FEDERAL RESERVE BANK of ST. LOUIS

ECONOMIC RESEARCH

FEDERAL RESERVE BANK OF ST. LOUIS

WORKING PAPER SERIES

Risky College Admissions and the (Mis)Allocation of Talent

Authors	Ricardo Marto; Yaacov Wittman
Working Paper Number	2026-026A
Release Date	November 2026
Citable Link	https://doi.org/10.20955/wp.2026.026

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

The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment.

RISKY COLLEGE ADMISSIONS AND THE (MIS)ALLOCATION OF TALENT

RICARDO MARTO AND YAACOV WITTMAN
Federal Reserve Bank of St. Louis

Why do high-achieving, low-income students not attend selective colleges? Leveraging restricted-use data, we show the gap arises at the application stage: lower-income students apply much less, while admission and enrollment differences are modest. We develop and estimate an equilibrium model in which heterogeneous students and colleges each hold noisy assessments of ability, borrowing-constrained students choose application portfolios under admission and financial-aid risk, and colleges with scarce seats compete through admission standards and tuition schedules. Because only the highest-ability low-income students apply, colleges offer them generous aid. Expanding aid raises tuition and standards, not enrollment. Better information, not subsidies, puts talent where it belongs.

KEYWORDS: college admissions, college applications, college enrollment, college market, credit constraints, financial aid, High School Longitudinal Study, information asymmetry, Pell grants, selective colleges, sequential search, sorting, tuition discrimination.

JEL CLASSIFICATION: E24, D82, I22, I23, I24.

Ricardo Marto and Yaacov Wittman: ricardo.marto@stls.frb.org; ywittman@sas.upenn.edu.

We thank Zachary Bleemer, Damien Capelle, Ciprian Domnisoru, Jesús Fernández-Villaverde, Grey Gordon, Jeremy Greenwood, Jonathan Heathcote, Lutz Hendricks, Dirk Krueger, Oksana Leukhina, Luca Mazzone, Alex Monge-Naranjo, Anusha Nath, Víctor Ríos-Rull, Todd Schoellman, and Holger Sieg for their insightful comments and suggestions. We are also grateful to seminar and conference participants at numerous institutions. The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

1. Introduction

How an economy allocates its talent shapes its stock of human capital and, ultimately, its aggregate productivity (Restuccia and Rogerson (2008); Hsieh et al. (2019)). In the United States, a central stage of that allocation runs through the college market. Selective colleges concentrate instructional resources and high-ability peers, and produce a sizable earnings premium for those they enroll. et attendance at these institutions varies sharply with parental income, even among students of similar measured ability. Students from the bottom quartile of family income are much less likely to complete college (Bailey and Dynarski (2011)) and far less likely to attend selective colleges (Chetty et al. (2020)) than students from the top. This paper asks how the college admissions system allocates talent across colleges. Which lever improves that allocation once the college market responds: the financial aid that subsidizes access, or the information on which admissions and pricing rest?

Building on Hoxby and Avery (2014), we revisit the empirical evidence on application and enrollment decisions of high-achieving low-income students. We leverage restricted-use data from the High School Longitudinal Study of 2009 (HSLS) that follow the same student through every stage of the admissions process (application, admission, and enrollment) so that the income gap in selective-college attendance can be decomposed stage by stage. To our knowledge, these data have not been used to study the three margins jointly.[1] The HSLS is a representative survey of high-school students followed through their postsecondary years, which includes detailed family, school, and financial characteristics. Students graduating high school in 2013 report the colleges they applied to, the admission decision at each, and where they enrolled.

Consider enrollment by parental income among top scorers. About 70% of students in the top SAT decile whose parents earn more than $195,000 enroll in a selective college, against 36% of those whose parents earn between $35,000 and $55,000. High-ability low-income students may be missing from selective colleges because they never applied, because they applied and were rejected, or because they were admitted and chose not to enroll. The HSLS lets us document each step, controlling for demographics, family background, and high-school characteristics.

