We acknowledge funding from the French National Research Agency (ANR) under the grant MATCHINEQ - ANR-22-CE26-0005-01 and the European Research Council (ERC) under the grant EDDYNCHOICE - 101162723.

†Toulouse School of Economics. guillem.foucault@tse-fr.eu. gfill.github.io.

‡PUC-Chile. juan.pal@uc.cl.

1 Introduction

Higher education is a central determinant of long-run earnings, inequality, and intergenerational mobility. Despite decades of expansion in attainment, the returns to a college degree have remained large (Card, 1999; Goldin and Katz, 2008), and access to selective institutions is a powerful lever for upward mobility (Chetty et al., 2020). Yet completion is unequally distributed, and the gap in degree attainment between high and low-income youth has widened in recent decades (Bailey and Dynarski, 2011; Dynarski et al., 2023). In response, governments have invested heavily in policies aimed at lowering financial barriers to college. As of today, more than 40 countries operate some form of tuition-free higher education (TF) policy, generally limited to public institutions, and recent debates in the United States over student-debt forgiveness, free community college, and need-based grants reflect a similar emphasis. Public spending on higher education in the OECD has risen accordingly (OECD, 2019), raising the question of how aid should be designed to balance equity, fiscal cost, and educational outcomes.

Higher education policies operate on two margins. On the extensive margin, lowering net tuition draws in marginal students who would otherwise have stayed out. On the intensive margin, the design of aid impacts field choice, persistence and completion among those who do enroll. Traditional instruments combine financial relief with performance requirements to keep incentives aligned. TF represents an extreme combination of both dimensions, eliminating out-of-pocket tuition for a broad eligible population without imposing progress conditions, and thus removing both a financial barrier and a performance lever simultaneously. The ex-ante effect on educational outcomes is therefore ambiguous. Eliminating tuition draws into the system a less-prepared pool of marginal entrants while, for infra-marginal students, it relaxes a binding budget constraint, frees up time previously devoted to part-time work, and reduces dropout after negative income shocks. Removing the performance requirement also impacts selection (drawing in students with a higher cost of effort) while reducing incentives for effort among infra-marginal students (Scott-Clayton, 2011). Countries with TF tend to exhibit lower aggregate graduation rates than comparable countries in which students face out-of-pocket fees, but aggregate measures cannot distinguish compositional changes from behavioral responses during college (Ferreyra et al., 2022).

This paper studies the effects of a large-scale TF reform on enrollment and progression in higher education, and develops a framework that disentangles selection from changes in effort. We exploit Chile's 2016 *Gratuidad* reform as a natural experiment, which granted tuition-free higher education to students from the lowest five income deciles in roughly half of universities and vocational centers (including some private institutions), and proceed

in three steps. First, we implement a two-way fixed effects (TWFE) design to estimate effects of TF on enrollment, persistence, dropout, and graduation. Second, we build and estimate a dynamic structural model of enrollment, program choice, and academic effort that endogenizes how students respond to changes in out-of-pocket tuition. Third, we use the estimated model to simulate counterfactual designs that re-introduce performance-based conditions on top of the tuition-free benefit, and assess their implications for fiscal cost and educational attainment. The structural exercise allows us to decompose aggregate impacts into a composition effect from marginal students induced to enroll and a behavioral effect on the effort choices of infra-marginal students.

We assemble administrative records from Chile's Ministry of Education that link the universe of national-exam takers to higher education enrollment, scholarship assignment, course credit completion, dropout, and graduation outcomes. The *Gratuidad* reform is a sharp, large-scale policy change with a clearly defined eligibility rule based on income decile, introduced on top of a well-developed pre-existing aid system that included subsidized loans (the State Guaranteed Loan) and need-based scholarships. Crucially, the pre-existing instruments were all performance-based, both at entry (scores at the National Entrance Exam) and during studies (course completion rates), and covered up to 80% of tuition. *Gratuidad* replaced these conditional schemes with full, unconditional tuition coverage for eligible students. The reform therefore varied both the level of net tuition and the performance conditioning of aid, providing the variation needed to identify selection and effort responses separately.

The TWFE estimates reveal responses in both the extensive and intensive margins. Enrollment among eligible students rose by 1.7 percentage points (p.p.) in 2016 and stabilized at around 4.3 p.p. from 2017 to 2019, after the policy was extended to vocational centers in 2017 and to the sixth decile in 2018. The increase was concentrated among lower-ability students who, prior to the reform, did not qualify for either subsidized loans or merit scholarships. Their enrollment rose by 9.6 p.p. in 2017 and by up to 16 p.p. in 2019, compared with about 1 p.p. for students already eligible for the most generous pre-existing aid. We find no evidence of crowding-out at non-eligible institutions, in contrast to the spillover patterns documented by Londoño-Vélez et al. (2020) in Colombia, suggesting that the bulk of the response comes from genuinely marginal students rather than substitution across institutions. As a result, the composition of incoming cohorts shifted to the left of the ability distribution, a pattern with first-order implications for downstream outcomes.

Estimating effects on educational outcomes is less straightforward, for two reasons. First, the change in composition confounds the response of infra-marginal students with that of new entrants. Second, the policy bundled lower out-of-pocket fees, which reduce the cost of

continuing, with the removal of performance requirements, which weakens the incentive to exert effort, so the two components may push outcomes in opposite directions. We address the first problem by exploiting a feature of the policy design. Cohorts that had already enrolled before 2016 became eligible for *Gratuidad* on their remaining tuition payments, with treatment intensity varying by cohort but with enrollment decisions already made. This holds composition fixed and isolates the response of infra-marginal students. Among these students, the policy reduced dropout by up to 2.2 p.p., raised graduation by up to 5.6 p.p., and increased on-time graduation by up to 6.1 p.p. These results suggest that, for these students, the lower cost of continuing dominates the weaker incentive to exert effort. For 2016-and-after cohorts, these effects diminish: the dropout reduction vanishes and graduation gains shrink, pointing to a negative composition shift that is roughly offset by the policy.

While the reduced-form analysis shows that the policy was successful in expanding access without lowering aggregate completion rates, three reasons motivate a structural approach. First, the TWFE estimates cannot decompose the relative contribution of marginal students induced to enroll and infra-marginal students who change their behavior. Second, the policy bundled two changes that the data cannot separately identify without further structure, namely the reduction in out-of-pocket tuition and the elimination of performance requirements. Third, evaluating alternative policy designs that preserve the equity gains of *Gratuidad* while reducing its fiscal cost requires a model of how students respond to changes in prices and financial aid rules. Between 2011 and 2019, public spending on higher education in Chile rose by 160% and reached 5.4% of GDP, against 2.9% on average in the rest of the OECD countries (OECD, 2019).

For these reasons, we develop a dynamic discrete-choice model of enrollment, field and institution choice, dropout, and academic effort. In the first stage, a student observes her type (gender, ability, socioeconomic status and unobserved trait) and her personalized choice set (Fack et al., 2019), and decides whether and where to enroll. Conditional on enrollment, in each subsequent period she chooses whether to continue and how much effort to exert, where effort shapes the distribution of credit completion. Performance feeds back into next-period out-of-pocket tuition through the rules of pre-existing scholarship programs, generating a dynamic link between current effort and future financial incentives. Assuming that students choose effort optimally, its marginal cost is recovered from the corresponding first-order condition, following De Groote (2025). The model thus recovers the joint distribution of fixed and marginal costs across students. Because it identifies which students enroll only because of the reform, it allows us to separate changes in aggregate outcomes driven by the compo-

sition of entrants from those driven by the behavioral responses of students who would have enrolled anyway.

We estimate the model by maximum likelihood on the 2016–2017 entry cohorts, for whom we observe academic progression for up to seven years. Continuation values are approximated from conditional choice probabilities using the inversion of Hotz and Miller (1993), so that estimation does not require solving the model over the full state space. We account for persistent unobserved heterogeneity with two latent types, following Arcidiacono and Miller (2011), who adapt the expectation-maximization (EM) algorithm to conditional choice probability estimators by updating the conditional choice probabilities jointly with each student's posterior type probabilities. We then use the estimated primitives to compute marginal costs and solve the model by backward induction. Because estimation uses post-reform cohorts only, the TWFE estimates provide an external benchmark for the model's simulated responses to *Gratuidad*.

Estimation reveals substantial heterogeneity in costs across the student distribution. Fixed costs increase with out-of-pocket tuition for low- and middle-income students, and less so for students with higher test scores. Beyond observed ability, which the model already controls for, we also uncover a latent type with a strong preference for being enrolled but slightly higher marginal costs of effort: these students persist in higher education because they value enrollment, not because studying is less costly for them. Marginal costs of effort decrease steeply with ability and, to a lesser extent, with income, in particular among students with low test scores, so that low-scoring, low-income students face the highest marginal costs of producing a unit of academic progress. The model also reproduces the enrollment response estimated by TWFE and its gradient across test-score groups, although it understates the 2017 increase among students below the scholarship threshold.

We find no evidence that the removal of performance requirements weakened effort among infra-marginal students. In the model, *Gratuidad* raises their graduation rate by about 1 p.p., and re-introducing an annual progress requirement leaves this gain almost unchanged. Selection, through marginal entrants and students who switch program, accounts for more than two thirds of the simulated increase in graduation, while the behavior of infra-marginal students changes little: aggregate completion is held in check by composition rather than by moral hazard.

We use the estimated model to evaluate alternative designs of TF: two that re-introduce the conditions of the pre-existing aid programs, namely an entry requirement (a test score of at least 500 for eligibility) and a continuation requirement (completing 70% of registered credits

each year to retain free tuition), and two that reallocate funding across income groups. In the model, TF raises enrollment by 3.0 p.p. and graduation by 1.9 p.p. at a public cost of $2,100 per student, two thirds of which is a transfer to students who would have enrolled anyway, or a net cost of $108,000 per additional graduate. The two requirements act on different margins and have very different returns. The entry requirement saves 61% of the outlay but retains only 14% of the graduation gain, because it excludes precisely the students whose outcomes aid changes most; its cost per additional graduate is nearly three times that of TF. The continuation requirement barely changes effort: for students at risk of falling behind, the threat of losing the benefit leads to dropout rather than to more study, so the requirement saves money by withdrawing aid from students who fall behind, not by strengthening incentives. Because it binds only after enrollment, it halves the outlay while retaining 77% of the graduation gain, lowering the cost per additional graduate to $70,000. These savings are not free: the graduation losses fall on low-ability, low-income students, and the requirement gives back about a quarter of TF's reduction in the graduation gap between high- and low-SES students.

Finally, TF is not on the efficiency-equity frontier. Phasing the benefit out over income deciles 5 and 6 and paying a stipend to students in the lowest decile delivers the same graduation gain as TF with 10% less outlay and a smaller SES gap, because the funding spent on middle-income students changes the behavior of few of them; extending the stipend to the three lowest deciles narrows the gap further, at a higher cost. Efficiency and equity are thus aligned when funding is reallocated from middle- to low-income students, and conflict when aid is made conditional on progress.

Related literature This paper builds upon several strands of literature. It first contributes to the literature on higher education financing and student outcomes. Broad reviews emphasize that aid design affects not only access, but also persistence, completion, and the timing at which students learn about prices and eligibility (Deming and Dynarski, 2010, Page and Scott-Clayton, 2016). There is substantial evidence that increasing grants or reducing net tuition raises enrollment and attainment (Dynarski, 2003, Castleman and Long, 2016, Angrist et al., 2015, Denning, 2017, Fack and Grenet, 2015, Londoño-Vélez et al., 2020, Dobbin et al., 2022), and that the effect of subsidies depends on the pre-reform financing environment and on whether aid takes the form of grants or loans (Solis, 2017, Molina and Rivadeneyra, 2021, Duryea et al., 2023, Bolhaar et al., 2024). Evidence on completion and performance is more mixed once composition effects are taken into account (Cohodes and Goodman, 2014, Denning, 2019, Denning et al., 2019, Ferreyra et al., 2024). Beyond partial equilibrium responses, Epple et al., 2006, Epple et al., 2017, and Abbott et al., 2019 show

that equilibrium supply responses and redistribution matter for the incidence of education-finance reforms, and Fu et al., 2024 show that state governments respond to free-college mandates by cutting per-student spending, so that enrollment rises while four-year graduation rates fall. This paper complements that literature by studying a large-scale tuition-free reform in a setting where the public aid system was already well developed, and by tracing both extensive-margin enrollment responses and post-enrollment outcomes.

It also contributes to the growing literature on tuition-free college and the design of free-college programs. Bucarey (2017) studies the crowding-out effects of Chile's tuition-free reform on low-income students at selective programs. Recent work emphasizes that design is consequential: unconditional or early-committed aid generates larger take-up responses than aid that remains uncertain or application-intensive (Burland et al., 2023, Harris and Mills, 2025), dynamic models imply that free college raises enrollment more than graduation because outcomes depend on the type of students induced to enroll and on their effort choices (Matsuda, 2020, Ferreyra et al., 2024), and promise programs and tuition reversals display meaningful heterogeneity across institutional settings (Carruthers and Fox, 2016, Schoonover et al., 2026, Bietenbeck et al., 2023). Our contribution is to bring this design question to the Chilean reform and to quantify how much of the observed effect operates through selection versus changes in effort among infra-marginal students.

Second, we contribute to the literature on the incentive structures and moral hazard inherent in educational funding. Subsidies lower the cost of entry, but they may also weaken the marginal cost of delay or failure. Scott-Clayton (2011) shows that academic incentives embedded in aid rules affect college achievement, Garibaldi et al. (2012) and Beneito et al. (2018) show that continuation tuition and retake pricing reduce delay and increase effort, and Montalbán (2022) finds that cash support is most effective when accompanied by demanding progress requirements. Students also respond to the composition of aid and to borrowing frictions (Marx and Turner, 2018, Falco and Reichlin, 2025). TF further releases students from debt accumulation, which affects major choice (Rothstein and Rouse, 2011), dropout (Stinebrickner and Stinebrickner, 2008), career decisions (Sieg and Wang, 2018), and post-college household formation (Black et al., 2023). This paper speaks directly to that mechanism because the Chilean reform simultaneously changed out-of-pocket tuition and removed academic-progress requirements for a large group of students.

Finally, we contribute to a growing structural literature on educational choices under dynamic incentives. Larroucau and Rios (2024) study retention and reapplication decisions in a centralized admissions market in which students learn about their match quality over time. Dynamic models have also been used to study the design of financial aid. Colas et al. (2021)

derive and quantify optimal need-based aid in a life-cycle model of enrollment, dropout, and parental transfers, while Joensen and Mattana (2024) compare grants and loans in a model of enrollment, work, and borrowing in which merit requirements condition aid on credit accumulation. Harris et al. (2026), who discipline their model with experimental and quasi-experimental estimates of aid effects, find that unconditional aid moves attendance far more than completion, and that only aid contingent on the degree raises completion among enrollees. In all three, academic progression is a stochastic process that does not respond to policy. A smaller set of papers model effort explicitly. Dubois et al. (2012) model endogenous learning effort with uncertain grade progression and use the experimental variation of Progresa to estimate the effect of conditional cash transfers on enrollment and performance. Ferreyra et al. (2022) treat in-college effort as an input to class completion in Colombia and show that omitting it overstates the role of ability, while Tincani et al. (2026) and Arslan (2026) study pre-college effort responses, to preferential admissions in Chile and to tuition subsidies in Turkey, the former allowing for biased beliefs about the returns to effort. Finally, De Groote (2025) studies track choice and grade retention in secondary school and shows that accounting for effort responses is key to predict changes in graduation and college enrollment; we borrow his identification of the marginal cost of effort and bring it to student financing in higher education, where effort determines eligibility for aid and the cost of continuing. We contribute by linking current effort to future educational costs through out-of-pocket fees and performance requirements, which allows the model to decompose adverse selection from moral hazard under a large-scale tuition-free reform.

Overview The rest of the paper is organized as follows. Section 2 describes the institutional details, the TF policy and introduces the data. Section 3 reports results from TWFE estimation, and Section 4 builds a model of enrollment and effort in higher education. Section 6 discusses the main results of the model and Section 7 performs counterfactual simulations. Section 8 concludes.

2 Background

This section describes the institutional setting, the policy of interest, and the data used in our analysis. Section 2.1 provides an overview of the Chilean higher education system and the pre-existing funding instruments. Section 2.2 describes the Tuition-Free Higher Education policy introduced in 2016 and the substitution patterns it generated across the previously available aid programs. Section 2.3 introduces the administrative records that we use to track enrollment, performance, and graduation outcomes, and defines the samples used in

the reduced-form and structural analyses.

2.1 Institutional Setting

Chilean higher education is well developed, with enrollment rising steadily since the General Law of Universities (*Ley General de Universidades*) of 1981, which encouraged the creation of new institutions by allowing the entry of private providers. In 2023, the system comprised a total of 138 institutions: 80 vocational schools (33 Professional Institutes and 47 Centers of Technical Formation) and 58 universities.[1]

Universities are divided between the so-called *traditional* universities (a mix of public and private institutions that receive direct support from the government) and *private* universities, which constitute the rest. The latter were created after 1981 and are mainly financed through tuition. Traditional universities are officially known as the Universities of the Rector's Council (*Consejo de Rectores de las Universidades Chilenas, CRUCH*) and are responsible for coordinating the higher education system. They comprise 30 universities and account for around half of total enrollment.[2]

Access to higher education is through a Centralized Admission Platform, compulsory for traditional universities and some private institutions. Vocational program admission is also possible, although most are conducted off-platform. After high-school completion, students applying to higher education take the centralized admission exam (*Prueba de Selección Universitaria*, PSU), which we refer to as the *test score* throughout the paper. This standardized test, similar to the SAT in the United States, ranges from 150 to 850 points, with a mean of 500 and a standard deviation of 110. Students then submit rank-ordered lists of up to 10 degree programs through the centralized system and are assigned via a deferred-acceptance algorithm (Gale and Shapley, 1962). In practice, the platform imposes a minimum test score of 450 to apply.

Higher education is costly in Chile. In 2009, the average tuition fee for a university program was equivalent to 47% of the median family income (Solis, 2017). Costs vary across institution types but remain high even in public institutions. Different funding instruments coexist. Students rely mainly on loans and grants from the Ministry of Education. Eligibility criteria are strongly related to test scores and family income, as summarized in Appendix Figure A4. The State Guaranteed Loan program (SGL), introduced in 2006, finances 80% of reference tuition and requires a test score of 475 or higher, with no socioeconomic requirement since

[1] https://www.ayudamineduc.cl/ficha/instituciones-vigentes-reconocidas-por-el-mineduc.

[2] In 2018, the government established a clear set of rules for being considered a traditional university with the promulgation of Law 21091, allowing the first three non-traditional universities to enter.

2014. Several scholarships also exist, with the *Beca Bicentenario* and *Beca Excelencia* being the most popular. On average, scholarships finance 80% of reference tuition. Eligibility requires a test score of 500, and excludes students in the top three income deciles. Both short-cycle programs (SCPs) and university degrees are covered by scholarships. Two-year programs typically require a high-school GPA above 5.0. Access to private loans is limited: in 2015, only 7.5% of student loans came from banks without a state guarantee (Bucarey et al., 2020).

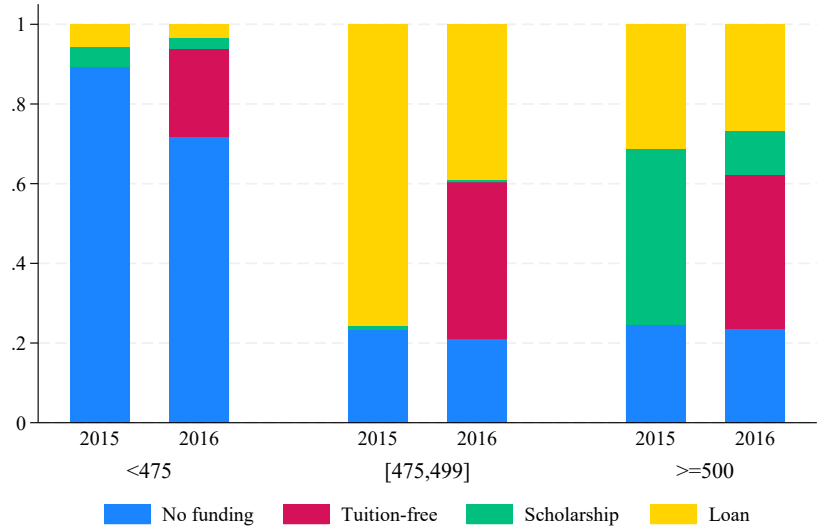
2.2 Tuition-Free Higher Education (TF) program

Since the implementation of SGL, student debt rose at a rate of 70% annually, and the number of recipients increased steadily from 15,800 in the first year to 652,000 students in 2016 (Bucarey et al., 2020). In 2011, students led mass protests demanding more affordable higher education. Michelle Bachelet was elected president in 2014 on a platform that included making college free by 2020. In 2015, the TF law was enacted, removing tuition fees for students in the bottom half of the income distribution. The policy was first implemented for the 2016 university cohort, expanded in 2017 to include vocational institutes, and further extended in 2018 to the sixth income decile. Applications for TF, loans, and scholarships are made through the Centralized Admission Platform by completing a short online form during the application period, regardless of whether the applicant applies to a platform degree or not. A schematic visualization of the policy is provided in Appendix Figure A5.

The introduction of TF policy generated different substitution patterns depending on students' test scores. Figure 1 shows the distribution of funding instruments before and after the policy, by test score segment. Students scoring below 475 were ineligible for subsidized loans or merit-based scholarships (except for a small scholarship targeting top students from low-SES schools). For this group, TF reduced the share of students with no funding. Students scoring between 475 and 500 qualified for subsidized loans but remained below the threshold for most scholarships; for them, substitution was mainly from loans to TF, relieving them of any debt upon graduation. Students scoring above 500 qualified for merit-based scholarships (conditional on income eligibility). For this group, substitution was mainly from scholarships to TF, which had two advantages: first, scholarships typically covered only 80% of tuition and fees, whereas TF provided full funding; second, scholarships required students to meet satisfactory progression standards (validating about 70% of annual course credits), while TF imposed no such requirements. Before the policy, about 20% of scholarship beneficiaries lost eligibility; of these, around 40% switched to subsidized loans, while the rest continued without public funding. After the policy, most substituted to TF, while those above the

income threshold continued either with loans or no funding, in roughly the same proportions as before.

Figure 1: Distribution of funding instruments, by test score



Notes: This figure shows the distribution of funding instruments by test score group, for 2015 and 2016. PSU is the national standardized test. Students with test scores below 475 are not eligible for public funding, students scoring between 475 and 500 are eligible for loans, and students scoring above 500 are eligible for scholarships.

2.3 Data

We use Chilean administrative records that provide detailed student and program data, including off-platform private universities and SCPs. We observe enrollment, admission test scores, scholarship assignment, socioeconomic information, and demographic characteristics. Throughout the paper, we measure a student's *ability* by her test score and her high-school GPA, both standardized in estimation. Because eligibility for loans and scholarships depends on the test score only, we group students by test score alone when comparing them across eligibility thresholds, and refer to these groups as *test-score groups*. Institutional data include type, tuition, location, and program length. In addition, we employ novel data on credit completion from the Information Service of Higher Education (SIES) covering the universe of enrolled students from 2016 to 2023. These data report the number of courses registered and whether they were completed.

Quasi-Experimental sample We restrict the sample to first-time test takers. The enrollment regressions use the 2013 to 2020 cohorts, and the regressions on educational outcomes

use students who enrolled in the 2009 to 2017 cohorts. Since we observe the data until 2023 and the standard university program lasts five years, we allow seven years for students to complete their program. Appendix Table B1 reports enrollment, controls, entry funding and outcomes for this sample by cohort. Between 100,000 and 130,000 students take the national exam for the first time each year, of whom about 60% enroll in higher education, split roughly half and half between the universities and SCPs. Panel B shows that around 55% of students are female. Panel C shows student financial aid. In 2013, around 30% of students received a subsidized loan at some point during their studies, and 40% received a scholarship. In 2016, these shares diminished to 19% and 25% respectively, while 34% of students received TF, which increased up to 59% in 2020. In 2012, scholarships were expanded to include the third income quintile, such that all eligible students for TF were already eligible for scholarships conditional on reaching a minimum test score (Bucarey, 2017).

Model sample For the structural model, we also exploit the rich data on credit completion. This dataset is only available from 2016. Since we allow students seven years to graduate, we estimate the model on a random sample of the 2016 and 2017 cohorts. Descriptive statistics for the model sample are reported in Appendix Table B2. It is a random sample of 7,000 students per cohort, which is representative of the full population (panels A, B and C). Panel D shows that students register almost all required credits per year (93%) which is consistent with the fact that tuition is paid annually irrespective of the registered courses. Students complete around 80% of the annual curriculum, which represents about 85% of the courses they register.

3 Policy Effects

This section provides reduced-form evidence on the effects of the TF policy on enrollment and educational outcomes. Section 3.1 presents the two-way fixed effects (TWFE) specification and discusses the identifying assumption. In Section 3.2, we estimate the extensive-margin response and document substantial heterogeneity across test-score groups and institution types. Section 3.3 addresses the composition problem that arises when estimating intensive-margin effects, and isolates the response of infra-marginal students by exploiting the differential exposure of pre-2016 cohorts to the policy. Section 3.4 discusses the main threats to identification and presents robustness checks.

3.1 Empirical Strategy

We are interested in causally estimating whether the TF policy had an impact on enrollment and educational outcomes. To do so, we exploit the exogenous variation in eligibility introduced by the policy and estimate its effect using a Two-Way Fixed Effects framework. We estimate the following regression model by OLS:

$$Y_{it} = \sum_{\substack{k=2013 \\ k\neq 2015}}^{2020} \beta_k \left(\mathbb{1}\{t = k\} \times \mathbb{1}\{dec_i < 6\} \right) + \gamma \, \mathbb{1}\{dec_i < 6\} + \delta_{\mathbf{t}} + \mathbf{X}_{\mathbf{i}}' \alpha_{\mathbf{x}} + \epsilon_{\mathbf{it}} \tag{1}$$

Y_{it} is an outcome of interest including enrollment and educational outcomes such as dropout, graduation, on-time graduation and persistence, defined as the time (in years) spent in college. $\mathbf{X}_{\mathbf{i}}'$ includes a constant, gender, test score, high school GPA, degree type, and mother's education. $\mathbb{1}\{dec_i < 6\}$ determines eligibility to TF, defined as belonging to the lower half of the income distribution. δ_t are cohort fixed effects. We report heteroskedasticity-robust standard errors. We use the year before the implementation (2015) as the reference category. β_t, starting 2016, is the effect of the policy reform.

The identifying assumption is that, in the absence of treatment, eligible and non-eligible students would have experienced similar trends in the outcomes of interest. This assumption is not testable, since the untreated outcomes of eligible students are never observed after 2015. We instead examine pre-treatment trends: parallel trends before the policy are suggestive, though not proof, that they would have continued after it. We also assess, in Appendix Section A.1, how sensitive the estimates are to violations of parallel trends (Rambachan and Roth, 2023; Roth, 2022). The design also requires no violation of the stable unit treatment value assumption (SUTVA): the outcomes of a student must not depend on the treatment status of other students, so there must be no spillovers between eligible and non-eligible students (for instance, through changes in the composition of applicants or in the availability of seats at institutions). Appendix Section A.3 reports descriptive evidence on these margins: admission cutoffs and tuition by institution tier (Appendix Figure A15 and Appendix Figure A16),[3] and vacancies (Appendix Figure A17). Admission cutoffs show no visible break around the introduction of TF, tuition keeps evolving in line with its pre-existing trend, and

[3]Elite: Pontificia Universidad Católica de Chile, Universidad de Chile, Universidad de Concepción, Universidad Técnica Federico Santa María, Pontificia Universidad Católica de Valparaíso and Universidad de Santiago de Chile. G8: Universidad Adolfo Ibáñez, Alberto Hurtado, Andrés Bello, de los Andes, del Desarrollo, Diego Portales, Finis Terrae and Mayor. Other CRUCH: the remaining universities of the Council of Rectors. Off-plat: universities outside CRUCH that are not in the G8. SCP: short-cycle programs at any institution.

vacancies did not expand much at adhering institutions, consistent with the 2016 Budget Law, which capped the growth of first-year enrollment at participating institutions at 2.7% unless authorized by the Ministry of Education (DIPRES, 2016).

Other papers have also estimated the causal effect of TF. To apply for any state grant, loan or TF, students must fill in the *Formulario Único de Acreditación Socioeconómica* (FUAS), an online form run by the Ministry of Education in which they report their household income. Since the official income decile is only observed for students who file this form, Martin (2024) uses the FUAS income decile and models aid applications jointly with college choice, as the policy increased applications. He also estimates a LATE with a fuzzy regression discontinuity at the eligibility cutoff, and Clasing-Manquian (2022) estimates a difference-in-differences with predicted eligibility. We complement these approaches by classifying students with the family income reported at the time of the PSU, before any aid application.

SUTVA cannot be tested directly, so the evidence in Appendix Section A.3 is suggestive rather than conclusive: it may not capture displacement across programs or changes within institutions. As an additional safeguard, we restrict the sample to the bottom and top income deciles, dropping students close to the eligibility threshold, who are the most exposed to displacement. This restriction also sidesteps the "disappearance of the seventh decile" (Gezer, 2026), the distortion in enrollment among students just above the threshold.

