ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS

WORKING PAPER SERIES

Raising College Access and Completion: How Much Can Free College Help?

Authors	Maria Marta Ferreyra, Carlos Garriga, Juan David Martin, and Angelica Maria Sanchez Diaz
Working Paper Number	2020-048A
Creation Date	October 2020
Citable Link	https://doi.org/10.20955/wp.2020.048
Suggested Citation	Ferreyra, M.M., Garriga, C., Martin, J.D., Sanchez Diaz, A.M., 2020; Raising College Access and Completion: How Much Can Free College Help?, Federal Reserve Bank of St. Louis Working Paper 2020-048. URL https://doi.org/10.20955/wp.2020.048

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

Raising College Access and Completion: How Much Can Free College Help? *

Maria Marta Ferreyra
The World Bank

Carlos Garriga
Federal Reserve Bank of St. Louis

Juan David Martin
The World Bank

Angelica Maria Sanchez Diaz
Georgetown University

July 8, 2020

Abstract

Free college proposals have become increasingly popular in many countries of the world. To evaluate their potential effects, we develop and estimate a dynamic model of college enrollment, performance, and graduation. A central piece of the model, student effort, has a direct effect on class completion, and an indirect effect in mitigating the risk of not completing a class or not remaining in college. We estimate the model using rich, student-level administrative data from Colombia, and use the estimates to simulate free college programs that differ in eligibility requirements. Among these, universal free college expands enrollment the most, but it does not affect graduation rates and has the highest per-graduate cost. Performance-based free college, in contrast, delivers a slightly lower enrollment expansion yet a greater graduation rate at a lower per-graduate cost. Relative to universal free college, performance-based free college places greater risk on students, but precisely for this reason leads them to better outcomes. Nonetheless, even performance-based free college fails to deliver a large increase in graduation rate, suggesting that additional, complementary policies might be required to elicit the large effort increase needed to raise graduation rates.

Keywords: Higher Education, free college, financial aid. *JEL codes*: E24, I21

*We thank Kartik Athreya, Raquel Bernal, Adriana Camacho, Stephanie Cellini, Doug Harris, Oksana Leukhina, CJ Libassi, Alexander Ludwig, Fabio Sanchez, James Thomas, and Susan Vroman for useful comments and suggestions. We thank Andrea Franco Hernandez for excellent research assistance. We benefited from seminars at The World Bank, the St. Louis Fed, the DC Economics of Education Working Group, and Universitat Autonoma de Barcelona, and by sessions at AEFP 2018, SED 2018, NASMES 2019, and LACEA 2018. Ferreyra acknowledges financial support from the World Bank. The views expressed herein do not necessarily reflect those of The World Bank, the Federal Reserve Bank of St. Louis, or the Federal Reserve System. Email contacts: mferreyra@worldbank.org, carlos.garriga@stls.frb.org, juandavid.0889@gmail.com, amsanchez456@gmail.com. Contact address: Maria Marta Ferreyra, The World Bank, 1818 H Street, NW, Washington, DC 20433.

1 Introduction

In modern economies, higher education is crucial role to the formation of skilled human capital. Not only can higher education raise a country's productivity; it can also lower income inequality. By subsidizing access to higher education, policymakers can contribute to these two roles. The question, of course, is how large a subsidy they should provide. Advocates of free college argue that policymakers should provide a full subsidy, resulting in zero tuition for students. While free college has existed for years in a number of countries,[1] free college proposals have sprouted recently in other countries, including the United States, Chile, and Colombia.

Free college advocates claim that only when college is free can it realize its promise, particularly that of lowering inequality. Given Latin America's distinction of being the most unequal region of the world (World Bank 2016), many in the region view free college as the ultimate solution to persistent, intergenerational inequality, and as an engine for social mobility. The existing evidence on government spending in higher education in the region, however, does not bode well for free college. As Figure 1 shows, countries in the region that finance a greater share of higher education spending do have higher enrollment rates (panel a), yet they fail to have higher graduation rates (panel b).

Figure 1: Higher Education Enrollment and Graduation by Public Spending Share in Latin America.

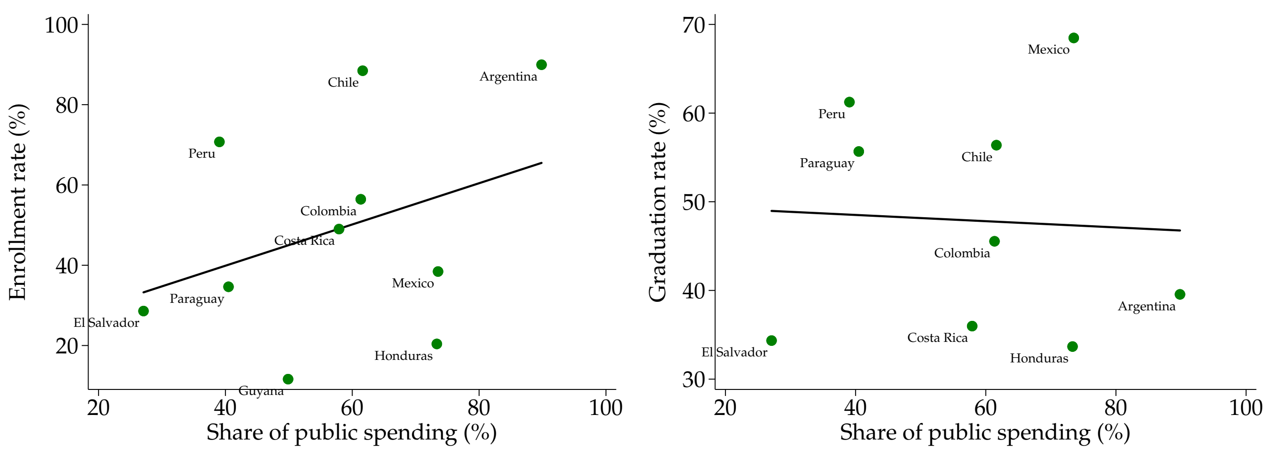

a. Enrollment

b. Graduation

Source: UNESCO for enrollment rates (panel a) in year 2017; own calculations based on SEDLAC (household surveys) for graduation rates in yearn 2017 (panel b); own calculations based on UNESCO for share of public spending (year 2017).
Notes: Panel a shows gross enrollment rate, defined as the ratio between higher ed enrollment and the number of individuals ages 18-24. Panel b shows graduation rate, defined as the ratio of individuals ages 25-29 who have graduated from higher education, and the number of individuals ages 25-29 who have ever started higher education. Share of public spending is defined as government spending in higher education relative to total spending in higher education.

To investigate the potential effects of free college, in this paper we develop a dynamic model of college enrollment, performance, and graduation. In the model, a college student faces risks

[1] Among the countries surveyed by OECD (2018), college is free at public institutions in Argentina, Brazil, Cuba, Czech Republic, Denmark, Ecuador, Estonia, Egypt, Finland, Germany, Greece, Iceland, Mexico, Norway, Panama, Poland, Slovenia, Sweden, Turkey, and Uruguay.

that may prevent her from completing a class or remaining in college. Through her effort, the student can mitigate these risks. We estimate the model with unique administrative data on the universe of higher education students in Colombia, a country which is highly representative of Latin America and the Caribbean, and use the parameter estimates to simulate free college programs differing in eligibility requirements. We study whether free college affects the risk facing students and perform a simple cost-benefit analysis of free college.

Understanding the trade-offs involved by free college is critical for any country considering it, particularly in developing economies. In Latin America, higher education enrollment rates rose from 21 to 40 percent between 2000 and 2010 (Ferreyra et al 2017). Since only 14 percent of the working age population has completed higher education, Mincerian returns are high (104 percent on average relative to a high school diploma). All Latin American countries subsidize public education at high rates, yet few provide student loans or funding for students at private institutions. While students in several of these countries have recently taken to the streets in demand of free college, the feasibility of free college has been called into question by the low growth and tight fiscal constraints of these countries in recent years - conditions, of course, that have only been aggravated by the COVID-19 pandemic.

For our analysis of Colombia, we focus on bachelor's programs, which capture 80 percent of higher education enrollment. In the high school class of 2005, 70 percent of students come from low-income families, and only 32.3 percent of students enroll in college within a five-year window. Among those who enroll in college, only 46 percent graduate, mostly late. Although student income and ability seem strongly related with enrollment, only ability seems strongly related with dropout rates – and this is through first-year dropouts, who account for most dropouts and are largely low-ability students. Meanwhile, the average number of classes completed by students per year (henceforth, performance) varies little across ability groups yet greatly within ability groups. The speed at which students complete the classes required for graduation (henceforth, "cumulative performance" or "cumulative classes completed") is highly persistent: the more classes a student completes in her first year, the more likely she is to keep the pace in subsequent years. Conversely, students with poor initial performance find it difficult to catch up later and are more likely to fall behind or drop out.

We write a model that captures these features of the data. In the model, high school graduates who are heterogeneous in income, ability, and idiosyncratic preferences decide whether to enroll in college. Depending on their final educational attainment, they will earn the wage corresponding to a high school graduate, college graduate, or college dropout. College graduation requires the completion of a set number of classes; students can take between 5 and 8 years to graduate. Each year, a college student chooses the number of classes she expects to complete (henceforth, her target), and this determines the effort she must make. Effort is costly, particularly for low-ability or low-income students. The number of classes completed by the student in a year is a function of her ability, effort, and a performance shock which depends partly on her cumulative performance. At the end of the year she receives a dropout shock, also dependent on cumulative performance, which may force her to drop out. Effort, then, has a direct effect on the number of classes completed in a year, and an indirect effect on the risk to class completion in subsequent years and to college continuity.

In order to graduate from college, the student must complete a predetermined number of classes, and each one must be completed in its totality. These indivisibilities limit policy impact: to affect class completion, a policy must induce a discrete, non-marginal effort change such that the student completes at least one additional class. And, to affect graduation rates, the policy must induce a large enough effort increase to complete all the required classes. These discrete effort changes may be simply too costly for some students.

We estimate the model using Simulated Method of Moments. We fit moments related to dropout, graduation, and cumulative performance (including patterns of persistence, catching up, and falling behind). The model evaluated at the parameter estimates (henceforth, the baseline) fits the data well. According to our estimates, effort has much greater impact than ability on the production of classes completed. If effort were not modeled as an input to classes completed, we would overestimate the role of ability by about 75 percent. From a policy standpoint, this would lead to an over-reliance on policies that promote selection of the most able students (positive selection) rather than policies which explicitly promote effort.

We simulate multiple free college programs differing in eligibility requirements: 1) universal (all students), 2) need-based (low-income students), 3) ability-based (high-ability students), and 4) performance-based (all students eligible in the first year; eligibility conditional on past cumulative performance in subsequent years). We also simulate a need-based version of (3) and (4). In each counterfactual we distinguish between existing students (who enroll in the baseline and the counterfactual) and new students (who do not enroll in the baseline but enroll in the counterfactual) to assess the impact of free college on graduation.

By lowering tuition to zero, free college raises consumption during college. This enhances the attractiveness of the being a college student (the "college experience") and has three effects on effort. First is the *loss-of-urgency effect*, whereby the student wishes to enjoy the enhanced college experience and loses the urgency to graduate. Second is the *substitution effect*, whereby the enhanced consumption compensates for greater effort. Other things equal, the loss-of-urgency effect leads to lower effort, whereas the substitution effect leads to more effort. Third is the *risk effect*, as the effort changes induced by the other two effects lead to performance changes which, in turn, affect the performance and dropout risks. Further, the longer the student stays in college, the more she exposes herself to risks. Which of these three effects prevails varies across students and eligibility requirements.

At the aggregate level, all free college programs expand enrollment. The largest expansion is for universal free college, followed by need-based and performance-based free college. Relative to the baseline, these programs increase enrollment by 70-85 percent. In our simulations, on average new students are of lower income and ability than existing students. The exception is ability- and ability-and-need-based free college, which induce positive selection of new students and attract new students who are more able, on average, than existing ones.

In contrast with these large enrollment rate effects, overall graduation rate effects are modest – between -2 and 7 percent relative to the baseline. For new students, the graduation rate effect depends on the policy-induced type of selection. For existing students, only performance-based programs accomplish an effect substantially different from zero, raising graduation rates between 9 and 14 percent relative to the baseline. This is because, by making free college contingent on performance, these programs incentivize effort and eliminate the loss-of-urgency effect, making students frontload effort in the early years. These results are consistent with the literature on college financial aid in the U.S., which has generally found positive and large effects on enrollment, small or null effects on graduation, and larger graduation effects for performance-based than unconditional financial aid.[2]

At the same time, these aggregate effects mask great heterogeneity across students. Consider, for instance, universal free college. Enrollment effects are largest for low- and middle-income students, and for mid-ability students. Hence, universal free college subsidizes many students who, by virtue of their ability or income, do not need the subsidy as they already enroll in the baseline. Graduation rate effects are similarly heterogeneous across students. They fall for high-ability or high-income students, who experience a strong loss of urgency, while they rise for low-ability or low-income students, who experience a strong substitution effect. In contrast to universal free college, performance-based free college induces greater effort on the part of *all* students and leads them all to higher graduation rates.

Our counterfactual findings provide an explanation for the enrollment and graduation rate patterns presented in Figure 1. Greater college funding substantially raises enrollment rates when a large fraction of high school graduates faces severe financial constraints, as they do in Latin America. It does not, however, raise graduation rates unless it is performance-based to incentivize effort, which is not the case in Latin America.

Interestingly, the higher effort and graduation rates induced by performance-based free college come at the cost of placing greater risk on students. We develop a measure of student anticipated risk in each year (equal to the coefficient of variation of the value of college), and compare it for the baseline and counterfactuals. We find that, in every scenario, anticipated risk falls when students exert greater effort or accumulate more completed classes. Anticipated risk is high in the initial two years and decreases rapidly afterwards, once students survive the initial attrition and settle on a performance path. In the initial years, universal free college lowers students' risk relative to the baseline by enhancing the college experience, whereas performance-based free college raises it by making the college experience contingent on performance. Facing the students with this greater risk, however, induces them to exert greater effort, which is their ultimate insurance mechanism. In other words, better college outcomes do not come from providing full but rather partial insurance to students.

At the same time, even the graduation rate increase from performance-based free college is relatively small. This is, in part, due to a composition effect, as the new college students in several programs are less likely to graduate than the existing ones. But, even among existing students, graduation rates rise relatively little. In other words, free college alone cannot substantially raise graduation rates. The indivisibilities discussed above help explain why: while free college might raise effort for some or even all students, it still fails to induce in many students the large effort increase needed to complete all graduation requirements.[3]

For a policymaker committed to providing free college, the question is how to choose among the programs presented here. We conduct a simple cost-benefit analysis to illuminate this question, and compare the per-graduate cost across programs. We find that all programs raise the per-graduate cost – if anything, because fewer students pay for college than in the baseline. For a policymaker who wishes to raise the fraction of high school graduates that finish college while limiting costs, the best option is performance-based or even need-based free college – but not universal free college. Nonetheless, the per-graduate cost of every scenario studied here is

[2]For recent reviews of this vast literature, see Avery et al (2019) and Dynarski and Scott-Clayton (2013).

[3]This result is reminiscent of Oreopoulos and Petronijevic (2019), who find that even when students realize that more effort is needed to improve outcomes, they adjust by lowering expectations rather than increasing effort.

far from low – ranging from one (baseline) to 2.5 (universal free college) times the per capita GDP – and should be considered with great care given current fiscal constraints.

By construction, our counterfactuals assume the most favorable scenario for free college. We assume that colleges have no capacity constraints; the average and marginal cost of educating new and existing students are the same; and free college does not crowd out parental transfers to their children in college. Further, we do not model taxation (which might be required to pay for free college programs), and assume that the wage of college graduates relative to high school graduates (henceforth, the college premium) does not fall with more college graduates.[4] Relaxing any of these assumptions would lead to less favorable free college outcomes.

The rest of the paper is organized as follows. Section 2 describes the related literature, and Section 3 describes our data. Section 4 presents our model, and Section 5 discusses its empirical implementation. Section 6 describes the estimation strategy and results. Section 7 presents the free college counterfactuals, including analyses of anticipated risk, fiscal costs, and potential general equilibrium effects on the labor market. Section 8 concludes.

2 Related Literature

Our paper relates to a large literature estimating sequential schooling models under uncertainty, with seminal contributions by Keane and Wolpin (2001), Eckstein and Wolpin (1998), and Keane (2002). This literature models college enrollment, performance, and college outcomes, and uncovers structural parameters based on students' observed choices during college. In one strand of this literature, researchers model students as acquiring information (learning) throughout college – regarding, for instance, their ability and preferences for college or specific majors, and their expected labor market performance. This literature includes, among others, Arcidiacono (2004), Arcidiacono et al (2016), Ozdagli and Trachter (2011), Stinebrickner and Stinebrickner (2014), and Trachter (2015). As in these papers, students in our model learn about their graduation probability based on their classes completed, and choose effort accordingly.

The idea that higher education is risky is not new (Levhari and Weiss 1974, Altonji 1993, Akyol and Athreya 2005), but the recent availability of college transcript data in the U.S. has helped estimate the role of risk in students' performance. These data reveal substantial and persistent heterogeneity in students' credit accumulation rates, which are strongly related to graduation probability. According to Hendricks and Leukhina (2017, 2018), based on their credit accumulation rates more than 50 percent of college entrants should be able to forecast whether they are at least 80 percent likely to graduate. According to Stange (2012), the large uncertainty faced by students makes them place a high value on the ability to drop out at any point in college rather than pre-commit to completing all graduation requirements. Our paper is similar to these in the use of administrative data to track students' performance, but different in that the risk associated to class completion or college continuity is not fully exogenous as in these two papers, but depends on an endogenous variable – student effort.

While the literature has placed much attention on the role of ability in performance and college outcomes, a growing line of research highlights the role of effort. Zamarro, Hitt, and

[4] In Section 7 we investigate potential general equilibrium effects associated with the greater supply of college graduates. Since we find them to be very small even in the medium run, we conclude that we can abstract away from them in our analysis.

Mendez (2019) use data from the Program for International Student Assessment (PISA) to show that different effort measures explain about a third of observed cross-country test score variation. Stinebrickner and Stinebrickner (2004) rely on time use surveys to estimate the effects of study time on grades. Ariely et al (2009) show that the use of incentives can help the average student improve her test performance, though the effect is more limited on high-ability students. Beneito et al (2018) provide evidence that the tuition increase implemented by Spanish colleges in 2012 boosted student effort. Ahn et al (2019) model effort in response to grading policies. We contribute to this line of research by explicitly modeling the role of effort and embedding it in a dynamic setting, where it affects class accumulation and risk mitigation.

In an efficient and equitable world, college enrollment would depend on student ability rather than parental resources (Cameron and Heckman 1998 and 1999, Carneiro and Heckman 2002). In Colombia, as in other countries, parental resources matter greatly to college enrollment even controlling for ability. This provides strong evidence for credit constraints limiting college access, as discussed in a large literature. Lochner and Monge-Naranjo (2011) develop a model that helps explain the rising importance of family income for college attendance in the U.S. even in the presence of credit. Solis (2017) finds that relaxing credit constraints in Chile had an immediate impact on enrollment and number of college years completed, particularly for low-income students. Parental resources and background, however, may be of limited importance. Hai and Heckman (2017) show that equalizing initial ability has larger effects on college outcomes and inequality than equalizing parental background. The importance of credit access weakens when students can supply work as a source of funding college, as suggested by Garriga and Keightley (2007). Although Colombia is a large developing economy, the market for student loans is very limited, covering only 7 percent of students in 2003 (ICETEX 2010). Lack of family resources, limited opportunities to work during college, and missing credit markets for student loans are clear impediments to college access in countries such as Colombia.

In this context, tuition subsidies appear as a simple tool to broaden college access. Our free college counterfactuals complement the literature on the recent free college policies in Chile (Bucarey 2018) and the elimination of free college in England (Murphy et al 2019). It also joins in the vast literature of college financial aid,[5] including the recent literature on free community college and the so-called "Promise" programs implemented in multiple U.S. states.[6]

3 Data and descriptive statistics

In this section we describe the salient features of our data. These shape our model, and give rise to the moments we match in estimation.

3.1 The 2005 cohort

Our data consists of student- and program-level information drawn from three different administrative datasets: Saber 11, SPADIES, and SNIES. The first one, Saber 11, contains students'

[5]For recent reviews of this vast literature, see Avery et al (2019) and Dynarski and Scott-Clayton (2013). Section 7.5 contains further references.

[6]These programs provide zero tuition to eligible students for state or local community colleges or four-year institutions. See, for instance, Carruthers et al (2018), Dynarski et al (2018), and Gurantz (2020).

test scores at the national mandatory high school exit exam (also named Saber 11), along with socio-economic information reported by the students when taking the test. Saber 11 is a standardized test that covers multiple academic fields and measures students' academic readiness for higher education. We average field scores and standardize the average by semester-year. We use the resulting standardized score as a measure of student ability, broadly understood as her preparedness for higher education – reflecting not only her innate ability but also her primary and secondary education quality. Family income is reported in brackets defined relative to the monthly legal minimum wage (MW), which is equal to 381,000 Colombian pesos (COP) in 2005 (US$ 1 = 2,321 COP in 2005.)

The second dataset, SPADIES, tracks college students. For each semester, it records the number of classes for which a student registers and the number of classes she passes, as well as her graduation or dropout date. It does not record the specific classes in which a student enrolls, how many times a class is taken until passing, or class grades. The third dataset, SNIES, contains program-level information including institution, field, and tuition.

We focus on the 2005 cohort, which is the group of approximately 415,000 students ages 15-22 who took Saber 11 in 2005. Since students typically graduate from high school the same year they take Saber 11, we can view this cohort as the high school graduates from 2005. We calculate deciles and quintiles of their ability distribution; in what follows, ability deciles and quintiles always refer to this distribution. For consistency with the model, in the statistics below we classify students into "student types" defined by combinations of student ability quintiles and family income brackets. Table 1 below shows the distribution of student types in the 2005 cohort. It shows that, while a remarkable 70 percent of high school graduates come from the lowest two income brackets, less than 5 percent come from the top one. It also shows that high-income students are more likely to belong to high-ability levels than their lower-income counterparts due to the strong, positive correlation between income and ability.

Table 1: Family income and Ability Distribution of High School Graduates.

Income	Ability quintile					
Bracket	1	2	3	4	5	Total
5+ MW	0.21	0.31	0.48	0.90	3.15	5.05
3-5 MW	0.88	1.08	1.37	1.94	3.43	8.69
2-3 MW	2.72	2.94	3.30	3.69	3.95	16.60
1-2 MW	8.47	8.99	9.16	8.69	6.64	41.95
<1 MW	7.95	6.89	5.80	4.58	2.49	27.71
Total	20.23	20.21	20.11	19.80	19.65	100.00

Source: Calculations based on Saber 11. The distribution refers to 415,269 high school graduates from 2005. *Notes*: Family income is reported in brackets; MW = monthly minimum wage. Ability is reported in quintiles of standardized Saber 11 scores. Quintile 1 is the lowest.