First, students submit few applications (2.4 on average and about 5 in the top decile of the SAT score distribution) and lower income students apply much less intensively to selective colleges. Controlling for a rich set of student, family, and high-school characteristics, a student in the 90th percentile of SAT whose parental income is between $35,000 and $55,000 has a 16 percentage point (p.p.) lower probability of applying to a selective college than a peer whose parents earn between $195,000 and $215,000 (31% versus 47%). Put differently, at the

[1] Hoxby and Avery (2014) use College Board records for every SAT and ACT taker and proxy applications with SAT score-sends, but do not include admission outcomes.

application stage family income substitutes for measured achievement at a steep rate, with that income gap worth roughly 130 SAT points (two-thirds of a standard deviation).

Second, differences in admission rates by income exist but are modest. Among applicants to selective colleges, the probability of admission rises by about 8 p.p. between those income groups (78% versus 86% at the 90th percentile of SAT). Third, among students admitted to both a selective and a nonselective college, lower income students are somewhat less likely to choose the selective one (but the difference is not statistically significant once we control for demographics). Almost all who turn down the selective offer enroll in a nonselective college instead. Finally, more than half of parents with income below $75,000 expected, a year before enrollment, not to be able to afford college, against less than 10% of parents earning $155,000 or more (even when the student is a top scorer).

Motivated by these facts, the second part of the paper presents a novel equilibrium model of the college market used to evaluate the macroeconomic impact of education policies. Students and colleges each hold their own assessment of a student's ability (the student's, from her academic record, and the college's, from the application file) and the two are noisy signals of each other, while parental resources are public knowledge through the FAFSA. Students choose costly application portfolios and which college to attend among the offers they receive. Colleges maximize the value added to their students and compete in admission standards and in tuition schedules that vary with parental income and their own assessment, taking into account that applicants may hold competing offers. Because each side acts on its own assessment, both admissions and financial aid are risky at the time of application.[2]

The model's distinguishing feature is that applications, admission standards, and student-specific prices are determined jointly in equilibrium, so who else applies shapes how much an applicant's file reveals about her ability and how much it reveals determines the aid she is offered. Our central lesson is that this joint determination matters for education policy as enrollment and welfare effects assuming tuition and admission standards fixed can be misleading.

The estimated model accounts quantitatively for the application and enrollment patterns in the data, with realistic tuition schedules that vary across and within colleges by parental income and test scores, and it rationalizes the coexistence of low application rates and generous aid. Selective colleges fund their low-income students generously because only the highest-ability among them apply. The model then serves as a laboratory to study college-access policies, namely expansions of need-based aid and of borrowing limits, cheaper applications, and changes in the correlation between the two sides' assessments. We judge every experiment by

[2]Experimental evidence supports this margin: Dynarski et al. (2021) find that removing the aid uncertainty sharply raises the applications and enrollment high-achieving, low-income students, whereas information without commitment has little effect (Bettinger et al., 2012).

two outcomes: students' welfare, measured in consumption-equivalent terms, and the allocation of talent, measured by a sorting efficiency index that measures the share of the earnings gap between randomly assigned and perfectly sorted college seats that the equilibrium closes.

We first study expansions of the federal Pell Grant program, which raise grant funding for low-income students and extend eligibility toward the middle of the income distribution. A large expansion considerably reduces the concentration of income at selective colleges and benefits high-ability students from the bottom half of the income distribution. But much of the aid is captured by the colleges. Evaluated at a fixed student mix, selective colleges raise the tuition charged to aid recipients by 80 cents for every grant dollar. Admission standards rise and total enrollment falls even as its composition improves. Average welfare declines, with higher-income students losing through the higher tax rate and the increased competition. If colleges' prices and standards were held fixed, the same policy would raise enrollment sharply and deliver sizable welfare gains even when fully financed by taxes.

The remaining experiments turn to the informativeness of applications. In a counterfactual economy in which low-income students apply to selective colleges at the same rate as their higher-income peers, the added low-ability applicants destroy the information that selective colleges rely on to sort. So, interventions that encourage applications without targeting ability dilute financial aid and selective admissions of lower-income students.

Motivated by the SAT-optional movement, we then lower the correlation between students' and colleges' ability assessments. Top-ability students lose 2 to 2.5% of lifetime consumption at every income level as ability-based discounts shrink and the value added of the colleges they attend falls. The policy is a tax on the students who benefit most from college. A scenario that combines this lower correlation with cheaper applications (the two margins of AI-assisted college applications) compounds both losses and pushes the bottom-quartile income share of selective enrollment to zero. Perfect sorting by students' own ability assessments would raise aggregate lifetime earnings by 1.2% and random assignment would lower them by 18%. This sorting efficiency becomes our summary measure of the allocation of talent: Every subsidy-based policy we evaluate lowers sorting efficiency, while better information alone raises it.