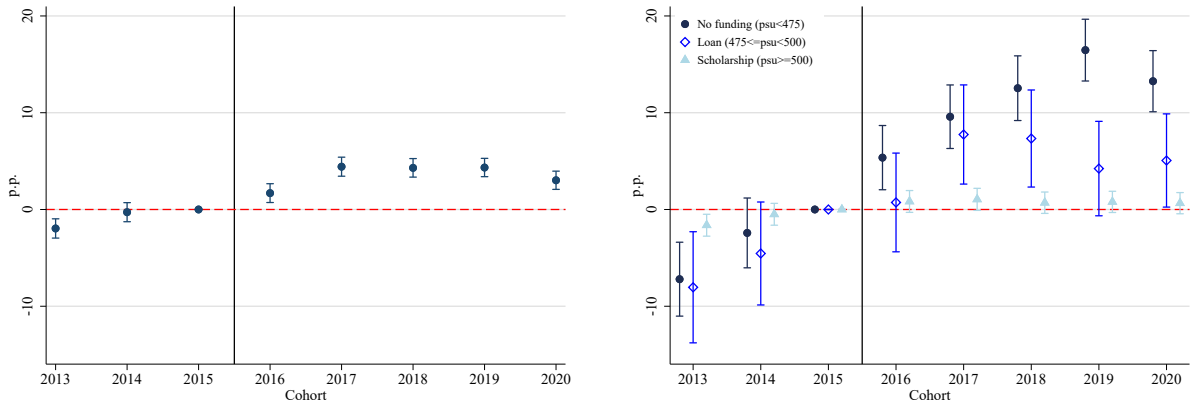
The TWFE estimator can be interpreted as the intent-to-treat effect (ITT) of the policy, given that not all eligible students actually receive TF (some do not enroll, while others enroll in non-eligible institutions).

3.2 Effects on enrollment

We first examine the impact of TF on the extensive margin, i.e., enrollment. The 2016 cohort is the first for which the policy could affect the enrollment decision. Figure 2 shows the results of estimating Equation 1. As shown in panel a), enrollment among students in the bottom three income deciles increased by 1.7 p.p. in 2016 and by around 4.3 p.p. from 2017 to 2019, relative to 2015, compared to students in the top two deciles. Results are similar, though slightly smaller, when the treated group is extended to the bottom five deciles (Appendix Figure A8). This increase in enrollment comes from students for whom the returns to education became positive or who were previously credit constrained. The larger effect, starting in 2017, is consistent with the expansion of free higher education to SCPs. As explained in Section 2, the control group are students in the upper income deciles. We then examine heterogeneity by test-score groups, which determine loan and scholarship eligibility.

In panel b), we observe that lower-ability students, who had limited access to public funding before 2016, are the group that increases enrollment the most. Differences between groups are significant in 2017, when lower-ability students increase enrollment by 9.6 p.p., compared to 1.0 p.p. for students who already had access to merit scholarships.

Figure 2: Enrollment



a) Effect on enrollment

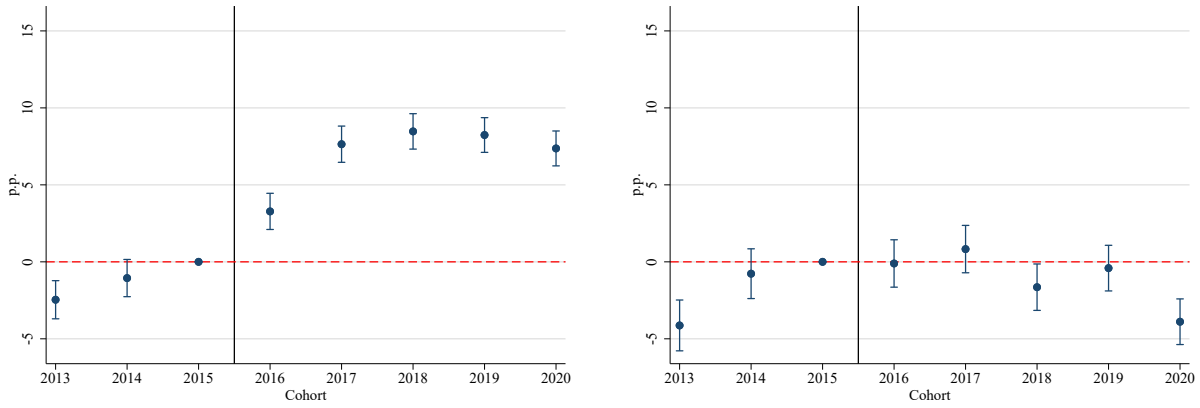
b) Effect on enrollment by funding eligibility

Notes: These figures show the results of estimating Equation 1. Sample is the universe of first time test takers. Panel (a) shows the aggregate results, comparing students in income deciles 1-3 with students in deciles 9-10, using family income self-reported at PSU registration, while panel (b) estimates the equation for students in deciles 1-3 by groups based on their test scores. PSU is the national standardized test. Students with test scores below 475 are not eligible for public funding, students scoring between 475 and 500 are eligible for loans, and students scoring above 500 are eligible for scholarships.

Substitution between eligible and non-eligible institutions Since TF affected around 50% of institutions and 70% of enrollment, it is also important to examine whether there was substitution from non-eligible to eligible centers, or whether enrollment is driven completely by marginal students. Figure 3 reports the estimation on enrollment for eligible and non-eligible institutions. Results show a significant increase in the enrollment of eligible students at institutions eligible for the policy. In panel a), we observe increases of 3.3 and 7.6 p.p. for years 2016 and 2017, respectively, and of 7.4 to 8.5 p.p. from 2018 onwards. In contrast to Londoño-Vélez et al. (2020), panel b) suggests no crowding-out. In fact, we find mostly a null effect on enrollment at non-eligible institutions up to 2019 (a 1.6 p.p. decrease in 2018 is the only estimate different from zero), with a decrease of 3.9 p.p. in 2020.[4]

[4]Appendix Figure A9 disaggregates both estimates by test-score group. The increase at eligible institutions is driven by students without access to public funding (7.8 p.p. in 2016, 14.8 p.p. in 2017 and above 20 p.p. from 2019, against 2 to 3 p.p. for scholarship-eligible students), while the null effect at non-eligible institutions combines a growing decrease among scholarship-eligible students (2 to 7 p.p.) with a 1 to 4 p.p. increase among students without public funding (significant only in 2019), consistent with the policy raising awareness of higher education among low-ability students with enough resources to enroll without TF.

Figure 3: Enrollment by institution eligibility



a) Eligible institutions

b) Non-eligible institutions

Notes: These figures show the results of estimating Equation 1 on enrollment at eligible and non-eligible institutions to the TF policy, comparing students in income deciles 1-3 with students in deciles 9-10, using family income self-reported at PSU registration, as in Figure 2. Sample is the universe of first time test takers; in panel a) students enrolled at non-eligible institutions are dropped, and in panel b) students enrolled at eligible institutions are dropped. Appendix Figure A9 disaggregates both panels by test score group.

Field Preferences We next look at whether students' preferences changed as a result of the policy. We use students' rank-ordered list and construct the dependent variable as the share of a field in a student's list. We restrict the sample to eligible institutions and estimate Equation 1 for each field. Results are presented in Appendix Figure A7. Panel a) shows that TF increased preferences for Health by 0.7 p.p. in 2016 and 1.2 p.p. in 2017 (1.5 p.p. in 2018), although these estimates are imprecise, with no clear effect on Arts and Humanities until 2019. STEM majors lose popularity starting in 2017 with respect to 2015, by 1.8 p.p. in 2017 and 3 to 5 p.p. from 2018 onwards. Panel b) decomposes the results into finer categories. In 2016 Education decreases by around 1 p.p., 10% with respect to the 2015 average, but the effect fades in later cohorts. These results match the evidence found by Castro-Zarzur et al. (2022), who find that TF eclipsed an incentive program to attract talent in Education majors.[5] The increase in Health preferences can be explained by the fact that they are the most expensive programs, along with law majors. Finally, the decrease in STEM majors is driven by Technology, which falls by 1.5 p.p. in 2017 and 4.3 p.p. by 2020.

[5]See Pal (2026) for a comprehensive analysis of this policy.

3.3 Effects on educational outcomes

Recovering treatment effects for the intensive margin is not as straightforward. If we were to estimate educational outcomes using Equation 1, composition and moral hazard could complicate interpretation: the distribution of students' test scores shifted to the left starting 2016, given that a substantial fraction of lower-ability students enrolled (marginal students), as can be seen in Appendix Figure A2. On top of that, from the 2016 cohort onwards eligible students were no longer subject to performance requirements. To keep the composition of students fixed, we exploit a particularity of the policy: cohorts prior to 2016 were eligible to TF for the remaining years of the nominal length of their degrees. For example, if they enrolled in 2015 in a five-year degree, they can be exempted from tuition for the remaining four years, two years for the 2015 cohorts, etc.

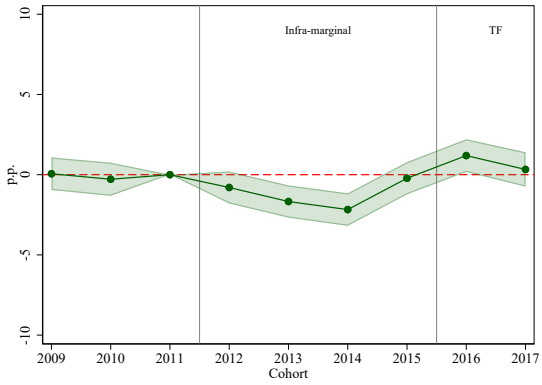
We can use this exogenous variation to partial out selection. 2012 to 2015 cohorts are potentially eligible to TF starting 2016 but cannot change their initial enrollment decision by construction. We adapt Equation 1 as

$$Y_{it} = \sum_{\substack{k=2009 \\ k \neq 2011}}^{2017} \beta_k^c \left(\mathbb{1}\{t = k\} \times \mathbb{1}\{dec_i < 6\} \right) + \gamma^c \, \mathbb{1}\{dec_i < 6\} + \delta_t^c + \mathbf{X_i}' \boldsymbol{\alpha}_x^c + \epsilon_{it}^c \tag{2}$$

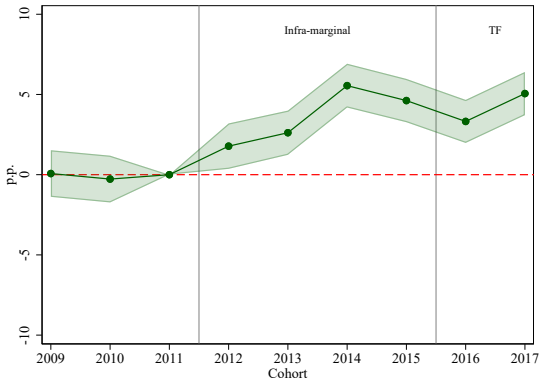
The above regression remains a standard TWFE regression, as the time dimension is the cohort. Treatment varies in intensity by cohort, but effects are always estimated with respect to a never treated cohort, so there are no issues of negative weighting (Callaway et al., 2024; de Chaisemartin and D'Haultfœuille, 2024). The variation exploited is purely cross-sectional.

Figure 4 shows the results of estimating Equation 2. Panels a) to d) show the results for dropout, graduation, field switch, and graduation on time, respectively, for students in deciles 1-3 relative to deciles 9-10. We find that dropout decreases by up to 2.2 p.p. for infra-marginal students (2014 cohort), and the effect reverses for the 2016 cohort (+1.2 p.p.), which is no longer different from zero in 2017. The composition of new entrants seems to dominate the positive effect from infra-marginal students. Graduation rises by up to 5.6 p.p. for students that receive the subsidy after their enrollment decision (2014 cohort), and remains 3.3 and 5.1 p.p. above the 2011 cohort for the 2016 and 2017 cohorts, despite their lower average ability. The probability of switching field falls steadily across cohorts, by about 4.2 p.p. for the 2015 and 2016 cohorts and 3.4 p.p. for 2017. Finally, as the subsidy gets larger (2014 and 2015 cohorts), the proportion of students who graduate on time rises by 5.4 and 6.1 p.p., and the effect declines to 4.5 and 3.2 p.p. for the 2016 and 2017 cohorts.

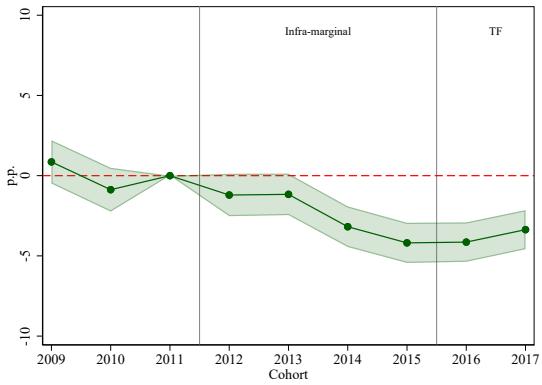
Figure 4: Educational outcomes



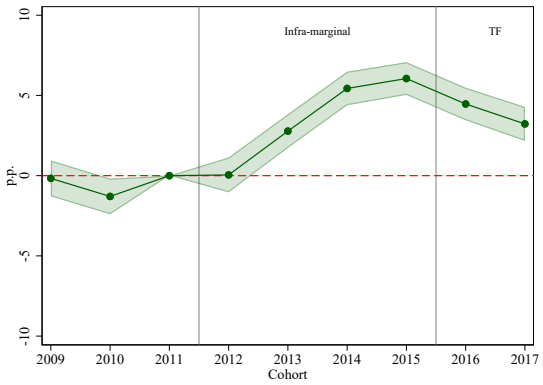
a) Dropout



b) Graduation



c) Field switch



d) Graduation on Time

Notes: These figures show the results of estimating Equation 2 on dropout, graduation, field switch, and graduation on time. Treated students are those in income deciles 1-3, and control students are those in deciles 9-10, using family income self-reported at PSU registration. Sample is enrolled students. Cohorts 2012 to 2015 are potentially eligible to TF starting 2016 but already made their initial enrollment decision. Cohort 2011 is the reference category.

For the 2012 to 2015 cohorts, these results are the overall effect of the policy on those who had already enrolled; for the 2016 and 2017 cohorts, they also reflect the change in the composition of entrants. TF lowers the cost of staying enrolled, which tends to reduce dropout and increase graduation. It does so by lowering out-of-pocket costs and by removing the performance requirements faced by students who previously benefited from scholarships. Conversely, students with lower test scores that now enter higher education face more difficulties in progressing and graduating. Our quasi-experimental findings suggest that both forces are at play: the composition effect partially offsets the gains for post-policy cohorts, which are smaller than for infra-marginal cohorts (and reversed for dropout in 2016) but remain positive and significant for graduation and on-time graduation. Since enrollment also increases, the total number of graduates rises. To validate reduced-form results and quantify both effects, we next build a model of enrollment and major choice with endogenous effort.

3.4 Discussion and robustness

Our quasi-experimental analysis suggests that TF policy had a positive effect on enrollment among eligible students, especially for low-ability students that previously had limited access to public funding. We find no evidence of crowding-out of non-eligible institutions, suggesting that most of the increase in enrollment comes from marginal students. Eligible institutions were able to absorb the increase in enrollment, especially lower-quality institutions. Regarding educational outcomes, we find that the policy reduces dropout and increases graduation for infra-marginal students. Effects are smaller for students that enrolled after the implementation of the policy, who are of lower observed ability on average: graduation and on-time graduation still increase, while dropout no longer decreases.

There are some limitations to our empirical strategy. We used eligibility to TF as a proxy for treatment, and compared eligible to non-eligible students as we do not have a continuous measure of income.

Crowd-out of average ability, high-income students If high-income students were to reduce their enrollment as a result of the policy, our estimates would be biased upwards. However, spillovers are unlikely as cut-offs scores were mostly unchanged, suggesting that the policy did not put additional stress on the capacity of institutions. As a robustness check, we estimate Equation 1 restricting the sample to income deciles 1, 2 and 9, 10, as crowd-out is more likely to occur around the eligibility threshold, and find similar results.

Parallel trends Appendix Section A.1 examines how sensitive the estimates of Figure 2 and Figure 4 are to violations of parallel trends. Following Roth (2022), we plot the coefficients against the linear trend fitted through the pre-period coefficients, and following Rambachan and Roth (2023) we compute confidence sets that remain valid when the post-treatment violation of parallel trends is at most M times the largest pre-treatment one. The enrollment effects of 2017 and 2018 remain different from zero for violations up to half the size of the largest pre-trend deviation ($M = 0.5$), as do the 2017 effect for students without public funding, the graduation gains of the 2012–2015 cohorts, the graduation-on-time gains of the 2014–2015 cohorts and the dropout reductions of the 2013–2014 cohorts. The 2016 enrollment effect, the field-switch estimates and the intensive-margin estimates for the 2016–2017 cohorts are robust at most to smaller violations ($M = 0.25$), and only the 2014 graduation effect survives at $M = 1$.

Increase in test-takers If the policy induced more low-income students to take the PSU test, our control group would be positively selected, biasing our results upwards. However, we find no evidence of an increase in test-takers as a result of the policy (Appendix Figure A14). Take-up rates were already high: around 80% of students in their last year of high school take the test, and close to 90% of those in income deciles 1–5. The number of low-income test takers is flat between the 2015 and 2016 cohorts (119,900 and 120,300 students), while their applications to higher education through the centralized system rise by 13% (from 47,000 to 53,000). What we see is an increase in applications from low-income students, not in test-taking.

Income manipulation Income is self-reported and verified by the government. Misreporting could lead to students that are actually ineligible being classified as eligible. If this is done on purpose, it could bias our estimates upwards. The clearest sign of such manipulation is the "disappearance of the seventh decile" documented by Gezer (2026): after the reform, fewer students report incomes just above the eligibility threshold. As noted in Section 3.1, our specification leaves out the deciles around the threshold and compares deciles 1–3 with deciles 9–10, so students who move across the threshold by misreporting belong to neither the treated nor the control group.

4 Model

The decision process of an individual considering attending higher education can be depicted in two stages. In the first stage, she chooses whether to enroll, and if so, which

field–institution combination to attend. To decide, she observes program characteristics and available funding instruments, and optimizes how much effort to exert in each program in her personalized choice set, determined by her score in the national exam (Fack et al., 2019).[6] Effort choice will determine her distribution of performance outcomes (credits). Conditional on enrollment, performance today endogenously determines out-of-pocket fees for the subsequent year. The student graduates when she accumulates the required number of credits of the enrolled program.

4.1 Timing and choices

Figure 5: Structural model layout

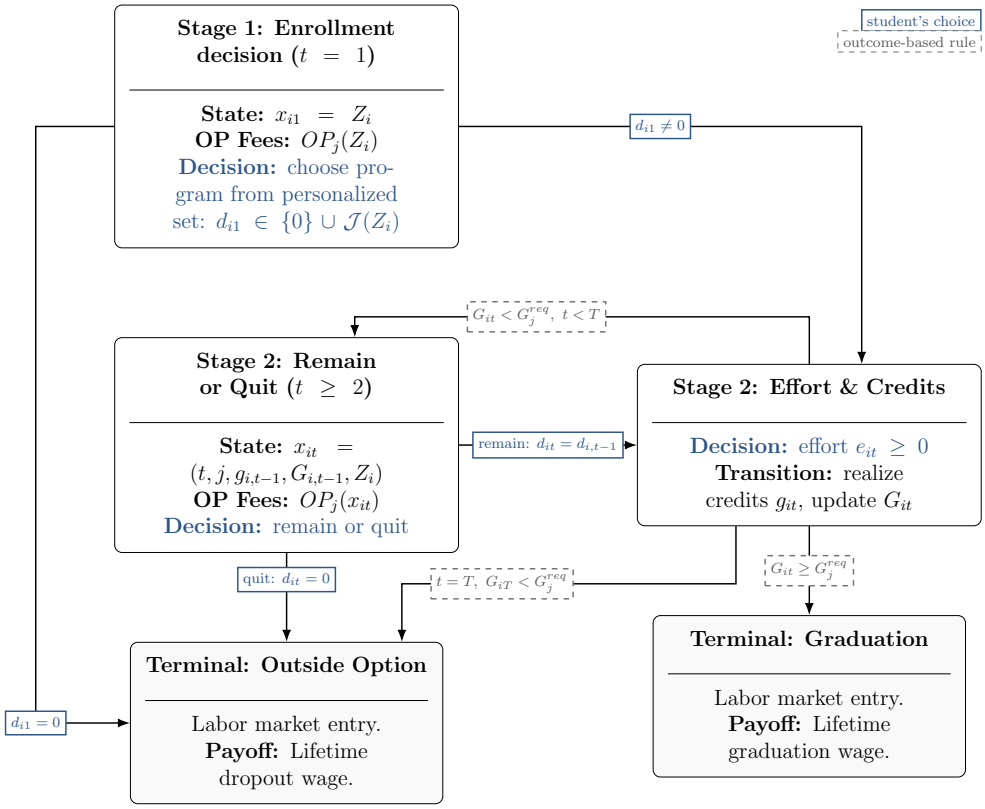


Figure 5 describes the timing and choices of the model. A student is characterized by her type Z_i, defined by sex, ability (test score and high-school GPA, see Section 2.3), socioeconomic status and unobserved trait m. In stage 1, the student decides whether to enroll in higher

[6]Choice sets are student specific: a field–institution combination is feasible for a student in a given year only if her test score is at least the program's admission cutoff for that year. For programs admitting through the centralized platform, the cutoff is the score of the last applicant admitted in the first round; for off-platform programs, it is the lowest test score among first-year enrollees. Short-cycle programs impose no score requirement and are therefore available to all students. Higher-scoring students thus face strictly larger choice sets, and the outside option is always available.

education. She observes her type Z_i, which characterizes her personalized choice set $\mathcal{J}(Z_i)$ and the funding instruments available. She decides between a field–institution combination $d_{i1} = j$, with $j \in \mathcal{J}(Z_i) = \{1, \ldots, J(Z_i)\}$ a program from her personalized choice set, and the outside option $j = 0$, where she would directly earn her expected lifetime non-enrollment wage.

Stage 2 is conditional on enrollment. It is an optimal stopping problem.[7] The student decides in each subsequent period whether to continue in the program, $d_{it} \in \{d_{i,t-1}, 0\}$, and how much effort $e_{it} \in [0, +\infty)$ to exert. By exerting effort, the student implicitly chooses the probabilities of performance outcomes $g_{i,t+1}$, which accumulate over time as $G_{it} := \sum_{s=1}^{t} g_{is}$.

The model allows for $T = 7$ academic periods.[8] They are indexed by $t = 1, \ldots, T$. The nominal length of program j is $\bar{t}_j$. A student can choose to drop out before the terminal boundary and earn the expected lifetime dropout wage. Dropout wages vary with the year of exit, capturing returns to incomplete higher education. In the final academic period she still exerts effort and can graduate; after that performance realization, a nongraduate exits mechanically because an eighth enrollment period is outside the model.

During periods before graduation ($t = 1, \ldots, \bar{t}_j - 1$), performance affects both the probability of dropout and, for scholarship holders, out-of-pocket fees $OP_j(x_{it})$: failing to meet the performance requirement ($g_{it} < 3$) implies paying the full tuition the following year. During the graduation period ($t = \bar{t}_j, \ldots, T$), the student can graduate and obtain the corresponding lifetime expected wage, postpone graduation when $t < T$, or choose to drop out. In period T, effort and graduation remain endogenous, but a nongraduate is forced to exit after performance is observed. Hence, the state space is $x_{it} = (t, d_{i,t-1}, g_{i,t-1}, G_{i,t-1}, Z_i)$.

[7] The model abstracts from switching programs after enrollment: a student who leaves her program is treated as a dropout. We do so for three reasons. First, data availability: separating switching from non-switching requires following students over more periods than our panel provides, so that a switch can be distinguished from a spell of non-enrollment and a student's post-switch outcomes can be compared with those of stayers. Second, our TWFE estimates show that the policy effect operates mainly through the extensive margin and through performance and persistence within the initial program, so switching is not a first-order margin for the effects we study. Third, switching is more naturally rationalized by models in which students learn about their match or ability over time, as in Larroucau and Rios (2024), whereas our model focuses on effort and financial incentives. Treating switches as dropouts therefore affects the level of measured attrition, but not the mechanism of interest.

[8] Since a typical degree program lasts five years, this horizon builds in a two-year margin, allowing completion to be measured with delay. The average program duration is 6.5 years for an undergraduate, 5 years for Professional Institutes and 3 years for technical programs (SIES, 2026).

4.2 Utility, wages and the dynamic program

Flow utility. Flow utility from being enrolled in program j during period t includes individual preferences for different programs, out-of-pocket fees sensitivity, and the psychological or physical "disutility" of academic work. Following De Groote (2025), we decompose it into a fixed cost (FC_j) that captures preferences for higher education, and a variable cost of effort:

$$u_j(x_{it}) + \varepsilon_{ijt} = -FC_j(x_{it}) - c_j(x_{it})e_{it} + \varepsilon_{ijt}, \tag{3}$$

with marginal cost $c_j(x_{it}) > 0$. Fixed costs $FC_j(x_{it})$ summarize all non-effort related preferences for programs including out-of-pocket fees $OP(Z_i, g_{i,t-1}, P_{j,t})$, defined as the fraction of program fees not covered by the state for student i in program j. These depend on tuition $P_{j,t}$, type Z_i and past performance $g_{i,t-1}$:

$$OP(Z_i, g_{i,t-1}, P_{j,t}) = \begin{cases} (1 - \lambda(Z_i, g_{i,t-1}))P_{j,t} & \text{if the student holds a scholarship,} \\ (1 - \lambda(Z_i))P_{j,t} & \text{otherwise.} \end{cases}$$

where $\lambda(\cdot) \in [0, 1]$ is the subsidy rate, i.e. the fraction of program fees $P_{j,t}$ covered by the government. For scholarship holders, $\lambda(\cdot)$ depends on past performance $g_{i,t-1}$ as well as on individual characteristics Z_i, since failure to meet credit requirements results in the loss of the subsidy. For loans and TF, $\lambda(\cdot)$ depends only on observable characteristics and does not vary with performance. Finally, the distribution of the idiosyncratic shock $\varepsilon_{ijt} \sim$ EV1 is iid and common knowledge. It captures what the student learns at the beginning of the period.

Wages. Leaving higher education, either by graduating or by dropping out, is a terminal outcome: the student enters the labor market and earns a stream of annual wages. We denote by $w_\ell(s, x_{it})$ the expected annual wage of a student with s years of labor market experience, where $\ell = j$ if she graduates from program j and $\ell = 0$ if she leaves without a degree. Annual wages depend on experience, gender, region, and the field and institution type of the program. Graduates earn the wage of their field and institution type, while the dropout wage increases with the credits accumulated before exit, capturing returns to incomplete higher education. Students who never enroll earn the non-enrollment wage. Students value their labor market prospects through the lifetime expected wage $W_\ell(x_{it})$, the discounted sum of expected annual wages over a working life of 40 years, with the common discount factor

β:

$$W_\ell(x_{it}) = \sum_{s=0}^{39} \beta^s \, w_\ell\big(s,\ \text{female}_i,\ \text{region}_\ell,\ field(\ell),\ \text{inst_type}(\ell)\big), \qquad \ell \in \{j, 0\}. \tag{4}$$

The same function applies to graduates ($\ell = j$) and to dropouts ($\ell = 0$); the two differ only in the field and institution-type arguments, which for dropouts take their nongraduate values.

Dynamic program. The dynamics of the optimization problem are the following: On the one hand, by exerting more effort today, the student i) increases her probability of accumulating credits and eventually graduating in subsequent periods and therefore raises her expected future wage, and ii) increases her probability to maintain public funding (passing the performance threshold) if eligible. On the other hand, exerting effort is costly and reduces the flow utility of attending higher education. The conditional value function $v_j(x_{it}, e_{it})$ can be written as

$$v_j(x_{it}, e_{it}) + \varepsilon_{ijt} = u_j(x_{it}, e_{it}) + \beta \sum_{\bar{g} \in \mathcal{G}} \phi^{\bar{g}}(e_{it}) \, \bar{V}(x_{i,t+1}(\bar{g})) + \varepsilon_{ijt}, \tag{5}$$

where $\bar{V}(x_{i,t+1}(\bar{g}))$ denotes the ex-ante (continuation) value function evaluated at the post-performance state $x_{i,t+1}(\bar{g})$. The first term corresponds to the current flow utility from enrolling in field–institution j. The second term captures the expected continuation value, discounted by β. Uncertainty in $x_{i,t+1}(\bar{g})$ arises solely from the realization of the performance outcome $\bar{g}$.