3.2 Enrollment rates

Although Colombia's higher education offers short-cycle and bachelor's programs (akin to two- and four-year programs in the U.S. respectively), we focus on bachelor's programs, which capture approximately 80 percent of the country's total higher education enrollment. In what follows, "college" refers to bachelor's programs, and "college outcomes" to the final outcomes -graduation and dropout, along with their timing (e.g., on-time graduation). We classify a student from the 2005 cohort as having enrolled in college if she did so between 2006 and 2010.[7]

In Colombia, as in our sample, enrollment in bachelor's programs is almost evenly split between public and private institutions. Since public institutions are heavily subsidized, they charge much less than private institutions. For an individual with an annual family income of twelve MWs, annual average tuition for a bachelor's program at a public and private higher education institution is equal to 24 and 135 percent of the familiy income, respectively.

Table 2 shows enrollment rates by student type for the 2005 cohort. Although the overall enrollment rate is 32 percent, enrollment rates vary widely among student types, from 9 to 84 percent.[8] Enrollment rates rise both with income and ability. On average, the enrollment gap between the highest and lowest income brackets is equal to 55 percentage points (pp) - similar to the 50-pp point gap between the highest and lowest ability. These gaps suggest that free college may have ample room to raise enrollment.

Table 2: Enrollment Rates by Income and Ability.

Income	Ability quintile					
Bracket	1	2	3	4	5	Total
5+ MW	32.85	44.14	58.87	69.23	83.85	73.38
3-5 MW	28.71	39.75	48.41	62.99	79.24	62.03
2-3 MW	20.34	28.72	36.96	48.03	67.88	43.50
1-2 MW	13.94	18.36	23.85	33.84	54.22	28.05
<1 MW	9.05	12.67	17.20	26.56	43.93	17.67
Total	13.43	19.15	26.20	38.93	63.74	32.29

Source: Calculations based on SPADIES and Saber 11, for 2005 high school graduates.
Notes: Each cell reports percent of high school graduates from a given income bracket and ability quintile who enrolled in a bachelor's program between 2006 and 2010. Income reported in brackets; MW = monthly minimum wage. Ability reported in quintiles of standardized Saber 11 scores; quintile 1 is the lowest.

3.3 Graduation and dropout rates

For the analysis of college outcomes and performance that follows, we focus on students from the 2006 college entry cohort that enroll in five-year bachelor's programs.[9] Our sample includes 27,344 students, of whom only 45.7 percent graduates - 15.1 percent graduates on time (in five years) and 30.6 percent graduates late (in 6-8 years).[10] The dropout risk is thus substantial.

[7] Among the 2005 high school graduates that enroll in college within that five year window, only 35 percent do so immediately following high school. A five-year window, then, provides a more accurate enrollment rate.

[8] For comparison, in the US the enrollment rate of individuals ages 16-24 who graduated high school in 2005 is 44.6 percent (Source: Digest of Education Statistics). If this enrollment rate allowed for a five-year window as in Colombia, it would clearly be higher.

[9] To analyze dropouts and cumulative performance, it is customary to focus on a group of students from the same entry cohort who study programs of the same length. Five-year programs in Colombia capture about three-quarters of the enrollment in bachelor's programs. Dropout rates correspond to the student's first higher education program.

[10] For comparison, in the US 59.2 percent students from the 2006 cohort graduate within six years - 39 percent on time (in four years), and 20.2 percent late (in five or six years). Source: Digest of Education Statistics.

Dropout rates are far from uniform over time. As Figure 2 shows, over a quarter of college students (or half of all dropouts) leave in year 1. Further, the first two years account for about 70 percent of all dropouts.

Figure 2: Dropout Timing.

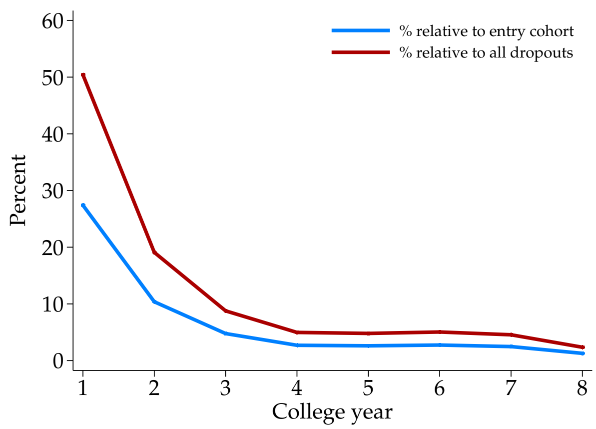

Source: Calculations based on SPADIES. Students belong to the 2006 entry cohort (first semester).
Notes: The blue line shows the percent of students who drop out in a given college year. The red line shows, among all dropouts, the percent of those who drop out in a given college year.

Figure 3: College Outcomes by Ability.

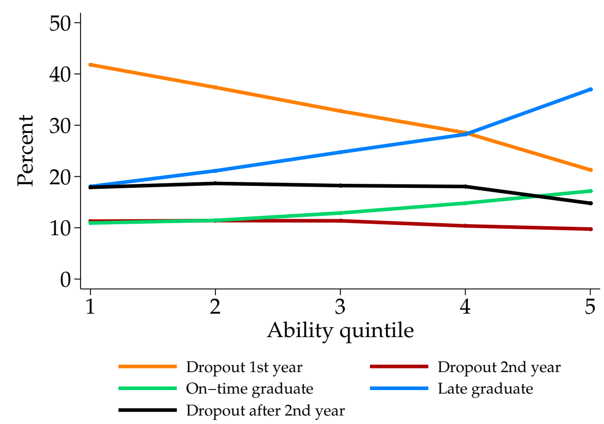

Source: Calculations based on SPADIES. Students belong to the 2006 entry cohort (first semester).
Notes: For each quintile, the graph shows the percent of students who attained each of the following outcomes: on-time graduate, late graduate, dropout in year 1, dropout in year 2, dropout after year 2. On-time and late graduates graduate in year 5, and years 6-8 respectively. For each ability quintile, percents add to 100.

Similar to enrollment rates, dropout rates vary widely across student types, ranging from 39 to 81 percent (Table 3). Conditional on income, higher ability students have lower dropout rates. On average, the dropout rate gap between the highest and lowest ability quintiles is equal to 25 pp. Dropout rates vary much less by income, as illustrated by the dropout rate gap of 14 pp between the highest and lowest income, and only 8 pp between the second highest and the lowest income. Free college, then, might affect enrollment more than dropout decisions.

Graduation rates are closely related to ability (Figure 3). Higher ability students are more likely to graduate, whether on time or late. Conversely, lower ability students are more likely to drop out, particularly in year 1. They are still more likely than other students to drop out afterwards, but the ability-based dropout rate gap becomes much smaller. While the fraction of graduates who finish on time is equal to one third, this fraction declines slightly with ability -from 38 percent for the lowest ability, down to 32 percent for the highest.

Taken together, these facts suggest that the dropout risk is asymmetric over time and across abilities, and is greatest for low-ability students in year 1. As a result, ability is a good predictor of dropout in year 1 but much less so in subsequent years - at which point, as explained below, the number of cumulative classes completed becomes a more powerful predictor.

Table 3: Dropout Rates by Income and Ability.

	Ability quintile					
	1	2	3	4	5	Total
5+ MW	81.36	65.83	61.48	52.13	39.04	44.73
3-5 MW	74.23	69.44	62.21	57.86	43.77	51.33
2-3 MW	68.54	67.58	63.73	57.68	46.53	55.09
1-2 MW	71.59	66.64	62.21	57.66	50.56	57.82
<1 MW	69.04	67.95	61.34	55.94	50.30	58.71
Total	70.99	67.44	62.37	56.96	45.84	54.36

Source: Calculations based on SPADIES. Students belong to the 2006 entry cohort (first semester).
Notes: Each cell reports percent of students from a given income bracket and ability quintile who drop out of their bachelor's program. A student is classified as a dropout if she does not graduate within eight years of having started her program. Income is reported in brackets; MW = monthly minimum wage. Ability is reported in quintiles of standardized Saber 11 scores. Quintile 1 is the lowest.

3.4 Cumulative classes completed

3.4.1 Cumulative classes completed and ability

Since every program requires a different number of classes for graduation, we normalize the requirement to 100 for all programs to facilitate exposition.[11] Lacking data on the number of classes that students must complete per year for on-time graduation, we assume the same number (20) for every year.[12] We use the term "classes completed" (or "performance") to denote the number of classes completed *in a given year*. We use "cumulative classes completed" (or "cumulative performance") for a given year as the total number of classes completed *over all years* up to (and including) that one. A student is on track for on-time graduation when she has completed her cumulative requirement up to that year - namely, when she has completed 20, 40, 60, 80, and 100 classes by the end of years 1 through 5 respectively.

It is useful to classify students into tiers based on cumulative performance relative to cumulative requirement. Tiers 1 through 4 correspond to students who complete the following *percent* of their cumulative requirement for the year: 95 percent or more for tier 1, (85, 95] percent for tier 2, (65, 85] percent for tier 3, and 65 percent or less for tier 4. To exemplify, consider a student who accumulates 16, 35, 42, 50, and 60 classes by the end of years 1 through 5 respectively, or 80 (=16/20*100), 88 (=35/40*100), 70, 62.5, and 60 percent of the cumulative requirements by year. This student falls in tiers 3, 2, 3, 4, and 4 in years 1 though 5 respectively. Appendix Table A.1 provides further details on tier classification. Importantly, a student can change tiers over time.

Figure 4's panel a shows the average number of classes completed by ability quintile in year 1. The thick black line depicts average over all students, whereas the color lines depict averages among the students who go on to attain the following outcomes: on-time graduation,

[11] For example, if a program's graduation requirement is 50 classes, completing 10 classes is equivalent to completing 20 percent of the program - or 20 classes, in our normalization. Since we do not observe each program's graduation requirement, we proxy for it by using the average number of classes completed by the program's graduates.

[12] In the data, we observe that students who graduate on time indeed complete classes at this pace. Our assumption of an equal number of classes per year is thus plausible.

late graduation, dropout in year 1, and dropout later.[13] While the figure focuses on year 1, a similar pattern holds for subsequent years. In other words, average classes completed varies little *across* abilities, both overall and conditional on outcomes. Nonetheless, it varies greatly *within* abilities: for a given ability, on-time graduates complete more classes than late graduates, who complete more than dropouts; among these, dropouts after year 1 complete more classes than year 1 dropouts. In other words, the number of classes completed -as early as in year 1- is a powerful predictor of college outcomes. This point is further illustrated in panel b, which classifies students into tiers by the end of year 1. Consistent with panel a, it shows the strong predictive power of early performance, suggesting that cumulative performance is highly persistent over time.

Figure 4: First-year Classes Completed and College Outcomes.

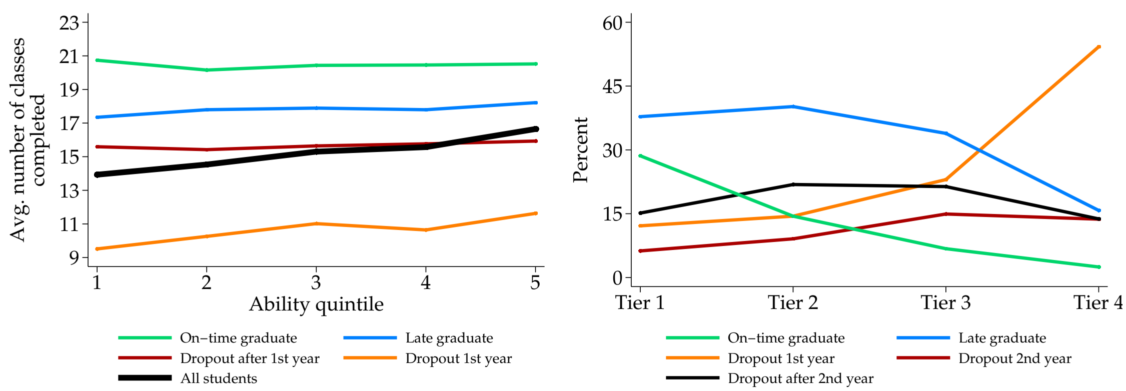

a. First-year classes completed by college outcome

b. College outcomes by first-year tier of classes completed

Source: Calculations based on SPADIES. Students belong to the 2006 entry cohort (first semester).
Notes: In panel a, each color represents a college outcome. The green line, for instance, shows the average number of classes completed by the end of the first year by students of each ability quintile who went on to graduate on time. The thick black line does the same for all students regardless of their college outcome. In panel b, students are classified by their tier of classes completed at the end of the first year. The graph shows the percent of students of a given tier who attain each college outcome. For the first year, tier 1 corresponds to 19+ classes completed; tier 2 to [17, 19); tier 3 to [13, 17); and tier 4 to [0, 13).

3.4.2 The persistence of cumulative performance

To explore this persistence, we consider whether students transition among performance tiers over time. Figure 5 depicts the probability of the following four outcomes for each year conditional on the previous year's tier: same-tier persistence, dropout, catch up, and fall behind. For example, a student who finished her first year in tier 2 has second-year probabilities of persistence, dropout, catch up, and fall behind equal to 29, 14, 20, and 39 respectively. If, for

[13] Some students drop out in the first semester of a given year. Since a period in our model is a school year (rather than a semester), for those students we impute a number of classes completed in their dropout year equal to twice the observed number for their last semester. This is a reasonable imputation, given that, in their dropout year, second-semester dropouts complete approximately twice as many classes, on average, as first-semester dropouts.

instance, she falls to tier 3 in her second year, then these probabilities for her third year are equal to 49, 12, 21, and 18 percent respectively.[14]

Panel a shows that same-tier persistence rises over time. This is, in part, associated to dropout rates that fall over time (panel b). Two issues related to dropout rates are worth noting. First, lower-performing students are more likely to drop out. Second, students from all tiers face a non-zero probability of dropping out. In other words, all students are subject to a dropout risk, although the risk is higher for low-performing students and, as we saw in Section 3.4.1, for low-ability students.

Figure 5: Tiers of Cumulative Classes Completed: Transitions Throughout College.

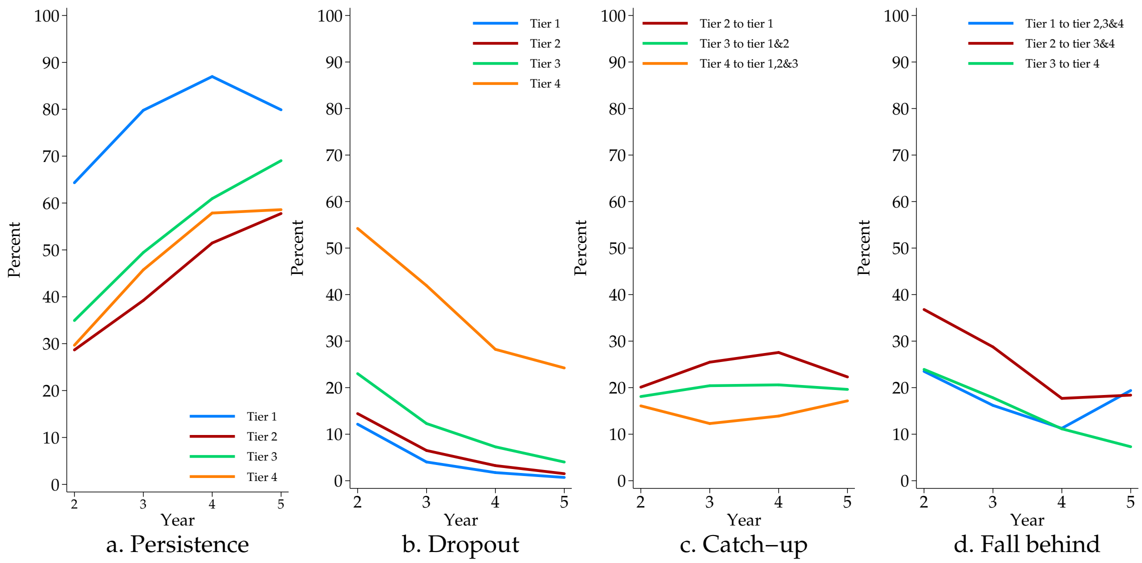

Source: Calculations based on SPADIES. Students belong to the 2006 entry cohort (first semester).
Notes: Each panel shows the probability that a student who ended the previous year in a given tier has one of the following outcomes in the current year: persist in that tier (panel a), drop out (panel b), catch up – namely, rise to a higher tier (panel c), or fall behind – namely, fall to a lower tier (panel d). For a given year and tier, probabilities add up to 100 across panels. For example, a student who finished year 1 in tier 3 is depicted in green. In year 2, she is 35, 23, 18, and 24 percent likely to persist in tier 3, drop out, catch up to tiers 1 or 2, and fall behind to tier 4 respectively.

Some students move across tiers (by catching up or falling behind), as shown in panels c and d. Higher-performing students are more likely to catch up and less likely to fall behind than others. At the same time, all students face a non-zero probability of falling behind. In other words, students are subject not only to a dropout but also a performance risk.

Taken together, these panels provide evidence of persistence in cumulative performance, explaining why first-year performance is a powerful predictor of college outcomes. A strong first year, however, is no guarantee of graduating or of not falling behind. Meanwhile, a poor beginning is hard to reverse and significantly raises the dropout risk. From the student's point of view, being on track with completed classes mitigates the risk of falling behind or dropping out, even if it does not fully eliminate it.

[14] These transitions hold for students as long as they are enrolled. Beginning in year 5, students can "transition" into graduation as well.

3.4.3 More on the role of ability, performance, and college outcomes

In Section 3.4.1, we established that, while there is little variation in academic progression across abilities, there is much more variation within abilities. We now explore this further by relying on our tiers.

Ability affects cumulative performance, particularly in year 1. Panel a of Figure 6 shows, for each ability quintile, the distribution of individuals across tiers at the end of year 1. Not surprisingly, high-ability students are most likely to belong to the top tier while low-ability students are most likely to belong to the bottom tier. Nonetheless, performance varies greatly within each ability quintile, and a sizable fraction of students from each quintile are concentrated in the middle tiers. Consistent with Section 3.4.1, this "thick middle" makes the cumulative performance vary little, on average, across abilities. A similar picture holds for year 5 (panel b), although by then the cumulative peformance distributions are more concentrated in tiers 1, 2, and 3 because most low-performing students have already dropped out. Note, also, that the distribution for higher ability students is more concentrated than that of less able students in both years, suggesting that lower-ability students face greater performance risk.

Figure 6: Tiers of Cumulative Classes Completed, by Ability.

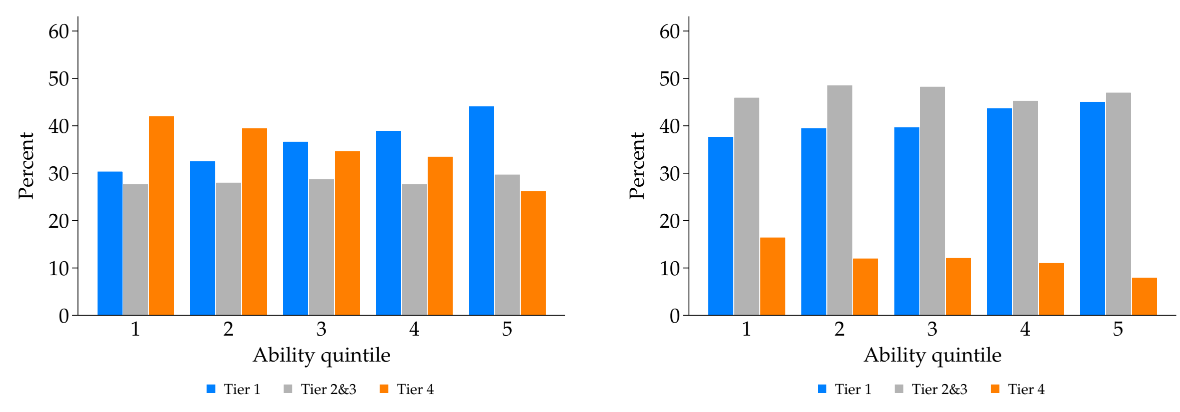

a. Year 1 b. Year 5

Source: Calculations based on SPADIES. Students are from the 2006 entry cohort (first semester).
Notes: Panel a (b) shows, for students of each ability who start year 1 (5), their classification into tiers of cumulative classes completed by the end of the year.

Not every student that reaches year 5 manages to graduate. Figure 7 shows that year 5's top tier students are more likely to graduate, and to do it on time. In contrast, bottom tier students are more likely to drop out than to graduate. By year 5, then, cumulative performance emerges as the main determinant of college outcomes.[15]

[15] The figure shows some patterns that may require additional explanation. Although students in tier 1 are closest to graduating, 40 percent of them graduate late - perhaps because they start working, or some graduation requirement (such as a thesis, or an English test) delays them. Some Tier 2 students graduate on time even though they have completed less than 95 percent of classes in year 5. Since we do not observe the actual number of classes required but rather the average number passed by graduates, these tier-2 students may have completed the required number of classes but not the graduates' average. A similar reasoning applies to on-time graduates from tiers 3 and 4. Of course, measurement error in number of classes passed or graduation date might also explain the on-time graduation rate in tiers 2, 3, and 4.

Figure 7: Graduation by Cumulative Classes Completed in Year 5.

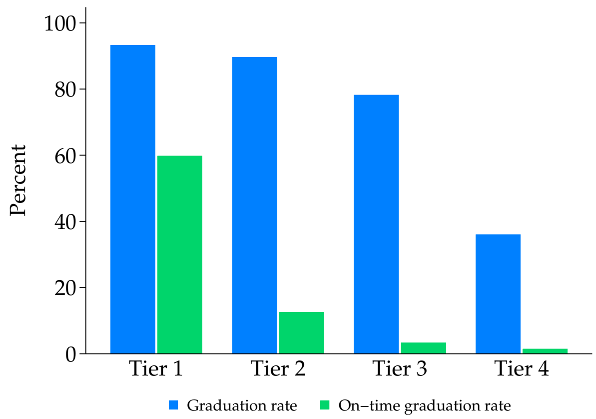

Source: SPADIES. Students belong to the 2006 entry cohort (first semester).
Notes: The figure refers to students who start year 5 and classifies them into tiers based on their cumulative classes completed by the end of the year. For each tier, the blue bar shows the percent of students who graduate, and the green bar shows the percent of students who graduate on time, relative to all graduates from that tier.

Figure 8: Average Number of Classes Taken in Year 1, by Ability.