A final section extends the model with graduation risk, estimated on completion rates by tier and SAT. Completion risk strengthens the sorting forces at the heart of the paper: perfectly informative assessments now widen the gap between the college tiers rather than closing it. Because students who expect not to graduate stay out when aid expands, aid feeds tuition and spending rather than enrollment, limiting the scale of aid expansions the market can absorb.

Related literature. This paper contributes primarily to the quantitative macroeconomic literature on human capital and education policy (Ionescu (2009), Lochner and Monge-Naranjo (2011), Herrington (2015), Ionescu and Simpson (2016), Krueger and Ludwig (2016), Kotera

and Seshadri (2017), Caucutt and Lochner (2017), Abbott et al. (2019), Colas et al. (2021)). These heterogeneous-agent economies evaluate college policies in general equilibrium but typically treat the market's supply side as passive. Our contribution is to endogenize prices, admissions, and college value added, and to show that the college market's response is first-order for the size and the sign of policy effects.[3] The mechanism also connects to the macroeconomic literature on misallocation (Restuccia and Rogerson (2008); Hsieh et al. (2019)) as the admission system's information frictions distort the assignment of talent and well-intentioned subsidies can amplify the distortion once prices respond.

A fast-growing quantitative literature examines specific margins of this market: financial-aid application frictions (Moschini and Raveendranathan (2026)), equilibrium admissions and affirmative action (Bodoh-Creed and Hickman (2018), Kapor (2020)), student debt and post-college choices (Folch and Mazzone (2026), Morazzoni (2023)), free college and school quality over the life cycle (Krueger et al. (2024), Wright and Zheng (2024)), the long-run evolution of college attainment and ability sorting (Donovan and Herrington (2019), Hendricks et al. (2021)), college access, quality, and intergenerational mobility (Hendricks et al. (2024), Capelle and Matsuda (2026), Wright (2025)), grant uncertainty and dropout risk (Vardishvili (2023)), competition among research universities (Alon et al. (2026)), and the drivers of tuition growth (Cai and Heathcote (2022); Gordon and Hedlund (2016, 2021)).

The framework builds on Epple et al. (2006), Epple et al. (2017), and Fu (2014), who study equilibrium models of the college market with quality-maximizing colleges. Because college quality in this tradition depends on the average ability of the student body, peer effects are built into the technology and the evidence supports them (Sacerdote, 2001, Zimmerman, 2003, Carrell et al., 2009). Our model reconciles the evidence finding little return to selectivity conditional on the set of colleges applied to (Dale and Krueger, 2002, 2014, Mountjoy and Hickman, 2020) and sizable returns for academically marginal students (Hoekstra, 2009, Zimmerman, 2014, Andrews et al., 2020, Bleemer, 2024) through the application margin. The empirical work motivating our question is broad. Chetty et al. (2020) document income segregation within and across U.S. colleges and Hoxby and Avery (2014), Dillon and Smith (2017), Belley and Lochner (2007), and Leukhina (2023) study the enrollment rates of high-ability, low-income students in other data. We explain the application gap as an informed response to admission and financial-aid risk, and evaluate in equilibrium the policies designed to close the gap.

[3] This response is empirically relevant: Lucca et al. (2018) and Turner (2017) provide evidence that college tuition increases in response to expansions in federal financial aid.

2. Empirical Motivation

High-ability, low-income students can end up outside selective colleges at three distinct stages: they may not apply, they may apply and be rejected, or they may be admitted and decide not to enroll. This section documents each stage using student-level data that record, for the same student, where she applied, whether she was admitted, and where she enrolled. We first describe the data (Section 2.1) and then present the evidence on enrollment (Section 2.2), applications (Section 2.3), admissions (Section 2.4), enrollment conditional on admission (Section 2.5), and expected affordability (Section 2.6). Throughout, the emphasis is on how outcomes vary with parental income *holding academic preparedness fixed*, so most figures focus on students in the top three deciles of the SAT distribution. Additional figures, the underlying regression tables, and the definitions of all HSLS variables are collected in Appendix A.