Given the Extreme-Value Type 1 assumption on the utility shock, we have a closed-form solution for the ex-ante value function, commonly known as the logsum expression, where γ is the Euler-Mascheroni constant. Instead, when the choice (the outside option) or the performance outcome leads to a deterministic state, the student earns the lifetime expected wage of either graduating or dropping out. Summarizing, the continuation value is given by:

$$\bar{V}(x_{i,t+1}(\bar{g})) = \begin{cases} \gamma + \ln\Big(\sum_{j' \in \mathcal{J}(Z_i)} \exp v_{j'}\big(x_{i,t+1}(\bar{g})\big) \Big) & \text{if } t < T+1 \text{ and } G_{it} < G_j^{req}, \\ \alpha_w \, W_j(x_{i,t+1}) & \text{if } G_{it} \geq G_j^{req}, \\ \alpha_w \, W_0(x_{i,t+1}) & \text{if } t = T+1 \text{ and } G_{it} < G_j^{req}, \\ & \text{or if } d_{it} = 0. \end{cases} \tag{6}$$

4.3 *Effective* effort and performance

Effective effort e_{it} is unobserved, while performance outcomes g_{it} are directly recovered from students' records. We model performance outcome $g_{i,t+1}$ as the result of effort e_{it} and a logistically distributed shock η_{it}:

$$g_{i,t+1} = \kappa \quad \text{if } \bar{g}_{jt}^{\kappa} < \ln(e_{it}) + \eta_{i,t+1} \leq \bar{g}_{jt}^{\kappa+1}, \tag{7}$$

where the cutoff values $\bar{g}_{jt}^{\kappa}$ and $\bar{g}_{jt}^{\kappa+1}$ are field–institution–year specific. In practice, performance is discretized into five categories, $\kappa \in \{0, 1, 2, 3, 4\}$. Students observe the realization of the shock $\eta_{i,t+1}$ only after choosing their effort e_{it}, but since they know the distribution of $\eta_{i,t+1}$ (standard logistic), they can compute the probability of each outcome conditional on effort:

$$\Pr\left(g_{i,t+1} = \kappa \mid e_{it}, x_{it}\right) = F\left(\ln(e_{it}) - \bar{g}_{jt}^{\kappa}\right) - F\left(\ln(e_{it}) - \bar{g}_{jt}^{\kappa+1}\right), \tag{8}$$

where $F(\cdot)$ denotes the logistic CDF. The higher the effort, the higher the expected performance. By choosing effort, students choose the distribution of performance outcomes. Making use of the logistic assumption of the performance shock $\eta_{i,t+1}$, for $\kappa = 0$ we can also rewrite Equation 8 as:

$$e_{it} = \frac{1 - \Pr(g_{i,t+1} = 0 \mid e_{it}, x_{it})}{\Pr(g_{i,t+1} = 0 \mid e_{it}, x_{it})}. \tag{9}$$

Effective effort can thus be interpreted as the odds of avoiding the lowest performance outcome (failing all registered credits).

Students observe realized performance outcomes $g_{i,t+1}$, which determine their state and therefore whether the student is forced out, is allowed to continue in the program or graduates. The model allows for heterogeneous performance thresholds $\bar{g}_{jt}^{\kappa}$ across programs j, capturing differences in grading standards and program requirements. Dropout or graduating entails receiving the lifetime expected wage of a nongraduate or graduate, W_0 or W_j in Equation 4, respectively.

4.4 Solution of the model

The model is solved by backward induction. Higher education is no longer feasible in $T + 1$ after the seventh performance realization, or once the student completes all required program credits, as detailed in Equation 6.

The terminal continuation value is used as an input in earlier periods. Importantly, performance in period T still changes graduation and therefore the relevant wage; only the subsequent enrollment/dropout choice is absent.

At each period t, the student first chooses an effort level e_{it} in her program j. The decision trades off a loss in flow utility with higher future expected gains. Assuming she is rational and behaving optimally, she is solving the following first-order condition for each program[9]:

$$\frac{\partial v_j(x_{it}, e_{it})}{\partial e_{it}} = \frac{\partial u_j(x_{it}, e_{it})}{\partial e_{it}} + \beta \sum_{g \in \mathcal{G}} \frac{\partial \phi^{\bar{g}}(e_{it})}{\partial e_{it}} \bar{V}(x_{it+1}(\bar{g})) = 0 \quad \text{if } e_{it} = e_j^*(x_{it}). \tag{10}$$

At the current-period state x_{it} the ex-ante value function $\bar{V}$ can always be computed, either by the logsum formula when a choice is involved, or directly the lifetime expected wage, as described in Equation 6, and the model is solved recursively by iterating backward to $t = 1$. Given the set of optimal effort levels $\{e_j^*(x_{it})\}_{j \in \mathcal{J}(x_{it})}$, the student chooses the program with the highest value $v_j(x_{it}, e_j^*(x_{it}))$. The resulting choice probabilities take the familiar logit form:

$$p_{jt} = \frac{\exp\big(v_j(x_{it}, e_j^*(x_{it}))\big)}{\sum_{j' \in \mathcal{J}(x_{it})} \exp\big(v_{j'}(x_{it}, e_{j'}^*(x_{it}))\big)}.$$

Here $v_j(x_{it}, e_j^*(x_{it}))$ is the conditional value function of Equation 5, evaluated at the optimal effort level $e_j^*(x_{it})$ implied by Equation 10.

5 Identification and Estimation

This section discusses the identification of the model primitives and describes the estimation procedure.

As is standard in the dynamic discrete choice literature, the discount factor is not separately identified from the per-period utility function (Magnac and Thesmar, 2002). We therefore fix β at 0.95 and normalize the utility of the outside option to the lifetime expected dropout wage.

We further assume that not enrolling in higher education is an absorbing state: individuals who drop out cannot resit the entrance exam or re-enter higher education in a later period.

[9] A sufficient condition for an interior solution is that students believe there is a non-zero probability to avoid the lowest performance outcome. Positive marginal costs ensure an infinite amount of effort is never exerted.

This restriction rules out option-value effects from delayed re-entry and lets the outside option be summarized by a single terminal payoff.

5.1 Identification of unobserved types

The unobserved component of Z_i is a latent student trait $m \in \{1, 2\}$. We assume that students are drawn from a mixture of two types with population probabilities π_m, independent of the observed characteristics Z_i^{obs}. The type is fixed over the student's life cycle, known to the student, and unobserved by the econometrician. It shifts the utility, effort and performance primitives, so that students with identical observed states may behave differently because of their type.

We draw the type independently of Z_i^{obs} because our setting does not raise an initial-conditions problem (Heckman, 1981). That problem arises when the initial state already contains outcomes of type-dependent choices made before the model starts: in Keane and Wolpin (1997), schooling at age 16 reflects earlier decisions of each type, so type probabilities must be conditioned on it. Our model starts at the first higher-education decision, and none of its initial state variables (entrance test score, high-school GPA, gender and socioeconomic status) is an outcome of the choices we model. These variables enter fixed costs, effort and dropout directly, and the type captures the persistent heterogeneity that remains after conditioning on them. Observed characteristics still shift each student's posterior type probabilities, through the likelihood of their history, but not the population prior.

We rely on the identification arguments for dynamic discrete choice models with time-invariant unobserved heterogeneity of Arcidiacono and Miller (2011); see also Arcidiacono and Ellickson (2011). First, conditional on the observed state x_{it} and the type m, choices, performance outcomes, and state transitions are independent across periods, so that the trait is the only source of persistence in unobservables. Second, the transition of the observed state depends on the type only through the observed state and the current choice, and the trait itself does not change over time. We have enough periods of choices and outcomes for the mixture to be identified from the within-student serial dependence in choices: the type-specific conditional choice probabilities and the type probabilities π_m are then jointly identified from the joint distribution of choice sequences, up to a labeling of traits. Under these conditions, the type enters the CCP-based representation of the continuation value exactly as an observed state would, which allows us to estimate the model with the conditional-choice-probability approach used below, replacing the unobserved type by its posterior probability given the observed history.

5.2 Identification of fixed and marginal costs

Marginal costs c_j are identified from the first-order condition (De Groote, 2025). The model attributes higher performance levels, conditional on the same continuation value, to lower marginal costs of effort. Fixed costs then rationalize the observed enrollment shares across programs. The sensitivity to out-of-pocket fees is pinned down by the variation in public funding eligibility, after controlling for observables, which serves as a price shifter.

Wage elasticity α_w is identified through an exclusion restriction. We use local labor market field-region variation. The underlying assumption is that conditional on region and field fixed effects, the interaction provides idiosyncratic local demand shocks that shift utility only through the wage. Region enters the model only through the wage, that is, it is excluded from utility. The region-by-field wage variation (estimated in Section 5.3) therefore shifts utility only through the wage and identifies α_w.

5.3 Estimation

We are able to circumvent full solution methods by representing the future value term as the conditional value function and a non-negative term that depends on the empirical conditional choice probabilities (Hotz and Miller, 1993). Given the Type-1 extreme value assumption on utility shocks, CCPs are of logit type. We specify flexible parametric functional forms for CCPs, effective effort, fixed costs and wage regression. Details on the estimation of these components are provided in Appendix C.

Because effort solves the first-order condition Equation 10, all students sharing the same state x_{it} exert the same effort $e_j^*(x_{it})$, so the index $\ln(e_j^*(x_{it}))$ in Equation 7 is a function of the state alone. We therefore estimate performance directly as a flexible ordered logit of credits (discretized into five categories, 0–4) on the state, with thresholds that vary by field, institution and year.

Annual wages $w_j(s, x_{it})$, which enter the lifetime expected wage $W_j(x_{it})$ of Equation 4, are estimated as

$$\begin{aligned}\log(w_j(s, x_{it})) = \beta_0^w + \beta_{f,field(j)}^w \,\text{female}_i + \beta_e^w \, s + \beta_{e2}^w \, s^2 \\ + \sum_{r=1}^{16} \beta_{r,field(j)}^w \, R_{ir} + \beta_{inst_type(j),\,field(j)}^w + \varepsilon_{ijt},\end{aligned} \tag{11}$$

where R_{ir} is a dummy for region r and $\text{inst_type}(j) \in \{\text{SCP}, \text{Private}, \text{CRUCH}\}$ denotes the institution type of program j. The gender gap and the institution-type premium are both

allowed to vary by field, and the region-field interactions $\{R_{ir}\,\beta^{w}_{r,field(j)}\}$ provide the variation used to identify α_w: conditional on region and field fixed effects, they act as local demand shocks that shift utility only through the wage. Because the wage equation is estimated on survey data, which is not linked to students' higher education history, the unobserved type cannot be accounted for in wages.

Exiting higher education is a terminal action, so $v_0(x) = \alpha_w W_0(x)$. Hotz and Miller (1993) show that the continuation value can be inverted as a function of the outside option continuation value and a log correction term, which corresponds to the conditional dropout probability

$$\bar{V}(x_{i,t+1}(\bar{g})) = \gamma + \alpha_w W_0(x_{i,t+1}) - \ln p_{0,t+1}(x_{i,t+1}(\bar{g})). \tag{12}$$

Replacing Equation 12 in Equation 6, we can write the continuation value as a function of the data, where the dropout probability $p_{0,t+1}$ is estimated as a flexible logistic regression of the observed dropout choice on the state.

From the first-order condition in Equation 10, we can directly recover the expression for marginal costs:

$$c_j(x_{it}) = \beta \sum_{\bar{g}\in\mathcal{G}} \frac{\partial \phi^{\bar{g}}(e_{it})}{\partial e_{it}} \bar{V}(x_{it+1}(\bar{g})) \quad \text{if } e_{it} = e^*_j(x_{it}), \tag{13}$$

where $\bar{V}(x_{it+1}(\bar{g}))$ is computed using the estimated CCPs and dropout wage, and $\frac{\partial \phi^{\bar{g}}(e_{it})}{\partial e_{it}}$ can be directly computed given the logistic assumption about η_{t+1} in Equation 7.

Replacing the expression for marginal costs in the flow utility, and defining

$$a^{\bar{g}}_{ijt} = \phi^{\bar{g}}(e^*_{ijt}) - e^*_{ijt} \left.\frac{\partial \phi^{\bar{g}}(e_{ijt})}{\partial e_{ijt}}\right|_{e_{ijt}=e^*_{ijt}}.$$

Note that the continuation component can be split as

$$\begin{aligned} EMAX_{ijt} &= \sum_{\bar{g}\in\mathcal{G}} a^{\bar{g}}_{ijt} \left[\gamma - \ln p_{0,t+1}(x_{i,t+1}(\bar{g}))\right], \\ EWAGE_{ijt} &= \beta \sum_{\bar{g}\in\mathcal{G}} a^{\bar{g}}_{ijt} W_{\bar{g}}(x_{i,t+1}) \end{aligned}$$

to estimate α_w. $W_{\bar{g}}$ equals the graduation wage on graduation branches and the dropout wage otherwise. At the terminal period T, $EMAX_{ijt} = 0$ by construction, while $EWAGE_{ijt}$

remains nonzero and preserves the effort incentive operating through the probability of graduation.

Finally, we reach the following expression for the conditional value function that we can directly estimate from the data

$$v_j(x_{it}, e^*_{ijt}) = -FC_j(x_{it}) + \beta EMAX_{ijt} + \alpha_w EWAGE_{ijt}. \tag{14}$$

Estimation algorithm: We estimate the model jointly with an expectation-maximization (EM) algorithm in the spirit of Arcidiacono and Miller (2011), in which the posterior probability q_{im} that student i has unobserved trait m is updated using both performance and choices. The trait distribution π_m is a population prior, common to all students and defined for every student in the structural sample, including those who never enroll, who inform it through their enrollment decision. The algorithm proceeds as follows; the full description is in Appendix Section C.4.

1. **Wages.** We estimate the wage equation Equation 11 on survey data and predict graduation and dropout wages for every program-state. Since the survey is not linked to students' higher education histories, wages do not depend on the unobserved trait and are estimated once, outside the EM loop.

2. **Starting values.** Following Arcidiacono and Miller (2011), the utility parameters start at the estimates of the model without unobserved heterogeneity, and the posteriors q_{im} at 1/2 for each type, perturbed by independent uniform draws so that the types can separate. We repeat the estimation from three such perturbations and from the two-step classification of Bonhomme et al. (2022); the converged estimates are stable across starts (Appendix Section C.4).

3. **EM loop.** Iterate over the following steps until the posteriors, the parameters and the objective stop changing:

   (a) *Reduced-form components.* Weighting each student by q_{im}, we estimate the ordered logit of credit completion and the dropout logit, both with trait shifters. The dropout probabilities deliver the continuation values $\bar{V}(x_{i,t+1}(\bar{g}))$ through Equation 12, and hence $EMAX^m_{ijt}$, $EWAGE^m_{ijt}$ and $c^m_j(x_{it})$ through Equation 13, separately for each trait.

   (b) *Fixed costs, utility parameters and trait prior.* Given the continuation values, we maximize the mixture likelihood of performance and choices over the fixed-cost and utility parameters, including FC_j and α_w, and over the type probabilities π_m.

(c) *Posterior updating.* We update q_{im} as the probability of trait m given the student's observed characteristics, performance and choices.

4. **Marginal costs.** Once the algorithm has converged, we compute $c_j(x_{it})$ from the FOC Equation 13 using the final dropout CCPs, performance ordered logit and posteriors.

Standard errors: The standard errors currently reported are second-step standard errors: they treat the first-stage estimates, the continuation values and the posteriors q_{im} as known, and therefore do not account for first-stage estimation error (the generated-regressor problem). Consistent standard errors are obtained by a nonparametric bootstrap clustered at the student level: we draw students with replacement from the structural sample, keeping each student's full history, and re-run steps 2–4 of the algorithm on each replicate. Re-estimating the whole procedure propagates both the uncertainty of the generated regressors $EMAX^m_{ijt}$ and $EWAGE^m_{ijt}$ and that of the posterior weights q_{im}. Wages are estimated on a separate survey and are held fixed across replicates. These bootstrap standard errors are being computed and will replace the conditional ones in the next version.

6 Results

This section first presents reduced-form estimates that predict dropout probabilities and performance outcomes. Then, structural parameters are discussed, which include students' fixed and marginal costs.

6.1 Dropout probabilities

Appendix Table B3 reports the results of a logistic regression model of dropout, and Appendix Figure A18 plots the implied dropout probability by year of enrollment and latent type, overall and by institution tier. Past performance is the strongest predictor of dropout: students who passed no credits in the previous year are far more likely to leave than those in the reference category, and the gradient is monotone across performance categories. Conditional on progress, higher ability lowers dropout, through both absolute ability (test score and high-school GPA) and the student's ability relative to their program peers; the protective effect of relative ability is stronger for low- and middle-SES students. Women are less likely to drop out.

Out-of-pocket fees raise the probability of dropping out. The effect is larger for women and for low- and middle-SES students, and smaller for students with higher test scores. Conditional on ability, progress and fees, low- and middle-SES students are slightly less likely

to drop out than high-SES students. Finally, the latent type shifts persistence substantially: at a given state, Type 2 students have roughly half the odds of dropping out of Type 1 students.

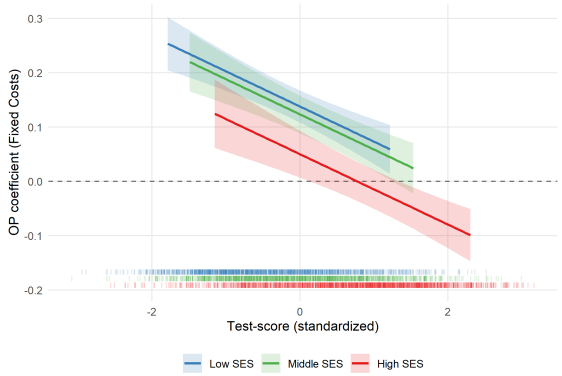
6.2 Performance

Appendix Table B4 presents the results of the ordered logit model of performance with five categories (0–4). Performance is highly persistent: credits passed in the previous year are the strongest predictor of current performance, and this persistence is weaker for students with higher test scores. Once product fixed effects are included, ability matters mainly through the student's position relative to their program peers: relative ability strongly raises performance, particularly for low- and middle-SES students, while the coefficients on absolute test score and GPA are small. Middle-SES students perform slightly better than high-SES students, and studying outside one's home region lowers performance.

Unlike dropout, performance responds little to out-of-pocket fees. The coefficient is small and positive for high-SES students and close to zero for low- and middle-SES students, so fees act mainly on the extensive margin rather than on credit accumulation. The latent type barely shifts performance at a given state (Appendix Table B4).

6.3 Fixed and Marginal Costs

Figure 6: Fee sensitivity by SES and test score



Notes: Coefficient on OP fees (thousand USD) in fixed costs as a function of the standardized test score, by SES group, with gender and GPA at the group means; positive values are costs. Shaded areas: 95% confidence bands. Ticks at the bottom: test-score distribution of each group. Each line spans the 5th–95th percentiles of its group.

Estimates of fixed costs and expected wage are presented in Table 1. Students' fixed costs increase with out-of-pocket fees for low- and middle-SES students, whereas the sensitivity of

Table 1: Fixed costs and marginal costs

	Fixed costs		Marginal costs (log)	
Female	-1.201***	(0.173)	0.409***	(0.070)
Low SES	0.007	(0.232)	-0.171**	(0.087)
Middle SES	0.328	(0.248)	-0.339***	(0.093)
PSU	-1.780***	(0.147)	-0.379***	(0.055)
GPA	-0.298***	(0.115)	-0.078	(0.048)
Type 2	-3.521***	(0.113)	0.144***	(0.008)
Enrol HE	2.086***	(0.166)		
OP fees (thousand USD)	0.031	(0.023)	-0.031***	(0.008)
OP × female	0.020	(0.016)	-0.023***	(0.006)
OP × low SES	0.098***	(0.026)	0.034***	(0.009)
OP × middle SES	0.083***	(0.025)	0.033***	(0.009)
OP × psu	-0.065***	(0.012)	0.014***	(0.004)
OP × gpa	0.022**	(0.011)	-0.011***	(0.004)
Credits (t-1)			-0.946***	(0.010)
Credits (t-1): 0	3.075***	(0.153)		
Credits (t-1): 1-2	1.178***	(0.118)		
Credits (t-1): 4	0.224*	(0.115)		
Credits (t-1) × female	-0.047*	(0.028)	-0.009	(0.006)
Credits (t-1) × low SES	-0.173***	(0.035)	0.035***	(0.007)
Credits (t-1) × middle SES	-0.213***	(0.039)	0.057***	(0.008)
Credits (t-1) × psu	-0.053***	(0.019)	0.094***	(0.004)
Credits (t-1) × gpa	0.007	(0.018)	-0.004	(0.004)
Continuation value				
EMax	0.950	(.)		
Expected wage (USD million)	10.232***	(1.038)		
Type 2 share (population)	0.596			
Type 2 share (enrollees)	0.822			
Product FE	Yes			
Field FE			Yes	
Institution FE			Yes	
Period FE			Yes	
Students	13,974		9,136	

Selected coefficients. The full set of estimates is reported in Appendix Table B5 and Appendix Table B6, and the functional forms in Equation 20 and Equation 13. Fixed costs are expressed as costs (all rows but EMax and the expected wage). Standard errors in parentheses; they do not account for the estimation error in the CCPs, continuation values and type posteriors.

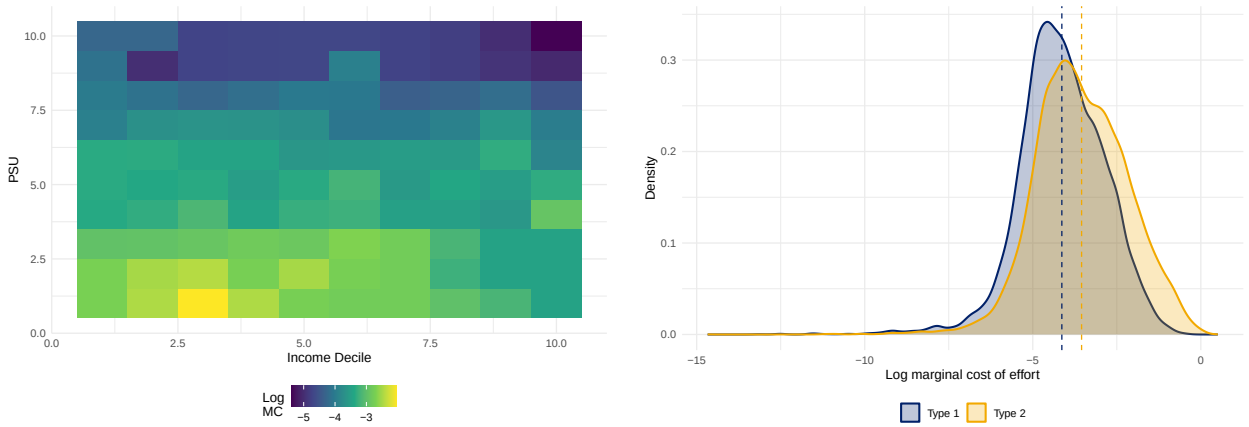
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

high-SES students is small and not statistically significant. Students with higher test scores are less sensitive to fees (Figure 6). Students also face higher fixed costs when the program is located in a different region, especially low- and middle-SES students and higher-ability students (Appendix Table B5). Female students and students with higher test scores and GPA face lower fixed costs on average.

Past progression significantly affects fixed costs: passing no credits in the previous year raises them sharply, and low performance (categories 1–2) raises them moderately. The cost reduction from passing more credits is larger for low- and middle-SES students and for students with higher test scores. Expected wage has a positive effect on the conditional value function, indicating that students value their future expected earnings.

The latent type, which is independent of observed characteristics, captures an intrinsic taste for higher education. Type 2 students face a fixed cost of being enrolled lower by 3.5 utility units, the largest coefficient in the table, and their odds of dropping out at a given state are half those of Type 1 (Appendix Table B3, Appendix Figure A18); accordingly, they make up 60 percent of the population but 82 percent of enrollees. They pass barely more credits at a given state (Appendix Table B4), so the first-order condition attributes to them a slightly higher marginal cost of effort (0.14 log points): they stay enrolled because they value being enrolled, not because studying is less costly for them.

Figure 7: Marginal Costs



Notes: Left: heatmap of log marginal costs by test score (y-axis) and income (x-axis) decile. Right: density of log marginal costs by latent type, weighted by posterior type probabilities; dashed lines mark type means.

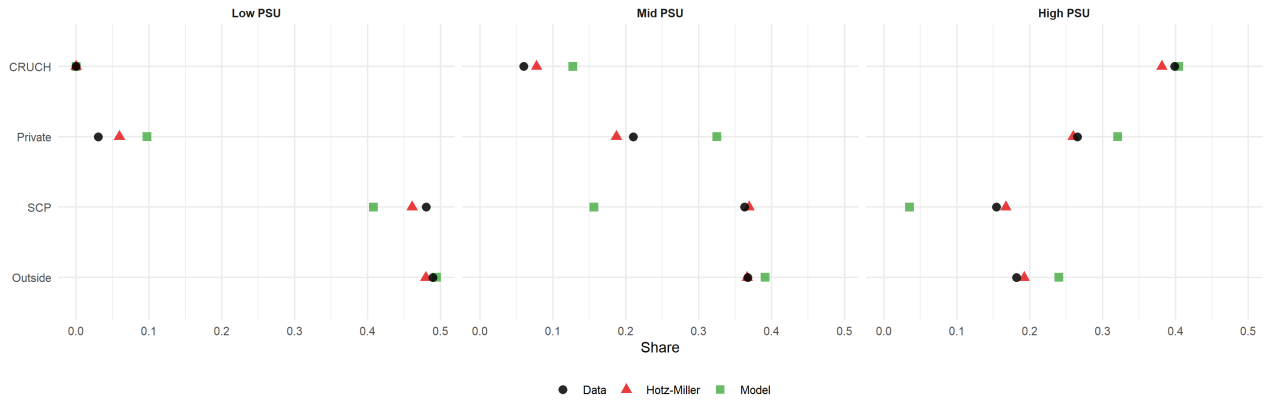
Figure 7 shows the distribution of log marginal costs by income and test score (left) and by latent type (right). The heatmap shows a steep gradient in ability and a weaker one in income, concentrated among students with low test scores.

The second column of Table 1 presents, for interpretational purposes, the unit-weighted projection of log marginal costs over the complete reachable state space. The fitted marginal cost decreases with test score and with credits passed in the previous year, is lower for low- and middle-SES students than for high-SES students, and is higher for female students. The interaction coefficients allow these gradients to vary with student characteristics.

6.4 Model Fit

6.4.1 In-sample Validation

Figure 8: Fit by test-score tercile



Notes: First-year choice shares by test-score tercile. Data: observed shares. Hotz–Miller: shares implied by the conditional choice probabilities estimated in the EM algorithm. Model: shares obtained by solving the model at the estimated parameters. CRUCH: universities of the Council of Rectors (traditional universities). Private: universities outside CRUCH. SCP: short-cycle programs at any institution. Outside: not enrolled in higher education.

Figure 8 compares first-year enrollment shares in the data with two sets of predictions, by test-score tercile. The Hotz–Miller predictions build continuation values from the estimated conditional choice probabilities and evaluate the choice model at the maximum likelihood estimates. They track the data closely in every tercile, as they should: with a full set of program fixed effects, the first-order conditions of the likelihood make predicted shares match observed shares almost by construction. They are therefore a check on the estimation, not a test of the model.

The model predictions are a more demanding test. They solve the dynamic program by backward induction with optimal effort, so the continuation values come from the model itself rather than from the estimated conditional choice probabilities. These probabilities are not used in the solution: they describe behavior under the policy environment observed in the data and would change in a counterfactual. The model therefore fits the data only if

the progression and dropout probabilities it generates are close to those estimated in reduced form. The model reproduces the main patterns: almost no student in the bottom test-score tercile enrolls in CRUCH institutions, and the share choosing the outside option falls with ability. The fit across institution types is less precise. The model tends to overpredict enrollment in private institutions and to underpredict enrollment in SCPs among students in the middle and top terciles, and it somewhat overpredicts the outside option in the top tercile. A plausible explanation for the SCP gap is that the model-solved continuation values imply fewer SCP graduations than observed, and since survey-based SCP graduation wages are quite low, the value of SCPs at entry is understated and fewer students choose them in the first period.

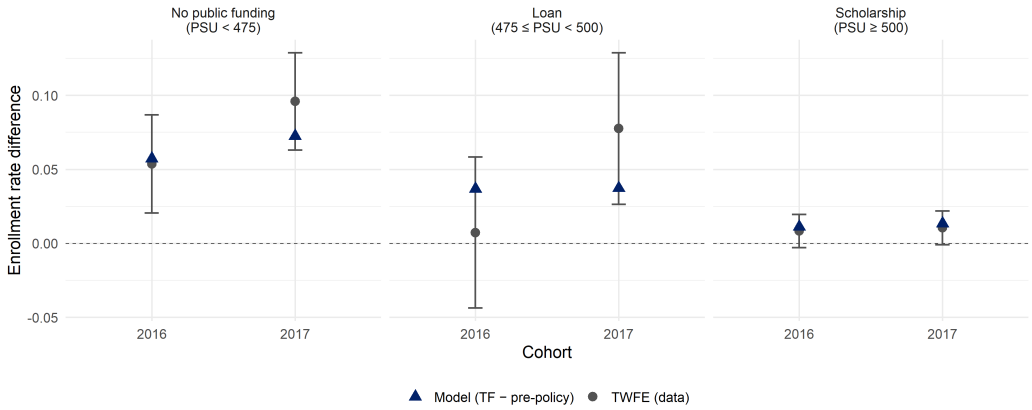
6.4.2 Out-of-sample Validation

The structural model is estimated entirely on post-reform (2016–2017) data, so its parameters are never directly disciplined by pre-reform behavior. To check that the model is not simply curve-fitting the post-reform scenario, we use it to simulate a *pre-policy* counterfactual that reverts the funding rule to the regime that prevailed before 2016: the *Gratuidad* coverage rate is set to zero, so eligible students face the same test-score- and income-conditional loan and performance-contingent merit-scholarship schedule as everyone else, with no unconditional tuition benefit.

We then compare the model-implied difference between the *Gratuidad* baseline and this simulated pre-policy regime to the reduced-form TWFE estimates of the enrollment response in Section 3.2, separately by cohort and by test-score group, matching the TWFE funding groups (no public funding, loan-only, and scholarship-eligible). Because the TWFE variation is never used to fit the structural parameters, agreement between the two designs indicates that the model captures the reform's actual effect on enrollment and is not an artifact of the post-reform sample it was estimated on.