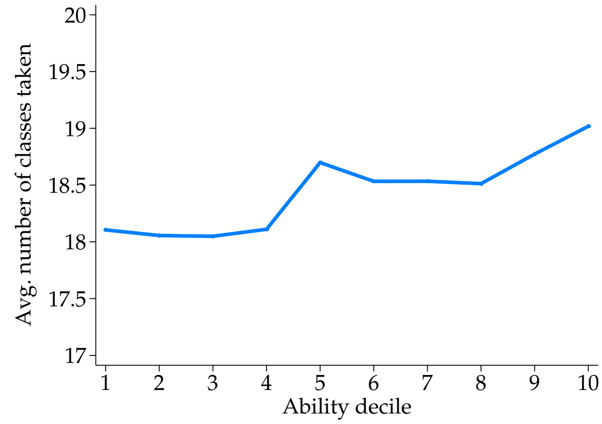

Source: Calculations based on SPADIES. Students belong to the 2006 entry cohort (first semester).
Notes: Figure shows the average number of classes for which students of a given ability register in year 1.

3.5 Number of classes taken

Although a student cannot fully control her performance, she can control the number of classes she takes in a given year. While the student may not pass all the classes for which she registers, this number is informative of her intended effort.

Figure 8 shows year 1's average number of classes taken, by ability decile, and depicts two important facts. First, on average students register for fewer than the 20 required classes. Second, on average higher ability students enroll in more classes. The first fact indicates that effort is costly - if it were not, students would register for 20 classes. The second fact indicates that effort is more costly for lower ability students - if it were not, they would enroll in the same number of classes as higher ability students. We will return to these facts when discussing the identification of effort in our empirical model.

Taking stock. The data shows that students of higher income or ability are more likely to enroll in college. Conditional on enrolling, lower ability students are much more likely to drop out, particularly in year 1. Through this channel, ability serves a strong predictor of graduation. Ability, however, is not a strong predictor of cumulative performance. Although cumulative performance varies little, on average, across abilities, it varies greatly within abilities. Further, cumulative performance is highly persistent over time - students who start on track are more likely to remain on track, less likely to fall behind, and more likely to catch up should they fall behind. As a result, cumulative performance is highly predictive -as early as in year 1- of final outcomes. Nonetheless, both performance and final outcomes are subject to shocks; even on-track students face the risk of falling behind or dropping out. The model presented below seeks to capture these data features.

4 Model

We model high school graduates who differ in ability, family income, and idiosyncratic preferences, and who choose whether to enroll in college or enter the labor force as high school graduates. College graduation requires the completion of a set number of classes. The combination of ability, effort, and performance and dropout risks determines class completion and college outcomes. Students can mitigate risk by being on track in class completion. After college, students enter the labor market, where wages depends on educational attainment (high school graduate, college graduate, or college dropout) and years of work experience, and remain there until retirement.

4.1 Student endowments and preferences

High school graduates differ in ability, $\theta \in \{\theta_1, \theta_2, \ldots, \theta_{n_\theta}\}$, and in the parental resources received if they enroll in college, $y \in \{y_1, y_2, \ldots, y_{n_y}\}$. The combination of ability and parental resources defines the student's type $j = 1, 2, \ldots, J$, with $J = n_\theta \times n_y$. For simplicity we refer to y as "income." Students receive y only if they enroll in college; if they do not, they receive their labor market wage as explained below.

A period, t, is equal to a year. For college students, instant utility depends on consumption, c, and study effort, e. Effort cost is heterogeneous across students as it depends on ability. The instant utility at t is $U(c_t, e_t, \theta) = u(c_t) - g(e_t, \theta)$. Function U satisfies $u' > 0$, $u'' < 0$, $g_1, g_{11} > 0$, $g_2 < 0$, and $g_{12} < 0$. Thus, effort cost is increasing in effort and convex, and total and marginal effort costs are lower for higher-ability students.

For workers, instant utility depends only consumption. Both for college students and workers, the discount factor is $\beta = 1/(1 + r)$, where $r > 0$ is the economy's risk-free interest rate.

4.2 College decisions

We model college as a multi-period, risky investment. To graduate, students must complete a required number of classes, h^{grad}. Students must spend at least five years in college in order to graduate, not to exceed eight. We use three concepts related to the number of classes per year. First is $\bar{x}$, the annual number of classes required to graduate on time. Second is q_t, the expectation formed by the student, at the beginning of the year, of how many classes she will complete (henceforth, "target.") The target can be greater or smaller than $\bar{x}$. Importantly, the student chooses effort for the year based on her target. Third is x_t, the actual number of classes she completes by the end of the year, which is a function of her effort, ability, and intervening luck. It can be greater or smaller than q_t because of "good" or "bad" luck, respectively.

Note that the model does not include the concept of "registering" for a class, and q_t should not be viewed as the number of classes for which the student registers. Instead, we assume that the student simply takes a class. This is because, in Colombia, there is no cost to add or withdraw from a class. Hence, students in Colombia register for the *maximum* number of classes they might complete, which is not necessarily the number they expect to complete.

For example, a student may start two classes but expect to complete only one ($q_t = 1$), perhaps because the other class is poorly taught. She may find out that that class is better taught than expected and complete the two classes ($x_t = 2$); or that both classes are poorly taught and complete none of them ($x_t = 0$); or that her expectation was correct and complete just one (x_t =1.) Since the student's behavior is dictated by the number of classes she expects to pass rather than the number she registers for, we model the former. In estimation we do use data on the number of classes for which students register, viewing it is an upper bound for q_t.

4.2.1 College technology

Let h_t denote the cumulative number of classes completed up to the end of t, or $h_t = \sum_{n=1}^{t} x_n$. Let $\bar{h}_t$ denote the average number of classes completed per year up to end of t, or $\bar{h}_t = h_t/t$. We assume that students start college with $h_0 = \bar{h}_0 = 0$, and by the end of year 1 attain $h_1 = \bar{h}_1 = x_1$.

While enrolled in college, students complete classes in year t according to the following production function:

$$x_t = H(\theta, e_t, z_t)\ \overline{x}. \tag{1}$$

The function describes completed classes, x_t, as a multiple of $\overline{x}$. The scalar $H(.)$ is a function of ability, effort, and a shock to classes completed (or performance shock), $z_t > 0$. We assume H is nonnegative and can be lower or greater than one. The shock, drawn from a continuous distribution known to the student, includes a random *i.i.d.* component, as well as a component that depends on the student's ability, cumulative classes completed up to beginning of t, and year. The dependence of the shock on past cumulative classes completed seeks to capture the observed persistence of classes completed, described in Section 3. Students can thus affect their "luck" next period by accumulating as many classes as possible this period. We allow the shock to depend on ability in order to capture the fact that ability may be correlated with other elements, not modeled, that systematically affect "luck."[16]

If the student knew z_t when choosing her effort, then choosing e_t would be equivalent to choosing classes completed, x_t. Since, as explained below, the student chooses e_t before z_t is realized, choosing effort is equivalent to choosing a target, q_t, where $q_t = E(x_t)$. Thus,

$$q_t = E[H(z_t, \theta, e_t)]\ \overline{x}. \tag{2}$$

Assuming that $H(\cdot)$ is linear on z_t so that it can be expressed as $H(z_t, \theta, e_t) = z_t\tilde{H}(\theta, e_t)$, target and effort are functions of $E(z_t)$:

$$q_t = E(z_t)\tilde{H}(\theta, e_t)\ \overline{x} \qquad \text{and} \qquad e_t = \tilde{H}_e^{-1}[\theta, q_t/(E(z_t)\ \overline{x})]. \tag{3}$$

Meanwhile, the actual number of classes completed, x_t, is a function of the effort chosen given the target, and of the realized z_t. Cumulative classes completed by the end of the year, h_t, is

$$h_t = h_{t-1} + x_t. \tag{4}$$

[16]For example, lower ability students may choose less selective programs than others, or may have lower levels of the non-cognitive skills necessary to succeed in college. These examples would lead to a negative and positive relationship between ability and the shock, respectively. In our estimation we let the data identify the sign of the relationship.

Finally, we assume that the production function in (1) is such that, when the student supplies zero effort, she completes zero classes: $H(z_t, \theta, 0) = 0$. For every student type, there always exists a level of effort, $\overline{e}_t$, that allows her to complete $\overline{x}$, or $H(z_t, \theta, \overline{e}_t) = 1$. Also, students of low ability can compensate for it, or for expected "bad luck", with high effort. For instance, consider students i and l , with $\theta_i < \theta_l$ and $E(z_i) < E(z_l)$. The low-ability student can compensate with higher effort, $e_i > e_l$, in order to have the same target as the other student, or $E[H(z_i, \theta_i, e_i)] = E[H(z_l, \theta_l, e_l)]$.

4.2.2 The student's optimization problem

The student faces a sequential problem. We differentiate between the pre-graduation years (when she cannot yet graduate) and the graduation years (when she is eligible to graduate depending on her cumulative number of classes completed). We divide each year into two subperiods. In the *first subperiod*, the student chooses her target number of classes and hence effort. At the end of it she receives the shock to the number of classes completed, which determines her actual (as opposed to target) number of classes completed and hence her cumulative classes completed. In the *second subperiod,* she graduates if she has accumulated the required number of classes; otherwise she draws a shock that determines whether she will remain in college next year or drop out ("dropout shock"). Thus, as long as she has not completed her graduation requirements, the student draws two shocks per year – one to classes completed in the year, and another to college continuity. The two shocks are endogenous in the sense that they depend on the student's cumulative performance, which she can affect through effort. In a given year, the state vector for a college student is (t, h_{t-1}, θ, y). Appendix Figure A.1 summarizes the timing of events and decisions, described in detail below.

Pre-Graduation Years (t = 1,...,4). During these years, students have not yet accumulated the required number of classes for graduation, or $h_t < h^{grad}$. In year 1, students start zero cumulative classes completed, $h_0 = 0$, and are heterogenous only in their type. Since students of a given type may vary in their first-year completed classes h_1 (depending, as we will see below, on their z_1 shock), from year 2 onwards students are heterogenous not only in their type but also in their cumulative classes completed at the beginning of the year, h_{t-1}.

At the beginning of the *first subperiod*, the student chooses e_t (and hence q_t) before z_t is realized. As in (2), the chosen target is a function of the $E(z_t)$. At the end of the first subperiod, z_t is realized and determines the number of cumulative classes completed, $h_t \geq h_{t-1}$.

In the *second subperiod*, the student receives her dropout shock, $d_t^{drop} = \{0, 1\}$, which determines whether she will remain in college next year or drop out, respectively. The probability that this shock leads her to drop out is a function of her cumulative classes completed after the realization of z_t, that is h_t, as well as her type and the year:

$$\Pr(d_t^{drop} = 1 \mid z_t) = \tilde{p}^d(t, h_t, \theta, y). \tag{5}$$

We assume that, by the end of t, if the student has accumulated less than a pre-specified number of classes for the year, h_t^{drop}, she must drop out: $\tilde{p}^d(t, h_t < h_t^{drop}, \theta, y) = 1$. If, in contrast, she completes $\bar{x}$ classes each year and is on track for on-time graduation, her dropout probability is very low: $\tilde{p}^d(t, h_t, \theta, y) \approx 0$. In general, $\tilde{p}^d$ is decreasing in h_t. Importantly, the student can lower $\tilde{p}^d$ by exerting effort, which raises h_t.

If the student drops out, she will receive the market wage of a college dropout from the following year onward; the value of dropping out is $V^{drop}(t{+}1)$.Meanwhile, the value of remainng in college is $V^{coll}(t+1, h_t, \theta, y)$.

Graduation Years ($t = 5,...,7$). These years are different from the previous ones in that college students become eligible to graduate depending on the number of cumulative classes completed. Those who have fulfilled graduation requirements, $h_t \geq h^{grad}$, will graduate and enter the labor market, whose value is $V^{grad}(t+1)$. For them, additional classes beyond h^{grad} yield zero marginal benefits. Remaining students draw the dropout shock to determine college continuity the following year.

Terminal year ($t = 8$). This is the last year that a student is allowed in college. At the end of it, only two outcomes are possible –the student graduates if $h_8 \geq h^{grad}$, or drops out otherwise- and continuation values are equal to $V^{grad}(9)$ and $V^{drop}(9)$ respectively.

We can now present the student's dynamic optimization problem from the first subperiod of each college year:

$$
\begin{aligned}
V^{coll}(t, h_{t-1}, \theta, y) = \max_{e_t} \Bigg\{ & U(c_t, e_t, \theta) + \beta E_z \Big[\mathbf{1}_{\{t \geq 5\}} \Pr\left(h_t \geq h^{grad}\right) V^{grad}(t+1) + \\
& \Pr\left(h_t < h^{grad}\right) \left[\tilde{p}^d(t, h_t, \theta, y) V^{drop}(t+1) + \right. \\
& \left. \left(1 - \tilde{p}^d(t, h_t, \theta, y)\right) V^{coll}(t+1, h_t, \theta, y) \right] \Big] \Bigg\},
\end{aligned} \tag{6}
$$

$$
\begin{aligned}
s.t. \quad & c_t = y - T(t, h_{t-1}, \theta, y) \\
& h_t = h_{t-1} + x_t \\
& x_t = H(z_t, \theta, e_t)\overline{x} \\
& c_t > 0.
\end{aligned}
$$

Here, the argument of $E_z[\cdot]$ is the continuation value function. Variable $T(\cdot)$ is tuition, constant regardless of the target, q_t.[17] To accomodate our counterfactuals, we write $T(\cdot)$ in general form so that it can vary by year, cumulative classes completed, ability, or income. In our baseline it varies only by y, as described in Section 5.2.1 below. For low-income students, tuition might exceed income, which would violate the $c_t > 0$ constraint and make enrollment unfeasible. Note the severe credit constraint: students cannot borrow to pay for tuition, nor can they save.[18] The policy function is the sequence of optimal efforts, $e^*(t, h_t, \theta, y)$, that solve the dynamic problem defined in (6).

[17] This is in keeping with the Colombian context, where students pay a fixed tuition regardless of the number of classes taken.

[18] We do not model student's decision to work while in college because our administrative data does not record this information. Further, data from Colombia's National Survey of Time Use (*ENUT*) reveals that high-income college students are more likely to work while in college than their lower-income counterparts, suggesting that the primary motivation to work is not necessarily to pay for college (details available upon request). For a model of student workers, see Garriga and Keightley (2007).

4.3 Workers

An individual can join the labor force after graduating from high school or college, or after dropping out from college.[19] The worker's optimization problem, written in recursive form, is

$$V^m(t) = \max_{c_t}\{u(c_t) + \beta V^m(t+1)\}, \tag{7}$$

$$s.t. \qquad c_t = w_t^m,$$

where $V^m(t)$ is the value function of a worker with educational attainment $m = \{hs, grad, drop\}$, denoting high school graduate, college graduate, and college dropout respectively. The worker's wage, w_t^m, is specific to educational attainment, and varies with t to allow for returns to experience. Note that V^m depends on t because of w, and because the value of working depends on the total number of years worked, given by the entry date into labor force.

4.4 Enrollment decision

In order to decide whether or not to enroll in college, a high school graduate compares the expected payoff of two choices - going to college, or joining the labor force as a high school graduate. The enrollment decision is a discrete choice problem, where the payoff associated to each option is the sum of three components. The first component is the expected value of going to college, $V^{coll}(t = 1, h_0 = 0, \theta, y)$ or of entering the labor force as a high school graduate, V^{hs}. The second component is a type-specific preference for college enrollment, $\xi_j = \tilde{\xi}(\theta_j, y_j)$, which captures type-related unobserved factors, such as parental education, that affect enrollment. We normalize the unobserved preference for joining the labor force as a high school graduate to zero for all types. The third component is an idiosyncratic choice-specific shock for each individual, ϵ^{hs} and ϵ^{coll}, corresponding to working as a high school graduate or enrolling in college, respectively. Thus, all individuals face the same V^{hs}, and individuals of a given type face the same V^{coll} and ξ_j; yet individuals within and across types differ in their idiosyncratic shocks. We assume that ϵ^{hs} and ϵ^{coll} are *iid* and distributed Type I Extreme Value with a scaling factor of σ_ϵ. The individual chooses to attend college if

$$\underbrace{V^{coll}(1, 0, \theta_j, y_j) + \xi_j + \sigma_\epsilon \epsilon^{coll}}_{\text{Value of going to college}} \geq \underbrace{V^{hs} + \sigma_\epsilon \epsilon^{hs}}_{\text{Value of working as a high school graduate}} \tag{8}$$

As a result, the probability of college enrollment for an individual of type j is

$$P^{coll}(\theta_j, y_j) = \frac{\exp\{(V^{coll}(1, 0, \theta_j, y_j) + \xi_j)/\sigma_\epsilon\}}{\exp\{(V^{coll}(1, 0, \theta_j, y_j) + \xi_j)/\sigma_\epsilon\} + \exp\{V^{hs}/\sigma_\epsilon\}}, \tag{9}$$

Its complement, $P^{hs}(\theta, y) = 1 - P^{coll}(\theta, y)$, is the probability of joining the labor force as a high school graduate.

[19] We assume that workers consume all their earnings and do not have access to credit markets, which is an accurate representation of developing economies. Since wages rise with experience and workers discount the future at the interest rate, they have no incentives to save.

5 Empirical implementation

In this section we describe the parameterization and computational version of the model. We also describe the algorithm to compute model predicted values for a given parameter point.

5.1 Functional forms

In the model, $t = 1$ corresponds to age 18. Retirement age is 65, or $t = 48$. Regardless of her educational attainment or when she joined the labor force, the individual accrues returns to experience (or becomes "experienced") from age 35 ($t = 28$) onwards.

The utility of college students is given by

$$U(c, e, \theta) = \frac{(c + \underline{c})^{1-\rho} - 1}{1 - \rho} - \mu \frac{e^{\gamma}}{(1 + \theta)^{k}}. \tag{10}$$

where the need to meet the minimum consumption level, $\underline{c}$, might limit low-income students' ability to enroll in college. To prevent this, we set $\underline{c}$ equal to one million COP.[20] The utility of workers is given by

$$u(c) = \frac{c^{1-\rho} - 1}{1 - \rho}. \tag{11}$$

We set $r = 0.04$, and assume $\sigma_{\epsilon} = 1$.

The production function to complete classes has constant returns to scale in ability and effort:

$$x_t = H(z_t, \theta, e_t)\overline{x} = z_t(\theta^{\alpha} e_t^{1-\alpha})\overline{x}, \tag{12}$$

where $\alpha \in (0, 1)$ is the elasticity of classes completed with respect to ability. Consistent with the model, we set $\bar{x} = 20$ classes. We set the minimum number of classes required to graduate, h^{grad}, equal to 98.[21]

The functional form for the z_t shock is as follows:

$$z_t = \exp\{-\exp\{-(\kappa_0 + \kappa_1 d_1 + \kappa_h \tilde{h}_{t-1} + \kappa_{\theta}\theta + (\sigma + \sigma_1 d_1 + \sigma_{\theta}\theta)\nu_t)\}\}, \tag{13}$$

where $\tilde{h}_{t-1}$ is a measure of past cumulative number of classes completed, with $\tilde{h}_{t-1} = \ln(\bar{h}_{t-1})$ for every $t > 1$, and $\tilde{h}_0 = 0$ for $t = 1$.The terms associated with d_1 allow the shock distribution to differ in year 1, when $d_1 = 1$.[22] The shock also depends on an *iid* component, ν_t, drawn from the uniform distribution $U(0, 1)$. The functional form in (13) ensures that $z_t \in (0, 1)$ for any combination of parameter values and for all $\tilde{h}, \theta \in \mathbb{R}$. Importantly, all the parameters in (13) affect the mean and variance of z_t. In Section 6.3 below we discuss the effect of $\tilde{h}_{t-1}$ and

[20] Our chosen value for $\underline{c}$ guarantees that, in our computational models, all students attain positive consumption if they enroll in college. We can think of $\underline{c}$ as the minimum consumption guaranteed to college students through student subsidies, such as those for food and transportation.

[21] We set this requirement to 98 rather than 100 because we observe students who graduate with slightly fewer than 100 classes - perhaps due to measurement error.

[22] In year 1, $\tilde{h}_0 = 0$, whereas $\tilde{h}$ is positive in subsequent years. This creates scaling problems in year 1, which we solve through κ_1. As documented in section 3.4.3, the variance of classes completed is higher in year 1 than in other years, which we capture with σ_1.

κ_θ on this mean and variance at our specific parameter estimates.

We parameterize the probability of dropping out as

$$\tilde{p}^d(t, h_t, \theta, y) = \frac{\exp\{\delta(t,\theta,y) + \pi \tilde{h}_t\}}{1 + \exp\{\delta(t,\theta,y) + \pi \tilde{h}_t\}}, \tag{14}$$

where $\delta(t, \theta, y)$ is a year-, ability- and income- specific fixed effect, and $\tilde{h}_t$ measures cumulative performance over all periods, including the current one. Evaluating $\tilde{p}^d(t, h_t, \theta, y)$ at $\pi = 0$ yields the "exogenous dropout probability" - namely, the dropout probability that students of a given type would have, in a given year, if they had accumulated no classes. It is "exogenous" because it is independent of effort. For example, low-income, low-ability students may have a high exogenous dropout probability in year 1 -perhaps because they lack parental guidance on how to navigate college- yet a lower one in subsequent years.

The model's full parameter vector is $\tilde{\Theta} = (\Theta, \boldsymbol{\xi}, \boldsymbol{\delta})$, where

$$\Theta = (\rho, \mu, \gamma, k, \alpha, \kappa_0, \kappa_{y_1}, \kappa_h, \kappa_\theta, \sigma, \sigma_{y_1}, \sigma_\theta, \pi) \tag{15}$$

is the vector of parameters common across individuals. Vector $\boldsymbol{\xi}_{J\times 1}$ contains type-specific unobserved preferences for college, ξ_j (see (9)) and $\boldsymbol{\delta}_{(J*8)\times 1}$ contains exogenous dropout probability fixed effects, $\delta(t, \theta_j, y_j)$, for the J types and 8 years (see (14)).

5.2 Computational representation

5.2.1 Student types

To build the empirical distribution of ability and income for school graduates, $\Phi(y, \theta)$, we start from Table 1, which classifies 2005 high school graduates by ability quintile and income bracket. We refine this table to work with ability deciles rather than quintiles, for a total of fifty student types. To construct values for θ, we start from the distribution of standardized Saber 11 test scores and normalize them between 0 and 1.[23] Our θ values are the 5th, 15th, ...95th percentiles from the normalized scores. We calculate the y value corresponding to each income bracket as the average annual per-capita income for that bracket, computed from Colombia's household survey data (SEDLAC) on family income and size. Lacking student-level data on tuition expenses, we estimate the tuition paid by students of a given y as the average annual tuition paid by students from the corresponding income bracket at public institutions, calculated from SNIES and SPADIES.[24]

Table 4 shows the resulting income and tuition corresponding to the underlying family income brackets. As the table shows, income varies greatly across income brackets. Although public institutions provide income-based tuition discounts, the highest income individuals do not pay proportionally to their income. While their per-capita income is about twenty times as large as that of the lowest-income individuals, their tuition is only 2.5 times as large.