2.1. *Data*

Student-level data. The analysis relies on the High School Longitudinal Study of 2009 (HSLS). Published by the National Center for Education Statistics (NCES) of the U.S. Department of Education, the HSLS is a nationally representative sample of about 25,000 ninth graders from 944 public and private high schools who are followed throughout their secondary and postsecondary education. Students and their parents were first interviewed in 2009, again in 2012 and 2013 as students were graduating from high school and applying to college, and once more in 2016. Math and science teachers, the school administrator, and the lead school counselor also completed surveys about students. We use the restricted-use version of the HSLS, which provides student SAT and ACT scores, the colleges students applied to and enrolled in, and more detailed household-level financial variables.

Parental income is reported by parents, who are asked for their household income from all sources in 2011 and choose among pre-set ranges (less than $15k, $15k to $35k, and $20k-wide bins up to more than $235k; panel (a) of Figure A.1 in the Appendix depicts the distribution). For college preparedness, we use the students' high-school GPA and SAT or ACT scores.[4] High-school GPA is reported directly by the high school attended and is honors-weighted by the NCES to make GPAs comparable across high schools. Panels (b)–(d) of Figure A.1 in the Appendix show the SAT and GPA distributions, and panels (e)–(f) their correlation with parental income. Students in the top decile of the SAT distribution scored 1,250 or above.

The HSLS also includes detailed information about each student's application portfolio. In the follow-up survey after completing high school (the 2013 Update), students were asked

[4]The NCES converts the ACT score into an equivalent SAT score for students who only took the ACT. Henceforth, whenever we mention the SAT, we refer to either the SAT or ACT score.

to identify the college they were currently attending and to list two other colleges they had applied to and seriously considered attending if given an offer. Students report whether they were admitted to each of these colleges. Students were also asked for the *total* number of applications they submitted. Most students report applying to three or fewer schools, so for the majority of the sample the three listed colleges cover the whole portfolio. Our sample consists of the 18,913 students (out of 25,206 interviewed) with non-missing parental income and SAT score who were not attending a foreign institution in November 2013 and answered the question about postsecondary enrollment. Appendix A.1 compares the demographic characteristics of this sample with those of the full HSLS and lists the variables and weights used throughout.

Identifying selective colleges. We focus on four-year public and not-for-profit colleges.[5] Following the literature (e.g., Hoxby and Avery (2014), Chetty et al. (2020)), we use Barron's selectivity index (*Profile of American Colleges*, 2015) to categorize colleges into selectivity groups.[6] We identify as *selective* the colleges in Barron's first and second tiers. This definition extends beyond the Ivy League to capture schools that offer strong academics and job prospects (liberal arts colleges, STEM-focused schools, and well-regarded public universities). All other four-year not-for-profit colleges are referred to as *nonselective*.

Table A.2 in the Appendix summarizes the differences between the two groups using 2013 data from the Integrated Postsecondary Education Data System (IPEDS). Out of 1,763 colleges, 186 are classified as selective. They enroll 16% of undergraduate students, whose average SAT score (1,292) places them in the top decile of the HSLS distribution. Weighting by enrollment, 55% of students at selective colleges attend a public institutions (73% at nonselective colleges); selective colleges have larger endowments ($131k per student vs. $13k), spend more per student ($21k vs. $9k), and post higher tuition ($25k vs. $13k, excluding room and board). They nonetheless charge very little to students at the bottom of the income distribution. The average net tuition for students from families earning less than $50k is negative at selective colleges, i.e., grants exceed sticker tuition (though not the full cost of attendance). Since most students in both tiers attend public institutions, the selectivity split cuts across the public–private line rather than along it. Appendix A.2 provides details on the IPEDS sample.

Measurement caveats. Three features of the data should be kept in mind when reading the evidence below. First, SAT scores are only observed for students who took the SAT or ACT, so the sample is tilted toward college-bound students, and every statistic broken down by SAT decile is conditional on having a score. Second, the application portfolio is censored at three colleges and admission decisions are self-reported. Censoring means that some selective-college appli-

[5]Two-year colleges are part of the outside option. They are nearly all open-admission, their tuition is low and uniform, and they offer little institutional aid, so the margins this paper studies are largely absent there.