Figure 9 shows that the model reproduces the 2016 enrollment response and its gradient across test-score groups: about 6 p.p. among students without public funding against 5 p.p. in the data, and close to zero for scholarship-eligible students in both. In 2017 the model preserves the gradient but understates its size, predicting 4 and 7 p.p. for the loan and no-funding groups against estimates of 8 and 10 p.p., although these gaps lie within the sampling uncertainty of the TWFE estimates. This out-of-sample response is the moment that matters most for the counterfactuals below, since it governs who is drawn into higher education by a change in tuition. The model gets its sign, magnitude in 2016, and gradient right, even if it misses part of the 2017 jump.

Figure 9: Out-of-sample fit: enrollment



Notes: TWFE estimates of the change in the enrollment rate (dots, with confidence intervals) and model-implied difference between the *Gratuidad* baseline and the simulated pre-policy regime (triangles, with simulation intervals), by test-score group and cohort.

7 Counterfactuals

We use the estimated model to simulate counterfactual funding policies, evaluating impacts on both enrollment and intensive (effort, credit completion, graduation) margins. Each policy is defined as a change to the funding rule $OP(x_{it})$ (income and test-score eligibility cutoffs, coverage rate, and performance conditioning), holding the estimated cost primitives $FC_j(\cdot)$, $c_j(\cdot)$, and wages fixed. For each scenario we resolve the full backward induction over the state space and forward-simulate outcomes by Monte Carlo, drawing performance and choice shocks many times per student to average out simulation noise, with effort chosen optimally each period from the FOC in Equation 10. Each student's latent type is drawn once from her posterior type probabilities and held fixed across scenarios (Appendix Section C.4). Throughout, the reference scenario is the simulated pre-TF regime; each counterfactual is reported as a change relative to this reference.

7.1 Decomposing the effects of each policy

A change in the funding rule affects outcomes through two distinct channels: it changes *who* enrolls and in which program (selection), and it changes how students who would have enrolled anyway behave once enrolled (behavior), i.e. their effort, progression and dropout. The reduced-form evidence in Section 3 has difficulty separating the two, since post-reform cohorts combine both channels; the structural model separates them for any policy, because

we simulate each student under every regime.[10]

Let $r \in \mathcal{R}$ index policy regimes, where $\mathcal{R}$ contains the reference (status-quo) regime PRE, the TF regime as implemented, TF, and the counterfactual designs studied below. Under regime r, $D_i^r \in \{0,1\}$ denotes the enrollment decision of student i and j_i^r her chosen program (institution-major combination). We classify students by comparing TF with PRE:

- **infra-marginal students (IM):** enroll in the *same program* under both regimes, $D_i^{\mathrm{TF}} = D_i^{\mathrm{PRE}} = 1$ and $j_i^{\mathrm{TF}} = j_i^{\mathrm{PRE}}$;
- **switchers:** enroll under both regimes but in a *different program*, $D_i^{\mathrm{TF}} = D_i^{\mathrm{PRE}} = 1$ and $j_i^{\mathrm{TF}} \neq j_i^{\mathrm{PRE}}$;
- **marginal students:** enroll only under TF, $D_i^{\mathrm{TF}} = 1$, $D_i^{\mathrm{PRE}} = 0$.

The groups are defined once, from TF versus the status quo, and held fixed across all counterfactual scenarios, so that each group is comparable across columns. Let Y_i^r denote a potential outcome (e.g. graduation) and π_{IM} the population share of infra-marginal students. For each scenario r, the total effect decomposes exactly as

$$\underbrace{\mathbb{E}[Y_i^r] - \mathbb{E}[Y_i^{\mathrm{PRE}}]}_{\Delta Y_{\mathrm{total}}} = \underbrace{\pi_{IM}\,\mathbb{E}[Y_i^r - Y_i^{\mathrm{PRE}} \mid IM]}_{\Delta Y_{\mathrm{behavior}}} + \underbrace{\Delta Y_{\mathrm{total}} - \Delta Y_{\mathrm{behavior}}}_{\Delta Y_{\mathrm{selection}}}.$$

$\Delta Y_{\mathrm{behavior}}$ is the change among students whose enrollment and program choice do not change, so it can only come from how funding alters their incentives once enrolled (lower out-of-pocket cost, and the performance conditions where present) and the resulting re-optimized effort and dropout; we do not separate the price and effort components of this response. $\Delta Y_{\mathrm{selection}}$ collects every change that comes through the enrollment margin: marginal students, whose outcome without the policy is mechanically that of a non-enrollee, and switchers, whose change in program is itself a change in the enrollment decision and whose outcome change cannot be attributed to funding alone. Selection can have either sign: marginal entrants are typically individually weaker than infra-marginal students (see Figure 10), but they raise the population graduation rate because they would not have graduated at all otherwise, while a restrictive design that keeps weaker marginal students out raises completion among those who remain. We apply the decomposition to each scenario below, and additionally report outcomes by test-score tercile and for income-eligible students.

[10] Simulations use common random numbers: the $t = 1$ taste shocks and the performance shocks are shared across scenarios, so a student's potential outcomes and choices are comparable across regimes.

7.2 Entry and continuation requirements

We saw that TF substitutes most of the pre-existing aid programs, which were conditional on performance. The removal of performance requirements may have weakened incentives for effort and progression. Pre-existing aid programs had performed well in attracting students (Falco and Reichlin, 2025; Solis, 2017, 2026), but the out-of-pocket cost could still be over 20% of tuition. From the Government's perspective, the current policy is very costly: between 2011 and 2019, spending on higher education increased by 160%, reaching 5.4% of GDP, compared with 2.9% in the rest of the OECD (OECD, 2019).

The counterfactuals below build on the conditions of the pre-existing aid programs described in Section 2.1: subsidized loans required a test score of at least 475, and merit scholarships required a test score of at least 500 and completion of at least 70% of registered credits each year (three of four credit-units in the model's discretized performance measure, $g_{t-1} \geq 3$). They are designed to be budget-improving while preserving the equity target of TF. We report four post-reference scenarios: TF as implemented, without requirements, and three designs that modify the funding rule $OP(x_{it})$ applying specifically to TF. The pre-existing loan and scholarship schedules are left untouched in all four scenarios. The designs differ in *when* the requirement binds: before higher education (at entry) or during higher education (to continue receiving the benefit). The three alternative designs are:

1. **Entry requirement (before higher education).** Impose a minimum test score of 500 for TF eligibility, harmonizing the threshold with the pre-existing merit scholarship. This changes the extensive margin only: no performance condition is added, and no fixed or marginal cost changes for students who remain eligible.
   Model prediction: mechanically reduces Government expenditure by shrinking the eligible pool, and attracts, on net, higher-ability students with lower fixed and marginal costs, so persistence and graduation among the enrolled population should improve through selection alone.

2. **Continuation requirement (during higher education).** Impose the same annual progress requirement used by the pre-existing scholarships ($g_{t-1} \geq 3$) to *retain* free tuition. This restores the performance-contingent component removed by TF, without changing who is initially eligible.
   Model prediction: first-period fixed and marginal costs are unaffected, since eligibility is unchanged. From the second year on, students at risk of falling below the threshold face a discrete jump in $OP(x_{it})$ if they fail to comply, so those with the highest marginal and fixed costs of effort should raise effort the most (or drop out if the increase is too

costly). If the expected utility of a first year is not compensated by this option value, some students may not enroll at all.

3. **Both requirements.** Combine (1) and (2): a minimum test score of 500 for eligibility, and the progress requirement to retain funding once enrolled.
   Model prediction: combines the selection and incentive channels above. Reduces the number of beneficiaries the most of the three scenarios, but the remaining recipients face both a more favorable cost distribution and stronger incentives to progress.

Mechanically, (1) affects selection into TF at $t = 1$; (2) affects $OP(x_{it})$ via g_{t-1} for *all* TF recipients, tightening incentives within a cohort; and (3) compounds both.

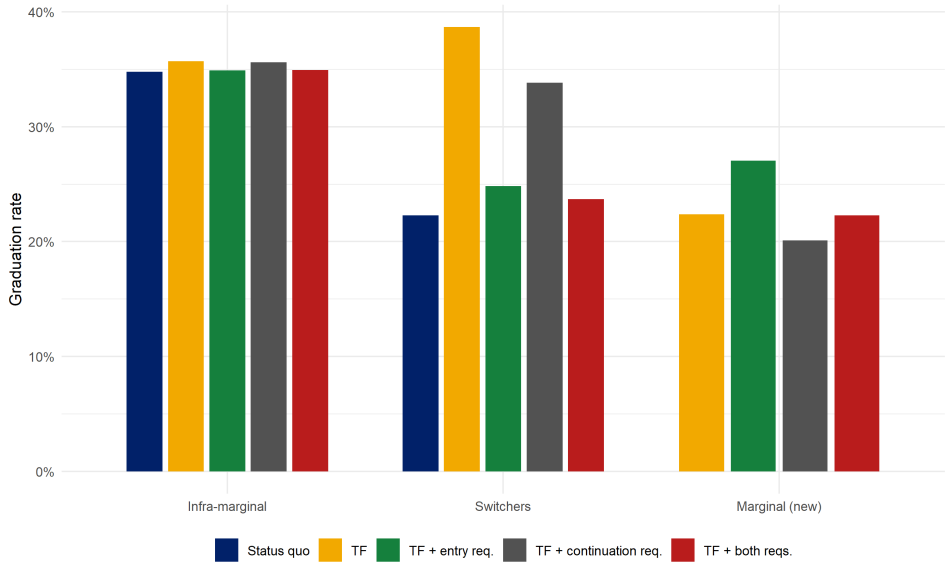
Table 2 summarizes the counterfactual results. The first column reports the status-quo level of each outcome under the simulated pre-policy regime, and the remaining columns report changes relative to it. The panels distinguish enrollment, graduation, graduation conditional on enrollment, and expected lifetime earnings. Changes in enrollment and graduation are in percentage points, and changes in earnings are in percent.

TF without requirements raises enrollment by 3.0 p.p. (from 59.5%) and graduation by 1.9 p.p. (from 20.1%), with the largest gains among low- and average-ability students. The continuation requirement retains most of these effects (+2.4 p.p. and +1.5 p.p.), whereas the entry requirement almost eliminates them (+0.3 p.p. each), and imposing both yields the smallest effects (+0.2 p.p.). Expected lifetime earnings follow the same ordering (+2.9%, +2.1% and +0.5%). The entry requirement is costly because it excludes precisely the students below the 500-point cutoff who respond most to TF. The continuation requirement instead applies only after enrollment, so it keeps most of the access gains and costs graduation mainly among low-ability students who later lose the benefit.

Figure 10 and Figure 11 show where these effects come from. Marginal entrants, drawn mostly from below 450 points, graduate at 22.4% under TF, well below the 35.7% of infra-marginal students. The entry requirement removes most of them, but those who still enroll graduate at a higher rate (27.1%) than under TF, because the requirement screens out the weakest entrants. The overall gain is therefore smaller even though completion among the remaining entrants is higher. Infra-marginal students, 55% of the cohort, account for only 0.5 of the 1.9 p.p. graduation gain under TF: their graduation rises from 34.8% to 35.7%. Marginal entrants contribute 0.7 p.p. The largest contribution, 0.8 p.p., comes from switchers, although they are only 5% of the cohort. They are not the strongest group, since their test scores are concentrated just below the sample mean, yet TF raises their graduation rate from 22.3% in their status-quo programs to 38.7%, the highest of any group, because

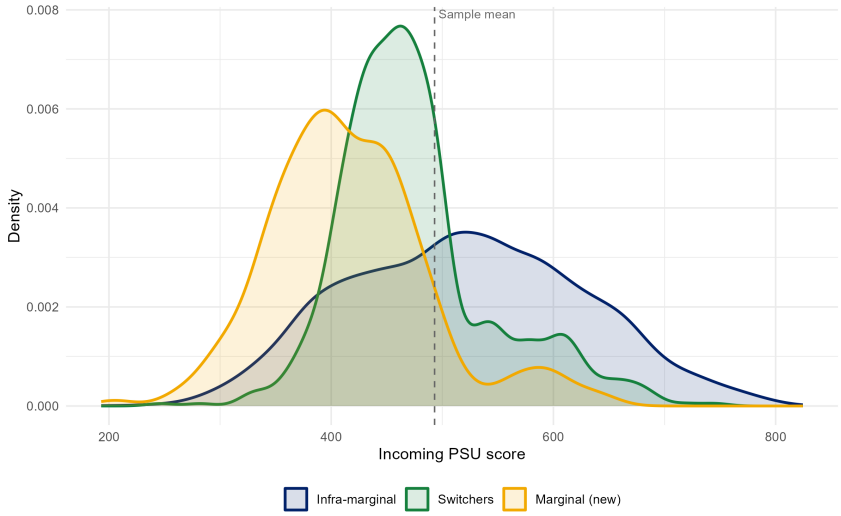
funding lets them move to programs in which they are more likely to graduate. Because they sit just below the 500-point cutoff, the entry requirement removes most of this gain (24.8%), whereas the continuation requirement keeps most of it (33.8%).

Figure 10: Graduation rate by margin group



Notes: Simulated graduation rate by margin group (infra-marginal students enrolled under both regimes, switchers, and new marginal entrants) and counterfactual scenario, as defined in Section 7.1. Rates are among the group's enrollees in each scenario. The status-quo bar for switchers is their graduation rate in the program they choose under the status quo; marginal entrants do not enroll under the status quo.

Figure 11: Incoming test score by margin group



Notes: Distribution of incoming test score by margin group. The dashed line marks the sample mean.

The two requirements also isolate the channels through which funding operates, and we use them to put a price on each. The entry requirement changes only who is eligible at

Table 2: Counterfactual effects on enrollment, graduation and expected earnings

	Status quo	Change relative to status quo			
	level	TF	Entry requirement	Continuation requirement	Both requirements
A. Enrollment (%, p.p.)					
All students	59.5	+2.97	+0.28	+2.40	+0.19
Infra-marginal students	100.0	0.00	0.00	0.00	0.00
Income-eligible students	56.3	+4.85	+0.45	+3.92	+0.31
Low ability	44.9	+5.71	0.00	+4.65	0.00
Average ability	58.3	+2.58	+0.20	+2.14	+0.15
High ability	75.4	+0.63	+0.63	+0.42	+0.42
B. Overall graduation (%, p.p.)					
All students	20.1	+1.94	+0.27	+1.49	+0.18
Infra-marginal students	34.8	+0.91	+0.13	+0.83	+0.14
Income-eligible students	16.9	+3.16	+0.43	+2.43	+0.30
Low ability	6.6	+2.74	0.00	+2.11	0.00
Average ability	17.3	+2.51	+0.24	+1.98	+0.19
High ability	36.5	+0.56	+0.56	+0.37	+0.37
C. Graduation conditional on enrollment (%, p.p.)					
All students	33.8	+1.49	+0.29	+1.09	+0.20
Infra-marginal students	34.8	+0.91	+0.13	+0.83	+0.14
Income-eligible students	30.1	+2.78	+0.52	+2.07	+0.36
Low ability	14.6	+3.77	0.00	+2.89	0.00
Average ability	29.7	+2.86	+0.30	+2.23	+0.24
High ability	48.4	+0.34	+0.34	+0.21	+0.21
D. Expected lifetime earnings (thousand 2015 USD, %)					
All students	166.5	+2.86	+0.52	+2.08	+0.34
Infra-marginal students	206.7	+0.60	+0.14	+0.52	+0.14
Income-eligible students	154.1	+5.04	+0.92	+3.66	+0.59
Low ability	127.0	+4.59	0.00	+3.37	0.00
Average ability	157.0	+4.19	+0.47	+3.11	+0.30
High ability	215.6	+0.87	+0.87	+0.56	+0.56

Notes: The first column reports the average level of each outcome for each group under the simulated pre-policy regime (status quo). The remaining columns are changes relative to that level. The units in parentheses in each panel title are those of the status-quo level and of the change: percentage points (p.p.) for enrollment and graduation, percent of the status-quo level for earnings. TF is unrestricted; the other columns add the named condition. The entry requirement (before higher education) is the minimum test score required to obtain TF; the continuation requirement (during higher education) is the academic-progress condition for keeping TF.

Overall graduation uses the full group and counts non-enrollees as non-graduates. Conditional graduation uses the enrollees in each scenario.

Expected lifetime earnings are the present value (5% discount rate, 40 years of experience) of model-predicted earnings, averaged over all students in the group. Non-enrollees earn the no-higher-education wage, graduates the degree wage of their program, and dropouts the dropout wage, which rises with credits earned. Predictions come from the wage regressions without shocks and measure earnings only: they exclude tuition, effort and other costs, so they are not a welfare measure.

Infra-marginal students are those who enroll in the same program both under the simulated pre-policy regime and under TF (the same program, given common random numbers across scenarios); membership is fixed across scenarios. Their enrollment change is zero by construction, since they enroll under both the status quo and TF and every scenario offers funding in between. Students who enroll under both regimes but in a different program are excluded. Ability groups are test-score terciles. Nearly all low-ability students are excluded from TF eligibility under the entry requirement, so their near-zero deltas there reflect that exclusion, not a lack of data.

$t = 1$ and adds no incentive, so its effect is pure selection; the continuation requirement leaves eligibility unchanged and adds only the progress condition, so its effect runs through the behavior of enrolled students and, through the option value of the first year, through entry. The entry requirement keeps 14% of TF's graduation gain (0.27 of 1.94 p.p.) and the continuation requirement 77% (1.49 p.p.): most of what TF achieves does not depend on admitting weaker students and can be preserved by conditioning on progress. Appendix D translates each scenario into a cost-benefit account per student of the cohort, in present value at entry, and lists its assumptions: the public outlay is split into a mechanical transfer to students whose enrollment path does not change and a behavioral response, the earnings gain is converted into income tax recovered, and students' willingness to pay is divided by the net public cost to obtain the marginal value of public funds (MVPF) of Hendren and Sprung-Keyser (2020). TF costs \$2,100 per student, two thirds of which goes to students who would have enrolled anyway, and buys 1.9 p.p. more graduates at a net cost of \$108,000 per additional graduate (MVPF between 1.1 and 2.2). The entry requirement saves 61% of the outlay but at \$308,000 per additional graduate, nearly three times the cost of TF, because the students it removes are those for whom aid changes the outcome most. The continuation requirement halves the outlay, keeps most of the graduation gain and lowers the net cost per additional graduate to \$70,000, with an MVPF of 1.2 to 3.0 (Table D1). The next subsection uses these two efficiency measures, together with the SES graduation gap, to compare alternative designs.

7.3 Equity and efficiency: redesigning TF

The motivation for free college is equity, but a universal or broadly targeted price cut pays most of its outlay to students whose decision it does not change. We therefore ask whether the same objectives can be reached at lower cost, and what the tradeoff between equity and efficiency looks like. We measure efficiency by the net public cost per additional graduate and the MVPF, and equity by the graduation gap between high- and low-SES students, i.e. the graduation rate of high-SES minus that of low-SES students in the cohort. Designs are compared with TF holding its outlay or its graduation gain fixed. Besides the continuation requirement, we consider two designs that reallocate funding toward low-income students:

1. **Phase-out and stipend for the poorest.** TF covers tuition fully up to income decile 4, half of tuition at decile 5 and nothing from decile 6 (a continuous phase-out in the spirit of Gezer, 2026), and students in income decile 1 receive a cash stipend of \$940 (2015 prices) for each year they are enrolled up to the nominal length of the

program.

2. **Phase-out and stipend for deciles 1–3.** Same phase-out, with the stipend extended to the three lowest income deciles.

The stipend is modeled as a negative out-of-pocket cost: it enters the fixed and marginal costs through the same (linear) price coefficient estimated from tuition variation, which assumes that a dollar of stipend is valued like a dollar of tuition.

Table 3 and Figure 12 summarize the results. TF is not on the efficiency-equity frontier. The first design, phase-out and stipend for decile 1, delivers the same graduation gain as TF (1.94 p.p.) with 10% less outlay, a smaller SES gap (11.9 against 12.1 p.p.) and a lower cost per additional graduate ($96,000 against $108,000): the outlay for deciles 5 and 6 buys few graduates because aid changes the behavior of few students there. The tradeoff appears at the ends of the frontier. The continuation requirement is the efficient end, with the lowest cost per additional graduate ($70,000) and the highest MVPF, but it widens the SES gap by 0.7 p.p. relative to TF, giving back about a quarter of TF's reduction in the gap (from 15.0 to 12.1 p.p.). The stipend for deciles 1–3 is the equitable end: it narrows the gap to 11.0 p.p. and raises low-SES graduation by 4.0 p.p. (against 2.9 p.p. under TF), but requires one third more outlay and a higher cost per additional graduate ($115,000), with an MVPF below 2. The entry requirement is dominated on both dimensions. In this model efficiency and equity are aligned in the reallocation from middle-income to low-income students, because the students whose behavior funding changes are almost all low-SES; they conflict when the outlay is varied, or when aid is made conditional on progress.

8 Conclusion

This paper investigates the impacts of Tuition-Free Higher Education policies on student enrollment, persistence, and academic performance, using the 2016 Chilean reform as a natural experiment. Our TWFE analysis reveals that the policy significantly boosted enrollment amongst eligible students by around 4 p.p., driven by lower-ability students, whose enrollment rose by up to 16 p.p. Effects on educational outcomes are estimated for infra-marginal students, with a decrease in dropout of up to 2.2 p.p. and an increase in on-time graduation of up to 6.1 p.p. These effects diminish for post-reform cohorts: graduation and on-time graduation gains persist but are smaller, and the dropout reduction vanishes.

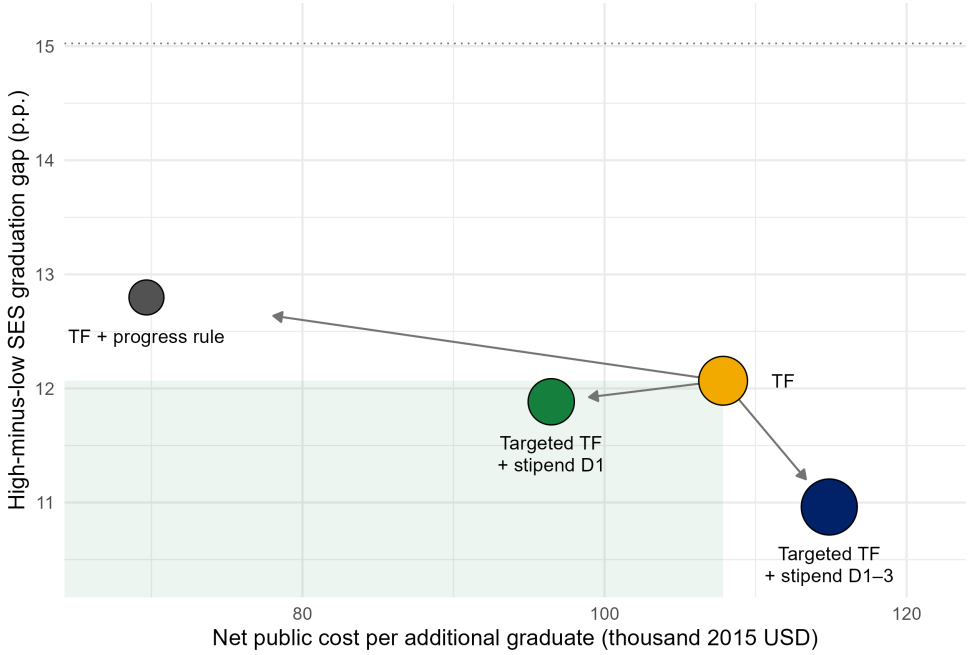
We build a structural model that endogenizes effort choices, capable of disentangling selection effects from behavioral responses, finding that the removal of performance-based

Table 3: Equity and efficiency of alternative funding designs

	Status quo	TF	Continuation requirement	Phase-out + stipend (decile 1)	Phase-out + stipend (deciles 1–3)	Entry requirement
A. Outlay and outcomes (status-quo level; change vs status quo)						
Public outlay ($000)		2.16	1.08	1.93	2.87	0.83
% of TF		100	50	90	133	39
Enrollment (%, p.p.)	59.5	+2.97	+2.40	+3.09	+4.17	+0.28
Graduation (%, p.p.)	20.1	+1.94	+1.49	+1.94	+2.43	+0.27
PV earnings ($000)	141.9	+2.45	+1.97	+2.40	+2.88	+0.45
B. Efficiency						
Net public cost ($000)		2.09	1.04	1.87	2.79	0.82
Net cost per extra grad. ($000)		108	70	96	115	308
Average outlay per grad. ($000)	11	20	15	18	22	15
MVPF, conservative		1.06	1.18	1.00	0.91	1.54
MVPF, upper bound		2.19	2.98	2.25	1.96	2.04
C. Equity						
SES graduation gap (p.p.)	15.02	12.07	12.80	11.88	10.96	14.55
change vs status quo		-2.96	-2.23	-3.14	-4.06	-0.48
Low-SES graduation (%, p.p.)	15.0	+2.89	+2.20	+3.12	+4.04	+0.40
Outlay share, low SES (%)	59	65	63	73	78	60

Notes: Per student of the simulated cohort, thousand 2015 USD, present value at entry (5% discount rate), each design relative to the simulated status quo. The phase-out keeps TF at full coverage up to income decile 4, covers half of tuition at decile 5 and nothing from decile 6. The stipend is a cash payment of $940 (2015 prices) per enrolled year, up to the nominal length of the program, to students in the indicated income deciles; it enters the model as a negative out-of-pocket cost. The continuation and entry requirements are those of Section 7.2. Public outlay counts TF at the regulated tariff (80% of list tuition), scholarships, Nuevo Milenio, BVP and stipends; loans are not counted. Net public cost subtracts the present value of the Chilean income tax on the earnings gain. Net cost per extra graduate divides net public cost by the additional graduates caused by the design relative to the status quo; it is an incremental measure that includes transfers to students whose behavior does not change. Average outlay per graduate divides all public aid under the design (status-quo aid plus the change in outlay) by the graduation rate of the cohort under the design. Conservative WTP is the out-of-pocket saving on status-quo enrollment paths, and the upper bound adds the after-tax earnings gain net of extra out-of-pocket spending; MVPF = WTP / net cost. The gap is the graduation rate of high-SES minus low-SES students in the cohort (status quo: 15.02 p.p.). Simulation noise in the gap differences between designs is below 0.03 p.p. (100 replications); it does not include parameter uncertainty.

Figure 12: Equity and efficiency of alternative funding designs



Notes: Each point is a funding design simulated with the estimated model. The horizontal axis is the net public cost per additional graduate (outlay net of the income tax on the earnings gain, divided by the additional graduates relative to the status quo) and the vertical axis the graduation gap between high- and low-SES students; lower is better on both. Arrows start at TF (as implemented). "TF + progress rule" is the continuation requirement, and "Targeted TF" the phase-out of TF between income deciles 4 and 6, combined with the stipend for decile 1 or for deciles 1–3. Bubble area is proportional to the public outlay per student: the continuation requirement costs 50% of TF's outlay, phase-out and stipend for decile 1 costs 90% and the stipend for deciles 1–3 costs 133%. The shaded area is better than TF on both dimensions and the dotted line is the status-quo gap (15.0 p.p.); the status quo has no position on the horizontal axis, which is measured relative to it. The entry requirement (cost per additional graduate of $308,000 and gap of 14.6 p.p.) is omitted to preserve the scale. Assumptions of the cost-benefit accounting are listed in Appendix D.

requirements did not lead to reduced effort or weakened outcomes overall. The policy induced a composition effect by attracting students with a higher cost of effective effort, while infra-marginal students responded to the policy by increasing effort, as lower out-of-pocket costs raise the value of remaining enrolled and hence the return to progressing.

Counterfactual simulations show that the design of TF matters for its cost more than for its effect on access. Two thirds of TF's outlay is a transfer to students who would have enrolled anyway, and each additional graduate costs about $108,000. Reinstating the eligibility conditions of earlier aid programs changes the cost per graduate considerably, and which condition is reinstated matters. An entry threshold on the test score is the least efficient design, since it excludes the students for whom aid matters most. A continuation requirement halves the outlay and keeps three quarters of the graduation gain, but the graduation it gives up is concentrated among low-ability students who later lose the benefit. Better designs reallocate rather than condition: phasing the benefit out for middle-income students and paying a stipend to the poorest achieves the same graduation gain as TF at lower cost and with a smaller socioeconomic gap.

Our analysis holds the supply side and the labor market fixed. Future research could incorporate supply responses: TF removed the main levers institutions used to adjust, setting regulated fees and capping vacancies, so institutions may respond through other margins, such as opening or closing programs. A second extension is the general equilibrium effect on the labor market, where a larger supply of graduates may lower the returns to a degree.

References

Abbott, B., Gallipoli, G., Meghir, C., and Violante, G. L. (2019). Education policy and intergenerational transfers in equilibrium. *Journal of Political Economy*, *127*, 2569–2624.

Angrist, J., Autor, D., Hudson, S., Pallais, A., Cantoni, E., Caldwell, S., Kim, O., Malone, B., Oyewole, K., and Scott, K. (2015). *Leveling up: Early results from a randomized evaluation of post-secondary aid*, National Bureau of Economic Research.

Arcidiacono, P., and Ellickson, P. B. (2011). Practical methods for estimation of dynamic discrete choice models. *Annual Review of Economics*, *3*, 363–394.