[23] Let sts denote the standardized test score. The normalized sts is equal to $(sts - \min(sts))/(\max(sts) - \min(sts))$.

[24] We use tuition at public HEIs because there is always a public HEI that the student can attend. Modeling the choice of college type (public or private) is beyond the scope of this paper.

Table 4: Income and Tuition.

Income Bracket	Avg. Per-Capita Household Income	Avg. Per-Capita Tuition
5+ MW	\$ 21,027,690	\$ 2,195,972
3-5 MW	\$ 9,191,642	\$ 1,826,386
2-3 MW	\$ 5,337,010	\$ 1,177,543
1-2 MW	\$ 2,952,288	\$ 978,690
<1 MW	\$ 1,119,633	\$ 855,493

Source: Calculations based on Saber 11 and SEDLAC (household surveys) for per-capita income; Ministry of Education of Colombia and SPADIES for tuition.
Notes: Since Saber 11 provides income brackets rather than actual income, we use SEDLAC (household surveys) data on household income and household size to calculate the average per-capita household income corresponding to households of a given bracket. To calculate the average tuition for a given bracket, we assign to each student the average tuition paid by students in her program, and average over students. Income is reported in brackets; MW = monthly minimum wage.

5.2.2 Workers

We use household surveys to compute the average wages earned by individuals with different educational attainment and experience in 2005. For workers aged 18-65, the average wage of a college graduate, a college dropout with at least one year of complete college, and a college dropout with less than one year of complete college is 160, 58, and 28 percent higher than the average wage of a high school graduate respectively.[25] Among college (high school) graduates, the average wage of experienced workers is 35 (29) percent higher than the average wage of inexperienced workers. Consistent with the data, we assume that the returns to experience of college dropouts are the same as those of high school graduates.

5.3 Computing predicted values

Since the model does not have a closed-form solution, we use a numerical algorithm to solve students' dynamic optimization problem for a given value of Θ. Appendix A.4 provides a full description of the algorithm. The estimation of $\boldsymbol{\delta}$ and $\boldsymbol{\xi}$ is nested within the model solution for a given of value of Θ, in the spirit of Berry, Levinsohn and Pakes (1995).

In anticipation of next section, a couple of remarks are in order. First, our model solution by construction replicates observed enrollment rates by type. Hence, we do not match enrollment rates in the estimation. Second, although our model solution attempts to replicate observed dropout rates at the (year, ability quintile, income) level, it does so with mixed success (see Appendix A.4.3 for further details). As a result, we are able to match dropout rates in the estimation.

[25]This creates, in effect, four college attainments - high school, college, some college (one year), some college (two or more years). The two "some college" categories correspond to college dropouts. We work with two rather than one dropout category because their wages are quite different from one another and hence provide different incentives to college students over time.

6 Estimation

In this section we describe the estimation strategy and identification. We also present parameter estimates, describe the model's fit, and address the role of effort given our estimates.

6.1 Estimation Strategy

We estimate the model parameters using Simulated Method of Moments (SMM). Our estimation searches for the value of Θ whose predicted moments, $\hat{\mathbf{M}}(\Theta)$, best match the observed ones, $\mathbf{M}$. The moments we match are listed in Table 5. They reflect the patterns of dropout, college outcomes, classes completed, and targets discussed in Section 3. Matching these 585 moments enables us to estimate our 13 parameters.

Table 5: Moments Matched in Estimation.

Data aspect	Moments	Number of Moments
Dropout rates	Dropout rate by year.	8
	Dropout rate by ability quintile and income.	25
	Dropout rate by ability decile.	10
College outcomes	College outcomes by ability quintile.	25
	Fraction of students that graduate by year (years 5-8).	4
Cumulative classes completed	Average number of cumulative classes completed by year, ability quintile, and college outcome.	140
	Average number of cumulative classes completed by year and ability decile (years 1-5).	50
	Distribution of students into tiers of cumulative classes completed, by year.	24
	Distribution of students into tiers of cumulative classes completed, by ability quintile and year (years 1-5).	75
Transition probabilities	Pr(tier Y in $t+1$|tier X in t) for years 1-7.	112
	$\Pr(d_t^{drop}=1 \vert \text{tier } X \text{ in } t)$ for years 1-8.	32
Target number of classes	Average target number of classes by ability decile and year.	80
Total		585

Source: Own estimation.
Notes: Moments per year are computed for years 1-8 unless otherwise specified. Tiers are 1-4, based on cumulative classes completed (see Subsection 3.4.1 and Appendix Table A.1 for further details). For "College outcomes," outcomes include on-time graduate, late graduate, drop out first year, drop out second year, drop out after second year. In "Cumulative classes completed," which are calculated by year, outcomes include on-time graduate (until year 5), late graduate, drop out this year, drop out later (until year 7); "this year" and "later" refer to the year under consideration. In "Transition probabilities", t refers to year; tiers X and Y are 1,...,4. Observed data for target number of classes is average number of classes for which the corresponding students registers.

Formally, our SMM parameter estimates solve the following problem:

$$\arg\min_{\Theta} \quad (\hat{\mathbf{M}}(\Theta)-\mathbf{M})'W^{-1}(\hat{\mathbf{M}}(\Theta)-\mathbf{M}), \tag{16}$$

where Θ is a 13×1 vector of parameters, $\mathbf{M}$ and $\hat{\mathbf{M}}$ are 585×1 vectors of sample and predicted moments, respectively, and W is a diagonal weighting matrix whose diagonal contains the standard error of the sample moments. We compute numerically the predicted values, $\hat{\mathbf{M}}$, for every value of Θ as explained in Appendix A.4.

6.2 Identification

A critical challenge is identifying the role of ability, effort, and performance shocks in the production of classes completed. Below we provide intuition for identification.

Effort. If effort had no role in the number of classes completed ($\alpha = 1$), or if it were costless ($\mu = 0$), then all students would take the required number of classes per year. The fact that most students take, on average, a lower number classes (see Section 3.5) indicates that effort does have a role in classes completed and helps identify μ. An increase in μ leads to lower targets, effort, and number of classes completed. An increase in μ also leads to lower college enrollment - particularly for low-income students, who have lower consumption than their wealthier counterparts to compensate for effort. The speed of accumulation of classes completed, as well as the transitions among tiers over time, helps identify γ. A high γ penalizes high effort levels and makes it costly to catch up. The fact that higher ability students take more classes, on average, than their lower-ability counterparts indicates that their effort cost is lower and identifies k. An increase in k raises the variance of average target, effort, and classes completed across abilities.

Performance shock. Despite the low variation of average classes completed across abilities, number of classes completed varies widely within abilities. Conditional on ability, effort varies by income. This explains some, but not all, of the within-ability variation of classes completed. The remainder of this variation, then, is explained by the performance shock, z. Since θ is between 0 and 1, we restrict z to be in this range as well. This helps us pin down the scale for effort. An increase in κ_0 makes shocks more favorable to all students, thus raising the number of classes completed and lowering dropout rates across the board. Parameter κ_1 is an intercept shifter that makes the scale of z comparable across years. An increase in κ_θ makes the expected shock relatively more favorable for low-ability students, and raises the dispersion in average classes completed and college outcomes across abilities. Parameter σ is identified by the overall variation of classes completed conditional on ability, whereas σ_θ is identified by the greater variation of classes completed among low- than high-ability students. The higher overall variation of classes completed in year 1 relative to other years identifies σ_1. After year 1, κ_h is identified by the persistence of students in their performance tiers.

Ability. Given the role of effort and performance shocks in the production of classes completed, the variation of average classes completed across abilities identifies α. This variation rises with an increase in α.

Other parameters. An increase in ρ raises the aversion to consumption variations over time and decreases the propensity to college enrollment. It also lowers the speed of class accumulation and increases time-to-degree. Finally, the sensitivity of dropout rates with respect to current classes completed, conditional on student income and ability, identifies π.

A sufficient condition for local identification is that the matrix of first derivatives of the moments' predicted values with respect to the parameter vector has full column rank when evaluated at the true parameter point. Evaluated at our parameter estimates, this matrix has full column rank in our sample.

6.3 Parameter Estimates

We now turn to our parameter estimates, shown in Table 6.

Table 6: Parameter Estimates.

Parameter	Symbol	Estimate
Utility function		
Consumption curvature	ρ	0.882
Effort weight	μ	0.062
Effort curvature	γ	4.727
Effort cost w.r.t. ability	k	1.225
Number of classes completed		
Elasticity w.r.t. ability	α	0.085
Performance shock		
Constant	κ_0	-4.207
Year 1 shifter	κ_1	3.534
Persistence component	κ_h	1.304
Ability component	κ_θ	0.407
Std. dev. of *iid* shock	σ	1.789
Std. dev. of *iid* shock - Year 1 shifter	σ_1	0.317
Std. dev. of *iid* shock - Ability shifter	σ_θ	-1.282
Dropout shock		
Cumulative performance component	π	-2.951

Source: Own estimation.

The estimated α is low, consistent with low variation of average classes completed across abilities. As a result, the estimated elasticity of credits completed with respect to effort (equal to $1-\alpha$) is high. We emphasize that our ability measure captures college academic readiness as measured by Saber 11 and not necesssarily "true" ability. Hence, our estimate indicates that college academic readiness plays a small role in the accumulation of classes completed relative to effort. Therefore, policies must affect effort in order to affect classes completed.

The estimated $\rho < 1$ indicates that students have low risk aversion to consumption changes over time and implies an elasticity of intertemporal substitution (equal to $1/\rho$) greater than one. This makes students quite willing to attend college and graduate on time.

The estimated γ, equal to 4.73, indicates a very high marginal cost of effort, exceeding typical quadratic costs, which prevents large catch-up efforts. The estimated k indicates that effort is negatively and strongly related to ability. When $k = 0$, effort cost is the same for all abilities, whereas when $k = 1$ effort cost falls with ability, at a decreasing rate. Our estimated value of 1.23 yields the same qualitative pattern as $k = 1$, but an even greater gap in effort cost among high and low ability students.

Given estimates for the shock-related parameters, we find that $E(z_t)$ is higher for students with higher past performance, which creates persistence. In addition, low-ability students have higher $E(z_t)$ and $Var(z_t)$ than high-ability students. Their (slightly) higher mean of z is consistent with the observed fact that they enroll in less selective (and presumably less demanding) programs than their abler counterparts.[26] Without this higher mean, it would be difficult to match the good performance and graduation of some low-ability students (see Sections 3.3 and 3.4). The more dispersed shock for lower-ability students, in turn, helps us match their higher variance of classes completed (see Section 3.4.3.)

To examine the relative impact of past performance and ability on z, consider students A and B. At the beginning of t, A has completed one more class than B and is more able, with $\Delta_\theta = \theta_A - \theta_B = 0.22$. This is a large ability difference, equal to the difference between the 55th and the 5th percentile, or between the 95th and the 75th percentile. Because A is abler than B, her $E(z_t)$ should be lower than B's, yet because she has completed more classes, her $E(z_t)$ should be higher. As it turns out, just having completed that one additional class gives her the same $E(z_t)$ as B's, even though B is much less able. In other words, h_{t-1} has a relatively larger impact than θ on $E(z_t)$.This makes z highly persistent and more dependent on something the student can control -her performance- than on ability, which she cannot control.

Finally, the estimated π indicates that an additional class completed by the end of the year, on average, decreases the probability of dropping out by about 5 pp. Since $\tilde{p}^d(\cdot)$ has a logistic functional form, this marginal effect -and therefore the incentive to accumulate classes completed- is stronger for students with intermediate values of the dropout probability rather than values close to zero or one.

Our full set of parameter estimates includes the dropout probability fixed effects in (14), $\hat{\delta}(t, \theta, y)$. To illustrate the relative magnitude of $\hat{\pi}$ and $\hat{\delta}(t, \theta, y)$, consider the average number of additional classes that a student from the second ability quintile ("Q2 student") must complete to attain the same dropout probability as a student from the top ability quintile ("Q5 student"). In year 1, she must complete more than 4 additional classes, or 25 percent of the annual requirements, reflecting a high exogenous dropout probability. In year 5 she only needs one additional class completed, as the Q2 students reaching year 5 are approximately on par with Q5 students. The important point is that, early on, low-ability students face a high exogenous dropout probability, which can only be reversed through very high initial effort or very favorable performance shocks.

6.4 Goodness of fit

The estimated model fits the data well, and replicates the patterns described in Section . The observed graduation rate is 45.64 percent, and the predicted one is 45.02 percent. The model captures the distribution of dropouts over time (Table 7) and across student types (Table 8). As Figure 9 shows, it also captures the distribution of college outcomes by ability. The on-time graduation rate (15.1 percent) is predicted perfectly. The observed fraction of 2005 high school graduates that complete college is 14.7 percent, and the predicted one is 14.5 percent.

As Appendix Tables A.3-A.10 show, we capture average classes completed by year and ability, and conditional on final outcome.[27] Therefore, the model replicates the wide variation

[26] We measure program selectivity as the average Saber 11 test score of the program's students. We find a negative correlation between student ability and program selectivity. In other words, less able students choose less selective programs.

[27] When assessing goodness of fit, it should be kept in mind that the use of a weighting matrix in (16) implies that

Figure 9: Goodness of Fit: College Outcomes.

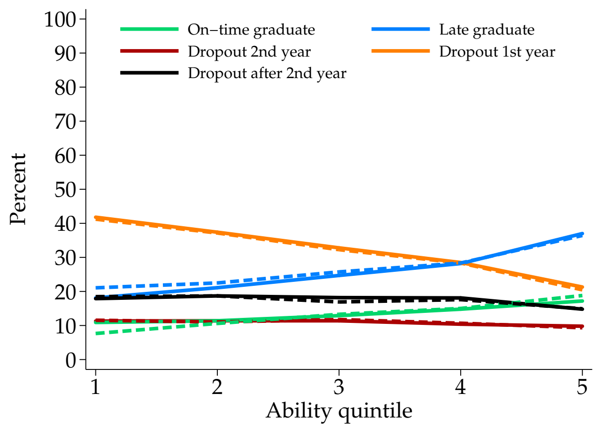

Source: SPADIES for observed data; fitted values for predicted data.
Notes: Whole lines show observed values; dashed lines show fitted values.

Table 7: Goodness of Fit: Dropout Year.

Year	Observed	Predicted
1st	27.41	27.92
2nd	10.38	10.43
3rd	4.77	4.94
4th	2.7	2.83
5th	2.61	2.81
6th	2.74	2.92
7th	2.48	1.94
8th	1.28	1.19
Total	54.36	54.98

Source: SPADIES for observed data; fitted values for predicted data.
Notes: Table shows the observed and predicted percent of students who drop out each year.

Table 8: Goodness of Fit: Dropout Rates by Income and Ability.

	Ability quintiles											
Income	Observed values						Predicted values					
Bracket	1	2	3	4	5	Total	1	2	3	4	5	Total
5+ MW	81.4	65.8	61.5	52.1	39.1	44.7	84.4	67.8	63.4	51.9	40.7	46.7
3-5 MW	74.2	69.4	62.2	57.9	43.8	51.3	81.6	68.3	62.2	57.3	44.4	52.5
2-3 MW	68.5	67.6	63.7	57.7	46.5	55.1	70.8	67.5	61.9	60.1	44.4	55.6
1-2 MW	71.6	66.6	62.2	57.7	50.6	57.8	69.9	66.5	60.9	56.2	46.9	56.7
<1 MW	69.0	67.9	61.3	55.9	50.3	58.7	69.6	66.3	58.6	53.7	46.2	57.5
Total	71.0	67.4	62.4	60.0	45.8	54.4	71.3	66.9	61.0	56.6	44.8	55.0

Source: Source: SPADIES for observed data; fitted values for predicted data.
Notes: Values are expressed in percentages (%). Income is reported in brackets; MW = monthly minimum wage. Ability is reported in quintiles of standardized Saber 11 scores. Quintile 1 is the lowest.

in classes completed within abilities, as well as the low variation across abilities (Appendix Figure A.2). Further, the model captures the qualitative patterns of persistence, drop out, catch-up and fall behind (Table 9), as well as the progressive concentration of students in the upper tiers over time (Figure 10). Finally, Appendix Figure A.3 shows average predicted target and average number of classes taken by the student. Recall that the latter is theoretically an upper bound for the former (see Section 4.2.) As such, our average predicted target is lower than the average observed number of classes.

moments with a greater number of underlying observations (such as those from early years, or for higher-ability students) will attain a better fit.

Table 9: Goodness of Fit: Transitions Among Tiers of Cumulative Classes Completed.

	Observed values				Predicted values			
	Year 2	Year 3	Year 4	Year 5	Year 2	Year 3	Year 4	Year 5
Persistence								
Tier 1	64.33	79.78	86.98	79.89	63.70	75.80	75.10	74.30
Tier 2	28.67	39.23	51.47	57.76	22.80	38.40	53.70	60.10
Tier 3	34.94	49.44	60.94	69.04	35.10	59.20	71.20	78.80
Tier 4	29.67	45.75	57.87	58.59	37.10	42.50	60.70	76.00
Dropout rate								
Tier 1	12.16	4.03	1.75	0.71	11.60	4.20	3.00	2.20
Tier 2	14.43	6.50	3.25	1.52	16.20	6.40	4.60	3.00
Tier 3	23.02	12.30	7.29	4.02	25.10	9.20	7.00	5.20
Tier 4	54.23	41.95	28.24	24.24	55.20	50.70	35.20	21.40
Prob. of Catch up								
Tier 3 to Tiers 1 & 2	18.10	20.42	20.61	19.63	26.60	20.30	13.10	9.40
Tier 4 to Tiers 1 & 2	4.26	0.55	0.22	0.14	0.82	0.00	0.00	0.00
Prob. of Fall behind								
Tier 1 to Tiers 3 & 4	11.01	5.41	2.54	1.66	4.30	1.20	0.30	0.00
Tier 2 to Tiers 3 & 4	36.8	28.77	17.70	18.40	26.00	25.60	22.80	24.40

Source: SPADIES for observed data; model simulations for predicted data.
Notes: Values are expressed in percentages (%).

Figure 10: Goodness of Fit: Tiers of Cumulative Classes Completed, by Year.

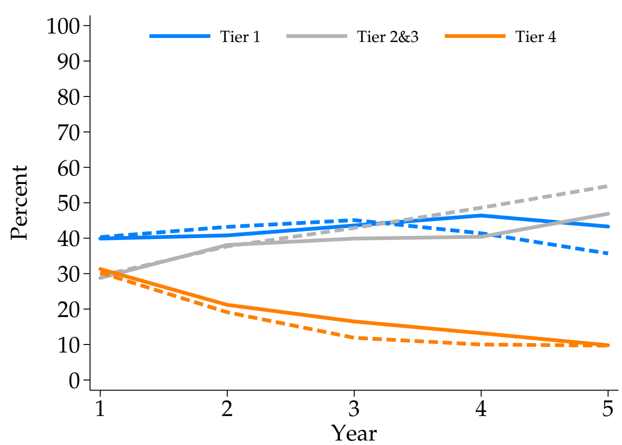

Source: SPADIES for observed data; fitted values for predicted data.
Notes: For each year, the figure shows the percent of students in each tier of cumulative classes completed. Whole and dashed lines show observed and predicted classification, respectively.

6.5 The Role of Effort

Our estimates show that class completion is much more responsive to effort than ability. Given the importance of effort, we begin by looking at average student effort by year (Figure 11). On average, effort is relatively high in years 1 and 4. In year 1, students work hard to complete classes and mitigate dropout risk. In year 4, they work hard in order to graduate. After year

5, average effort falls as the remaining students have few classes left.

Figure 11: Predicted Average Effort by Year.

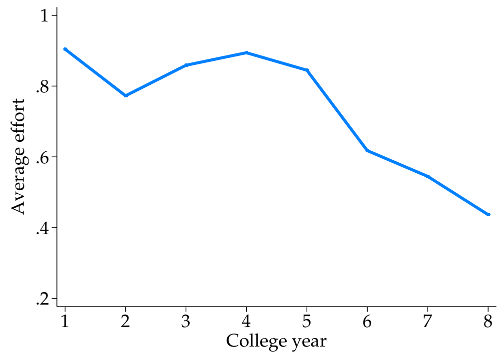

Source: Model baseline predictions.
Notes: The figure shows average student effort for each year in the simulated baseline.For a given year, individual efforts are normalized by the 95th percentile of the year's effort distribution.

In general, administrative datasets do not provide effort measures. When these datasets are used to estimate the relationship between classes completed and ability, the lack of effort measures might bias the estimates. We investigate this issue in Table 10. We begin with the regression of log classes completed, $\ln(x_t)$, on log ability for the observed data (column 1) and the simulated baseline (column 2). The coefficients on log ability in columns (1) and (2) are very close, which is additional evidence of our good fit. In both cases, log ability explains about 20 percent of the variation in log classes completed.

Table 10: Classes Completed Per Year.

	Actual data	Simulated Data		
	(1)	(2)	(3)	(4)
ln(ability)	0.166***	0.156***	0.090***	0.085***
	(0.015)	(0.005)	(0.005)	(0.000)
ln(effort)			0.854***	0.915***
			(0.004)	(0.000)
ln(shock to classes completed)				1.000***
				(0.000)
Constant	2.060***	2.748***	2.197***	2.996***
	(0.012)	(0.004)	(0.004)	(0.000)
R^2	0.213	0.204	0.518	1.000
Num. Obs.	123,101	127,044	127,044	127,044

Source: OLS estimation using SPADIES for actual data; model's baseline predictions for simulated data.
Notes: Dependent variable is ln(classes completed per year), or $\ln(x_t)$ in the model. An observation is a student-year; years 1-8 are included. Ability is θ, effort is e_t^*, and shock to classes completed is z_t. All regressions include year fixed effects (not shown). Standard errors (in parentheses) are clustered by student.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

One advantage of our structural model is that we can recover variables -effort and shocks- not observed in the administrative data. This allows us to expand regression (2) in order to gauge the relative roles of effort, performance shocks, and ability. Adding effort (column 3), we explain an additional 30 percent of the variation in classes completed. As expected, we explain all the variation in classes completed when we add the shock as well (column 4), with a coefficient on log ability equal to our point estimate for α. Further, in column 4 effort and the shock acount for 47 and 52 percent of the variation in classes completed, respectively, leaving a mere 1 percent explained by log ability. In other words, in our model almost the whole variation in classes completed is due to effort and shocks.[28]

Importantly, when effort is not controlled for, the coefficient on ability is overestimated by about 75 percent. Therefore, when we do not control for effort, which is typically unobserved in administrative data, we attribute to ability an effect that is actually mediated through effort, since higher ability students have lower effort cost and exert more effort). By underestimating the role of effort, policies might place too much weight on a student trait -ability- that cannot be changed rather than on a student choice -effort- that could, in principle be changed.