[6]The index divides colleges into six categories: (1) most competitive, (2) highly competitive, (3) very competitive, (4) competitive, (5) less competitive and (6) noncompetitive.

cations go unrecorded, so application rates to selective colleges are understated, more so for students with long portfolios. Third, several statistics condition on an endogenous choice (having applied to a selective college, or having been admitted to both types of colleges). Because low-income applicants to selective colleges are likely to be positively selected on unobservables, which is in fact the mechanism at the heart of the model in Section 3, the conditional statistics should be interpreted as descriptions of who reaches each stage rather than as causal effects of income. The regressions we report control for a rich set of student, family, and high-school characteristics, but are subject to the same caveat.

2.2. *Who attends selective colleges?*

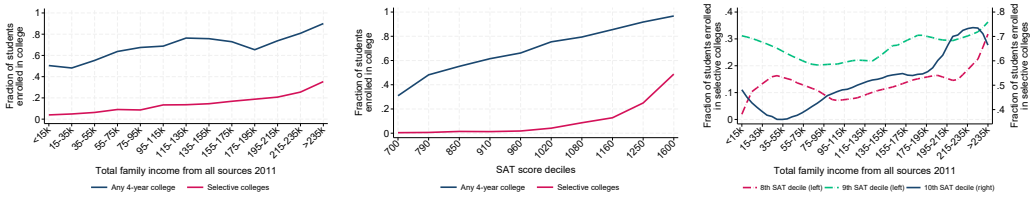
We start with enrollment patterns by parental income and SAT scores. Lower income students are less likely to attend selective colleges than their wealthier peers, even among top test scorers. Panel (a) of Figure 2.1 shows the fraction of students enrolled in any four-year college and in a selective college across the parental income distribution. Higher-income students are much more likely to attend a selective college: about 21% of students whose parental income is between \$195 and \$215k attend a selective college, agains about 6% of students in the \$35-55k range. Panel (b) shows the same fractions by SAT decile. Students who scored highly on standardized tests are far more likely to attend a selective college. Sorting by high-school GPA gives a similar picture (see panel (a) of Figure A.2 in the Appendix).

Since parental income and SAT scores are positively correlated (the correlation coefficient is 0.33), panel (c) displays enrollment rates by parental income for the top three SAT deciles. Among the best test takers, those with higher-income parents are more likely to attend a selective college, although the relation is not monotone at the bottom of the income distribution. About 70% of students in the top SAT decile whose parental income is above \$195k enrolled in a selective college; the fraction is 36% for students whose parents earn between \$35k and \$55k, and above 40% for students whose parents earn less than \$35k.[7]

Many demographic characteristics correlate with both parental income and attendance. The HSLS allows us to control for a rich set of them, including high-school GPA, gender, race, parents' education, family structure, and high-school characteristics. Since we observe where students applied, we can also restrict attention to students who applied to and enrolled in college. Table A.3 in the Appendix reports logit estimates in which the dependent variable is an indicator for enrolling in any four-year not-for-profit college *conditional on applying to college* (left columns), for enrolling in a selective college *conditional on applying to any college* (middle columns), and for enrolling in a selective college *conditional on applying to a selec-*

[7]The top three deciles of the GPA distribution show similar differences across the income distribution; see panel (b) of Figure A.2 in the Appendix.

FIGURE 2.1.—College enrollment



(a) By parental income (b) By SAT score (c) By parental income and SAT

Note: Panel (a) shows the fraction of students enrolled in any four-year not-for-profit college (blue) and in a selective college (pink) by parental income. Panel (b) shows the same fractions by SAT score decile. Panel (c) shows the fraction of students enrolled in a selective college by parental income for the top three deciles of the SAT score distribution (8th decile in pink, 9th in green, 10th in blue). Figure A.2 in the Appendix replicates panel (b) and (c) using high-school GPA. *Source:* HSLS and authors' calculations.

tive college (right columns). Figure A.2 in the Appendix plots the expected probabilities across the parental income distribution for students at the 50th, 70th, and 90th percentiles of the SAT distribution, holding the other covariates at their means.