Arcidiacono, P., and Miller, R. A. (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, *79*, 1823–1867.

Arslan, H. A. (2026). Pay or prepare? how tuition shapes student effort in college admissions. *SSRN Electronic Journal*.

Bailey, M. J., and Dynarski, S. M. (2011). Gains and gaps: Changing inequality in u.s. college entry and completion. In R. S. Foundation (Ed.), *Whither opportunity? rising inequality, schools, and children's life chances* (pp. 117–132).

Beneito, P., Boscá, J. E., and Ferri, J. (2018). Tuition fees and student effort at university. *Economics of Education Review*, *64*, 114–128.

Bietenbeck, J., Leibing, A., Marcus, J., and Weinhardt, F. (2023). Tuition fees and educational attainment. *European Economic Review*, *154*, 104431.

Black, S. E., Denning, J., Dettling, L. J., Goodman, S., and Turner, L. J. (2023). Taking it to the limit: Effects of increased student loan availability on attainment, earnings, and financial well-being. *SSRN Electronic Journal*.

Bolhaar, J., Kuijpers, S., Webbink, D., and Zumbuehl, M. (2024). Does replacing grants by income-contingent loans harm enrolment? new evidence from a reform in dutch higher education. *Economics of Education Review*, *101*, 102546.

Bonhomme, S., Lamadon, T., and Manresa, E. (2022). Discretizing unobserved heterogeneity. *Econometrica*, *90*, 625–643.

Bucarey, A. (2017). *Who pays for free college? crowding out on campus*.

Bucarey, A., Contreras, D., and Muñoz, P. (2020). Labor market returns to student loans for university: Evidence from chile. *Journal of Labor Economics*, *38*, 959–1007.

Burland, E., Dynarski, S., Michelmore, K., Owen, S., and Raghuraman, S. (2023). The power of certainty: Experimental evidence on the effective design of free tuition programs. *American Economic Review: Insights*, *5*, 293–310.

Callaway, B., Goodman-Bacon, A., and Sant'anna, P. (2024). *Difference-in-differences with a continuous treatment*.

Card, D. (1999). The causal effect of education on earnings. *Handbook of Labor Economics*, *3*, 1801–1863.

Carruthers, C. K., and Fox, W. F. (2016). Aid for all: College coaching, financial aid, and post-secondary persistence in tennessee. *Economics of Education Review*, *51*, 97–112.

Castleman, B. L., and Long, B. T. (2016). Looking beyond enrollment: The causal effect of need-based grants on college access, persistence, and graduation. *Journal of Labor Economics*, *34*, 1023–1073.

Castro-Zarzur, R., Espinoza, R., and Sarzosa, M. (2022). Unintended consequences of free college: Self-selection into the teaching profession. *Economics of Education Review*, *89*, 102260.

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.

Clasing-Manquian, P. (2022). *The promise of free tuition: The case of chile* [Doctoral dissertation].

Cohodes, S. R., and Goodman, J. S. (2014). Merit aid, college quality, and college completion: Massachusetts adams scholarship as an in-kind subsidy. *American Economic Journal: Applied Economics*, *6*, 251–85.

Colas, M., Findeisen, S., and Sachs, D. (2021). Optimal need-based financial aid. *Journal of Political Economy*, *129*, 492–533.

De Groote, O. (2025). Dynamic effort choice in high school: Costs and benefits of an academic track. *Journal of Labor Economics*.

de Chaisemartin, C., and D'Haultfœuille, X. (2024). Difference-in-differences estimators of intertemporal treatment effects. *Review of Economics and Statistics*, 1–45.

Deming, D., and Dynarski, S. (2010). College aid. In *Targeting investments in children: Fighting poverty when resources are limited* (pp. 283–302).

Denning, J. T. (2017). College on the cheap: Consequences of community college tuition reductions. *American Economic Journal: Economic Policy*, *9*, 155–188.

Denning, J. T. (2019). Born under a lucky star. *Journal of Human Resources*, *54*, 760–784.

Denning, J. T., Marx, B. M., and Turner, L. J. (2019). Propelled: The effects of grants on graduation, earnings, and welfare. *American Economic Journal: Applied Economics*, *11*, 193–224.

DIPRES. (2016). *Ley de presupuestos* (tech. rep.).

Dobbin, C., Barahona, N., and Otero, S. (2022). *The equilibrium effects of subsidized student loans*.

Dubois, P., de Janvry, A., and Sadoulet, E. (2012). Effects on school enrollment and performance of a conditional cash transfer program in mexico. *Journal of Labor Economics*, *30*, 555–589.

Duryea, S., Ribas, R. P., Sampaio, B., Sampaio, G. R., and Trevisan, G. (2023). Who benefits from tuition-free, top-quality universities? evidence from brazil. *Economics of Education Review*, *95*, 102423.

Dynarski, S., Page, L., and Scott-Clayton, J. (2023). College costs, financial aid, and student decisions. In *Handbook of the economics of education* (pp. 227–285, Vol. 7). Elsevier.

Dynarski, S. M. (2003). Does aid matter? measuring the effect of student aid on college attendance and completion. *American Economic Review*, *93*, 279–288.

Epple, D., Romano, R., Sarpça, S., and Sieg, H. (2017). A general equilibrium analysis of state and private colleges and access to higher education in the u.s. *Journal of Public Economics*, *155*, 164–178.

Epple, D., Romano, R., and Sieg, H. (2006). Admission, tuition, and financial aid policies in the market for higher education. *Econometrica*, *74*, 885–928.

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

Fack, G., Grenet, J., and He, Y. (2019). Beyond truth-telling: Preference estimation with centralized school choice and college admissions. *American Economic Review*, *109*, 1486–1529.

Falco, A. D., and Reichlin, Y. (2025). Grants vs. loans: The role of financial aid in college major choice. *The Economic Journal*.

Ferreyra, M. M., Garriga, C., Martin-Ocampo, J. D., and Sanchez-Diaz, A. M. (2024). The limited impact of free college policies. *European Economic Review*, *168*, 104800.

Ferreyra, M., Garriga, C., Martin-Ocampo, J. D., and Diaz, A. M. S. (2022). Cows don't give milk: An effort model of college graduation.

Fu, C., Ishimaru, S., and Kennan, J. (2024). Government expenditure on the public education system. *International Economic Review*, *65*, 43–73.

Gale, D., and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, *69*, 9.

Garibaldi, P., Giavazzi, F., Ichino, A., and Rettore, E. (2012). College cost and time to complete a degree: Evidence from tuition discontinuities. *The Review of Economics and Statistics*, *94*, 699–711.

Gezer, F. (2026). Non-distortionary financial aid design in college admissions: Theory and practice from chile.

Goldin, C. D., and Katz, L. F. (2008). *The race between education and technology*. The Belknap Press of Harvard University Press.

Harris, D. N., and Mills, J. (2025). Should college be "free"? evidence on free college, early commitment, and merit aid from an eight-year randomized trial. *American Economic Journal: Economic Policy*, *17*, 373–406.

Harris, D. N., Taber, C. R., and Yang, Y. (2026). *College success: Understanding the roles of financial aid, loans, preparation, and incentives*.

Heckman, J. J. (1981). The incidental parameters problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process. In *Structural analysis of discrete data with econometric applications*. Cambridge, MA: MIT Press.

Hendren, N., and Sprung-Keyser, B. (2020). A unified welfare analysis of government policies. *The Quarterly Journal of Economics*, *135*, 1209–1318.

Hotz, V. J., and Miller, R. A. (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, *60*, 497–529.

Joensen, J. S., and Mattana, E. (2024). *Student aid design, academic achievement, and labor market behavior: Grants or loans?*

Keane, M. P., and Wolpin, K. I. (1997). The career decisions of young men. *Journal of Political Economy*, *105*, 473–522.

Larroucau, T., and Rios, I. (2024). *Dynamic college admissions*.

Londoño-Vélez, J., Rodríguez, C., and Sánchez, F. (2020). Upstream and downstream impacts of college merit-based financial aid for low-income students: Ser pilo paga in colombia. *American Economic Journal: Economic Policy*, *12*, 193–227.

Magnac, T., and Thesmar, D. (2002). Identifying dynamic discrete decision processes. *Econometrica*, *70*, 801–816.

Martin, A. (2024). *Does free tuition help or hinder the poor?*

Marx, B. M., and Turner, L. J. (2018). Borrowing trouble? human capital investment with opt-in costs and implications for the effectiveness of grant aid. *American Economic Journal: Applied Economics*, *10*, 163–201.

Matsuda, K. (2020). Optimal timing of college subsidies: Enrollment, graduation, and the skill premium. *European Economic Review*, *129*, 103549.

Molina, T., and Rivadeneyra, I. (2021). The schooling and labor market effects of eliminating university tuition in ecuador. *Journal of Public Economics*, *196*, 104383.

Montalbán, J. (2022). Countering moral hazard in higher education: The role of performance incentives in need-based grants. *The Economic Journal*, *133*, 355–389.

OECD. (2019). *Education at a glance* (tech. rep.).

Page, L. C., and Scott-Clayton, J. (2016). Improving college access in the united states: Barriers and policy responses. *Economics of Education Review*, *51*, 4–22.

Pal, J. (2026). Education policy and the quality of public servants.

Rambachan, A., and Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, *90*, 2555–2591.

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, *4*, 305–22.

Rothstein, J., and Rouse, C. E. (2011). Constrained after college: Student loans and early-career occupational choices. *Journal of Public Economics*, *95*, 149–163.

Schoonover, P., Attridge, J., Carruthers, C. K., Welch, J. G., Conrad, C., Ferman, B., Klafehn, A., Cadena, B., Mckinnish, T., Barham, T., and Mansfield, R. (2026). A promise worth keeping? impacts of free community college on degrees and earnings.

Scott-Clayton, J. (2011). On money and motivation: A quasi-experimental analysis of financial incentives for college achievement. *Journal of Human Resources*, *46*, 614–646.

Sieg, H., and Wang, Y. (2018). The impact of student debt on education, career, and marriage choices of female lawyers. *European Economic Review*, *109*, 124–147.

SIES. (2026). *Duración real y en exceso de los programas de educación superior* (tech. rep.). Subsecretaría de Educación Superior.

Solis, A. (2017). Credit access and college enrollment. *Journal of Political Economy*, *125*, 562–622.

Solis, A. (2026). University loans and grants: Effects on educational and labor market outcomes. *Journal of Labor Economics*, *44*.

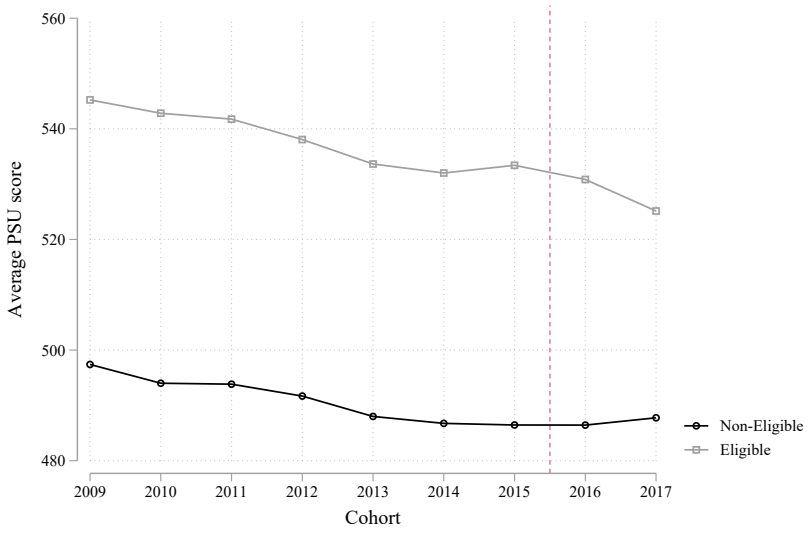
Stinebrickner, R., and Stinebrickner, T. (2008). The effect of credit constraints on the college drop-out decision: A direct approach using a new panel study. *American Economic Review*, *98*, 2163–84.

Tincani, M. M., Kosse, F., and Miglino, E. (2026). Beliefs and the incentive effects of preferential college admissions: Evidence from an experiment and a structural model *.

Appendices

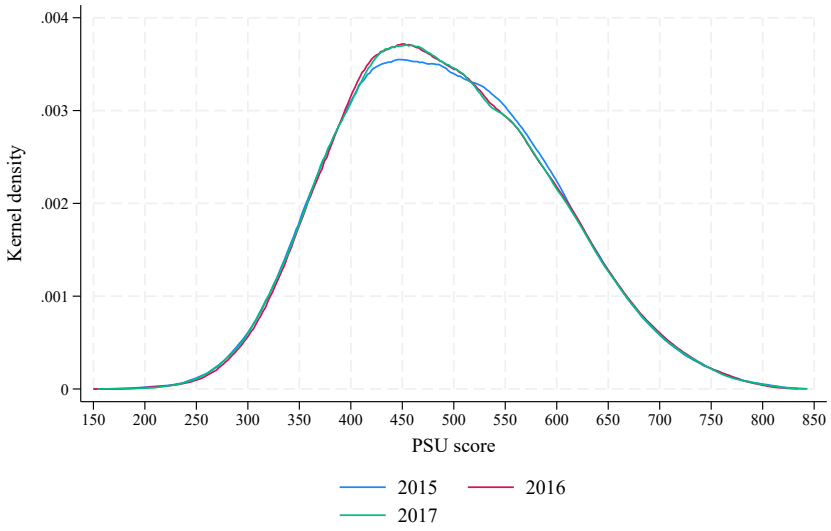
A Figures

Figure A1: Average PSU score over time for participating and non-participating institutions



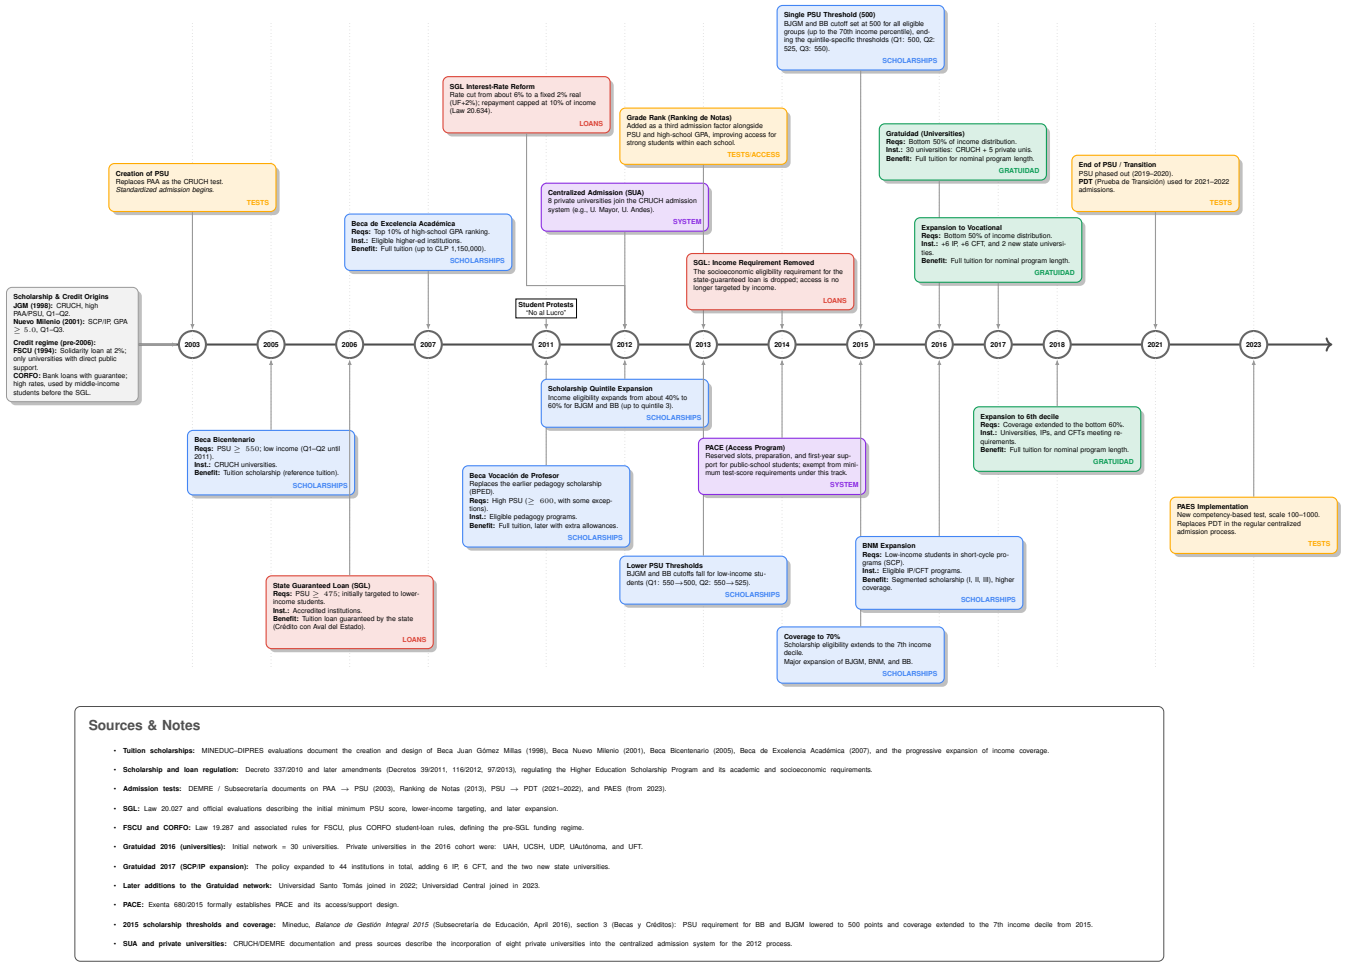
Notes: Annual average PSU score by groups of eligible and non-eligible institutions to free higher education.

Figure A2: Distribution of test scores of enrolled students, 2015–2017



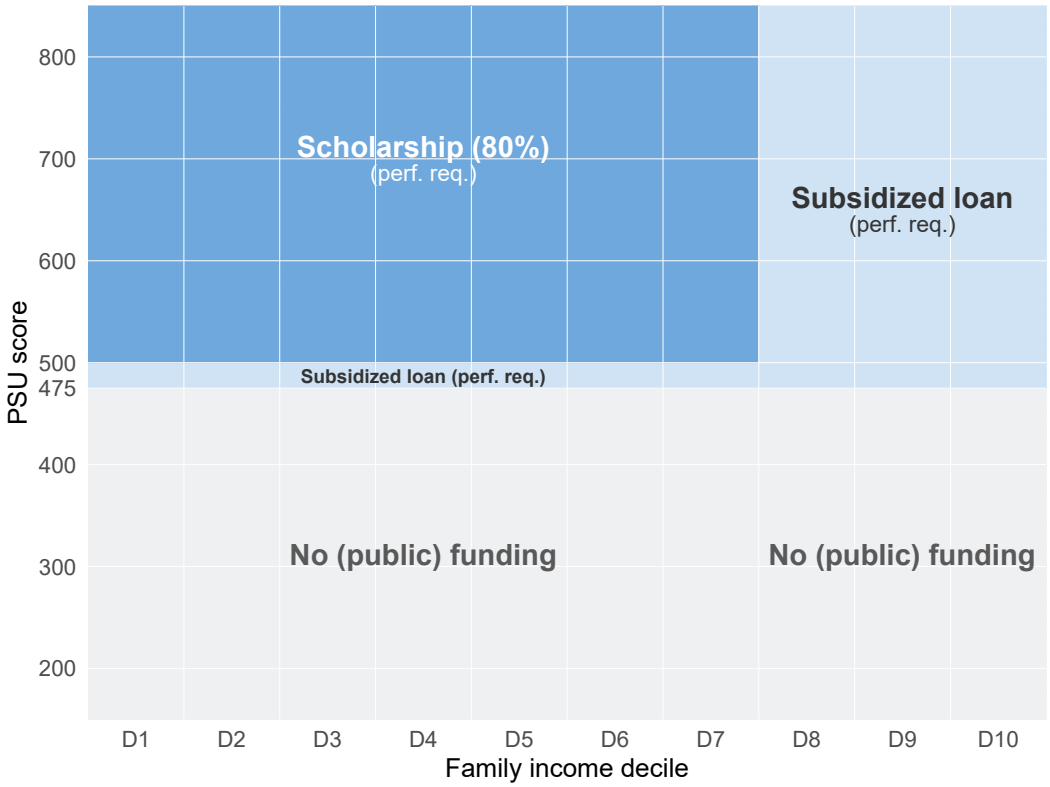
Notes: Kernel density of the PSU score of students enrolled in higher education, by entry cohort (2015, 2016 and 2017). PSU is the national standardized test.

Figure A3: Education reforms 2003-2024



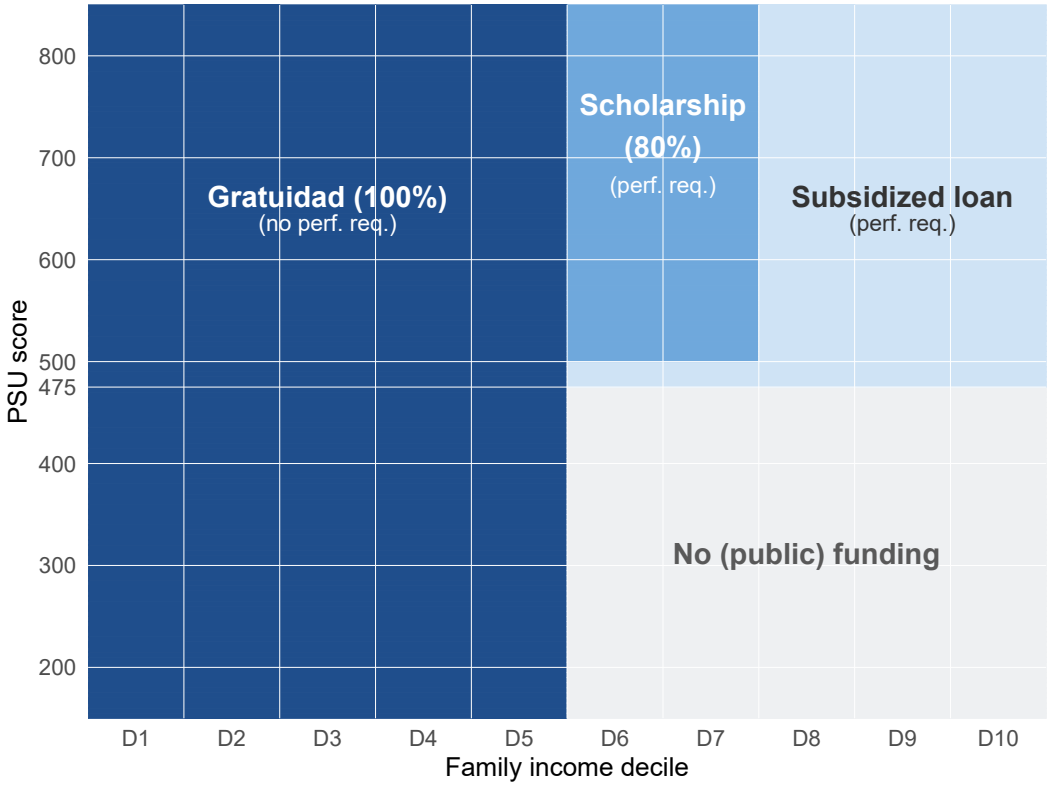
Notes: CRUCH: universities of the Council of Rectors (traditional universities); private universities are all others. IP: professional institutes; CFT: technical training centers; SCP: short-cycle programs. PAA, PSU, PDT and PAES: successive national admission tests; SUA: centralized admission system. SGL: State Guaranteed Loan; FSCU: university solidarity loan fund; CORFO: loans guaranteed by the state development agency. BJGM, BB and BNM: *Juan Gómez Millas*, *Bicentenario* and *Nuevo Milenio* scholarships. PACE: access program for public-school students. Q1–Q3: income quintiles; UF: inflation-indexed unit of account.

Figure A4: Funding instruments by PSU score and family income, 2015



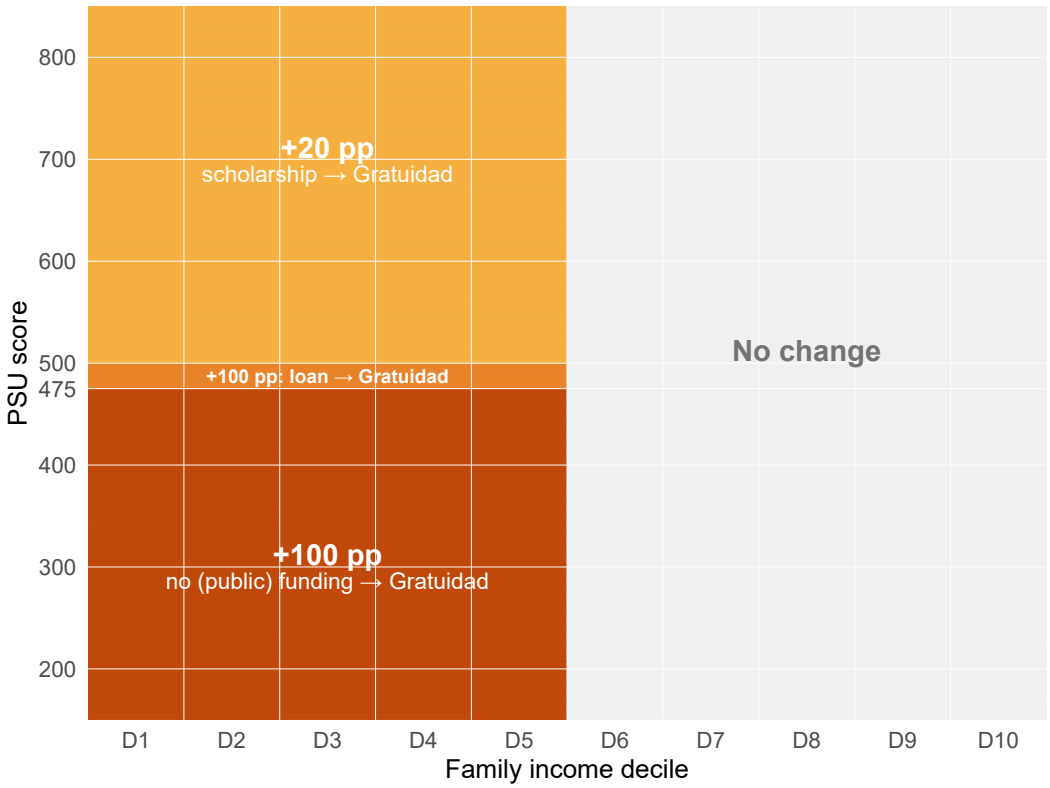
Notes: Funding instrument available to students by PSU score (vertical axis) and family income decile (horizontal axis) before the TF policy. Scholarships finance 80% of reference tuition and loans are subsidized; all instruments require meeting academic performance requirements.

Figure A5: Funding instruments by PSU score and family income, 2016



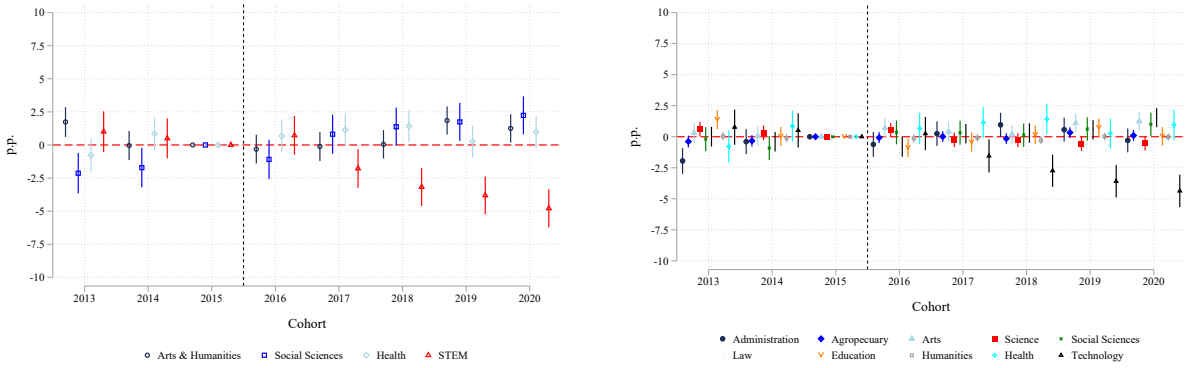
Notes: Funding instrument available to students by PSU score (vertical axis) and family income decile (horizontal axis) after the introduction of TF (*Gratuidad*) for the 2016 cohort. *Gratuidad* covers 100% of reference tuition with no performance requirement and applies to students in income deciles 1 to 5 attending eligible institutions.

Figure A6: Change in funding instruments after TF



Notes: Change in the share of students receiving *Gratuidad* between 2015 and 2016, by PSU score (vertical axis) and family income decile (horizontal axis), and the instrument it substitutes for. No change for deciles 6 to 10.