These findings elicit two important questions. The first is whether policies could actually raise the number of classes completed. Combining equations (3) and (12), we arrive at the effort associated with a particular target:

$$e_t = \left(\frac{q_t}{\bar{x} E(z_t) \theta^{\alpha}} \right)^{\frac{1}{1-\alpha}}. \tag{17}$$

As the expression shows, effort is greater the higher the target, the lower the expected shock, and the lower the ability. If a policy aims at raising classes completed, it must raise a student's optimal target, q_t^*, and its corresponding effort. Since the number of classes completed is discrete (for instance, a student can raise q_t^* from 14 to 15, but not to 14.5), a target increase requires a non-marginal effort increase. The answer to the first question, then, is that a policy can raise the number of classes completed provided it incentivizes a non-marginal effort increase.

Assuming a policy can accomplish this goal, the second question is which students are most likely to respond to it. A student chooses effort as the solution of the dynamic optimization problem in (6). In a given period, her chosen effort is related to her ability, income, target number of classes, cumulative classes completed (which may create the need to catch up), expected shock to classes completed, and graduation probability. In principle, the relationship between ability and effort is ambiguous. On the one hand, higher-ability students have lower effort costs, which induces them to exert greater effort. On the other hand, they need less effort to attain a given target, which induces them to exert less effort.

Using the simulated data from our model, Table 11 shows the correlates of effort based on some reduced-form regressions. In addition to ability, column 1 controls for target number of classes, whereas columns 2, 3, and 4 control for other variables that determine target and effort – namely, past cumulative classes completed (proxied by average classes completed per year in past years), expected dropout probability, and expected shock to classes completed. Column 1 shows that higher ability students exert lower effort, whereas column 2, 3 and 4 show the opposite. The reason for these opossing results is that column 1 shows the effect of ability controlling for target (i.e., the effect indicated in equation 17), by which higher ability students

[28] These results are based on a R-squared decomposition of the regression in column 4, whcih quantifies the fraction of R-squared that is attributable to each independent variable. Note that the fraction attributable to the shock is partly attributable to *past* effort, since the shock is a function of past classes completed, which depends on past effort.

need less effort for a given target, whereas the other columns show ability effects including those on target choice, by which higher ability individuals choose higher targets and effort.

Table 11: Effort.

	(1)	(2)	(3)	(4)
ln(ability)	-0.015^{***}	0.076^{***}	0.107^{***}	0.067^{***}
	(0.003)	(0.003)	(0.003)	(0.003)
ln(target)	0.787^{***}			
	(0.004)			
ln(average classes completed)		-0.173^{***}		
		(0.008)		
ln(expected dropout probability)			0.022^{***}	
			(0.000)	
ln(expected shock to classes completed)				-0.159^{***}
				(0.006)
Constant	-1.537^{***}	1.019^{***}	0.764^{***}	0.519^{***}
	(0.011)	(0.018)	(0.003)	(0.006)
Adj. R^2	0.766	0.285	0.557	0.286
Num. obs.	127,044	127,044	127,044	127,044

Source: OLS estimation using model's simulated baseline values.
Notes: The dependent variable is ln(effort), or $\ln(e_t^*)$ in the model. An observation is a student-year. Ability is θ, target is q_t, and cumulative classes completed is $\bar{h}_{t-1}$. Expected shock to classes completed is $E[z_t]$; it varies across students and over time. Expected dropout probability is $E[\tilde{p}^d(t, \cdot)]$, calculated by the individual prior to the realization of z_t. All regressions include year and income fixed effects (not shown). Standard errors (in parentheses) are clustered by student. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column (2) shows that students with better past perfomance exert less effort. This is the outcome of a direct and an indirect effect of past performance on effort, both of which move in the same direction. Because of the direct effect, students with better past performance are closer to completing graduation requirements and need less effort. Because of the indirect effect, they expect better shocks to classes completed and a lower dropout probability, both of which make effort less necessary (columns 3 and 4). The direct and indirect effects of past performance illustrate the importance of a strong beginning: students who perform well in the early years can expect better "luck" in the future and a higher graduation probability, thereby eliminating the need for a costly catch-up.

We can now return to our second question. Policies that seek to affect effort might be more successful among high-ability students, whose effort cost is lower. But, importantly, they should aim at raising effort in the early years so that students have less need of a costly catch-up later on. As a result, free college targeted to specific students (e.g., of high ability), or contingent on student's cumulative classes completed (e.g., performance-based policies) might be highly effective at raising effort and number of classes completed. These conclusions motivate some of the counterfactuals presented below.

7 Free college simulations

In this section we present counterfactual simulations of alternative free college programs. We begin by describing their setup as well as the theoretical impact of free college on student decisions. Then we describe free college's aggregate effects on enrollment and graduation, and examine effort in order to understand graduation rate effects. We then turn to the heterogeneity of free college effects by student type, and relate our counterfactual results with findings from the empirical literature on college financial aid. We analyze the role of risk, both in the baseline and under free college. We conduct a simple cost-benefit analysis of free college, and conclude with an examination of the possible long-run effects of free college.

7.1 Counterfactuals' setup

We simulate the following free college programs, which provide free tuition to the students listed below:

- **Universal free college:** all college students.
- **Need-based free college:** low-income students, defined as those from the lowest two income brackets.
- **Ability-based free college:** high-ability students, defined as those from the highest ability quintile.
- **Ability- and need-based free college**: students who are both low income and high ability – namely, students from the lowest two income brackets and the highest ability quintile.
- **Performance-based free college:** all students in year 1, but only high-performing students afterwards. In a given year, a high-performing student is one who finished the previous year in tiers 1 or 2 based on her cumulative performance. Note that a student may be high-performing in one year but not in another, depending on her cumulative classes completed.
- **Performance- and need-based free college:** all low-income students (from the lowest two income brackets) in year 1, but only to students who are both low-income and high-performing afterwards.

Note that, by design, performance-based and performance- and need-based free college entail a risk for the student, as the zero tuition is not guaranteed but must be attained through performance. In contrast, the other free college programs guarantee zero tuition based on ex-ante student traits, which are not contingent on performance.

As we analyze these counterfactuals, it is convenient to distinguish between two groups of students–*existing* and *new.* For a given counterfactual, we define existing students as those who enroll both in the baseline and at least in year 1 of the counterfactual. In contrast, new students are those who do not enroll in the baseline but enroll in the counterfactual.

Three important assumptions hold in our counterfactuals. First, we assume a perfectly elastic higher education supply, whereby institutions can adjust capacity as needed in order to absorb additional students at a constant marginal cost, and tuition does not rise in response to greater demand. If there were capacity constraints, institutions would need to ration free college. Second, we assume that parents continue to transfer the same y to their children in college even when college becomes free. In other words, public college funding does not crowd out private funding at all. If, in contrast, parents reduce their transfers one-for-one (full crowd-out), free college has no effects and is similar to the baseline. More generally, crowding out of any degree is analogous to a reduction in tuition subsidy.[29] Third, we assume that the average cost of educating new and existing students is the same. This does not hold, for instance, when new students need additional services, such as remedial education or non-academic supports. If new students are more costly, then a government with fixed resources cannot fund as many students. Since these assumptions provide the most favorable setting possible for free college, the results presented below are best viewed as an upper bound on free college effects.

7.2 Free college and student choices

Free college affects students on the extensive and intensive margins. On the extensive margin it raises college entry by adding new students, thereby affecting the size and composition of the student body. On the intensive margin, it affects student effort via the following channels:

1. **Loss-of-urgency effect**. By raising consumption during college, free college enhances the "college experience," as it raises the value of being a student relative to joining the labor force and makes the student less eager to leave college. Other things equal, this effect leads to lower effort.

2. **Substitution effect**. Since free college raises consumption while in college, it allows students to exert more effort without losing utility – in other words, it provides additional consumption to compensate for greater effort. Other things equal, this effect leads to higher effort.

3. **Risk effect.** The performance and dropout shocks, z_t and d_t^{drop} respectively, depend on the number of cumulative classes completed, which in turn depends on effort. To the extent that free college raises (lowers) effort, it also lowers (raises) risk. Further, the longer (shorter) a student spends in college, the greater (lower) her exposure to risks.

At the extensive margin, free college affects enrollment. At the intensive margin, it affects graduation; whether graduation rises or falls depends on the net effect on effort. Importantly, performance- and performance-and-need based free college provide a direct incentive to student effort by making tuition contingent on performance. As a result, they eliminate the loss-of-urgency effect.

[29]There is a large literature endogenizing parental transfers as a function of education cost and labor market returns. See, for instance, Keane and Wolpin (2001), Restuccia and Urrutia (2004), and Abbot et al (2013).

7.3 Aggregate Outcomes

7.3.1 Enrollment and graduation

Table 12 shows the aggregate effects of the free college programs on enrollment and college outcomes. All these programs raise enrollment relative to the baseline, albeit in different magnitudes. Universal free college delivers the largest enrollment increase (28 pp), followed by need-based free college (23 pp) and performance-based free college (21.3 pp). In contrast, ability-based free college (whether need-based or not) has the smallest effect on enrollment rates (3 or 4 pp). The magnitude of the enrollment effect is related to the percent of high school students who are eligible for free college – equal to 100 for universal and performance-based free college, 71 percent for need-based and performance- and need-based free college, and only 20 percent for ability-based free college.

New students account for almost half of the student body under universal free college but just about 10 percent under ability-based free college. The entry of new students changes not only the size but also the composition of the student body. The fraction of low-income students rises in all programs, particularly those that are need-based. As for the fraction of high-ability students, it rises with ability- and ability-and-need based free college because these programs induce a positive selection of new students, who are of higher ability than the existing ones on average. The other programs, in contrast, induce a negative selection of new students. Nonetheless, performance- and performance-and-need based free college attract new high-ability students at a higher rate than universal- and need-based free college.

While four out of six programs raise enrollment rates by more than 15 pp, no policy affects graduation rate by more than 3 pp, and two programs (universal and need-based free college) actually lower it. These aggregate effects are small yet mask considerable differences between existing and new students. For existing students, graduation rate remains almost constant in all counterfactuals except for performance and performance-and-need based free college, which raise it by 4 to 6 pp. For new students, graduation rate depends on the policy-induced selection of new students. In programs with positive selection, the graduation rate of new students is almost 10 pp higher than that of existing students in the baseline, consistent with the greater effort exerted by higher-ability students. In the other programs, new students graduate at lower rates than existing ones. Still, performance and performance-and-need based free college graduate new students at higher rates than universal and need-based free college. In terms of on-time graduation rate, aggregate effects are very small (between -1 and 1 pp.), though they are large among the positively-selected new students under ability-based free college.

The ultimate goal of free college programs is raising the fraction of high school graduates completing college. By definition, this fraction is the product of enrollment and graduation rates. All programs raise this fraction relative to the baseline. Universal free college delivers the greatest increase (about 12 pp), followed very closely by performance-based free college. Relative to universal free college, performance-based free college delivers a lower increase in enrollment rates, as many students seek to avoid the risk involved in it. However, since performance-based free college graduates students at a higher rate than universal free college, it delivers almost the same increase in the fraction of high school graduates completing college. The trade-off between risk and outcomes illustrated by this comparison is a theme in the analysis that follows.

Table 12: **Free College Counterfactuals: Aggregate Outcomes.**

Attribute	Data	Baseline	Universal	Need	Ability	Ability & need	Perf.	Perf. & need
Enrollment rate (%)	32.3	32.3	59.9	54.9	36.8	35.4	53.6	49.7
Eligible students (%)			100.0	71.0	20.0	9.8	100.0	71.0
Student body composition (%)								
New students			46.1	41.2	12.2	8.9	39.8	35.0
Low income	52.4	52.4	66.0	72.1	54.5	56.6	64.0	69.0
High ability	39.5	39.5	28.8	29.0	46.9	44.9	30.9	31.0
Graduation rate (%)	45.6	45.0	43.5	43.6	45.6	45.7	47.8	44.9
Existing students			45.8	45.6	44.4	44.8	51.3	48.7
New students			40.9	40.8	54.3	54.1	42.7	42.6
On-time graduation rate (%)	15.1	15.1	13.9	14.4	15.2	15.6	15.9	16.1
Existing students			14.0	14.5	14.3	14.8	16.0	15.9
New students			13.8	14.3	22.2	24.2	15.7	16.4
High school graduates that complete college (%)	14.7	14.5	26.1	24.0	16.8	16.2	25.6	22.3

Source: Model's predictions for baseline and counterfactuals.
Notes: In these simulations, college is free for the following: all students (universal); students from the two lowest income brackets (need-based); students from the top ability quintile (ability-based), students from the two lowest income brackets and the top ability quintile (ability-and-need based); students classified in Tiers 1 or 2 the previous period (performance-based), students classified in Tiers 1 or 2 the previous period from the two lowest income brackets (performance- and need- based). Eligible students are those who could, in principle, make use of free college. Existing students are those who enroll both in baseline and counterfactual. New students are those who do not enroll in baseline but enroll in counterfactual. Low income=two lowest income brackets; high ability=top ability quintile.

7.3.2 Graduation and effort

With the goal of understanding the graduation rate difference between universal and performance-based free college, we focus on existing students and examine how their average effort and accumulation of classes completed differs between baseline and counterfactuals.

Under universal free college, on average students increase effort in the first two years but backload it afterwards. The substitution and risk effects prevail in years 1 and 2, leading students to work harder (Figure 12 panel a) and accumulate more classes (panel b) than in the baseline. Nonetheless, this greater accumulation of classes, together with the loss-of-urgency effect, leads students to work less and accumulate fewer classes than in the baseline in years 3-5. To compensate for their slower pace and graduate, students raise effort in years 6-8. Thus, students conduct an intertemporal re-allocation of effort which, on average, delivers almost the same graduation rate as the baseline, yet a lower on-time-graduation rate.

In contrast, under performance-based free college students frontload effort. Absent the loss-of-urgency effect, the substitution and risk effect lead students to work harder than in baseline in years 1-3. Since this makes them accumulate more classes than in the baseline in those years, they can work less later while still accumulating more classes than in the baseline. These effort

Figure 12: Free college: Changes in effort and cumulative classes completed

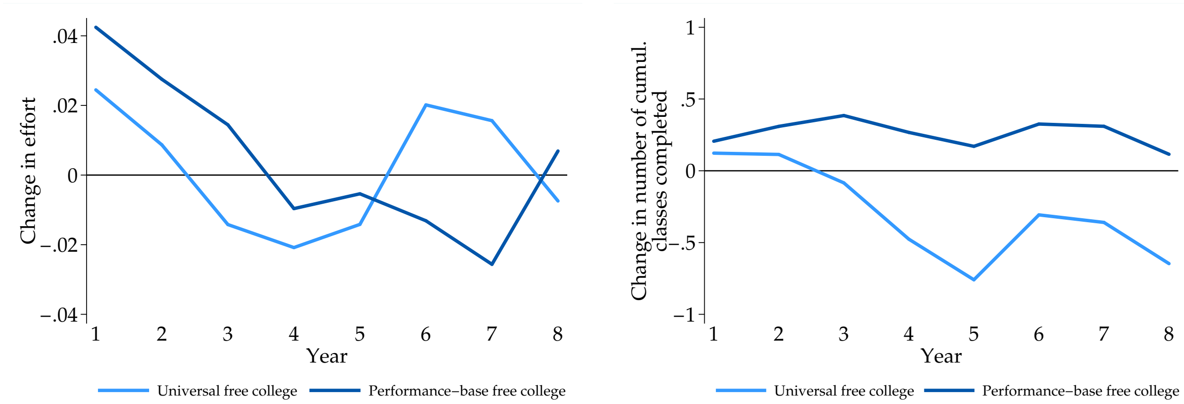

a. Change in effort w.r.t. to baseline   b. Change in number of cumulative classes

Source: Model's simulations.
Notes: For existing students in the counterfactuals, panel a shows average effort in the corresponding counterfactual minus average effort in the baseline. Panel b does the same for average number of cumulative classes completed.

changes imply not just an intertemporal reallocation of effort but also a net effort increase, which raises both graduation and on-time graduation rates relative to the baseline. In sum, the effort incentives contained in performance-based free college are the key to their success with graduation rates.

7.4 Outcomes by student type

The aggregate outcomes shown above mask great heterogeneity among student types. To explore it, we classify students into nine groups based on income and ability. Income groups include high, middle and low-income–corresponding to the top, two middle, and two bottom income brackets respectively. Ability groups include high, middle, and low-ability–corresponding to the top, two middle, and two bottom ability quintiles respectively.

7.4.1 Enrollment

Table 13 shows the enrollment rate for each student group in the baseline (panel a) and the enrollment rate increase relative to the baseline under universal and performance-based free college (panels b and c, respectively). We focus on these two counterfactuals because the remaining four are special cases of these. For example, the "low income" row of panel b displays the effects of need-based free college, and the "high ability" column displays the effects of ability-based free college.

Both free college programs raise enrollment rates for all student groups. Under universal free college, the greatest effects are for low- and middle-income students, who are most budget-constrained in the baseline. As a result, enrollment rates become more equal among student groups. For instance, the baseline enrollment gap between high-income, high-ability and low-income, low-ability students is equal to $70.2 = 83.8 - 13.6$ pp, but shrinks to $49.5 = (83.8 +$

Table 13: Free College Counterfactuals: Enrollment Effects.

	a. Baseline Enrollment Rate			b. Universal Free College: Change			c. Performance-Based Free College: Change		
				Ability					
Income	High	Mid	Low	High	Mid	Low	High	Mid	Low
High	83.8	65.5	39.3	7.1	11.4	12.0	5.8	9.0	9.2
Mid	73.1	47.4	27.3	15.0	21.1	18.3	12.5	16.8	13.9
Low	51.4	26.0	13.6	32.2	36.5	27.8	27.1	28.5	20.1

Source: Model's predictions for baseline and counterfactuals.
Notes: For each student type, panel a shows predicted enrollment rate (%) in the baseline; panel b shows the enrollment rate difference (pp) between universal free college and baseline; and panel c shows the enrollment rate difference (pp) between performance-based free college and baseline.

$7.1) - (13.6 + 27.8)$ pp under universal free college. From the point of view of policy design, the greater responsiveness of low- and middle-income students to free tuition justifies targeting it to them, as in our need-based free college simulation. In contrast, the lower responsiveness of high- than mid-ability students to free tuition indicates that ability-based free college cannot be justified by enrollment considerations. Instead, it can be justified on the grounds that it raises overall graduation rate through the positive selection of new students, and it lowers the income-based enrollment gap among high-ability students (from 32.4 to 7.3 pp.)

As for performance-based free college (panel c), its enrollment effects are qualitatively similar to those of universal free college, albeit smaller. Effects are similar because performance-based free college does provide free college to all students in year 1. Effects are smaller because some potential students do not enroll in order to avoid the effort necessary to remain in tiers 1 or 2.

Our enrollment results are consistent with the small free college program run in Colombia between 2015 and 2018, *Ser Pilo Paga* ("being diligent pays off.") This was an ability-, need-, and performance-based program that offered free tuition to high-ability, low-income students for the theoretical duration of their program; students who dropped out were required to pay the tuition back. Among the eligible population, the program raised college enrollment rates by 32 pp (Londoño et al 2020), in line with the 27-32 pp increase depicted in Table 13 (panels b and c) for these students.

7.4.2 Graduation

To compare graduation rates in the baseline and counterfactuals we focus on existing students, whose graduation rate changes are entirely driven by the incentives created by free college.[30] Following the same logic as with enrollment, Table 14 shows baseline graduation rates (panel a), as well as the graduation rate changes relative to baseline under universal and performance-based free college (panels b and c, respectively.)

Universal free college. Although it raises graduation rates for some groups, it lowers them for others. The reason is that the prevailing effort effect –loss of urgency, substitution, or risk– varies among student groups. In general, loss of urgency is stronger among higher-income

[30] Free college generates incentives both for existing and new students. We focus our analysis on existing students because they allow us to compare behavioral changes between the baseline and the counterfactuals. Since new students do not enroll in the baseline, they do not allow for this comparison.

Table 14: Free College Counterfactuals: Graduation Rate Effects for Existing Students.

	a. Baseline Graduation Rate			b. Universal Free College: Change			c. Performance-Based Free College: Change		
				Ability					
Income	High	Mid	Low	High	Mid	Low	High	Mid	Low
High	59.3	44.4	26.3	-1.8	-1.9	3.9	4.4	3.2	5.5
Mid	55.6	39.9	29.9	-1.9	3.5	1.7	4.9	6.6	6.5
Low	53.3	42.5	32.1	-1.3	4.4	-1.5	8.1	8.8	2.6

Source: Model's predictions for baseline and counterfactuals.
Notes: For existing students of each type, panel a shows predicted graduation rate (%) in the baseline; panel b shows graduation rate difference (pp) between universal free college and baseline; and panel c shows graduation rate difference (pp) between performance-based free college and baseline.

students, whose high baseline consumption only gets higher with free college. The substitution effect is stronger among lower-ability students, for whom effort is more costly and for whom free college provides greater compensation for additional effort. The risk effect is stronger for students with fewer cumulative classes completed, particularly when the loss of urgency has led them to delay graduation and has exposed them to risk for a longer time.

As panel b shows, graduation rates fall in the upper triangle (students of high ability or high income) and generally rise in the lower triangle (students of low ability or low income). In the upper triangle, baseline effort is already high because effort cost is low and consumption high. Since there is little room for additional effort, the substitution effect is weak. And, since consumption is already high, the loss-of-urgency effect is strong and prevails, leading to lower effort, greater risk, and lower graduation rates.

The story is reversed in the lower triangle. Baseline effort is low because effort cost is high and consumption low, which gives room for a strong substitution effect. Since consumption remains low even with free college, the loss-of-urgency effect is weak and the substitution effect prevails, leading to higher effort, lower risk, and higher graduation rates.

The exception in the lower triangle is the students with the lowest income and ability, for whom graduation rates fall. The combination of lowest ability (which renders effort very costly) and lowest income (which makes consumption as a college dropout higher than as a college student, even with free college) leads these students to lower effort and drop out at higher rates.

Performance-based free college. In contrast to universal free college, performance-based free college raises graduation rates for all student groups by eliminating the loss-of-urgency effect. The substitution effect prevails and leads to greater effort, which in turn mitigates risk. In the aggregate, performance-based free college raises the graduation rate of existing students from 45 to 51.3 percent (see Table 12), which is the greatest graduation rate increase for existing students among all programs considered here. Even new students attain a higher graduation rate under performance-based than universal free college (42.7 v. 40.9 percent).