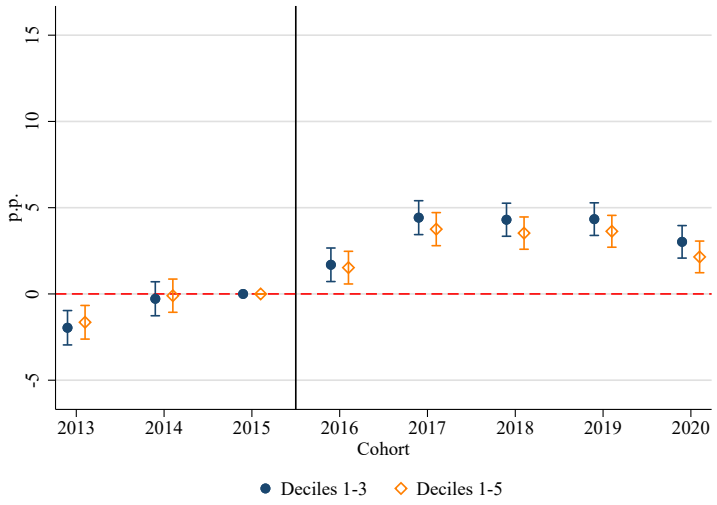
Figure A7: TWFE: Preferences



a) 4 fields b) 10 fields

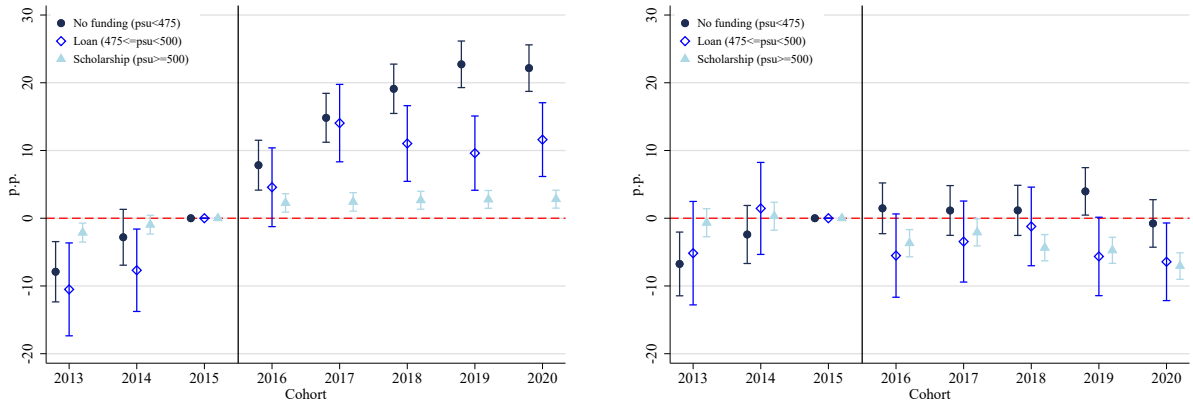
Notes: These figures show the results of estimating Equation 1 on the share of each field in students' rank-ordered lists. Treated students are eligible students in income deciles 1-3, and control students are ineligible students in income deciles 9-10 (self-reported at PSU registration). Panel (a) groups the results in four fields, while panel (b) uses the original ten-field classification. STEM includes Agriculture, Science and Technology. Social Sciences includes Administration, Social Sciences and Law. Arts and Humanities includes Education, Arts and Humanities. Health includes Health Sciences.

Figure A8: Enrollment: alternative treated groups



Notes: Results of estimating Equation 1 on enrollment, comparing students in income deciles 1-3 (main specification) and in deciles 1-5 with students in deciles 9-10, using family income self-reported at PSU registration. Sample is the universe of first time test takers. 95% confidence intervals with heteroskedasticity-robust standard errors.

Figure A9: Enrollment by institution eligibility and funding eligibility



a) Eligible institutions

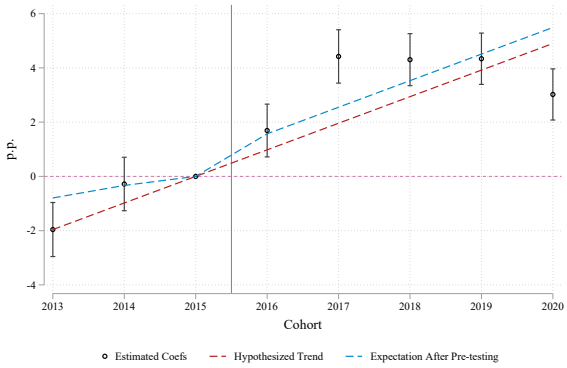
b) Non-eligible institutions

Notes: These figures disaggregate Figure 3 by test score group. Results of estimating Equation 1 on enrollment at eligible and non-eligible institutions to the TF policy, comparing students in income deciles 1-3 with students in deciles 9-10, using family income self-reported at PSU registration, by groups based on PSU test scores: students scoring below 475 are not eligible for public funding, students scoring between 475 and 500 are eligible for loans, and students scoring above 500 are eligible for scholarships. Sample is the universe of first time test takers; in panel a) students enrolled at non-eligible institutions are dropped, and in panel b) students enrolled at eligible institutions are dropped. 95% confidence intervals with heteroskedasticity-robust standard errors.

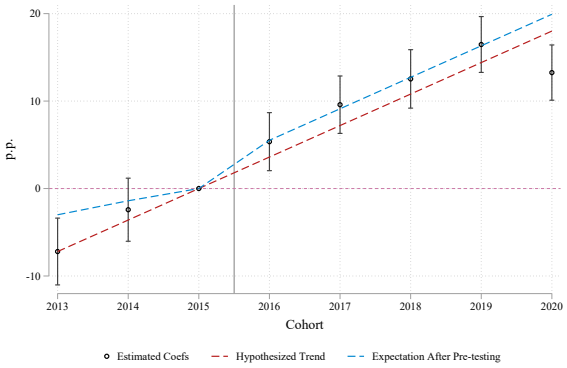
A.1 Sensitivity of the TWFE estimates to violations of parallel trends

This section reports two checks on the parallel-trends assumption behind Figure 2 and Figure 4, on the same samples and specifications. First, following Roth (2022), Appendix Figure A10 and Appendix Figure A11 plot the estimated coefficients against a hypothesized linear violation of parallel trends, the line fitted through the pre-period coefficients and the reference cohort, together with the coefficients one would expect to estimate under that trend conditional on a pre-trend test not rejecting. Second, Appendix Figure A12 and Appendix Figure A13 report the robust confidence sets of Rambachan and Roth (2023) under the relative-magnitudes restriction, which bounds the change in the differential trend between consecutive post-treatment periods by M times the largest change between consecutive pre-treatment periods. $M = 0$ is the original confidence interval, and a coefficient is robust at M if its confidence set excludes zero.

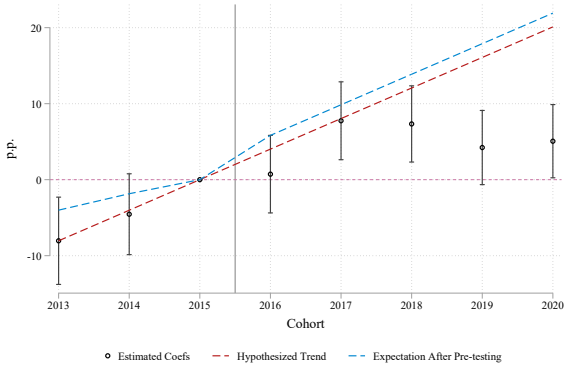
Figure A10: Hypothesized pre-trends: enrollment



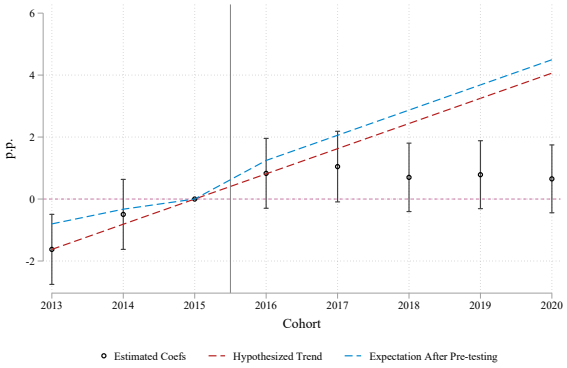
a) All students



b) No public funding (PSU $<$ 475)



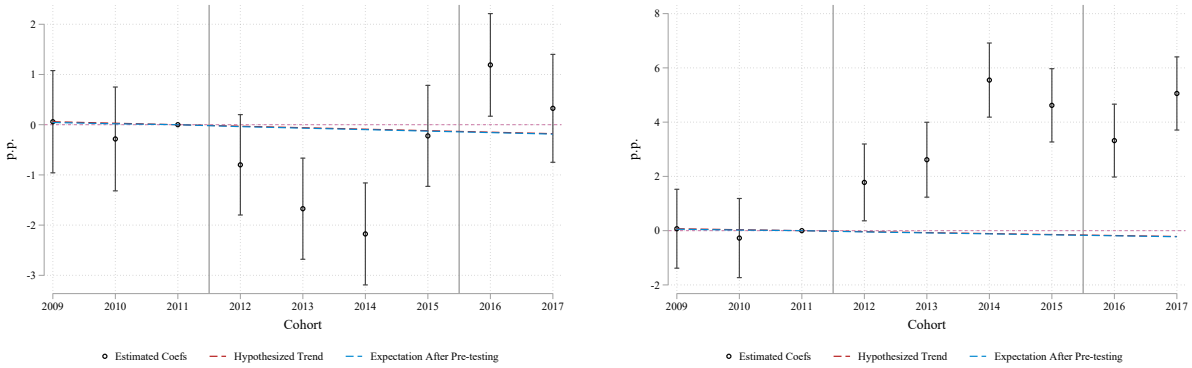
c) Loans (PSU 475–499)



d) Scholarships (PSU $\geq$ 500)

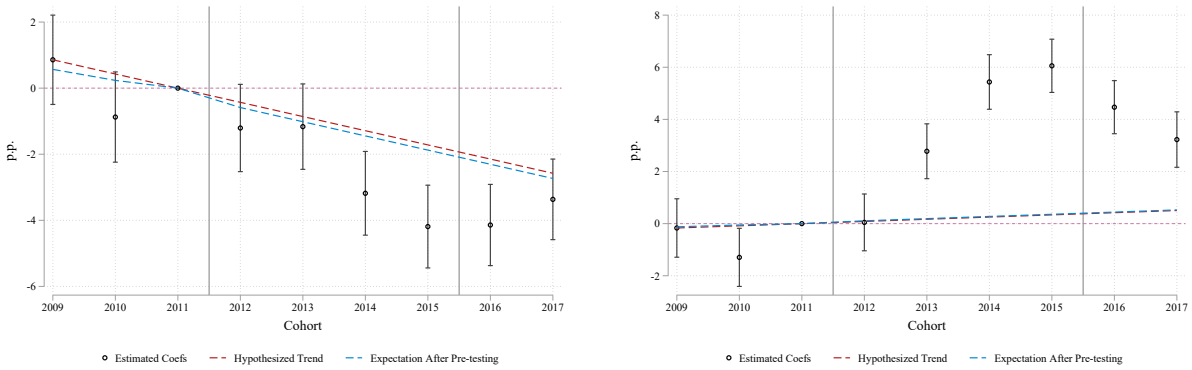
Notes: Event-study coefficients of Equation 1 on enrollment (percentage points, 95% confidence intervals with heteroskedasticity-robust standard errors), as in Figure 2: students in income deciles 1-3 relative to deciles 9-10, reference cohort 2015, for all students (panel a) and by PSU funding group (panels b-d). The hypothesized trend is the linear trend fitted through the pre-period coefficients and the reference cohort, extended to the post-period (Roth, 2022). "Expectation after pre-testing" is the expected value of the coefficients under that trend, conditional on a pre-trend test not rejecting parallel trends. Trends of the fitted slopes (0.98, 3.60, 4.02 and 0.81 p.p. per cohort in panels a-d) would be detected by the pre-trend test with probability 0.97, 0.96, 0.81 and 0.82.

Figure A11: Hypothesized pre-trends: educational outcomes



a) Dropout

b) Graduation

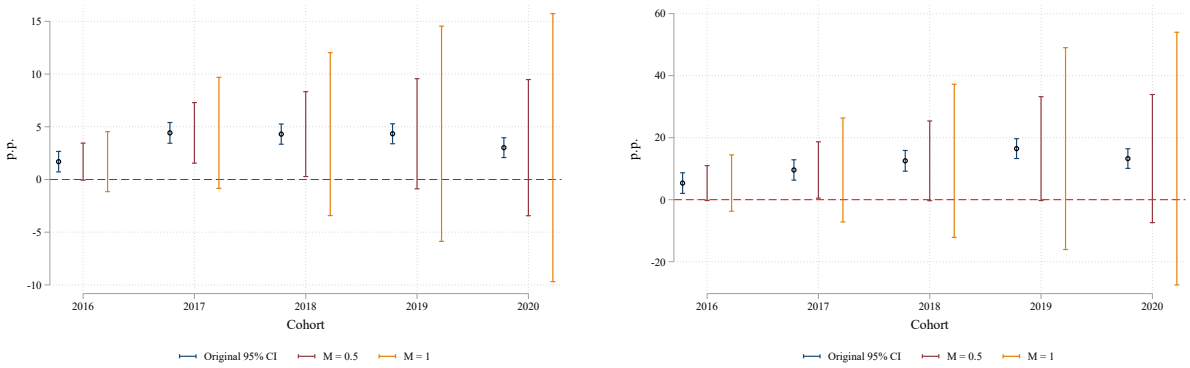


c) Field switch

d) Graduation on time

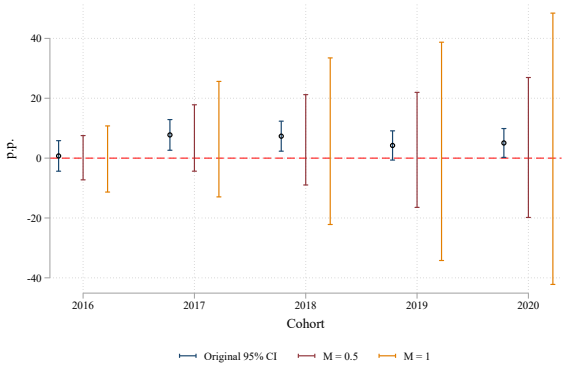
Notes: Event-study coefficients of Equation 2 (percentage points, 95% confidence intervals with heteroskedasticity-robust standard errors), as in Figure 4: enrolled students in income deciles 1-3 relative to deciles 9-10, reference cohort 2011. The hypothesized trend is the linear trend fitted through the pre-period coefficients and the reference cohort (Roth, 2022), and "Expectation after pre-testing" the expected coefficients under that trend conditional on a pre-trend test not rejecting. The pre-period coefficients are close to zero, so the fitted trends are flat (-0.03, -0.04, -0.43 and 0.08 p.p. per cohort in panels a-d) and would be detected with probability of only 0.09, 0.09, 0.28 and 0.10; the confidence sets in Appendix Figure A13 are the more informative check.

Figure A12: Robustness to violations of parallel trends: enrollment

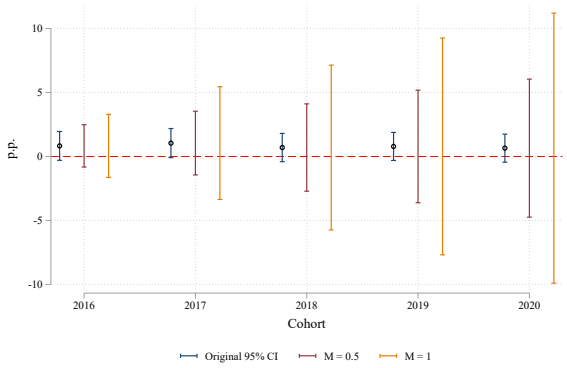


a) All students

b) No public funding (PSU $<$ 475)



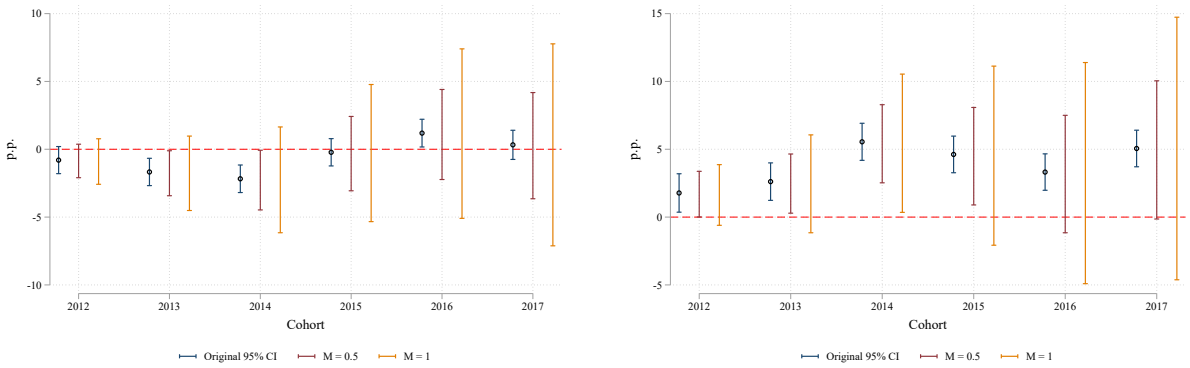
c) Loans (PSU 475–499)



d) Scholarships (PSU $\geq$ 500)

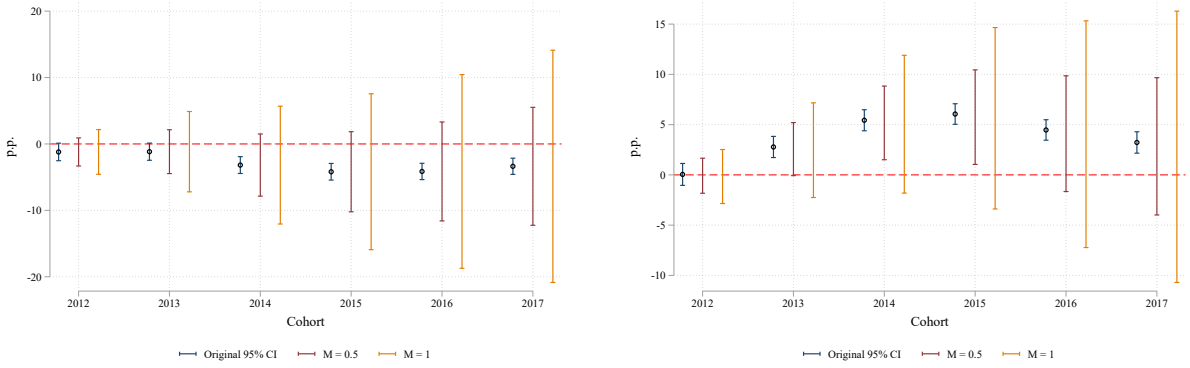
Notes: Robust 95% confidence sets of Rambachan & Roth (2023) for each post-period coefficient of Appendix Figure A10 (percentage points) under the relative-magnitudes restriction: the change in the differential trend between consecutive post-treatment cohorts is at most M times the largest change between consecutive pre-treatment cohorts. $M = 0$ is the original 95% confidence interval, $M = 1$ allows post-period violations as large as the largest pre-period one. Conditional least-favorable hybrid confidence sets computed with the `honestdid` Stata package. The 2017 and 2018 effects for all students and the 2017 effect for students without public funding remain different from zero at $M = 0.5$; no coefficient does at $M = 1$.

Figure A13: Robustness to violations of parallel trends: educational outcomes



a) Dropout

b) Graduation



c) Field switch

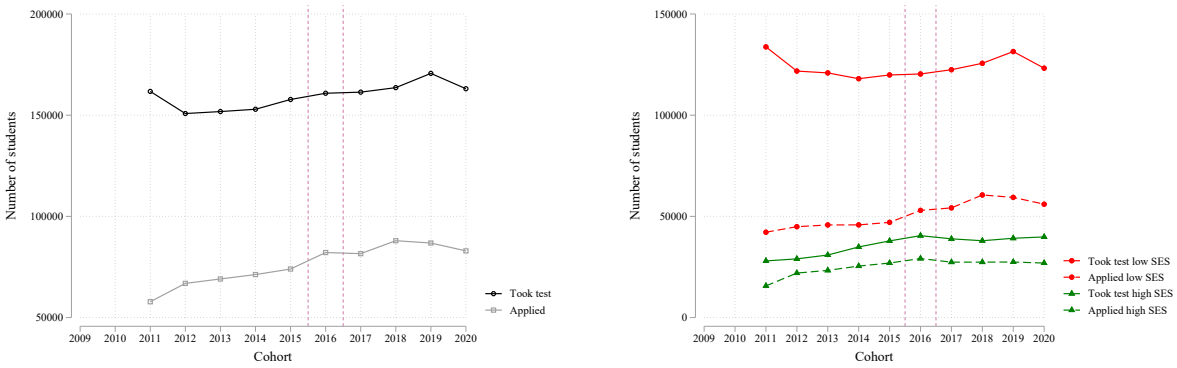
d) Graduation on time

Notes: Robust 95% confidence sets of Rambachan & Roth (2023) for each post-period coefficient of Appendix Figure A11 (percentage points) under the relative-magnitudes restriction, as in Appendix Figure A12; $M = 0$ is the original 95% confidence interval. The graduation gains of the 2012–2015 cohorts, the graduation-on-time gains of the 2014–2015 cohorts and the dropout reductions of the 2013–2014 cohorts remain different from zero at $M = 0.5$, and the 2014 graduation effect also at $M = 1$. The field-switch estimates and the 2016–2017 cohort estimates do not survive $M = 0.5$.

A.2 Test-taking and applications around the reform

The DiD of Section 3 is estimated on test takers. If TF had induced more low-income students to sit the PSU, the composition of the treated group would change with the reform. Appendix Figure A14 shows the number of students in their last year of high school who took the test and who applied through the centralized admission system, overall and by family income.

Figure A14: PSU test takers and applicants by cohort



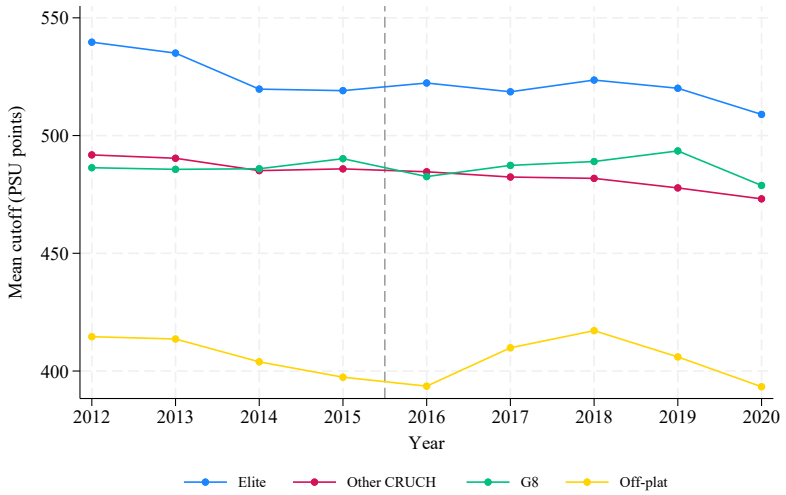
a) All students b) By family income

Notes: Students in their last year of high school, by the year in which they would sit the PSU (cohort), 2013–2020. "Test-takers" counts students with a valid PSU language score; "Applied" counts students who submitted at least one application through the centralized admission system. Panel b splits students by family income as defined in the test-takers dataset (reported at PSU registration): low SES is family income deciles 1–5, high SES deciles 6–10. The dashed vertical lines mark the 2016 cohort, the first under TF, and the 2017 cohort, when TF was extended to short-cycle programs.

A.3 Threats to SUTVA: spillovers to non-eligible students

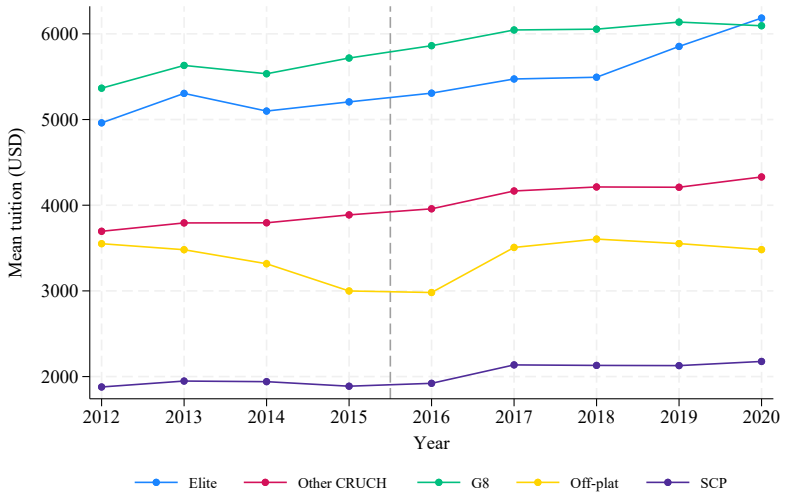
The DiD design in Section 3 requires that the policy does not affect non-eligible students through changes in institutions' supply, selectivity or prices. Appendix Figure A15, Appendix Figure A16 and Appendix Figure A17 summarize the supply-side margins through which such spillovers could operate: admission cutoffs, prices, and vacancies.

Figure A15: Mean admission cutoff by institution tier



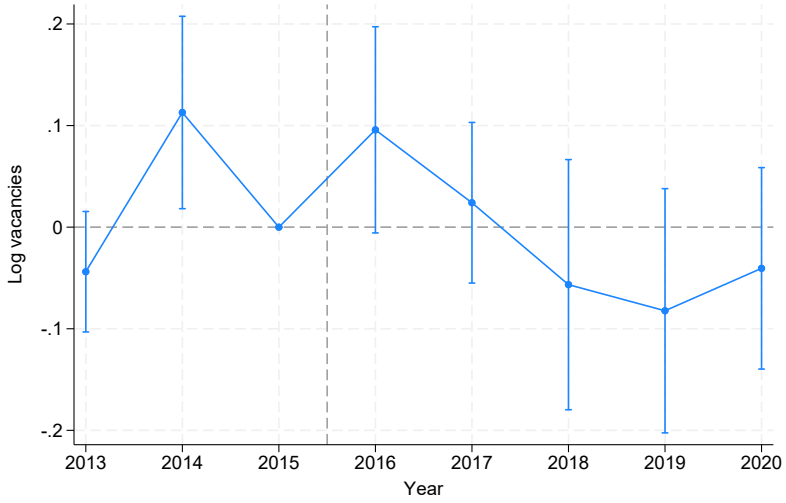
Notes: Mean cutoff of bachelor programs, by tier, 2013–2020. Elite: Pontificia Universidad Católica de Chile, Universidad de Chile, Universidad de Concepción, Universidad Técnica Federico Santa María, Pontificia Universidad Católica de Valparaíso and Universidad de Santiago de Chile. G8: Universidad Adolfo Ibáñez, Alberto Hurtado, Andrés Bello, de los Andes, del Desarrollo, Diego Portales, Finis Terrae and Mayor. Other CRUCH: the remaining universities of the Council of Rectors (CRUCH). Off-plat: universities outside CRUCH that are not in the G8. The dashed line marks the introduction of TF.

Figure A16: Mean tuition by institution tier



Notes: Mean annual tuition of programs, by tier, in real December 2009 USD (pesos deflated by the December CPI and converted at 500 CLP per USD). Tiers as in Appendix Figure A15: Elite (PUC, U. de Chile, U. de Concepción, UTFSM, PUCV, USACH); G8 (UAI, UAH, UNAB, U. de los Andes, UDD, UDP, U. Finis Terrae, U. Mayor); Other CRUCH (remaining CRUCH universities); Off-plat (non-CRUCH, non-G8 universities); SCP (short-cycle programs at any institution). The dashed line marks the introduction of TF.

Figure A17: TWFE event study: log vacancies of bachelor programs



Notes: Coefficients of an event study of log vacancies on the interaction between year and an indicator for programs at institutions adhering to TF (relative to non-adhering universities), with program and year fixed effects and 2015 as the omitted year. Standard errors clustered by institution; bars show 95% confidence intervals. Only a few non-adhering universities have a 2015 cutoff, so inference rests on few control clusters.