Ultimately, the policymaker is interested in raising the fraction of high school graduates who complete college. Table 15 shows this fraction in the baseline (panel a) as well as the differences with respect to it under universal and performance-based free college (panels b and c, respectively.) This fraction rises for all student groups in both counterfactuals. Performance-based free college is at least as effective as universal free college for every student group except low-income students of middle or low ability. For these students, enrollment rates rise substantially more under universal than performance-based free college given their inherent difficulty in meeting the zero-tuition performance requirements.

Table 15: Free College Counterfactuals: Percent of High School Graduates That Complete College.

	a. Baseline Percent of High School Graduates			b. Universal Free College: Change			c. Performance-Based Free College: Change		
	Ability								
Income	High	Mid	Low	High	Mid	Low	High	Mid	Low
High	49.7	29.1	10.3	2.9	3.1	5.1	7.1	5.9	5.3
Mid	40.7	18.9	8.2	6.7	10.1	6.4	10.7	10.1	6.4
Low	27.4	11.0	4.4	16.7	17.5	8.5	19.1	15.5	7.0

Source: Model's predictions for baseline and counterfactuals.
Notes: For existing students of each type, panel a shows the percent of high school graduates that graduate from college; panel b shows difference (pp) in this variable between universal free college and baseline; and panel c shows difference (pp) in this variable between performance-based free college and baseline.

7.5 Discussion

Our predicted effects are consistent with some empirical regularities as well as the literature on higher education financial aid. For Latin America, they provide an explanation for the cross-country evidence depicted in Figure 1, which shows that government funding for higher education has a large, positive relationship with enrollment rates but a weak (actually negative) relationship with graduation rates. This evidence is consistent with our counterfactuals, which show a greater impact of free college on enrollment than graduation rates. The graduation rate evidence is also consistent with the financial aid regimes prevailing in those countries, which are typically universal, need-based or ability-based – but not performance-based (Ferreyra et al 2017). Indeed, for the most common regimes in Latin America, namely universal and need-based financial aid, our counterfactuals predict a decline in graduation rate (see Table 14), consistent with Figure 1.

Our findings are also consistent with the literature on U.S. higher education financial aid, which has found positive effects of financial aid on enrollment and graduation.[31] In recent years, states have implemented a variety of financial aid programs, based on merit and/or aid,[32] as well as "Promise" programs, which provide zero tuition to eligible students for local community colleges or state four-year institutions.[33] Overall, these programs have had positive effects on

[31] See, for instance, Bettinger (2004), Dynarski (2003), Hoxby and Turner (2013) and the references therein, as well as the surveys by Avery et al (2019), Deming and Dynarski (2009), Dynarski and Scott-Clayton (2013), Long (2008), Page and Scott-Clayton (2016).

[32] See, for instance, Bettinger et al (2019), Castleman and Long (2016), Cornwell et al (2006), Dynarski (2000, 2004, 2008), Scott-Clayton (2011), Scott-Clayton and Zafar (2016), and the references therein.

[33] These programs vary across states in terms of eligible institutions (community colleges v. four-year institutions) and students. See, for instance, Carruthers et al (2018), Dynarski et al (2018), Gurantz (2020).

enrollment and graduation. Also consistent with our counterfactuals, the literature has found greater effects of financial aid on enrollment than on graduation, with the latter ranging from 0 to 6-7 pp.[34] As in our counterfactuals, the literature has found that performance-based financial aid improves college outcomes more than unconditional aid.[35]

Even though performance-based free college raises graduation rates more than other free college programs in our simulations, it still delivers a relatively small graduation rate increase (between 6 and 14 percent of the baseline graduation rate depending on whether we consider existing or new students, and whether it is need-based). Further, graduation rate effects are small not only in our counterfactuals but also in the literature, as discussed above. Such small effects beg the question of why free college (or financial aid, in general) fails to substantially raise graduation rates.

The answer emerging from our model is that free college fails to deliver a large, non-marginal effort increase. As Table 11 shows, students who need to catch up exert more effort. Particularly when the student has fallen far behind, catching up is difficult (see Section 3.4.2), which renders performance incentives of limited use. This suggests that additional supports promoting class completion, particularly in year 1, may be needed.[36] These include remedial education, advising, mentoring, and tutoring.[37] Recent evidence (Deming and Walters 2017) indicates the effectiveness of directing funding to institutions for these supports. Indeed, in their examination of possible policies to raise graduation rates in the U.S., Avery et al (2019) conclude that need-based free college combined with higher funding to institutions might be the most cost-effective policy. A cautionary tale, however, comes from Oreopoulos and Petronijevic (2019), who find that these supports help students realize the need for greater effort but have little effect on college continuity or graduation. The reason is that students respond not by raising effort, but by expecting less of themselves.

Moreover, our counterfactuals have assumed two highly favorable – perhaps unrealistic – conditions. First, we have assumed no capacity constraints in higher education. Free college, however, would likely meet capacity constraints – in which case institutions would ration access. Bucarey (2018) explores the potential effects of free college (*gratuidad*) in Chile, where it was introduced in 2016 for the bottom 60 percent of the income distribution. Data prior to 2016 shows that when low-income students are given additional financial aid, institutions respond by becoming more selective, thereby leading to a lower share of low-income students. Using his structural model, Bucarey (2018) predicts that capacity-constrained institutions would follow a similar strategy in the presence of free college, with similar effects. Second, we have implicitly assumed that college quality remains the same after college becomes free. This might not happen, for instance, if the policymaker reimburses colleges for less than the full tuition. Murphy et al (2019) provide evidence that higher education quality rose in England when free college ended. The reason is that, under free college, the institutions were not receiving enough funding. The abolition of free college has been implemented as an income-contingent loan from the government, which covers full per-student costs throughout college. Murphy et al (2019)

[34] For recent, well-identified studies, see Bettinger et al (2019), Denning (2017), Mayer et al (2015), and Scott-Clayton (2011).

[35] See Dynarski and Scott Clayton (2013).

[36] In the context of our model, these policies would amount to changing parameters from the distribution of z in order to reach higher $E(z)$ and lower $V(z)$.

[37] Clotfelter et al (2018), Evans et al (2017), Scrivener et al (2015) and Sommo et al (2018) document positive effects of these supports on graduation rates.

show that the new system has led to higher enrollment, greater participation of lower-income students, and higher per-student funding.

7.6 Free college and anticipated risk

For students, college is a risky investment. They are uncertain about the number of classes they will be able to complete in a given year as well as their chances to remain enrolled from year to year. In this section we investigate whether free college can mitigate risk. Free college affects risk through two channels. First, it raises the value of remaining in college relative to dropping out by raising consumption during college and enhancing the "college experience." Second, it affects effort through the loss-of-urgency, substitution, and risk effects described above. Effort changes affect the number of classes completed, which in turn affects the future shocks to classes completed and the dropout probability.

These consumption and effort changes are distinctive of our setting. Consumption changes are of great importance in Colombia's severely credit-constrained economy. In an economy with credit, consumption is more closely related to lifetime income (as in standard Ben-Porath models), and this channel is hence more muted. Effort changes, in turn, are not present in models with exogenous risk, in which risk is unrelated to effort (e.g. Hendricks and Leukhina 2017).

In our model, these consumption and effort changes affect the distribution of the value of college, $V^{coll}(t, h_{t-1}, \theta, y)$, defined in (6). Whereas $V^{coll}(t, h_{t-1}, \theta, y)$ is calculated for the student's optimal effort, e_t^*, while taking expectation over z_t, we now define the college payoff from *any* given effort and realization of z_t:

$$\begin{aligned}\tilde{V}^{coll}(t, h_{t-1}, \theta, y; z_t, e_t) = U(c_t, e_t, \theta) \ + \beta \Bigg[& \mathbf{1}_{\{t \geq 5\}} \Pr\left(h_{t-1} + H(z_t, \theta, e_t)\, \bar{x} \geq h^{grad}\right) V^{grad}(t+1) + \\ & \Pr\left(h_{t-1} + H(z_t, \theta, e_t)\, \bar{x} < h^{grad}\right) \Big(\tilde{p}^d(t, h_{t-1} + H(z_t, \theta, e_t)\, \bar{x}, \theta, y) V^{drop}(t+1) + \\ & \left(1 - \tilde{p}^d(t, h_{t-1} + H(z_t, \theta, e_t)\, \bar{x}, \theta, y)\right) V^{coll}(t+1, h_{t-1} + H(z_t, \theta, e_t)\, \bar{x}, \theta, y) \Big) \Bigg].\end{aligned} \tag{18}$$

From here, it follows that $E_z[\tilde{V}^{coll}(t, h_{t-1}, \theta, y; z_t, e_t^*)] = V^{coll}(t, h_{t-1}, \theta, y)$.

Our proposed measure of anticipated risk for student i, at the beginning of period t under scenario p (baseline or counterfactual) is the coefficient of variation of her college payoffs:

$$CV_{it}^p = \frac{\sqrt{\mathrm{Var}_z \left[\tilde{V}^{coll}(t, h_{it-1}, \theta, y; z_{it}, e_{it}^*) \,|\, p \right]}}{E_z \left[\tilde{V}^{coll}(t, h_{it-1}, \theta, y; z_{it}, e_{it}^*) \,|\, p \right]}, \tag{19}$$

where the right-hand side is the ratio between the standard deviation and expected value of college payoffs, calculated for the student's optimal effort, e_{it}^*, chosen under scenario p. Given e_{it}^*, the randomness in $\tilde{V}^{coll}(\cdot)$ comes from the performance and dropout shocks, z and d^{drop} respectively, which are associated with cumulative classes completed and therefore effort. Note that anticipated risk varies across students and programs, and over time.

An important question is whether, in our baseline, effort lowers anticipated risk. To investigate this matter, in principle we might think of using our simulated baseline data to regress anticipated risk on effort, controlling for cumulative past performance. However, this regression would suffer from endogeneity, as optimal effort would appear in both sides of the regression. Since e^*_{it} is a function of the state variables, t, h_{it-1}, θ_i, and y_i, we instrument for effort using the state variables, and run the second-stage regression reported in Table 16.[38] The regression indicates that, in our baseline model, effort does lower anticipated risk. Further, the effect is large. Students, indeed, use effort very effectively to mitigate performance and dropout risk.

Table 16: Anticipated Risk.

	Dependent variable ln(Anticipated risk)
ln(effort)	-2.340^{***}
	(0.089)
ln(average classes completed)	-1.767^{***}
	(0.043)
Constant	-0.772^{***}
	(0.068)
Num. obs.	116,761

Source: 2SLS estimation using model's simulated baseline values.
Notes: The dependent variable is ln(anticipated risk), or $\ln(CV_{it})$ for the baseline. An observation is a student-year. Included students are "existing" in both counterfactuals; upper 5% tail of risk has been trimmed. Effort is e^*_{it}, instrumented with by classes completed, as well as year, income, and ability fixed effects. ln(average classes completed) is $\tilde{h}_{it-1}$, as defined in the model. The regression includes year fixed effects (not shown). Standard errors (in parentheses) are clustered by student. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 13 shows the average anticipated risk in the baseline, and under universal and performance-based free college. To facilitate comparisons, we focus on the set of students who enroll in college in all three scenarios (akin to our previous "existing students" concept).[39]

In the three scenarios, average anticipated risk is U-shaped with respect to time. In year 1, it is high for two reasons. The first is student body composition, which includes a large share of students with a high exogenous dropout probability (many of which, indeed, drop out in the initial years). The second is that individual students are uncertain about their future performance, and their exogenous dropout probability is higher in year 1 than in subsequent years. In years 2 through 5, anticipated risk falls as many high-risk students drop out while the remaining students settle on a performance path. Anticipated risk rises after year 6, as the students who remain at that point face the risk of not graduating at all.

Relative to the baseline, free college raises $E_z[\tilde{V}^{coll}(\cdot)]$ by raising consumption while in college and, in some cases, effort (see Figure 12). Free college also raises $\mathrm{Var}_z[\tilde{V}^{coll}(\cdot)]$ at least in year 1, although at some point $\mathrm{Var}_z[\tilde{V}^{coll}(\cdot)]$ becomes lower than in the baseline. The net result is that, relative to the baseline, universal free college raises anticipated risk in year 1 but lowers it afterwards. Meanwhile, performance-based free college raises it in years 1 through 3, and

[38]First-stage results can be found in Appendix Table A.11.

[39]While students are the same in year 1, they might not be in subsequent years because free college affects dropout rates. All measures reported in this section eliminate the 5-percent upper tail, which contains extreme values.

Figure 13: Anticipated Risk.

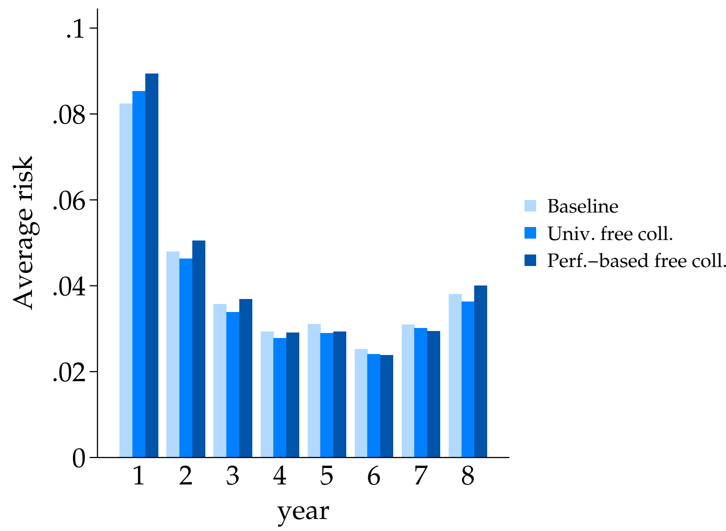

Source: Own calculations based on model's predictions for baseline and counterfactuals.
Notes: The figure shows, for each year, average anticipated risk for baseline and counterfactuals, calculated for students who begin college in year 1 in all three scenarios. Upper 5% tail has been trimmed.

lowers it afterwards. Note, however, that free college changes anticipated risk little relative to the baseline because effort changes little, as discussed in Section 7.5.

It is interesting to compare anticipated risk under universal and performance-based free college. On the one hand, performance-based free college elicits greater effort than universal free college. On the other hand, it makes consumption while in college contingent on performance. These two forces have opposing effects on anticipated risk. When the consumption force prevails (years 2-5), anticipated risk is higher under performance-based than universal free college. In other words, performance-based free college exposes students to greater risk than universal free college, but also elicits greater effort from them.

To summarize, college is a risky enterprise. Effort, however, can substantially lower anticipated risk. Risk is frontloaded and decreases dramatically after the first two years, once students survive the initial attrition and settle on a performance path. Universal free college lowers risk relative to the baseline after year 1, albeit very little. By introducing uncertainty in free college availability, performance-based free college subjects students to greater anticipated risk than the baseline or universal free college, but this very uncertainty leads them to higher effort and graduation rates.

7.7 A simple cost-benefit analysis of free college

As we saw in Section 7.4, free college programs vary in their ultimate outcome, namely the fraction of high school graduates who finish college. Nonetheless, the policymaker would be mistaken in simply choosing the policy that maximizes this fraction, as it might also be the most costly. Hence, in this section we conduct a simple cost-benefit analysis of the simulated free-college programs. The benefit is the fraction of high school graduates who finish college, and the cost is the net public spending incurred to produce them.

Lacking data on the annual cost of educating a college student, C, we assume it is equal to the tuition paid by students from the highest income bracket at public universities (equal to 2.2 million COP on average; see Table 4). This is a lower bound of the true cost, as public colleges subsidize tuition even for the highest income students, and private universities (whose tuition might be closer to actual costs) charge more than 6 million COP on average. We continue to assume a perfectly elastic college supply and a constant annual cost per student. We calculate the average net cost per graduate under policy p as follows:

$$anc_p = \frac{\sum_{t=1}^{8} \sum_{i=1}^{N_t^p} (C - T^p(t, h_{it}, \theta_i, y_i))}{G^p}, \tag{20}$$

where N_t^p is the number of students enrolled under policy p in year t, C is the annual cost of educating a college student (constant across programs and students, and over time), and $T^p(\cdot)$ is the tuition paid by student i under policy p. In the baseline, tuition varies across students based on income, as described in Table 4. In the counterfactuals, it varies as described in Table 4, with the modifications created by the corresponding free college policy. Equation (20) takes into account that the policymaker incurs a cost C to educate every student but may collect tuition from some of them. By adding over all students enrolled (rather than just those who graduate) in the numerator, the net cost per graduate incorporates the cost of dropouts.

For each policy, Figure 14 shows the net cost per graduate (panel a). For convenience, it also shows the fraction of high school graduates who finish college (panel b), which provides a measure of G. We measure the net cost per graduate relative to per capita GDP.

Figure 14: Free College Cost

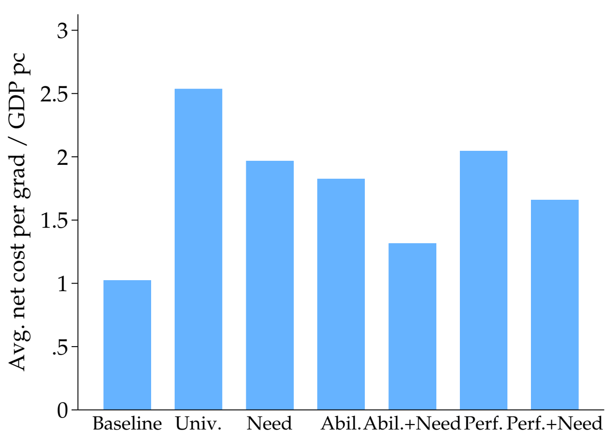

a. Average net cost per college graduate (relative to GDP per capita)

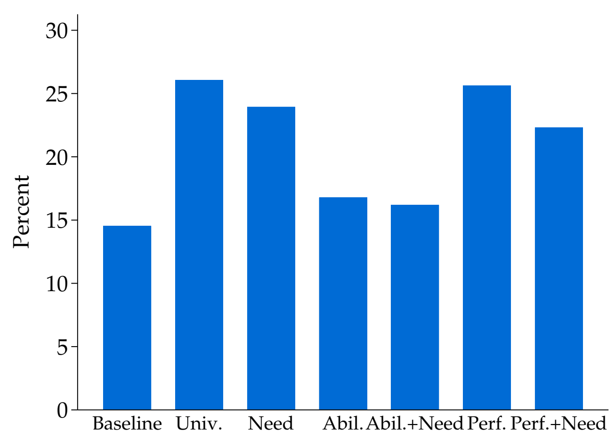

b. Percent of high school graduates that complete college

Source: Model simulations for enrollment and graduates. See the text for (gross) cost assumptions. GDP per capita is from the National Administrative Department of Statistics (DANE).
Notes: Average net cost per college graduate is computed as described in the text.

Net cost per graduate is already high in the baseline – about the same as per capita GDP. Even if enrollment did not rise, free college would raise net cost per graduate just because fewer students would pay tuition. The combination of lower tuition revenues and higher enrollment raises net cost per graduate under all free college regimes. The cost rises the least under ability-and-need based free college, and the most under universal free college.This is, of course, a reflection of the number of graduates from each policy: while universal free college adds the most graduates relative to the baseline, ability-and-need based free college adds the least.

Consider, now, a policymaker who is interested in maximizing the number of graduates subject to a budget constraint. Universal and performance-based free college maximize the number of graduates, yet performance-based free college does it at a lower per-graduate cost than universal free college. This is because the incentives embedded in performance-based free college produce fewer dropouts and graduate students faster than universal free college. By adopting performance-based rather than universal free college, the policymaker would save no less than 50 percent of per capita GDP in each graduate. Further, he would raise efficiency relative to the baseline by incentivizing effort and focusing public spending on the students who, given their performance, are most likely to graduate.

Despite these public savings and efficiency gains, the policymaker might be reluctant to adopt performance-based free college due, perhaps, to political considerations. Indeed, students throughout Latin America have been clamoring for unconditional, universal free college (*gratuidad*). As Figure 14 shows, universal free college raises the number of graduates only slightly more than need-based free college yet costs an additional 50 percent of per capita GDP per graduate. The reason is that universal free college provides a transfer to medium- and high-income students who would attend college even anyway. Therefore, when having to choose between universal and need-based free college, the latter is preferrable.

One caveat is in order before finishing this section. College dropouts seem to acquire human capital during their time in college, as given by their wage premium relative to high school graduates (see Section 5.2.2). This may motivate the policymaker to raise enrollment regardless of graduation. The question, then, is whether starting but not completing a bachelor's program is the most efficient way of acquiring human capital beyond high school. A short-cycle program, lasting two or three years, might be a more cost-effective option, and one that would avoid the psychic cost for students derived from starting but not finishing a bachelor's program. It would also avoid the overcrowding of bachelor's programs during the initial years, which are critical to final outcomes. The policymaker might consider providing free college for short-cycle programs, in the spirit of the relativley higher tuition subsidies provided for community colleges rather than bachelor's programs at public institutions in the United States (Denning 2017).

7.8 Long-run effects

One limitation of the current analysis is the assumption that, despite the greater supply of college graduates due to free college, the college premium does not change. In reality, however, the college premium might fall, thereby reducing the incentives to attend college. In addition, financing free college might require taxes, which would further depress after-tax wages for college graduates. In other words, free college might unleash general equilibrium effects that could undermine some of the effects previously described.[40]

To analyze the equilibrium relationship between the college premium and the number of college graduates, we postulate a production function following Katz and Murphy (1992),

[40] Heckman et al (1998) conclude that ignoring general equilibrium effects can bias the estimated effects of tuition subsidies.

Heckman et al (1998), and Card and Lemieux (2001). Ignoring for the moment the connection between capital and technology, consider a CES production function, $F(N_t^h, N_t^g) = \left(A_t^h(N_t^h)^\lambda + A_t^g(N_t^g)^\lambda\right)^{\frac{1}{\lambda}}$, that combines high school and college graduates, whose amounts are N_t^h and N_t^g respectively. In this function, A_t^h and A_t^g, are efficiency parameters, and $\omega = 1/(1-\lambda)$ is the elasticity of substitution between college and high school graduates. In a competitive labor market, the college premium is equal to

$$\frac{w_t^g}{w_t^h} = \frac{A_t^g}{A_t^h}\left(\frac{N_t^g}{N_t^h}\right)^{\lambda-1}, \tag{21}$$

where w_t^g and w_t^h are the wages of college and high school graduates, respectively. Taking logs on both sides and defining a time difference ($\triangle$) yields an expression commonly used to study changes in wage inequality:

$$\triangle \ln\left(\frac{w_t^g}{w_t^h}\right) = \triangle \ln\left(\frac{A_t^g}{A_t^h}\right) - \frac{1}{\omega}\triangle \ln\left(\frac{N_t^g}{N_t^h}\right). \tag{22}$$

In this equation, changes in relative wages depend on two components. The first is technical changes in the relative productivity of the two labor inputs. Skill-biased technical change, or $\triangle \ln(\frac{A_t^g}{A_t^h}) > 0$, increases the college premium. The second is changes in relative supply, adjusted by the elasticity of substitution. An increase in the supply of college graduates, N_t^g, lowers the college premium by an amount inversely related to ω. Over time, the evolution of the college premium depends on the relative strength of these two opposite forces.