B Tables

Table B1: Enrollment and outcomes, quasi-experimental sample

Extensive margin	2013	2014	2015	2016	2017	2018	2019	2020
A. Enrollment								
Enrolled in higher education	0.626	0.633	0.631	0.640	0.658	0.659	0.646	0.662
University enrollment	0.308	0.314	0.313	0.343	0.337	0.343	0.338	0.369
SCP enrollment	0.316	0.319	0.318	0.298	0.321	0.316	0.308	0.293
Not enrolled	0.374	0.367	0.369	0.360	0.342	0.341	0.354	0.338
B. Controls								
Female	0.553	0.551	0.551	0.548	0.553	0.549	0.544	0.547
Family income bracket (1–10)	3.33	3.59	3.75	3.95	4.16	4.04	4.06	4.40
PSU	486.7	487.8	490.3	495.2	495.8	499.9	494.4	497.6
GPA	541.2	542.1	544.0	550.5	555.1	559.0	563.6	576.7
C. Funding								
Gratuidad	0.000	0.000	0.000	0.339	0.558	0.596	0.585	0.587
Subsidized loan	0.308	0.310	0.292	0.188	0.120	0.110	0.117	0.108
Merit scholarship	0.390	0.388	0.442	0.247	0.118	0.111	0.121	0.120
No (public) funding	0.302	0.302	0.266	0.226	0.204	0.182	0.177	0.186
Observations	102,480	104,348	106,956	108,695	103,951	119,227	127,146	118,717

Intensive margin	2009	2010	2011	2012	2013	2014	2015	2016	2017
D. Outcomes									
Graduation	0.358	0.353	0.344	0.364	0.366	0.392	0.385	0.379	0.365
On-time graduation	0.168	0.151	0.154	0.169	0.177	0.199	0.188	0.179	0.206
Persistence (years enrolled)	3.51	3.42	3.36	3.39	3.33	3.42	3.42	3.50	3.51
Field switch	0.267	0.270	0.268	0.264	0.263	0.246	0.235	0.217	0.195
Dropout	0.160	0.168	0.164	0.157	0.170	0.182	0.196	0.226	0.269
Observations	51,872	56,332	58,333	55,962	64,015	66,059	67,490	69,609	68,391

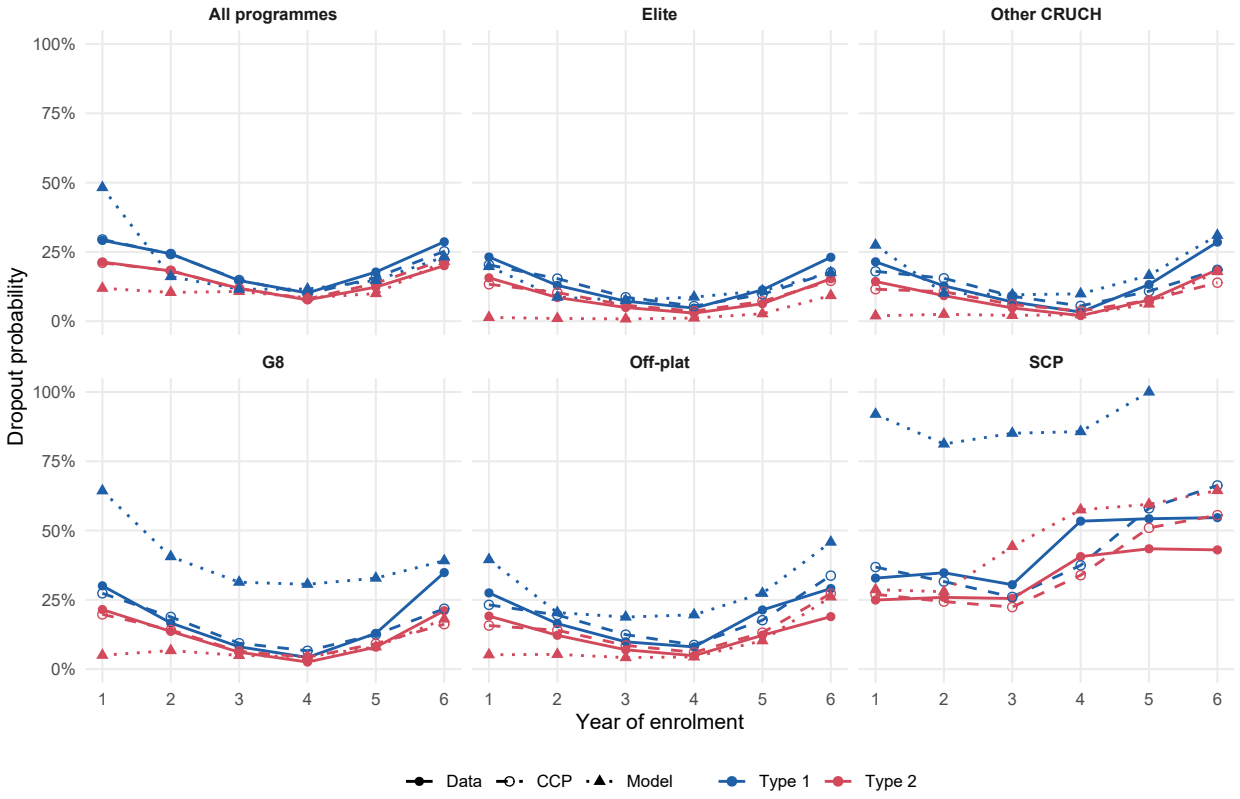
Notes: The extensive sample follows the pooled enrollment TWFE regression; the intensive sample follows the pooled enrolled-student TWFE regression. University/SCP rows are unconditional extensive-sample shares. SCP: short-cycle programs at any institution. Funding shares condition on enrollment and use the priority Gratuidad, loan, then merit scholarship. Characteristics are the reduced-form controls (family income defines the control group). Observations are students in the extensive sample (Panels A–C) and enrolled students in the intensive sample (Panel D). On-time graduation means within nominal duration.

Table B2: Descriptive statistics, model sample

	2016	2017
A. Students and entry choices		
University enrollment	0.359	0.284
SCP enrollment	0.349	0.316
Not enrolled	0.292	0.401
B. Student characteristics		
Female	0.512	0.527
Family income bracket (1–10)	4.02	4.45
PSU	493.7	492.3
GPA	549.6	547.5
C. Funding		
Gratuidad	0.253	0.414
Subsidized loan	0.222	0.160
Merit scholarship	0.267	0.177
None of these funding sources	0.259	0.250
D. Academic progress		
Registered / annual curriculum	0.935	0.932
Completed / annual curriculum	0.810	0.815
Completed / registered	0.844	0.854
Observations (student-years)	15,470	13,162
Observations (students)	6,981	6,993

Notes: Cohorts 2016–2017 in the current model-estimation input. SCP: short-cycle programs at any institution. Panels A–C count each student once (Observations (students)) at entry, including the outside option; funding conditions on entry enrollment and uses the priority Gratuidad, loan, then merit scholarship. PSU and GPA are original scores, not standardized model regressors. Panel D weights each chosen enrolled student-year equally over the observed model horizon. Credit rates are normalized by the annual curriculum, and pool credit- and course-reporting programs.

Figure A18: Dropout probability by year of enrollment and latent type



Notes: Probability of leaving higher education at the end of each year of enrollment, among students still at risk (enrolled, credits left to graduate, not forced out). Type 1 and Type 2 are the latent types of the model. *Data*: observed dropout rate weighted by each student's posterior type probability from the EM estimation; because types are latent, this series is observed only conditional on the estimated posteriors. *CCP*: fitted probability of the dropout logit (Appendix Table B3), weighted the same way. *Model*: dropout rate in the baseline simulation of the structural model, with one type drawn per student and replication from the posterior. Tiers as in Appendix Figure A15: Elite (six selective CRUCH universities), Other CRUCH (remaining universities of the Council of Rectors, CRUCH), G8 (eight private universities outside CRUCH), Off-plat (non-CRUCH, non-G8 universities) and SCP (short-cycle programs).

Table B3: Dropout logit: all coefficients

OP fees (t+1, thousand USD)	0.018**
	(0.007)
Relative ability	-0.096***
	(0.027)
Diff. Region	0.005
	(0.023)
Credits left	-0.080***
	(0.006)
Credits: 0	3.744***
	(0.037)
Credits: 1	2.448***
	(0.030)
Credits: 2	1.114***
	(0.023)
Credits: 4	-0.498***
	(0.022)
Female	-0.744***
	(0.044)
Low SES	-0.179***
	(0.060)
Middle SES	-0.237***
	(0.062)
PSU	-0.456***
	(0.047)
GPA	-0.232***
	(0.029)
OP (t+1) × female	0.026***
	(0.005)
OP (t+1) × low SES	0.017**
	(0.009)
OP (t+1) × middle SES	0.017**
	(0.008)
OP (t+1) × psu	-0.010**
	(0.004)
OP (t+1) × gpa	0.005
	(0.003)
Credits × female	-0.066***
	(0.007)
Credits × low SES	-0.058***
	(0.010)
Credits × middle SES	-0.041***
	(0.010)
Credits × psu	0.016***
	(0.005)

continued

Credits × gpa	-0.038***
	(0.005)
Credits left × female	0.045***
	(0.002)
Credits left × low SES	0.005
	(0.003)
Credits left × middle SES	0.003
	(0.003)
Credits left × psu	0.033***
	(0.002)
Credits left × gpa	0.006***
	(0.002)
Rel. Abil. × female	0.057***
	(0.010)
Rel. Abil. × low SES	-0.094***
	(0.014)
Rel. Abil. × middle SES	-0.067***
	(0.015)
Diff. Region × female	0.007
	(0.020)
Diff. Region × low SES	-0.137***
	(0.026)
Diff. Region × middle SES	-0.060**
	(0.028)
Diff. Region × psu	0.064***
	(0.013)
Diff. Region × gpa	0.033**
	(0.013)
Elite × female	-0.063
	(0.038)
Elite × low SES	-0.113**
	(0.048)
Elite × middle SES	-0.068
	(0.049)
Elite × psu	0.117***
	(0.037)
Elite × gpa	0.102***
	(0.030)
Other CRUCH × female	-0.049
	(0.043)
Other CRUCH × low SES	0.012
	(0.053)
Other CRUCH × middle SES	0.072
	(0.056)
Other CRUCH × psu	0.112***
	(0.043)

continued

Other CRUCH × gpa	0.049
	(0.031)
G8 × female	-0.043
	(0.054)
G8 × low SES	0.116*
	(0.066)
G8 × middle SES	0.068
	(0.064)
G8 × psu	0.212***
	(0.056)
G8 × gpa	0.056
	(0.039)
SCP × female	0.204***
	(0.030)
SCP × low SES	0.060
	(0.039)
SCP × middle SES	0.074*
	(0.041)
SCP × psu	0.501***
	(0.023)
SCP × gpa	0.225***
	(0.020)
Type 2	-0.673***
	(0.009)
t = 2	-0.133***
	(0.020)
t = 3	-0.676***
	(0.036)
t = 4	-0.896***
	(0.054)
t = 5	-0.639***
	(0.069)
t = 6	-0.391***
	(0.089)
Constant	-1.497***
	(0.155)

Product FE and period FE included. Tiers: Elite (six selective CRUCH universities), Other CRUCH (remaining universities of the Council of Rectors, CRUCH), G8 (eight private universities outside CRUCH) and SCP (short-cycle programs at any institution); the omitted tier is the remaining private universities (see Appendix Figure A15). Standard errors in parentheses (naive: CCPs and continuation values held fixed).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B4: Performance ordered logit: all coefficients

Female	-0.017
	(0.026)
Low SES	0.025
	(0.035)
Middle SES	0.101***
	(0.037)
PSU	-0.094***
	(0.030)
GPA	0.053***
	(0.018)
OP fees (thousand USD)	0.016***
	(0.004)
Relative ability	0.357***
	(0.017)
Diff. Region	-0.131***
	(0.014)
Credits left (t-1)	-0.076***
	(0.004)
Credits (t-1): 0	-1.397***
	(0.029)
Credits (t-1): 1	-1.195***
	(0.023)
Credits (t-1): 2	-0.525***
	(0.014)
Credits (t-1): 4	0.947***
	(0.013)
OP × female	0.004
	(0.003)
OP × low SES	-0.015***
	(0.005)
OP × middle SES	-0.013***
	(0.004)
OP × psu	-0.006**
	(0.002)
OP × gpa	0.005**
	(0.002)
Credits (t-1) × female	-0.017***
	(0.004)
Credits (t-1) × low SES	-0.013**
	(0.006)
Credits (t-1) × middle SES	-0.035***
	(0.006)
Credits (t-1) × psu	-0.065***
	(0.003)

continued

Credits (t-1) × gpa	-0.007**
	(0.003)
Credits left (t-1) × female	0.013***
	(0.002)
Credits left (t-1) × low SES	-0.012***
	(0.002)
Credits left (t-1) × middle SES	-0.009***
	(0.002)
Credits left (t-1) × psu	-0.020***
	(0.001)
Credits left (t-1) × gpa	0.016***
	(0.001)
Rel. Abil. × female	-0.027***
	(0.006)
Rel. Abil. × low SES	0.073***
	(0.009)
Rel. Abil. × middle SES	0.036***
	(0.009)
Diff. Region × female	0.025**
	(0.012)
Diff. Region × low SES	0.064***
	(0.016)
Diff. Region × middle SES	0.020
	(0.017)
Diff. Region × psu	-0.025***
	(0.008)
Diff. Region × gpa	0.018**
	(0.008)
Elite × female	-0.106***
	(0.021)
Elite × low SES	0.104***
	(0.026)
Elite × middle SES	0.013
	(0.026)
Elite × psu	0.065***
	(0.020)
Elite × gpa	0.009
	(0.016)
Other CRUCH × female	0.014
	(0.022)
Other CRUCH × low SES	0.048*
	(0.027)
Other CRUCH × middle SES	0.035
	(0.029)
Other CRUCH × psu	-0.036*
	(0.022)

continued

Other CRUCH × gpa	-0.006
	(0.016)
G8 × female	-0.037
	(0.029)
G8 × low SES	0.106***
	(0.036)
G8 × middle SES	0.041
	(0.035)
G8 × psu	-0.106***
	(0.030)
G8 × gpa	0.028
	(0.021)
SCP × female	0.136***
	(0.018)
SCP × low SES	0.030
	(0.023)
SCP × middle SES	-0.017
	(0.025)
SCP × psu	-0.051***
	(0.014)
SCP × gpa	0.100***
	(0.012)
Type 2	0.012**
	(0.006)
t = 2	-1.883***
	(0.022)
t = 3	-2.161***
	(0.025)
t = 4	-2.053***
	(0.031)
t = 5	-2.153***
	(0.039)
t = 6	-2.883***
	(0.046)
t = 7	-3.483***
	(0.055)
Cut point 1	-5.386***
	(0.091)
Cut point 2	-4.848***
	(0.091)
Cut point 3	-3.893***
	(0.091)
Cut point 4	-3.540***
	(0.091)

continued

Product FE and period FE included. Tiers: Elite (six selective CRUCH universities), Other CRUCH (remaining universities of the Council of Rectors, CRUCH), G8 (eight private universities outside CRUCH) and SCP (short-cycle programs at any institution); the omitted tier is the remaining private universities (see Appendix Figure A15). Standard errors in parentheses (naive: CCPs and continuation values held fixed).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B5: Fixed costs: all coefficients

EMax	0.950
	(.)
Female	-1.201***
	(0.173)
Middle SES	0.328
	(0.248)
Low SES	0.007
	(0.232)
PSU	-1.780***
	(0.147)
GPA	-0.298***
	(0.115)
OP fees (thousand USD)	0.031
	(0.023)
Diff. Region	0.573***
	(0.081)
Relative ability	1.671***
	(0.052)
Credits (t-1): 0	3.075***
	(0.153)
Credits (t-1): 1-2	1.178***
	(0.118)
Credits (t-1): 4	0.224*
	(0.115)
Credits left (t-1)	-0.005
	(0.011)
OP × female	0.020
	(0.016)
OP × middle SES	0.083***
	(0.025)
OP × low SES	0.098***
	(0.026)
OP × psu	-0.065***
	(0.012)
OP × gpa	0.022**
	(0.011)
Diff. Region × female	-0.000
	(0.068)
Diff. Region × middle SES	0.406***
	(0.092)
Diff. Region × low SES	0.629***
	(0.086)
Diff. Region × psu	0.446***
	(0.050)

continued

Diff. Region × gpa	0.368***
	(0.045)
Credits (t-1) × female	-0.047*
	(0.028)
Credits (t-1) × middle SES	-0.213***
	(0.039)
Credits (t-1) × low SES	-0.173***
	(0.035)
Credits (t-1) × psu	-0.053***
	(0.019)
Credits (t-1) × gpa	0.007
	(0.018)
Rel. Abil. × female	-0.032
	(0.026)
Rel. Abil. × middle SES	-0.435***
	(0.038)
Rel. Abil. × low SES	-0.573***
	(0.036)
Credits left (t-1) × female	0.057***
	(0.008)
Credits left (t-1) × middle SES	-0.040***
	(0.011)
Credits left (t-1) × low SES	-0.032***
	(0.010)
Credits left (t-1) × psu	-0.059***
	(0.006)
Credits left (t-1) × gpa	0.014**
	(0.005)
Enrol HE × Agricultural	0.837***
	(0.261)
Enrol HE × Arts and Architecture	0.381
	(0.242)
Enrol HE × Science	-0.497*
	(0.260)
Enrol HE × Social Sciences	-0.220
	(0.175)
Enrol HE × Law	-0.297
	(0.214)
Enrol HE × Education	-0.406**
	(0.189)
Enrol HE × Humanities	0.985***
	(0.379)
Enrol HE × SCP	-2.796***
	(0.200)
Enrol HE × SCP-Gratuidad	-3.639***
	(0.148)

continued

Enrol HE × Health	-0.635***
	(0.162)
Enrol HE × Technology	-0.385**
	(0.156)
Enrol HE	2.086***
	(0.166)
Enrol HE × Elite	-0.384***
	(0.119)
Enrol HE × Other CRUCH	1.025***
	(0.126)
Enrol HE × G8	0.654***
	(0.153)
Elite × female	0.257***
	(0.085)
Elite × middle SES	-0.709***
	(0.112)
Elite × low SES	-0.851***
	(0.109)
Elite × psu	-0.791***
	(0.078)
Elite × gpa	-0.947***
	(0.058)
Other CRUCH × female	0.418***
	(0.085)
Other CRUCH × middle SES	-0.735***
	(0.116)
Other CRUCH × low SES	-0.770***
	(0.112)
Other CRUCH × psu	-0.384***
	(0.082)
Other CRUCH × gpa	-0.521***
	(0.056)
G8 × female	0.184*
	(0.110)
G8 × middle SES	0.241*
	(0.138)
G8 × low SES	0.315**
	(0.134)
G8 × psu	-0.752***
	(0.104)
G8 × gpa	0.114
	(0.074)
SCP × female	1.059***
	(0.075)
SCP × middle SES	0.177*
	(0.106)

continued

SCP × low SES	0.371***
	(0.100)
SCP × psu	0.481***
	(0.056)
SCP × gpa	-0.111**
	(0.049)
Expected wage (USD million)	10.232***
	(1.038)
Type 2	-3.521***
	(0.113)

Product FE included; fixed costs in the cost sign. Tiers: Elite (six selective CRUCH universities), Other CRUCH (remaining universities of the Council of Rectors, CRUCH), G8 (eight private universities outside CRUCH) and SCP (short-cycle programs at any institution); the omitted tier is the remaining private universities (see Appendix Figure A15). Standard errors in parentheses (naive: CCPs and continuation values held fixed).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B6: Log marginal cost: all coefficients

Female	0.409***
	(0.070)
Low SES	-0.171**
	(0.087)
Middle SES	-0.339***
	(0.093)
PSU	-0.379***
	(0.055)
GPA	-0.078
	(0.048)
Type 2	0.144***
	(0.008)
OP fees (thousand USD)	-0.031***
	(0.008)
Diff. Region	0.185***
	(0.025)
Relative ability	-0.265***
	(0.021)
Credits (t-1)	-0.946***
	(0.010)
Credits left (t-1)	0.065***
	(0.006)
t = 2	2.943***
	(0.027)
t = 3	3.214***
	(0.033)
t = 4	2.793***
	(0.050)
t = 5	2.354***
	(0.066)
t = 6	2.849***
	(0.099)
t = 7	3.363***
	(0.119)
Credits (t-1) × female	-0.009
	(0.006)
Credits (t-1) × low SES	0.035***
	(0.007)
Credits (t-1) × middle SES	0.057***
	(0.008)
Credits (t-1) × psu	0.094***
	(0.004)
Credits (t-1) × gpa	-0.004
	(0.004)

continued

Credits left (t-1) × female	-0.034***
	(0.003)
Credits left (t-1) × low SES	0.020***
	(0.004)
Credits left (t-1) × middle SES	0.014***
	(0.004)
Credits left (t-1) × psu	0.023***
	(0.002)
Credits left (t-1) × gpa	-0.021***
	(0.002)
OP × female	-0.023***
	(0.006)
OP × low SES	0.034***
	(0.009)
OP × middle SES	0.033***
	(0.009)
OP × psu	0.014***
	(0.004)
OP × gpa	-0.011***
	(0.004)
Rel. Abil. × female	0.028***
	(0.010)
Rel. Abil. × low SES	-0.089***
	(0.014)
Rel. Abil. × middle SES	-0.055***
	(0.015)
Rel. Abil. × psu	0.000
	(0.005)
Rel. Abil. × gpa	-0.007
	(0.007)
Diff. Region × female	-0.045**
	(0.021)
Diff. Region × low SES	-0.044
	(0.028)
Diff. Region × middle SES	0.007
	(0.031)
Diff. Region × psu	0.030**
	(0.013)
Diff. Region × gpa	-0.028**
	(0.014)
Elite × female	0.175***
	(0.042)
Elite × low SES	-0.124**
	(0.053)
Elite × middle SES	0.065
	(0.054)

continued

Elite × psu	-0.346***
	(0.041)
Elite × gpa	-0.039
	(0.034)
Other CRUCH × female	-0.032
	(0.045)
Other CRUCH × low SES	-0.011
	(0.056)
Other CRUCH × middle SES	-0.055
	(0.064)
Other CRUCH × psu	-0.072
	(0.045)
Other CRUCH × gpa	0.022
	(0.032)
G8 × female	0.131**
	(0.055)
G8 × low SES	-0.147**
	(0.069)
G8 × middle SES	-0.100
	(0.067)
G8 × psu	-0.127**
	(0.056)
G8 × gpa	-0.033
	(0.042)
SCP × female	-0.253***
	(0.029)
SCP × low SES	0.020
	(0.035)
SCP × middle SES	0.087**
	(0.037)
SCP × psu	0.012
	(0.021)
SCP × gpa	-0.146***
	(0.019)
Constant	-4.390***
	(0.135)

Field FE, institution FE and period FE included. Tiers: Elite (six selective CRUCH universities), Other CRUCH (remaining universities of the Council of Rectors, CRUCH), G8 (eight private universities outside CRUCH) and SCP (short-cycle programs at any institution); the omitted tier is the remaining private universities (see Appendix Figure A15). Standard errors in parentheses (naive: CCPs and continuation values held fixed).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C Estimation

C.1 Two-step classification of types: an alternative starting value

The EM algorithm of Appendix Section C.4 is started from perturbed uninformative posteriors and, as a check, from a preliminary two-step classification of enrolled students into types. This subsection describes the classification, which is used only as a starting value and does not enter the estimates reported in the paper. Following the two-step grouped fixed-effects approach of Bonhomme et al. (2022), as implemented empirically by Joensen and Mattana (2024), we classify enrolled students into M latent types via k-means on two standardized, individual-average residual moments: completion conditional on registered credits and completion relative to the intended annual course load. Both moments are obtained as OLS residuals from

$$y_{it} = \mathbf{x}_i'\alpha_x + \mathbf{z}_{it}'\alpha_z + \lambda_{field(j)} + \lambda_{inst(j)} + \lambda_t + \varepsilon_{it}, \tag{15}$$

where y_{it} is, in turn, the registered-credit completion rate and the intended-annual-load completion rate, $\mathbf{x}_i$ collects individual controls (sex, standardized PSU and GPA, socioeconomic status), and $\mathbf{z}_{it}$ collects state controls (previous-period and cumulative credit completion, delay, and remaining-credits-to-graduate indicators). Fixed effects for field, institution (HEI), and period are included, allowed to differ between SCP and university programs. Residuals are collapsed to individual-level averages and standardized to unit variance before clustering.

Clustering uses k-means++ initialization: cluster centers are seeded sequentially with probability proportional to squared distance from existing centers, and the procedure is repeated over 20 independent seeds to guard against convergence to a local optimum, keeping the run with the lowest within-cluster sum of squares (WCSS). Resulting clusters are relabeled from highest to lowest average of the two standardized moments, so that Type 1 always denotes the highest-productivity group. The $M = 2$ classification is the one used as a starting value of the EM algorithm, which has two types; $M = 3$ is estimated on the same moments purely as a robustness check on this choice of M and is not used anywhere else.

Appendix Table C1 compares the fit and stability of the $M = 2$ and $M = 3$ classifications, and Appendix Table C2 compares the resulting type profiles.

Moving from $M = 2$ to $M = 3$ mechanically lowers WCSS and raises the share of variance explained (Appendix Table C1), since additional clusters can only improve within-sample fit. The improvement is not merely mechanical curve-fitting, however: the Calinski-Harabasz index also favors $M = 2$ over $M = 3$ once the loss of degrees of freedom is taken into

Table C1: Comparison of two- and three-type classifications

K	Students	Moments	WCSS	Q(K)	Explained (%)	CH	Best seed	Optimum starts	Min. share (%)	Max. share (%)
2	191,058	2	130,990	0.686	65.7	366,277	2	20	24.0	76.0
3	191,058	2	81,543	0.427	78.7	352,121	1	20	9.6	67.0

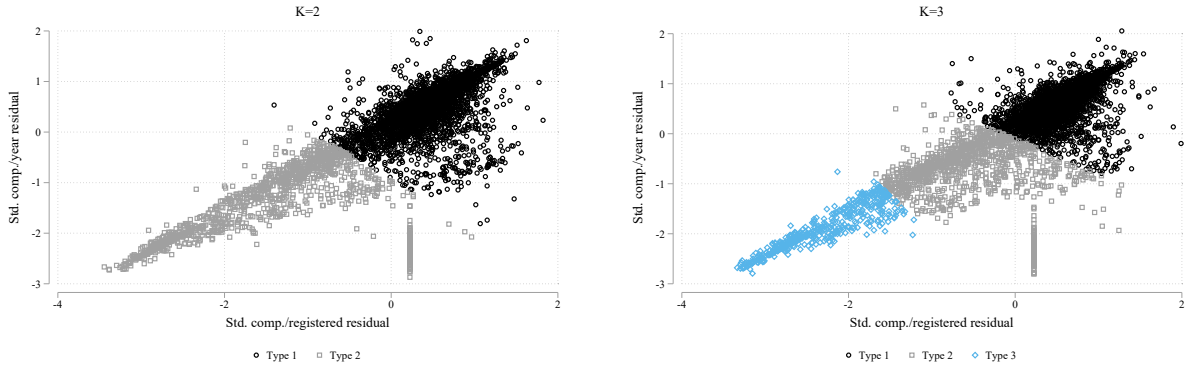
Notes: Both classifications use the same two standardized individual-average residual moments and 20 k-means++ starts. WCSS mechanically decreases with K. K=2 is the production classification; K=3 is appendix-diagnostic only.

Table C2: Profiles of two- and three-type classifications

	K=2		K=3		
Statistic	Type 1	Type 2	Type 1	Type 2	Type 3
Students	145,202	45,856	128,019	44,621	18,418
Share (%)	76.0	24.0	67.0	23.4	9.6
Std. comp./registered residual	0.448	-1.417	0.525	-0.508	-2.421
Std. comp./year residual	0.463	-1.467	0.576	-0.836	-1.979
Completed/registered ratio	0.902	0.299	0.930	0.560	0.106
Completed/intended-year ratio	0.830	0.211	0.870	0.389	0.078
Graduated by current credit state (%)	38.5	0.0	43.2	1.6	0.0
Behavioral dropout before t=6 (%)	53.6	94.7	49.8	87.9	99.1
Forced terminal exit at t=6 (%)	7.8	5.3	7.0	10.5	0.9
Observed academic periods	3.69	1.47	3.81	2.12	1.11

Notes: Types are ordered from highest to lowest mean standardized academic-productivity residual. Behavioral dropout excludes forced sixth-year nongraduate exit. K=2 is used in all structural estimates; K=3 is diagnostic only.

Figure C1: Classification moments by type: K=2 vs. K=3



a) K=2 b) K=3 (appendix diagnostic)

Notes: Standardized comp./registered residual (x-axis) against standardized comp./year residual (y-axis), by assigned type. A random subsample of 5,000 students is plotted for legibility; the k-means fit itself uses the full classification sample.

account, and the additional cluster comes from splitting the low-productivity $M = 2$ group (Type 2, 25.1% of students) into a moderate-productivity group and a very small, very-low-productivity group (Type 3 under $M = 3$, 10.5% of students, completed/registered ratio of 0.107, and essentially universal behavioral dropout before the sixth period). This third group is small enough, and separated enough from the rest of the distribution, that Appendix Figure C1 shows the additional structure is visually apparent but concentrated in a thin tail rather than reflecting a qualitatively distinct third mode. We therefore use the $M = 2$ classification, for parsimony and cluster-size balance, as the alternative starting value of the EM algorithm, treating $M = 3$ as a robustness check on the sensitivity of the classification to M.

C.2 Functional forms

This subsection gives the empirical specification of the two reduced-form objects that enter the estimation, the performance ordered logit, which delivers effective effort, and the dropout logit, which delivers the continuation values, as well as the structural fixed costs. Throughout, $Z_i = (Z_i^{obs}, m_i)$, where Z_i^{obs} collects the observed characteristics (a female dummy, low- and middle-SES dummies with high SES as the reference, and standardized PSU and high-school GPA scores) and $m_i \in \{1, 2\}$ is the unobserved trait of Section 5.1. Program j enters through its field $field(j)$, its institution tier $tier(j)$ (elite, other CRUCH, G8, off-platform and SCP) and a product fixed effect α_j, where a product is an institution–field pair. The outside option $j = 0$ is the reference product and every regressor below is zero for it. Two student–program match variables complete the program side: D_{ij} indicates that the program is located in a different region from the student's school, and $A_{ij} = (psu_i - \overline{psu}_j)/sd_j$ is the student's PSU score relative to the program's peers. The state x_{it} enters through the period t, previous-period performance $g_{i,t-1}$ (dummies $\mathbb{1}\{g_{i,t-1} = k\}$ for $k \in \{0, 1, 2, 4\}$, with $k = 3$ as the reference, and the linear value in interactions), the credits remaining to graduate $L_{i,t-1} = G_j^{req} - G_{i,t-1}$, and out-of-pocket fees $OP_j(x_{it})$, which depend on Z_i, $g_{i,t-1}$ and tuition as defined in the main text.

The unobserved trait enters each equation as an additive shifter $\delta\, \mathbb{1}\{m_i = 2\}$, with Type 1 as the reference: it shifts the level of effort, the dropout hazard and the utility of enrolling, but not their slopes with respect to fees or the other regressors. Because m_i is part of x_{it}, every object derived from these equations (effort, dropout probabilities, $EMAX$, $EWAGE$ and marginal costs) is type specific; the superscript m used in the main text is omitted here. Within the EM algorithm of Appendix Section C.4, each student contributes one observation per type, weighted by the posterior q_{im}.

C.2.1 Performance and effective effort

We estimate performance as an ordered logit of $g_{i,t+1} \in \{0, \dots, 4\}$ on a linear index $\tilde{e}_j(x_{it})$ with thresholds $\hat{c}_1 < \cdots < \hat{c}_4$,

$$\Pr\big(g_{i,t+1} = \kappa \mid x_{it}\big) = F\big(\hat{c}_{\kappa+1} - \tilde{e}_j(x_{it})\big) - F\big(\hat{c}_{\kappa} - \tilde{e}_j(x_{it})\big), \qquad \hat{c}_0 = -\infty,\ \hat{c}_5 = +\infty, \tag{16}$$

where F is the logistic CDF and

$$\begin{aligned}
\tilde{e}_j(x_{it}) = {} & Z_i^{obs\prime}\alpha_Z^e + \delta^e\,\mathbb{1}\{m_i = 2\} + \alpha_p^e\,OP_j(x_{it}) + \sum_{k\in\{0,1,2,4\}} \alpha_{g,k}^e\,\mathbb{1}\{g_{i,t-1} = k\} \\
& + \alpha_L^e\,L_{i,t-1} + \alpha_D^e\,D_{ij} + \alpha_A^e\,A_{ij} \\
& + \sum_n \Big(\alpha_{pn}^e\,OP_j(x_{it}) + \alpha_{gn}^e\,g_{i,t-1} + \alpha_{Ln}^e\,L_{i,t-1} + \alpha_{Dn}^e\,D_{ij} + \alpha_{An}^e\,A_{ij}\Big)\,Z_{in}^{obs} \\
& + \sum_\tau \sum_n \alpha_{\tau n}^e\,\mathbb{1}\{tier(j) = \tau\}\,Z_{in}^{obs} + \lambda_t^e + \alpha_j^e.
\end{aligned} \tag{17}$$

Only differences between the index and the thresholds are identified. We map the estimates to the model of Equation 7 by setting

$$\ln e_j^*(x_{it}) = \tilde{e}_j(x_{it}) - \hat{c}_1, \qquad \bar{g}^\kappa = \hat{c}_\kappa - \hat{c}_1, \tag{18}$$

so that the first threshold is zero and $\Pr(g_{i,t+1} = 0 \mid x_{it}) = F(-\ln e_j^*(x_{it})) = 1/(1 + e_j^*(x_{it}))$: effective effort is the odds of avoiding the lowest performance outcome, as in the main text. The field–institution–period variation in the thresholds of Equation 7 is captured by the product and period effects in the index, which move the four thresholds in parallel.

Three details on the specification. First, previous performance enters the main effects as categories and the interactions linearly (centered at the reference category), so that the interactions with observed characteristics are parsimonious while the main effects remain flexible. Second, A_{ij} is interacted with the female and SES dummies only, since it is itself built from the PSU score. Third, the tier interactions omit the off-platform tier and the period effects omit the first period. Out-of-pocket fees enter both directly and interacted with every observed characteristic, so that the funding instruments affect the effort margin as well as the choice margin. The equation is estimated on all enrolled student-periods of the panel sample, including the final academic period.

C.2.2 Dropout probability

The dropout probability $p_{0,t+1}$ of Equation 12 is a logit of the decision to leave the program once $g_{i,t+1} = \bar{g}$ is realized. It is estimated on the post-performance states at which a continue-or-exit choice exists: the student has not completed the program, is not mechanically forced out, and the post-performance state is realized before $T+1$. Its index has the same structure as Equation 17, evaluated at the post-performance state $x^+ \equiv x_{i,t+1}(\bar{g})$:

$$\begin{aligned}\text{logit } p_{0,t+1}(x^+) = {} & Z_i^{obs\prime}\alpha_Z^d + \delta^d \mathbb{1}\{m_i = 2\} + \alpha_p^d OP_j(x^+) + \sum_{k\in\{0,1,2,4\}} \alpha_{g,k}^d \mathbb{1}\{\bar{g} = k\} \\ & + \alpha_L^d L_{it}(\bar{g}) + \alpha_D^d D_{ij} + \alpha_A^d A_{ij} \\ & + \sum_n \Big(\alpha_{pn}^d OP_j(x^+) + \alpha_{gn}^d \bar{g} + \alpha_{Ln}^d L_{it}(\bar{g}) + \alpha_{Dn}^d D_{ij} + \alpha_{An}^d A_{ij}\Big) Z_{in}^{obs} \\ & + \sum_\tau \sum_n \alpha_{\tau n}^d \mathbb{1}\{tier(j) = \tau\} Z_{in}^{obs} + \lambda_t^d + \alpha_j^d,\end{aligned} \tag{19}$$

where $L_{it}(\bar{g}) = G_j^{req} - G_{i,t-1} - \bar{g}$ are the credits still missing after the outcome and $OP_j(x^+)$ are the fees the student will face next period given $\bar{g}$, including the loss of the scholarship subsidy when $\bar{g} < 3$. This is the channel through which the performance requirements of the funding instruments affect persistence. The same restrictions on the interactions apply as in the performance equation.

C.2.3 Fixed costs

Fixed costs collect every non-effort component of the utility of enrolling in program j, relative to the outside option:

$$\begin{aligned}FC_j(x_{it}) = {} & \mathbb{1}\{j \neq 0\}\Big(Z_i^{obs\prime}\alpha_Z + \delta \mathbb{1}\{m_i = 2\}\Big) + \alpha_p OP_j(x_{it}) + \sum_n \alpha_{pn} OP_j(x_{it}) Z_{in}^{obs} \\ & + \alpha_D D_{ij} + \alpha_A A_{ij} + \sum_{k\in\{0,\{1,2\},4\}} \alpha_{g,k} \mathbb{1}\{g_{i,t-1} \in k\} + \alpha_L L_{i,t-1} \\ & + \sum_n \Big(\alpha_{gn} g_{i,t-1} + \alpha_{Ln} L_{i,t-1} + \alpha_{Dn} D_{ij} + \alpha_{An} A_{ij}\Big) Z_{in}^{obs} \\ & + \sum_\tau \sum_n \alpha_{\tau n} \mathbb{1}\{tier(j) = \tau\} Z_{in}^{obs} \\ & + \mathbb{1}\{\text{first period}\}\Big(\alpha_E + \alpha_{E,field(j)} + \alpha_{E,tier(j)}\Big) + \alpha_j.\end{aligned} \tag{20}$$

The first term is the utility of being enrolled, which varies with observed characteristics and with the trait; this enrollment shifter is the only place where the trait enters fixed costs, so that the trait moves the level of utility from higher education but not the sensitivity to fees. The fee terms give the price sensitivity and its heterogeneity across observed characteristics; their coefficients are identified from the variation in funding eligibility (Section 5.2). The entry cost in the first period, which varies by field and institution tier, separates the one-off cost of entering higher education from the per-period cost of staying, which matters because continuation choices after the first period are between staying and exiting only. The product fixed effects α_j rationalize the observed enrollment shares across programs. Previous-performance categories 1 and 2 are pooled, as the corresponding cells are sparse in the choice sample. The parameters are estimated from the conditional value function of Equation 14, $v_j(x_{it}, e^*_{ijt}) = -FC_j(x_{it}) + \beta EMAX_{ijt} + \alpha_w EWAGE_{ijt}$, with β fixed at 0.95 and a wage coefficient α_w that is common across types and characteristics.

C.2.4 Labor market

The wage equation is Equation 11 and lifetime wages are built from it through Equation 4; we only document the sample and the implementation choices here. The equation is estimated on the pooled 2015, 2017 and 2022 waves of the CASEN household survey, on employed respondents aged 25 to 65, using the survey expansion weights, with survey-year effects (2015 as the base) so that all predictions are expressed in real 2015 US dollars of annual labor income. Potential experience is age minus years of schooling minus six, and the quadratic experience profile is allowed to differ between respondents who never enrolled, SCP graduates and university graduates; Equation 11 displays the common profile for brevity. The survey identifies nine broad fields of study, with SCP degrees pooled into a single field, so that SCP earnings are identified through the institution-type terms. Institution type distinguishes technical and professional SCP programs, private universities and CRUCH universities, and no institution-level premium is included beyond the type. Because predictions are used in levels, the intercept is shifted by the log of the weighted Duan smearing factor so that the exponentiated index is the conditional mean wage. The region R_{ir} is the region of the program for enrolled alternatives and the region of the student's school for the outside option.

The annual dropout wage $w_0(s, x_{it})$ after enrollment is predicted from a separate equation estimated on respondents who attended but did not complete higher education, with the same gender, field, region, institution-type and experience terms, survey-year effects, and the number of completed years of higher education (capped at four). In the model it is evaluated

at the program's field and institution type and at the student's accumulated credits converted into completed-equivalent years, so that the dropout wage increases with the year of exit. The never-enrollment wage of the outside option at the enrollment stage uses Equation 11 with the non-enrollment values of its arguments. Finally, lifetime wages are the discounted sum of the predicted annual wages over the working life in Equation 4, at an annual discount rate of 5 percent, expressed in millions of real 2015 US dollars.

C.3 Conditional value function

This subsection details the continuation-value and envelope expressions summarized in Section 5.3 and gives the closed forms implied by the ordered logit. The unobserved trait is part of Z_i and hence of x_{it}, so all objects below are evaluated separately for each type, with the type-specific effort and dropout probabilities of Appendix Section C.2.

Branches of the continuation value. For an enrolled alternative, Equation 5 gives $v_j(x_{it}, e_{it}) = -FC_j(x_{it}) - c_j(x_{it})e_{it} + \beta \sum_{\bar{g}\in\mathcal{G}} \phi^{\bar{g}}(e_{it})\bar{V}(x_{i,t+1}(\bar{g}))$. By Equation 6, after the performance draw $\bar{g}$ the student is in one of three situations. Either she has completed the program and earns the graduate wage; or she has not and, if $t < T+1$, chooses between staying and exiting, in which case the logsum is replaced by the Hotz–Miller inversion of Equation 12 with the dropout probability of Equation 19; or she has not completed the program and the outcome is realized in $t = T+1$, after the final academic period, in which case she exits mechanically. Writing $G_{it} = G_{i,t-1} + \bar{g}$,

$$\bar{V}\big(x_{i,t+1}(\bar{g})\big) = \begin{cases} \alpha_w\, W_j(x_{i,t+1}), & G_{it} \geq G_j^{req}, \\ \gamma + \alpha_w\, W_0(x_{i,t+1}) - \ln p_{0,t+1}\big(x_{i,t+1}(\bar{g})\big), & G_{it} < G_j^{req} \text{ and } t < T+1, \\ \alpha_w\, W_0(x_{i,t+1}), & G_{it} < G_j^{req} \text{ and } t = T+1. \end{cases} \tag{21}$$

The final academic period still contains effort and a performance draw; what is absent is a choice after that draw. The terminal nongraduation branch therefore contains the dropout wage but neither γ nor $-\ln p_0$, both of which arise only from a genuine stay-or-exit choice. The same wage-only branch applies to states from which completion is no longer feasible within the horizon, which we treat as mechanical exits. The outside option is terminal in every period, $v_0(x_{it}) = \alpha_w\, W_0(x_{it})$, with no discount factor multiplying this contemporaneous payoff.

Envelope representation. At an interior optimum the first-order condition Equation 10 gives $c_j(x_{it})\, e^*_{ijt} = \beta \sum_{\bar g} e^*_{ijt}\, \partial\phi^{\bar g}(e^*_{ijt})/\partial e\; \bar V(x_{i,t+1}(\bar g))$, so that

$$v_j(x_{it}, e^*_{ijt}) = -FC_j(x_{it}) + \beta \sum_{\bar g \in \mathcal{G}} a^{\bar g}_{ijt}\, \bar V\big(x_{i,t+1}(\bar g)\big), \tag{22}$$

with the weights $a^{\bar g}_{ijt}$ defined in Section 5.3. Inserting Equation 21 and separating the terms whose coefficient is fixed at β from those multiplying α_w yields Equation 14, with

$$EMAX_{ijt} = \mathbb{1}\{t < T\} \sum_{\bar g:\, G_{it} < G_j^{req}} a^{\bar g}_{ijt} \Big[\gamma - \ln p_{0,t+1}\big(x_{i,t+1}(\bar g)\big)\Big], \tag{23}$$

$$EWAGE_{ijt} = \beta \sum_{\bar g \in \mathcal{G}} a^{\bar g}_{ijt}\, W_{\bar g}(x_{i,t+1}), \qquad W_{\bar g} = \begin{cases} W_j & \text{if } G_{it} \geq G_j^{req}, \\ W_0 & \text{otherwise.} \end{cases} \tag{24}$$

Graduation branches contribute to $EWAGE$ only. In the final academic period $EMAX_{ijT} = 0$, whereas $EWAGE_{ijT}$ still depends on effort through the graduation probabilities, so that effort remains incentivized at $t = T$.

Ordered-logit closed forms. Let $e = e^*_{ijt}$ and, for the thresholds of Equation 18, define $F_\kappa \equiv F(\bar g^\kappa - \ln e)$ and $f_\kappa \equiv F_\kappa(1-F_\kappa)$, the logistic density at the same point, with $F_0 = f_0 = 0$ and $F_5 = 1$, $f_5 = 0$. The performance probabilities of Equation 8 are, using the symmetry $F(-z) = 1 - F(z)$,

$$\phi^\kappa(e) = F_{\kappa+1} - F_\kappa, \qquad \frac{\partial \phi^\kappa(e)}{\partial e} = -\frac{f_{\kappa+1} - f_\kappa}{e}, \qquad \kappa = 0, \ldots, 4, \tag{25}$$

where the derivative uses $\partial F(\bar g^\kappa - \ln e)/\partial e = -f_\kappa / e$. The envelope weights are therefore

$$a^\kappa_{ijt} = \phi^\kappa(e) - e\, \frac{\partial \phi^\kappa(e)}{\partial e} = F_{\kappa+1} - F_\kappa + f_{\kappa+1} - f_\kappa. \tag{26}$$

Since $\sum_\kappa \phi^\kappa(e) = 1$ and $\sum_\kappa e\, \partial\phi^\kappa(e)/\partial e = -(f_5 - f_0) = 0$, the weights sum to one, $\sum_\kappa a^\kappa_{ijt} = 1$: the envelope correction redistributes probability mass across performance outcomes without changing its total, although individual weights need not be positive. With the normalization $\bar g^1 = 0$, $F_1 = 1/(1+e)$ and $f_1 = e/(1+e)^2$, so that for instance $a^0_{ijt} = \frac{1}{1+e} + \frac{e}{(1+e)^2}$ and $a^4_{ijt} = 1 - F_4 - f_4$.

Finally, substituting Equation 25 in Equation 13 and summing by parts gives the marginal

cost in closed form,

$$c_j(x_{it}) = \beta \sum_{\kappa=0}^{4} \frac{\partial \phi^{\kappa}(e)}{\partial e} \bar{V}\big(x_{i,t+1}(\kappa)\big) = \frac{\beta}{e} \sum_{\kappa=1}^{4} f_{\kappa} \Big[\bar{V}\big(x_{i,t+1}(\kappa)\big) - \bar{V}\big(x_{i,t+1}(\kappa-1)\big) \Big], \tag{27}$$

a weighted sum of the gains in continuation value from moving up one performance category, with weights f_κ/e. Marginal costs are thus positive whenever continuation values increase in performance. This is the expression used to recover $c_j(x_{it})$ once the EM algorithm has converged.

C.4 Estimation algorithm

We estimate the model jointly using an expectation-maximization (EM) algorithm in the spirit of Arcidiacono and Miller (2011), in which the type posteriors are updated using both performance and choices. Two features matter for the sample. First, the trait distribution π_m is a *population* prior, independent of observed characteristics (Section 5.1), and it is defined for every student in the structural sample, including those who never enroll. Individuals who choose the outside option have no performance or dropout history, so they inform the type distribution only through their enrollment decision, which is what allows the type shares to be identified in the population rather than among enrollees. Second, students observed in the panel of enrolled students but outside the structural sample have no enrollment choice set; for them the posterior is updated from performance and dropout alone, with a prior conditional on enrollment. The algorithm proceeds as follows.

1. **Wages.** We estimate the wage equation Equation 11 on survey data and predict graduation and dropout wages for every program-state. Since the survey is not linked to students' higher education histories, wages do not depend on the unobserved trait and can be estimated once, outside the EM loop.

2. **Starting values.** The EM algorithm requires starting values for the utility parameters θ and for the posteriors q_{im}. Following Arcidiacono and Miller (2011), θ starts at the estimates of the model without unobserved heterogeneity, the type shifters at zero (a small offset separates the two types), and the type probabilities π_m at the mean of the starting posteriors. For the posteriors we start from an uninformative 0.5/0.5 split, perturbed by independent uniform draws for each student and type and renormalized to sum to one, so that the starting posterior is centered on one half but differs across students. Because the mixture likelihood may have several local maxima, we repeat the estimation from three such perturbations, with different seeds, and from the two-

step classification of Appendix Section C.1, in which case the start also includes the type coefficient estimated with the classified types. Appendix Table C3 compares the converged estimates.

3. **EM loop.** Iterate over the following steps until the posteriors, the parameters and the value of the objective stop changing:

   (a) *Reduced-form components (performance ordered logit and dropout CCPs), M-step.* Weighting each student by the current posterior q_{im}, we estimate the ordered logit of credit completion (0–4), which delivers the index $\ln e_j^*(x_{it})$, and the dropout logit $p_{0,t+1}$. Both are specified in Appendix Section C.2 and include product fixed effects, period fixed effects, out-of-pocket fees and trait shifters. They are estimated on the *panel sample*, which contains all enrolled students and is considerably larger than the structural sample, so they are estimated more precisely. This step delivers, for each student and type, the likelihood L_{im}^{perf} of the observed sequence of performance outcomes. The dropout logit is used differently: it is not a likelihood term for students in the structural sample, whose observed continuation decisions are explained by the structural model below. Its role there is to supply the predicted dropout probabilities $p_{0,t+1}$ at the states of the structural data, for each type.

   (b) *Continuation values.* Evaluating the estimated dropout probabilities at every state of the structural data, we obtain $\bar{V}(x_{i,t+1}(\bar{g}))$ through Equation 12 separately for each type, and hence $EMAX_{ijt}^m$, $EWAGE_{ijt}^m$ and the marginal costs $c_j^m(x_{it})$ through Equation 13.

   (c) *Fixed costs, utility parameters and type prior, M-step.* Taking the continuation values as given, we maximize over the fixed-cost and utility parameters θ (including FC_j and α_w) and over the type probabilities π_m, the mixture likelihood of all students in the *structural choice sample*, enrolled or not:

$$\sum_i \ln \sum_m \pi_m \, L_{im}^{perf} \prod_t P_{imt}(\theta),$$

   where $P_{imt}(\theta)$ is the probability of student i's observed choice at t given type m: the program or outside-option choice at $t = 1$, based on v_j in Equation 14, and the continue-or-exit decision afterwards. The dropout CCP is not included in L_{im}^{perf}, because the continue-or-exit decision already enters through P_{imt} and would otherwise be counted twice. Program fixed effects are included in FC_j,

so that fixed costs rationalize the observed enrollment shares across programs, while the sensitivity to out-of-pocket fees is identified from the variation in public funding eligibility. Since the type prior is a constant, its maximizer is the average of the current posteriors over the structural sample.

(d) *Type updating, E-step.* For students in the structural sample, the posterior is

$$q_{im} = \frac{\pi_m \, L_{im}^{perf} \prod_t P_{imt}(\theta)}{\sum_{m'} \pi_{m'} \, L_{im'}^{perf} \prod_t P_{im't}(\theta)},$$

so that it uses the program choice as well as performance, and never-enrollees are assigned a posterior that reflects their characteristics and their decision not to enroll. For students in the panel who are not in the structural sample, we use

$$q_{im} \propto \pi_m^E \, L_{im}^{perf} \, L_{im}^{drop},$$

where L_{im}^{drop} is the likelihood of the observed continuation or dropout decisions and π_m^E is the average posterior among these students at the current iteration. Since these students are enrollees, π_m^E describes the type mix conditional on enrollment, not the population prior. Students observed for more periods have more informative posteriors, and the posteriors are carried over as weights to the next iteration.

4. **Marginal costs.** Once the algorithm has converged, we use the final dropout CCPs, performance ordered logit and posteriors to compute $c_j(x_{it})$ from the FOC Equation 13, stacking each realized higher education episode once per type and weighting it by q_{im}.

Stability across starting values. Appendix Table C3 reports the converged estimates from the four starts. The three perturbed uninformative starts converge to the same point: their log likelihoods differ by less than one unit, and the type shares and the coefficients agree to the third decimal. The start from the two-step classification converges to a nearby point with a log likelihood lower by about eight units, a slightly larger wage weight and type shift, and the same type shares and fee coefficients. The classification initially assigns a posterior of zero or one to every enrolled student, with a type 2 share of one third in the population, but the posteriors move to the same split of about 0.6 within the first three outer iterations, so that the type is pinned down by the data rather than by the start. We retain the estimates of the first perturbed start, which attains the highest likelihood, for the results reported in the paper.

Table C3: Stability of the EM estimates across starting values

	Perturbed 0.5/0.5 posterior			Classification
Starting values	Random 1	Random 2	Random 3	Two-step k-means
Fit and convergence				
Log likelihood	-75,725.3	-75,725.8	-75,725.6	-75,733.4
Outer EM iterations	41	41	41	33
Converged	Yes	Yes	Yes	Yes
Type distribution				
Type 2 share (population)	0.596	0.596	0.596	0.597
Type 2 share (enrollees)	0.822	0.822	0.822	0.821
Fixed costs (cost sign)				
Type 2	-3.521	-3.522	-3.521	-3.577
Enrol HE	2.086	2.087	2.086	2.119
Female	-1.201	-1.201	-1.201	-1.207
PSU	-1.780	-1.780	-1.780	-1.789
OP fees (thousand USD)	0.031	0.031	0.031	0.032
OP × low SES	0.098	0.098	0.098	0.097
OP × middle SES	0.083	0.083	0.083	0.083
Expected wage (USD million)	10.232	10.232	10.232	10.762
Performance ordered logit				
Type 2	0.012	0.008	0.010	-0.042
OP fees	0.016	0.016	0.016	0.016
Dropout logit				
Type 2	-0.673	-0.673	-0.673	-0.544
OP fees $(t+1)$	0.018	0.018	0.018	0.018

Notes: Converged joint-EM estimates by starting value. Random 1–3: posteriors start at 0.5/0.5 perturbed by independent uniform draws (three seeds), utility parameters at the estimates of the model without unobserved heterogeneity. Two-step k-means: posteriors and type coefficient start from the classification of Appendix Section C.1. Type 1 is the reference type. Fixed costs are reported in the cost sign; the performance and dropout coefficients keep their estimated sign. Random 1 is the estimate used in the paper (Table 1).

Latent types in the simulations. The counterfactual simulations of Section 7 require a type for each student. Solving the dynamic problem once per type for every student and weighting the two solutions by the posterior doubles the size of an already large state space. Instead, we assign each student a permanent type with a single uniform draw: student i is of type 2 if $u_i < q_{i2}$, with $u_i \sim U[0,1]$ and q_{i2} her converged posterior, and of type 1 otherwise. The draw is made once and reused in every policy scenario, so that a student's type is part of the common random numbers of Section 7.1. Because the posterior is the probability of each type given the student's observed history, the drawn types reproduce the posterior type distribution in expectation, and for a sample of our size the simulated aggregates coincide, up to simulation noise, with those of the posterior-weighted mixture.

D Cost-benefit accounting of the counterfactual policies

This appendix details the cost-benefit accounting summarized at the end of Section 7.2 and used to compare funding designs in Section 7.3. We translate each scenario into costs and benefits per student of the simulated cohort, in thousand 2015 USD and present value at entry (5% discount rate), relative to the pre-policy regime. The accounting rests on the following assumptions.

1. *Public outlay.* TF is paid at the regulated tariff, set to 80% of list tuition, plus scholarships, Nuevo Milenio, BVP and, where present, stipends, which are paid in full. Loans are not counted as public spending and program-funded grants are treated as private. The *mechanical* component re-prices the status-quo enrollment paths under the new rule; the *behavioral* component is the remaining change in aid caused by changes in enrollment and persistence.

2. *Earnings and the fiscal externality.* Expected lifetime earnings are the present value of model-predicted wages over 40 years of experience, without shocks, with separate profiles for graduates, dropouts (whose wage rises with credits earned) and non-enrollees. The fiscal externality is the change in the present value of the Chilean income tax, applying the standard progressive schedule (exempt below 13.5 monthly UTM and rising to 40%) year by year to each student's earnings profile. We do not count VAT, social contributions or any labor-supply response. Because projected earnings are close to the exempt threshold, the tax recovers only about 3% of TF's outlay.

3. *Willingness to pay and MVPF.* The conservative measure of willingness to pay is the

reduction in out-of-pocket spending on the status-quo enrollment paths (students who change their behavior value the price change at the margin at about zero, by the envelope theorem). The upper bound adds the after-tax earnings gain, net of the extra out-of-pocket spending of those students. The marginal value of public funds is willingness to pay divided by net public cost, i.e. outlay minus the fiscal externality (Hendren and Sprung-Keyser, 2020).

4. *Scope.* The analysis is partial equilibrium: wages, tuition, capacity and institutional adherence are held fixed. Simulation noise (100 replications with common random numbers) is small relative to the differences we report; we do not account for parameter uncertainty.

Table D1 reports the results. TF costs $2,100 per student, two thirds of which is the mechanical transfer to students who would have enrolled anyway, and buys 1.9 p.p. more graduates, at a net cost of $108,000 per additional graduate (MVPF between 1.1 and 2.2). The entry requirement saves 61% of the outlay but at $308,000 per additional graduate, nearly three times the cost of TF: the students it removes are those for whom aid changes the outcome most. The continuation requirement halves the outlay, keeps most of the graduation gain and lowers the net cost per additional graduate to $70,000, with an MVPF of 1.2 to 3.0.

Table D1: Back-of-the-envelope cost-benefit of the counterfactual funding policies

	Status quo	TF	Entry requirement	Continuation requirement	Both requirements
A. Outcomes (status-quo level; change vs status quo)					
Enrollment (%, p.p.)	59.5	+2.97	+0.28	+2.40	+0.19
Graduation (%, p.p.)	20.1	+1.94	+0.27	+1.49	+0.18
PV earnings ($000)	141.9	+2.45	+0.45	+1.97	+0.36
B. Public cost ($000 per cohort student)					
Public student aid (level)	2.14	4.30	2.98	3.23	2.34
Mechanical transfer		+1.46	+0.72	+0.68	+0.12
Behavioral response		+0.69	+0.11	+0.41	+0.07
Tax on earnings gain		-0.07	-0.02	-0.05	-0.01
Net public cost		+2.09	+0.82	+1.04	+0.18
C. Willingness to pay ($000) and MVPF					
Willingness to pay, conservative		+2.21	+1.26	+1.22	+0.51
Willingness to pay, upper bound		+4.58	+1.67	+3.09	+0.86
MVPF, conservative		1.06	1.54	1.18	2.86
MVPF, upper bound		2.19	2.04	2.98	4.77
D. Cost per additional student ($000)					
Gross cost per extra enrollee		73	303	45	101
Gross cost per extra graduate		111	315	73	104
Net cost per extra graduate		108	308	70	97

Notes: Per student of the simulated cohort, thousand 2015 USD, present value at entry (5% discount rate). The first column is the level under the simulated status quo (enrollment and graduation in %); the other columns are changes relative to it. Public aid counts TF at the regulated tariff (80% of list tuition) plus scholarships, Nuevo Milenio and BVP; loans, direct funding to institutions and other student support are not counted, so public student aid is not total public spending on higher education. The change in public cost has three parts. The *mechanical transfer* applies the policy's funding rule to the enrollment paths that students would have followed without the policy, holding their behavior fixed. The *behavioral response* is the further change in aid because the policy changes who enrolls and for how long. The *tax on earnings gain* is the change in the present value of the Chilean income tax (standard progressive schedule, 0–40%, applied year by year to the experience profile of each student's earnings, discounted back by years enrolled); a negative number is a saving. Net public cost is the sum of the three parts. Willingness to pay (WTP) is what students gain from the policy: the conservative measure is the reduction in their out-of-pocket spending on the status-quo enrollment paths (students who change behavior are valued at about zero at the margin), and the upper bound adds the after-tax earnings gain net of the extra out-of-pocket spending of those students. The marginal value of public funds is MVPF = WTP / net public cost (Hendren and Sprung-Keyser, 2020): the dollars of benefit per dollar the government spends. It is infinite when the net public cost is zero or negative, i.e. the policy pays for itself. Cost per additional student divides the outlay by the additional enrollees or graduates relative to the status quo; the net version subtracts the income tax. These are incremental measures and include transfers to students whose behavior does not change.