To investigate the extent to which the increased supply of college graduates might depress the college premium, we focus on universal free college, which increases supply the most and provides an upper bound to the potential college premium decline. We examine short- and long-run effects. For the latter, we assume that universal free college becomes a permanent policy beginning with our 2006 cohort, and simulate the evolution of the share of college graduates and college premium. We assume no skill-biased technical change, which delivers an upper bound to the college premium decline.

The critical parameter for this exercise is ω, which we estimate as 3.2 for Colombia.[41] Using this estimate, the top and bottom panel of Figure 15 show the evolution of the share of college educated workers and the college premium, respectively. The former rises over time as young, more educated cohorts replace older, less educated ones. The replacement is complete around the year 2045, when the economy reaches steady state.

As in our model, the 2006 entry cohort starts to graduate in 2010. In 2006, the initial share of college graduates in the workforce equals 27 percent.[42] Universal free college raises

[41] In the literature there is a large number of estimates for this parameter. For example, Katz and Murphy (1992) estimate it to be 1.4 for the U.S. for the period 1963-1987, but the low range of estimates also include 0.5. Card and Lemieux (2001) estimate it for the United States (1959 to 1996) and the United Kingdom (1974-1996) and find estimates in the (2- 2.5) range. Following Card and Lemieux (2001), we estimate the prodution function parameters using Colombia's household surveys (SEDLAC) for 2001-12. Due to data limitations, college graduates include those from 2- and 4-year programs.

[42] Notice that the initial share of 27 percent is higher than the initial fraction of high school graduates that finish college, equal to 14.5 in the baseline. The reason is that 27 percent includes graduates from all bachelor's programs (rather than only five-year programs) as well as graduates from two-year programs. This is because

Figure 15: Free College and and the College Premium

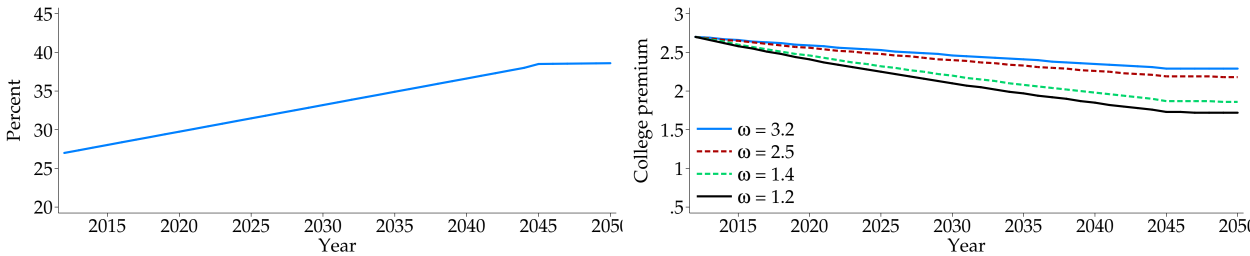

a. Share College in Labor Force

b.College Premium

Source: Own calculations based on estimates.
Notes: Panel a shows predicted share of college graduates in the labor force, calculated as the share of workers with a college degree (of all ages) relative to all workers (all ages and educational attainment) in the economy. Panel b shows predicted college premium, calculated as the ratio between wage for college graduates and wage for high school graduates.

this share for a particular cohort by 12 pp, to a long-run value of 39 percent (panel a). The college premium falls relatively little (panel b), from 2.6 to 2.28. Further, the decline is slow. The reason is that, in the short-run, the increase in college graduates from the first few cohorts represent a small addition to the stock of college-educated workers, which encompasses 35 to 40 cohorts. Holding other things constant, the impact on the share of college-educated workers and the college premium is negligible. Over time, the continuous inflow of a greater number of college graduates produces a larger impact. This, however, is quite small for our estimated elasticity, and unlikely to dissuade many high school graduates from college enrollment. As a robustness check, panel b depicts the implied path of the college premium for a range of ω estimates for developed economies.[43] The college premium declines more for these values, yet even the greatest decline -down to 1.70- still leaves the economy with a substantial college premium, comparable to that in the U.S. Skill-biased technical change, of course, would lead to a smaller decline for any value of ω. While forecasting technical change is clearly difficult, our working asumption of no skill-biased technical change should probably be viewed as extreme.

In sum, both the short- and long-run decline of the college premium induced by free college would likely be small in Colombia. As a result, abstracting away from the general equilibrium implications of free college might not substantially bias our predicted effects. College, thus, would remain a worthy investment even if more college graduates entered the labor market.[44]

the college premium is calculated for all higher education graduates, based on data that does not distinguish between graduates from bachelor's and two-year programs.

[43] We use the point estimates of 1.4 from Katz and Murphy (1992) and Heckman et al (1998), 2.5 from Card and Lemieux (2001), and include lower value of 1.2. For each elasticity value, the intercept has been adjusted to generate the same college premium for the year 2012.

[44] These results are consistent with Garriga and Keightley (2007), who conduct a similar analysis for the U.S. in a model that endogenizes the college premium, government budget constraint, after-tax earnings, and labor supply decisions. They show that tuition subsidies have small effects on the college premium even when accounting for taxes. Their results are partly due to the limited effects of the policy on graduation rates, and to the high elasticity of substitution between college and high school graduates.

8 Conclusions

In this paper we have developed and estimated a dynamic model of college enrollment, performance, and graduation. A central piece of the model, student effort, has a direct effect on the completion of classes, and an indirect effect mitigating the risk of not completing a class or not remaining enrolled. We have estimated the model using rich administrative data for Colombia. According to our estimates, effort has much greater impact than ability on class completion. Failing to model effort as an input to class completion leads to overestimating the role of ability by about 75 percent, and to favoring policies that promote positive selection of new students rather than effort. We have used our parameter estimates to simulate free college programs differing in eligibility. According to our simulations, universal free college expands enrollment the most but has the highest per-graduate cost and does not raise graduation rates. Performance-based free college, in contrast, delivers a slightly lower enrollment expansion but has a greater graduation rate and a lower per-graduate cost. Performance-based free college faces students with greater risks than the current baseline or universal free college, but precisely for this reason elicits greater effort and graduation rates.

For existing students, free college has little impact on effort and graduation rates, even when it is performance-based. This suggests that additional, complementary policies might be needed. Given the high persistence of performance, improving first-year performance and retention is critical. Supports such as tutoring, remedial education, mentoring, and advising might be helpful, particularly for the most disadvantaged students. Further, most college programs in Latin America start with classes specific to the major or field, and rarely include general education requirements that transfer easily across majors. As a result, changing majors typically implies starting from scratch, which leads many students to abandon higher education. Shortening bachelor's programs (which last five or six years) and connecting them with the labor market might raise effort and graduation rates as well.

Finally, there might be an additional reason why free college fails to elicit large effort changes: particularly for low-income students, free college might not be enough. For them, the effort necessary to complete college might require not only free tuition but also an additional, generous stipend. In countries with limited fiscal resources, providing this additional stipend might require lowering the existing college subsidies for affluent students – a redistribution that, while politically costly, would nonetheless enhance equity and efficiency.

References

[1] Abbot, B., Gallipoli, G., Meghir, C., & Violante, G. (2013). Education decisions and intergenerational transfers in equilibrium (Working Paper). New York: New York University.

[2] Ahn, T., Arcidiacono, P., Hopson, A., & Thomas, J. R. (2019). Equilibrium grade inflation with implications for female interest in stem majors. NBER Working Paper No. w26556.

[3] Akyol, A. & Athreya, K. (2005). Risky higher education and subsidies. *Journal of Economic Dynamics and Control, 29(6)*, pp. 979-1023.

[4] Altonji, J. (1993). The demand for and return to education when education outcomes are uncertain. *Journal of Labor Economics, 11(1),* pp. 48-83.

[5] Arcidiacono, P. (2004). Ability sorting and the returns to college major. *Journal of Econo-*

metrics, 121(1-2), pp.343-375

[6] Arcidiacono, P., E. Aucejo, A. Maurel, & T. Ransom (2016). College attrition and the dynamics of information revelation. NBER Working Paper No. 22325.

[7] Ariely, D., U. Gneezy, G. Loewenstein, & N. Mazar (2009). Large stakes and big Mistakes. *Review of Economic Studies, 76(2),* pp. 451-469.

[8] Avery, C., Howell, J., Pender, M., & Sacerdote, B. (2019). Policies and payoffs to addressing America's college graduation deficit. In BPEA Conference Draft, Fall.

[9] Beneito, P., J. Bosca & J. Ferri (2018). Tuition fees and student effort at university. *Economics of Education Review, 64(C),* pp. 114-128.

[10] Berry, S., J. Levinsohn, A. Pakes (1995). Automobile prices in market equilibrium, *Econometrica,* pp: 841–890

[11] Bettinger, E. (2004). How financial aid affects persistence. In College choices: The economics of where to go, when to go, and how to pay for it (pp. 207-238). University of Chicago Press.

[12] Bettinger, E., Gurantz, O., Kawano, L., Sacerdote, B., & Stevens, M. (2019). The long-run impacts of financial aid: Evidence from California's Cal Grant. *American Economic Journal: Economic Policy*, 11(1), 64-94.

[13] Bucarey, A. (2018). Who pays for college? Crowding out on campus. Working Paper, Massachusetts Institute of Technology.

[14] Cameron S.V. & Heckman, J. J. (1998). Life cycle schooling and dynamic selection bias: Models and evidence for five cohorts of American males. *The Journal of Political Economy, 106(2).* pp. 262-333.

[15] Cameron S.V. & Heckman, J. J. (1999).The dynamics of educational attainment for blacks, hispanics and whites (Working Paper 7249). Cambridge, MA: National Bureau of Economic Research.

[16] Card, D. & Lemieux, T. (2001). Can falling supply explain the rising return to college for younger men? A cohort-based analysis (Working Paper 7655). Cambridge, MA: National Bureau of Economic Research.

[17] Carruthers, C. K., Fox, W. F., & Jepsen, C. (2018). Promise kept? Free community college and attainment in Tennessee. Working Paper.

[18] Castleman, B. L., & 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(4),* pp. 1023-1073.

[19] Clotfelter, C. T., Hemelt, S. W., & Ladd, H. F. (2018). Multifaceted aid for low-income students and college outcomes: Evidence from North Carolina. *Economic Inquiry, 56(1),* pp. 278-303.

[20] Cornwell, C., Mustard, D. B., & Sridhar, D. J. (2006). The enrollment effects of merit-based financial aid: Evidence from Georgia's HOPE program. *Journal of Labor Economics, 24(4),* pp. 761-786.

[21] Deming, D., & Dynarski, S. (2009). Into college, out of poverty? Policies to increase the postsecondary attainment of the poor. NBER Working Paper, 15387.

[22] Deming, D., & Walters, C. (2017). The impacts of price and spending subsidies on US post-secondary attainment. NBER Working Paper, 23736.

[23] Denning, J. T. (2017). College on the cheap: Consequences of community college tuition reductions. *American Economic Journal: Economic Policy, 9(2),* pp. 155-88.

[24] Dynarski, S. (2000). Hope for whom? Financial aid for the middle class and its impact on

college attendance (No. w7756). National Bureau of Economic Research.

[25] Dynarski, S. M. (2003). Does aid matter? Measuring the effect of student aid on college attendance and completion. *American Economic Review, 93(1),* pp. 279-288.

[26] Dynarski, S. (2004). The new merit aid. In College choices: The economics of where to go, when to go, and how to pay for it (pp. 63-100). University of Chicago Press.

[27] Dynarski, S. (2008). Building the stock of college-educated labor. *Journal of human resources, 43(3),* pp. 576-610.

[28] Dynarski, S., & Scott-Clayton, J. (2013). Financial aid policy: Lessons from research (No. w18710). National Bureau of Economic Research.

[29] Dynarski, S., Libassi, C. J., Michelmore, K., & Owen, S. (2018). Closing the gap: The effect of a targeted, tuition-free promise on college choices of high-achieving, low-income students (No. w25349). National Bureau of Economic Research.

[30] Eckstein, Z., & Wolpin, K. I. (1998). Youth employment and academic performance in high school. CEPR Discussion Paper Series No. 1861

[31] Evans, W. N., Kearney, M. S., Perry, B. C., & Sullivan, J. X. (2017). Increasing community college completion rates among low-income students: Evidence from a randomized controlled trial evaluation of a case management intervention (No. w24150). National Bureau of Economic Research.

[32] Ferreyra, M., Avitabile, C., Botero Alvarez, J., Haimovich Paz, F., & Urzua, S. (2017). At a crossroads: higher education in Latin America and the Caribbean. The World Bank.

[33] Garriga, C. & M.P. Keightley (2007). A general equilibrium theory of college with education subsidies, in-school labor supply, and borrowing constraints. Federal Reserve Bank of St. Louis Working Paper 2007-051A.

[34] Gurantz, O. (2020). What does free community college buy? Early impacts from the Oregon Promise. *Journal of Policy Analysis and Management, 39(1),* pp. 11-35.

[35] Hai, R., & J.J. Heckman (2017). Inequality in human capital and endogenous credit constraints. NBER Working Paper No. 22999.

[36] Heckman, J. J., Lochner, L. & Taber, C. (1998). General equilibrium treatment effects: A study of tuition policy. *The American Economic Review, 88(2)*, pp. 381-386.

[37] Hendricks, L. & Leukhina, O. (2017). How risky is college investment? *Review of Economic Dynamics, 26,* pp. 140-163.

[38] Hendricks, L., & Leukhina, O. (2018). The return to college: selection and dropout risk. *International Economic Review, 59(3),* pp. 1077-1102.

[39] Hoxby, C., & Turner, S. (2013). Expanding college opportunities for high-achieving, low income students. Stanford Institute for Economic Policy Research Discussion Paper, 12, 014.

[40] ICETEX (2010). Informe de gestión 2003-2010. Instituto Colombiano de Crédito Educativo y Estudios Técnicos en el Exterior (ICETEX).

[41] Katz, L. F., & Murphy, K. M. (1992). Changes in relative wages, 1963-1987: supply and demand factors. *The Quarterly Journal of Economics, 107(1),* pp. 35-78.

[42] Keane, M. P. & Wolpin, K. I. (2001). The effect of parental transfers and borrowing constraints on educational attainment. *International Economic Review*, *42(4),* pp. 1051-1103.

[43] Keane, M. P. (2002). Financial aid, borrowing constraints, and college attendance: Evidence from structural estimates. *The American Economic Review*, *92(2)*, pp. 293-297.

[44] Levhari, D. & Weiss, Y (1974). The effect of risk on the investment in human capital. *The*

American Economic Review, 64(6), pp. 950-963.

[45] Lochner, L.J. & A. Monge-Naranjo (2011). The nature of credit constraints and human capital. *American Economic Review, 101(6),* pp. 2487-2529.

[46] Londoño, J., Rodriguez, C., & Sanchez, 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(2),* pp. 193-227.

[47] Long, B. T. (2008). What is known about the impact of financial aid? Implications for policy. An NCPR Working Paper. National Center for Post-secondary Research.

[48] Mayer, A., Patel, R., Rudd, T., & Ratledge, A. (2015). Designing scholarships to improve college success: Final report on the Performance-based scholarship demonstration. New York: MDRC.

[49] Murphy, R., S. Clayton, & G. Wyness (2019). The end of free college in England: Implications for quality, enrollment, and equity. *Economics of Education Review, 71,* pp. 7-22.

[50] OECD (2018). Education at a glance 2018: OECD indicators. Organisation for Economic Co-operation and Development.

[51] Page, L. C., & Scott-Clayton, J. (2016). Improving college access in the United States: Barriers and policy responses. Economics of Education Review, 51, 4-22.

[52] Oreopoulos, P. & U. Petronijevic (2019). The remarkable unresponsiveness of college students to nudging and what we can learn from it. NBER Working Paper No. 26059.

[53] Ozdagli, A. K., & Trachter, N. (2011). On the distribution of college dropouts: Household wealth and uninsurable idiosyncratic risk. Working Papers No. 11-8.

[54] Restuccia, D. & Urrutia, C. (2004). Intergenerational persistence of earnings: The role of early and college education. *American Economic Review, 94(5)*, pp. 1354-1378.

[55] Scott-Clayton, J. (2011). On money and motivation a quasi-experimental analysis of financial incentives for college achievement. *Journal of Human resources, 46(3),* pp. 614-646.

[56] Scott-Clayton, J., & Zafar, B. (2016). Financial aid, debt management, and socioeconomic outcomes: Post-college effects of merit-based aid. Working Paper 22574, National Bureau of Economic Research.

[57] Scrivener, S., Weiss, M. J., Ratledge, A., Rudd, T., Sommo, C., & Fresques, H. (2015). Doubling graduation rates: Three-year effects of CUNY's Accelerated Study in Associate Programs (ASAP) for developmental education students. New York: MDRC.

[58] Solis, A. (2017). Credit access and college enrollment. *Journal of Political Economy, 125(2),* pp. 562-622.

[59] Sommo, C., Cullinan, D., Manno, M., Blake, S., & Alonzo, E. (2018). Doubling graduation rates in a new state: Two-year findings from the ASAP Ohio demonstration. New York: MDRC, Policy Brief.

[60] Stange, K.M. (2012). An empirical investigation of the option value of college enrollment. *American Economic Journal: Applied Economics 4 (1),* pp. 49-84

[61] Stinebrickner, R. & T. R. Stinebrickner (2004). Time-Use and college outcomes. *Journal of Econometrics, 121(1-2),* pp 243-269.

[62] Stinebrickner, R. & T. R. Stinebrickner (2014). Academic performance and college dropout: Using longitudinal expectations data to estimate a learning model. *Journal of Labor Economics, 32(3),* pp. 601 - 644.

[63] Trachter, N. (2015). Stepping stone and option value in a model of post-secondary education. *Quantitative Economics, 6(1),* pp. 223-256.

[64] World Bank (2016). Poverty and shared prosperity 2016: Taking on inequality. World Bank Publications.

[65] Zamarro, G., C. Hitt, & I. Mendez (2019). When students don't care: Reexamining international differences in achievement and student Effort. *Journal of Human Capital, 13(2),* pp. 519-552.

For Online Publication: Appendices

A Appendix

A.1 Tiers of cumulative classes completed

Table A.1: Tiers: Lower Bound by Year.

Tier	Year 1	Year 2	Year 3	Year 4	Year 5
Tier 1	19	38	57	76	95
Tier 2	17	34	51	68	85
Tier 3	13	26	39	52	65
Tier 4	0	0	0	0	0
Recommended cumulative classes	20	40	60	80	100

Source: Own classification.
Notes: This table shows the lower bound for each tier, by year, expressed in number of cumulative classes completed. For a given year, tiers are defined by the percent of the year's recommended cumulative classes that have been completed by the student. Tier 1: 95 percent or more; Tier 2: (85, 95] percent; Tier 3: (65, 85] percent; Tier 4: 65 percent or less. For example, in year 2 the lower bounds (expressed in cumulative classes completed) are as follows: 38 classes = 0.95×40 for Tier 1; 34 classes = 0.85×40 for Tier 2; and 26 classes = 0.65×40 for Tier 3.

A.2 Wages in Colombia by Educational Attainment and Age

Table A.2: Average hourly wage by age bracket and educational attainment.

	Age bracket			
	18-60	18-22	23-35	36-60
College graduates	6,308	3,636	5,305	7,171
HS graduates	2,424	1,845	2,213	2,864
College dropouts; completed 1 year or less	3,091	2,349	2,897	3,619
College dropouts, completed 2 years or more	3,824	2,459	3,340	4,451

Source: Household surveys for Colombia (SEDLAC); year 2005.
Notes: Wages are expressed in Colombian pesos (COP) of 2005. Calculations include males and females who work. Attainment reflects an individual's highest completed level of schooling.

A.3 Model timeline

The figure below summarizes the timing of events and student decisions:

Figure A.1: Summary of Timing of Events and Individuals' Decisions

Enrollment Decision ($t = 0$) $\begin{cases} \text{College} \\ \text{Labor force} \end{cases}$

College ($t = 1$:4)

State	i. Choice	ii. Classes	iii. Accumulated	iv. Dropout	v. Payoffs
(t, h_{t-1}, θ, y)	e_t	classes	Completed	shock d_t^{drop}	$V^{drop}(t+1)$
		shock z_t	completed h_t		$V^{coll}(t+1, h_t, \theta, y)$

College ($t = 5$:7)

State	i. Choice	ii. Classes	iii. Accumulated	iv. Dropout	v. Payoffs
(t, h_{t-1}, θ, y)	e_t	completed	classes	shock d_t^{drop}	$V^{drop}(t+1)$
		shock z_t	completed h_t		$V^{coll}(t+1, h_t, \theta, y)$
					$V^{grad}(t+1)$

College ($t = 8$)

State	i. Choice	ii. Classes	iii. Accumulated	iv. Dropout	v. Payoffs
$(8, h_7, \theta, y)$	e_8	completed	classes	shock d_8^{drop}	$V^{drop}(9)$
		shock z_8	completed h_8		$V^{grad}(9)$

A.4 Computation and Estimation

In this appendix we first summarize the computational solution of the model, and then provide detail on its main steps.

A.4.1 Solving the model: a summary

Recall that the state vector is (t, h_{t-1}, θ, y).We discretize the state space for a total of 40,400 points. We simulate $N = 100{,}000$ high school graduates from the empirical distribution of ability and income, $\Phi(\theta, y)$.[45] From a given student type, a fraction of the simulated high school graduates receives a college enrollment shock equal to 1 and enrolls in college (thus becoming the "simulated college students"); the fraction is equal to the type's observed college enrollment rate (see Appendix A.4).[46]

To compute the model's predictions for a given value of Θ, the algorithm proceeds as follows:

1. For each point in the state space, use backward induction to solve for the sequence of optimal efforts (i.e., the policy function) and the value function, $e^*(t, h_{t-1}, \theta, y)$ and $V^{coll}(t, h_{t-1}, \theta, y)$ respectively.

2. For each simulated college student, and for every year she is enrolled, combine her optimal effort with the corresponding ν_t shock to determine the year's classes completed and

[45] For each simulated high school graduate, we draw one *i.i.d.* shock per possible year, $\{\nu_{it}\}_{t=1}^{8}$. For a given simulated high school graduate, these shocks are the same for parameter vector values, and at baseline and counterfactuals.

[46] In the free-college counterfactuals, the fraction of enrolled students is given by equation ((9)), where the value of college changes in response to free college.

probability of dropping out, $\tilde{p}^d(t, h_t, \theta, y)$. Draw the binary dropout shock; the shock is equal to 1 with probability $\tilde{p}^d(t, h_t, \theta, y)$.

3. Based on step 2, aggregate the simulated dropout decisions to obtain a predicted dropout rate for each of the 400 (t, θ, y)-combinations.

4. Find the vector $\boldsymbol{\delta}$ that minimizes the distance between the predicted and observed dropout rate for each (t, θ, y)-combination, using the contraction mapping algorithm described in Appendix A.4.3.

5. By comparing the value of going and not going to college for each type, $V^{coll}(1, 0, \theta, y)$ and V^{hs}, respectively, find the type-specific college enrollment shock, ξ_j, that renders the type indifferent between going and not going to college. Further details are provided in Appendix A.4.4.

Solving steps 1-5 of the dynamic optimization problem for $100,000$ simulated high school graduates and $40,400$ states takes approximately 8 minutes in a 1.4 GHz Intel Core i5 processor. Since the model does not have a closed-form solution, in estimation the problem must be solved anew for each value of Θ. Note that the estimation of $\boldsymbol{\delta}$ and $\boldsymbol{\xi}$ is nested within the model solution for a given of value of Θ, in the spirit of Berry et al (1995), as described in Appendix A.4.3.

A.4.2 Further details on solving the model

The solution of the dynamic problem and computation of predicted outcomes for a given value of Θ involve three steps: calculating the value of working by educational attainment, solving for the policy and value functions, and simulating college students.

Calculating the value of working by educational attainment Since we solve the student's dynamic programming problem by backward induction, we begin by calculating the final value of the individuals' finite horizon problem. We calculate the value of future discounted payoffs of working as a college graduate, high school graduate, college dropout with one year of college, and college dropout with two or more years of college. In the timing of the model, $t = 1$ when the individual is 18 years old, and either starts college or joins the labor force as a high school graduate. The value of working as a high school graduate since $t = 1$ onward is

$$V^{hs} = \sum_{t=1}^{L} \beta^{t-1} u(w_t^{hs}), \tag{23}$$

where w_t^{hs} is the average wage for a high school graduate in year t, and L is retirement age (65 years old, or $L = 48$). For this and the other educational attainments, we allow the wage to vary over time to incorporate returns to experience (which accrue after age 35, that is for $t > 17$). Similarly, the value of working as a college dropout who has completed n years of college is

$$V^{drop}(n+1) = \sum_{t=n+1}^{L} \beta^{t-n-1} u(w_t^{drop}), \tag{24}$$

where w_t^{drop} is the wage an individual who dropped out of college receives in year t. Finally, the value of working as a college graduate who took n years to graduate is:

$$V^{grad}(n+1) = \sum_{t=n+1}^{L} \beta^{t-n-1} u(w_t^{grad}), \tag{25}$$

Solving for the policy and value functions Since the state, (t, h_{t-1}, θ, y), has four dimensions, we build two-four dimensional grids - one for the policy function, $e^*(t, h_{t-1}, \theta, y)$, which contains the optimal effort choice for each state, and another for the optimal payoffs, $V^{coll}(t, h_{t-1}, \theta, y)$ associated with $e^*(t, h_{t-1}, \theta, y)$. The grid includes 8 points (years) for t, 101 points for h (to represent $0, 1, 2, \ldots, 100$ credits completed), ten points (ability deciles) for θ, and five points for y for a total of 40,400 points. Each ability decile is represented by its median. In other words, θ takes on values corresponding to the 5th, 15th, ... 95th percentile of the ability distribution.

We use backward induction to solve the Bellman equation for each period. Starting from the last year of college, $t = 8$, when students must either graduate or drop out, the Bellman equation is:

$$V^{coll}(8, h_7, \theta, y) = U(c_8, e_8, \theta) + \beta E_z \Big[\Pr\left(h_8 \geq h^{grad}\right) V^{grad}(9) + \Pr\left(h_8 < h^{grad}\right) V^{drop}(9) \Big]. \tag{26}$$

By choosing e_8 to maximize V^{coll} for every state, $(8, h_7, \theta, y)$, we find both the policy function and the value function, $e^*(8, h_7, \theta, y)$ and $V^{coll}(8, h_7, \theta, y)$.

Moving backwards to $t = 7$, we proceed analogously

$$\begin{aligned} V^{coll}(7, h_6, \theta, y) = U(c_7, e_7, \theta) \ + \beta E_z \Bigg[& \Pr\left(h_7 \geq h^{grad}\right) V^{grad}(8) + \\ & \Pr\left(h_7 < h^{grad}\right) \Big(\tilde{p}^d(7, h_7, \theta, y) V^{drop}(8) + \\ & \left(1 - \tilde{p}^d(7, h_7, \theta, y)\right) V^{coll}(8, h_7, \theta, y) \Big) \Bigg], \end{aligned} \tag{27}$$

where all outcomes $V^{grad}(8)$, $V^{drop}(8)$ and $V^{grad}(8, \cdot)$ are already known. We continue this procedure for $t = 6, \ldots, 1$ in order to complete the calculation of $e^*(t, h_{t-1}, \theta, y)$ and $V^{coll}(t, h_{t-1}, \theta, y)$ for all possible states.

We use the resulting policy function, $e^*(t, h_{t-1}, \theta, y)$, whenever we need to compute a student's optimal effort during estimation or counterfactuals. In addition, we use the resulting value function at $t = 1$, that is, $V^{coll}(1, 0, \theta, y)$, to calculate the value of enrolling in college, which is to be compared with the value of joining the workforce as a high school graduate, V^{hs}, following equation (8).

Simulating college students We simulate $N = 100{,}000$ high school graduates. We make draws from the joint distribution of ability of parental transfers of high school graduates in 2005, $\Phi(\theta_j, y_j)$. Recall that a student type j is given by a (θ_j, y_j) combination. We have $J = 50$ types.

For each type, let $P^{coll}(\theta_j, y_j)$, equal to the actual, observed share of individuals *of that type* that enrolls in college. Note that $P^{coll}(\theta_j, y_j)$ varies across types, as illustrated by Table 1. Consider individual i who belongs to type j. For each simulated individual, we draw a binary variable, d_i^{enr}, to determine whether the student goes to college or not. More specifically,

$$d_i^{enr} = \begin{cases} 1, & i \text{ goes to college, with probability } P^{coll}(\theta_j, y_j) \\ 0, & i \text{ does not goes to college, with probability } 1 - P^{coll}(\theta_j, y_j) \end{cases} \tag{28}$$

Simulated students who receive $d_i^{enr} = 1$ are those who enroll in college. In other words, the proportion of simulated students of a given type who receive $d_i^{enr} = 1$ is the same as the proportion of actual students of that type who enroll in college.[47] For students who do not enroll in college, the value function is V^{hs}. For those who enroll in college, we simulate classes completed and dropout shocks as described below.

For $t = 1$, we use the policy function $e^*(1, 0, \theta_j, y_j)$ corresponding to every student type j. Since all students start at $t = 1$ with zero accumulated credits, $h_0 = 0$, the policy function assigns the same effort to all students of a given type, j. Then, we draw the *iid* shock ν_{i1} for each student; this, in turn, yields a value for the z_{i1} shock. The combination of the student's ability, effort, and z_{i1} shock yields the number of completed credits, h_{i1}. Because of the z shock, individuals of a given type attain different values of h by the end of the first period.

For student i, we use the realized h_{it} to establish whether the student drops out before the second period. The student receives a draw of the binary variable d_{it}^{drop}; if the draw is equal to 1, she drops out. The probability of $d_{it}^{drop} = 1$ is a function of student type, year, and average performance up to (and including) the corresponding year:

$$d_{it}^{drop} = \begin{cases} 1, & i \text{ drops out of college, with probability } \tilde{p}^d(t, h_{it}, \theta_j, y_j) \\ 0, & i \text{ continues in college, with probability } 1 - \tilde{p}^d(t, h_{it}, \theta_j, y_j) \end{cases} \tag{29}$$

where $\tilde{p}^d$ is defined as in (14). The binary variable d_{it}^{grad} indicates whether a student graduates; the graduation requirement is $h^{grad} = 98$.Whenever $t \geq 5$ and $h_{it} \geq h^{grad}$, we set $d_{it}^{grad} = 1$ and $d_{it}^{drop} = 0$. In other words, a student in year 5 or beyond who has completed at least 98 credits is no longer subject to the risk of dropping out, and automatically graduates. In addition, a student who reaches $t = 8$ without having completed at least 98 classes cannot graduate ($d_{it}^{grad} = 0$) and must drop out ($d_{it}^{drop} = 1$).

The final outcome of the simulation is a "dataset" with $N = 100{,}000$ simulated high school graduates, some of whom enroll in college. For those who enroll, we obtain their number of classes completed by year, final outcome (graduation or drop out), along with the period in which they either drop out or graduate. This dataset mimics our observed student-level administrative data.

[47] For a large number of simulations such as ours, this is asymptotically equivalent to simply assigning $d_i^{enr} = 1$ to a fraction of simulated students from a given type equal to the type's observed enrollment rate.

A.4.3 Estimation of fixed effects in the dropout probability

We now describe the estimation of the time- and type-specific fixed effects that enter in the dropout probability, $\delta(t,\theta,y)$. This estimation is nested within the estimation Θ,as it must take place for every possible value of Θ.

From the simulation of college students described above, we compute the predicted dropout rates by year, ability quintile, and income. This is the predicted fraction of students of type j who drop out in every t, or $\hat{p}_{jt}^{drop} = f^d(\delta(t,\theta,y);\Theta)$. We compare these predicted rates with the observed ones, denoted by p_{jt}^{drop}, and compute a measure of distance between them.

For each pair of predicted and observed dropout rates, we calculate the fixed effects $\delta(t,\theta,y)$ that minimize this distance. We do so through an iterative contraction mapping, in the spirit of Berry et al (1995). While Berry et al (1995) uses a contraction mapping to find the unobserved product characteristics that make predicted market shares for each product equal to their observed counterparts, we search for the time- and type-fixed effects that make observed dropout rates as close as possible to their observed counterparts for each period and student type.

Formally, we use a contraction mapping algorithm to find the vector $\boldsymbol{\delta} = [\delta(t,\theta,y)]_{J(8)\times 1}$ that fulfills the following condition:

$$\|\mathbf{f}^d(\boldsymbol{\delta};\Theta) - \mathbf{p}^{drop}\| \leq \epsilon^d, \tag{30}$$

where ϵ^d is our chosen tolerance level. Below are the algorithm steps; recall that they are conditional on a given parameter point, Θ:

1. Define an initial guess for the fixed effects vector, $\boldsymbol{\delta}^{(0)}$.

2. Solve the dynamic optimization problem (see Appendix A.4.)

3. With the resultant panel data of simulated outcomes, compute the predicted vector of drop out rates $\mathbf{f}^d(\boldsymbol{\delta}^{(0)};\Theta) = \hat{\mathbf{p}}^{drop}$.

4. Using the observed drop out rates, compute the components of the updated fixed effects vector, $\boldsymbol{\delta}^{(1)}$, as follows:

$$\delta^{(1)}(t,\theta_j,y_j) = \ln\left(\frac{p_{jt}^{drop}}{f^d(\delta^{(0)}(t,\theta_j,y_j);\Theta)}\right). \tag{31}$$

5. Using $\boldsymbol{\delta}^{(1)}$ as the new initial guess, repeat steps 1 through 4 until condition (30) is satisfied or a predetermined maximum number iterations is reached.

The algorithm may not be able to meet (30) due to non-convexities in the model. For instance, in $t = 8$, dropping out in the model is a deterministic function of the number of credits completed, whereas in the data we observe some individuals graduate without having completed all credits (due, perhaps, to measurement errors in number of classes completed.) Another non-convexity arises because, in the Berymodel, the student must meet a minimum number of cumulative credits per period, h_t^{drop}. If she does not complete them, she must drop out. In the data, in contrast, some students remain enrolled even though they do not meet that requirement.

A.4.4 Recovering type-specific preferences for college enrollment

Recall that $\xi_j = \tilde{\xi}(\theta_j, y_j)$ is the type-specific unobserved preference shock for enrolling in college. For a given value of Θ, we recover it as follows. From the computation of equilibrium algorithm described in Section A.4 we compute the value function, $V^{coll}(\cdot)$, for every state (t, h_t, θ, y), which allows us to compare the value of going to college, $V^{coll}(1, 0, \theta_j, y_j)$, with the value of working as a high school graduate, V^{hs}. Thus, the value of ξ_j is such that the predicted probability of enrolling to college is equal to the observed one. Under the assumption that $\sigma_\epsilon = 1$ (see Section 5.1), we solve for ξ_j in equation (9):

$$\xi_j = \ln\left(\frac{P^{coll}(\theta_j, y_j)}{1 - P^{coll}(\theta_j, y_j)}\right) - (V^{coll}(1, 0, \theta_j, y_j) - V^{hs}). \tag{32}$$

During counterfactuals, we hold these shocks at their baseline values, since they are preference parameters.

A.5 Goodness of fit: additional evidence

In this section we provide additional evidence regarding the model's fit of the data.

Table A.3: Goodness of Fit: Cumulative Classes Completed by Year 1.

Classification	Observed values	Predicted values
Ability quintile		
1	13.9	15.7
2	14.5	16.3
3	15.3	16.5
4	15.6	16.0
5	16.6	17.4
On-time graduate	20.5	20.7
Late graduate	18.0	19.4
Dropout later	15.8	14.9
Dropout this year	10.9	13.2
Total	15.8	16.7

Source: SPADIES for observed data; model simulations for predicted data.
Notes: Values are expressed in percentages (%).

Table A.4: Goodness of Fit: Cumulative Classes Completed by Year 2.

Classification	Observed values	Predicted values
Ability quintile		
1	31.7	32.7
2	32.3	33.3
3	33	33.6
4	33.6	33.1
5	34.5	35.5
On-time graduate	41.2	41.5
Late graduate	35.4	37.8
Dropout later	30.4	29.0
Dropout this year	23.8	21.7
Total	33.8	34.2

Source: SPADIES for observed data; model simulations for predicted data.
Notes: Values are expressed in percentages (%).

Table A.5: Goodness of Fit: Cumulative Classes Completed by Year 3.

Classification	Observed values	Predicted values
Ability quintile		
1	49.3	51.4
2	49.8	52.0
3	51.2	52.8
4	52.0	52.3
5	52.9	54.3
On-time graduate	62.0	62.0
Late graduate	52.5	55.4
Dropout later	44.5	42.9
Dropout this year	36.4	37.5
Total	52.1	53.2

Source: SPADIES for observed data; model simulations for predicted data.
Notes: Values are expressed in percentages (%).

Table A.6: Goodness of Fit: Cumulative Classes Completed by Year 4.

Classification	Observed values	Predicted values
Ability quintile		
1	67.2	68.3
2	68.2	69.9
3	69.2	70.6
4	70.3	70.6
5	71.8	72.0
On-time graduate	82.8	82.1
Late graduate	70.1	72.0
Dropout later	59.3	53.9
Dropout this year	46.4	55.6
Total	70.6	71.0

Source: SPADIES for observed data; model simulations for predicted data.
Notes: Values are expressed in percentages (%).

Table A.7: Goodness of Fit: Cumulative Classes Completed by Year 5.

Classification	Observed values	Predicted values
Ability quintile		
1	83.9	83.1
2	85.7	85.7
3	85.9	86.6
4	87.1	86.9
5	88.6	87.6
On-time graduate	99.0	99.5
Late graduate	86.6	87.1
Dropout later	72.4	64.4
Dropout this year	69.7	64.5
Total	87.5	86.9

Source: SPADIES for observed data; model simulations for predicted data.
Notes: Values are expressed in percentages (%).

Table A.8: Goodness of Fit: Cumulative Classes Completed by Year 6.

Classification	Observed values	Predicted values
Ability quintile		
1	88.8	90.4
2	91.3	92.0
3	91.5	92.9
4	92.7	92.6
5	94.0	93.3
Late graduate	95.6	96.8
Dropout later	78.5	68.8
Dropout this year	84.5	77.3
Total	93.0	92.8

Source: SPADIES for observed data; model simulations for predicted data.
Notes: Values are expressed in percentages (%).

Table A.9: Goodness of Fit: Cumulative Classes Completed by Year 7.

Classification	Observed values	Predicted values
Ability quintile		
1	93.6	91.3
2	95.3	92.3
3	95.8	93.1
4	96.6	92.3
5	97.6	93.5
Late graduate	98.3	98.4
Dropout later	82.7	75.7
Dropout this year	86.2	76.7
Total	96.8	92.9

Source: SPADIES for observed data; model simulations for predicted data.
Notes: Values are expressed in percentages (%).

Table A.10: Goodness of Fit: Cumulative Classes Completed by Year 8.

Classification	Observed values	Predicted values
Ability quintile		
1	96.0	92.2
2	97.2	91.6
3	98.0	94.3
4	98.4	90.5
5	99.0	91.4
Late graduate	98.9	99.3
Dropout this year	88.2	81.5
Total	98.5	91.6

Source: SPADIES for observed data; model simulations for predicted data.
Notes: Values are expressed in percentages (%).

Figure A.2: Goodness of Fit: Cumulative Classes Completed by Ability Decile and Year.

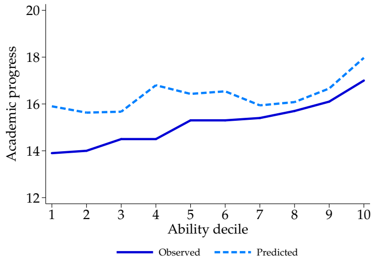

a. Year 1

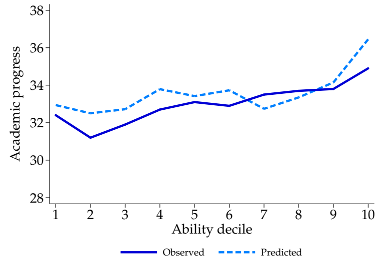

b. Year 2

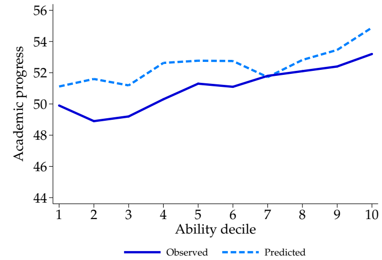

c. Year 3

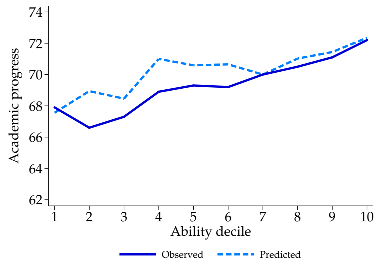

d. Year 4

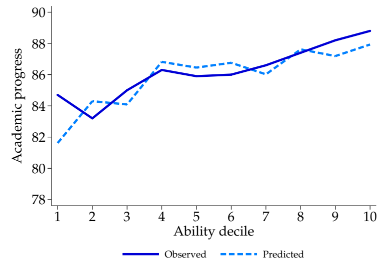

e. Year 5

Source: SPADIES for observed data; model simulations for predicted data.
Note: For each year, the panels depict observed and predicted cumulative number of classes completed by ability decile. Figures correspond to students who begin each year.

Figure A.3: Goodness of Fit: Target Number of Classes by Ability Decile and Year

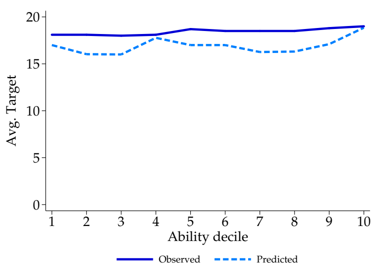

a. Year 1

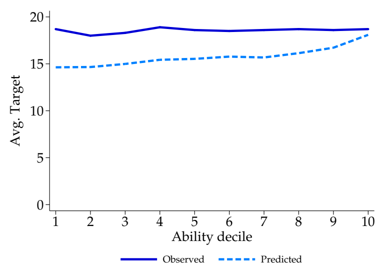

b. Year 2

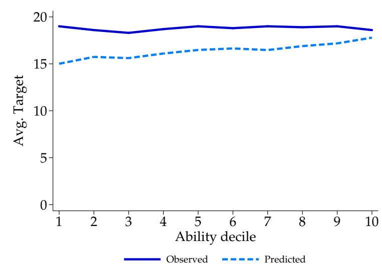

c. Year 3

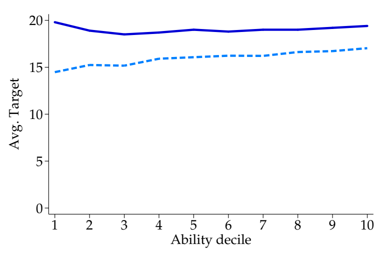

d. Year 4

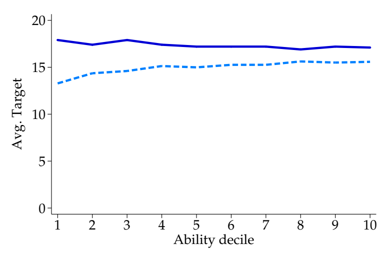

e. Year 5

Source: SPADIES for observed values; model's own simulations for predicted values.
Note: Observed values correspond to the average number of classes for which students register; predicted values correspond to target number of classes as defined in the model.

A.6 Anticipated Risk

Table A.11: Determinants of Effort.

	Dependent variable ln(Optimal Effort)
ln(average classes completed)	-0.060^{***}
	(0.013)
Income 1-2 MW	-0.001
	(0.003)
Income 2-3 MW	-0.008^{**}
	(0.003)
Income 3-5 MW	0.001
	(0.003)
Income 5+ MW	-0.031^{***}
	(0.003)
Year=2	0.090^{**}
	(0.039)
Year=3	0.099^{***}
	(0.037)
Year=4	0.036
	(0.038)
Year=5	-0.014
	(0.037)
Year=6	-0.256^{***}
	(0.037)
Year=7	-0.289^{***}
	(0.034)
Year=8	-0.321^{***}
	(0.039)
Ability Q2	0.007
	(0.006)
Ability Q3	0.035^{***}
	(0.005)
Ability Q4	0.055^{***}
	(0.005)
Ability Q5	0.107^{***}
	(0.005)
Constant	0.707^{***}
	(0.005)
R^2	0.226
Num. obs.	116,761

Source: First-stage of 2SLS estimation based on model's predicted baseline values.
Notes: The dependent variable is ln(optimal effort), or $\ln(e_t^*)$. An observation is a student-year. Included students are "existing" in universal and performance-based free-college counterfactuals. Upper 5% tail of risk has been trimmed. Independent variables are state variables at t. Ability is θ; income is y; year is t; average classes completed is $\bar{h}_{t-1}$. The regression includes year fixed effects (not shown). Standard errors (in parentheses) are clustered by student. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.