Two results stand out. First, conditional on applying, the probability of enrolling in *some* four-year college is flat in parental income at all SAT levels. Second, among applicants to any college, the probability of enrolling in a selective college is significantly increasing in income (the income coefficient is 0.071 with a standard error of 0.026). The point estimates imply that, among applicants to any college, a student at the 90th percentile of the SAT in the highest income category is about 7 p.p. more likely to enroll in a selective college than a peer with the same SAT score whose parents earn less than $35k. The evidence below locates where it arises.

These patterns are line with previous evidence on the role of parental income in college attendance using earlier cohorts from the NLSY79 and NSLY97 (e.g. Belley and Lochner (2007) and Leukhina (2023)). We now go one step further and separate the roles of applications, admissions, and enrollment conditional on admission.

2.3. *Who applies to selective colleges?*

Students must apply before they can enroll. Two facts about applications stand out: students send few applications and lower income students apply less intensively to selective colleges.

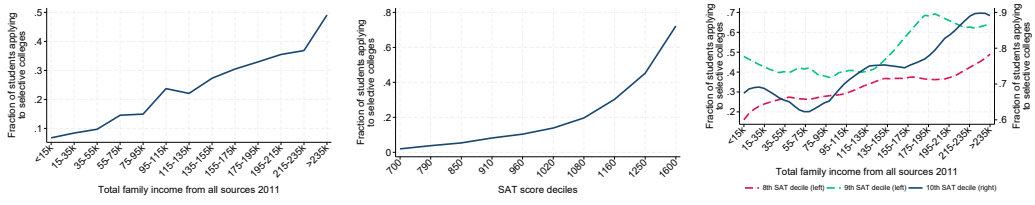
Number of applications. Although the data do not allow us to count with certainty the applications submitted to *selective* colleges, we observe the total number of applications sent to any four-year college. That number is low: on average 2.4, with 90% of high-school students submitting six or fewer applications. Figure A.3 in the Appendix shows how it varies with parental income and SAT scores. Students whose parents earn less than $55k submit fewer than two applications, about half as many as their higher-income peers, but the gap narrows considerably once we condition on submitting at least one application (as many low-income students do not apply at all). Applications rise steeply with test scores. Students in the top SAT decile submit roughly five applications, against fewer than three for the median student. Among top scor-

ers, differences by parental income are modest. Application counts are U-shaped in income for students in the top SAT decile, who submit between four and six applications in nearly every income bin, with the most at the bottom and the top of the income distribution.[8]

Applications to selective colleges. We now turn to the three colleges students list. Panel (a) of Figure 2.2 displays the fraction of students who applied to a selective college by parental income. About 49% of students whose parents earn more than $235k applied to a selective college, against less than 10% of students whose parents earn between $35k and $55k. Panel (b) shows the fraction by SAT decile. About 72% of the very top scorers applied, compared to 10% of students with the median SAT score. Panel (c) focuses on the top three SAT deciles. Among top students, those with lower incomes are less likely to apply to a selective college.

Among students in the top SAT decile, more than 85% of those whose parents earn more than $195k apply to a selective college, compared with about 66% of those whose parents earn between $35k and $55k. The differences across parental income are even more pronounced in the 8th and 9th SAT deciles.[9] Expressed as an exchange rate, the logit estimates in Table A.4 imply that, among top scorers, moving a student from the $35k–$55k parental-income group to the $195k+ group raises the probability of applying to a selective college by as much as an increase of roughly 130 points on the SAT (about two-thirds of a standard deviation).

FIGURE 2.2.—Applications to selective colleges



(a) By parental income (b) By SAT score (c) By parental income and SAT

Note: Panel (a) shows the fraction of students who applied to a selective college by parental income. Panel (b) shows the fraction of students who applied to a selective college by SAT decile. Panel (c) shows the fraction of students who applied to a selective college by parental income for the top three deciles of the SAT distribution (8th decile in pink, 9th in green, 10th in blue). Figure A.3 in the Appendix replicates panels (b) and (c) using high-school GPA. *Source:* HSLS and authors' calculations.

The shares above mask important differences in demographic characteristics. Yet, controlling for them still reveals that parental income plays a major role in whether students apply to selective colleges—in particular among top students. Table A.4 in the Appendix reports logit estimates in which the dependent variable is an indicator for applying to any four-year, not-for-profit college (left columns) and for including a selective college in the application portfolio *conditional on applying to college* (right columns). Figure A.3 in the Appendix plots the ex-

[8]Figure A.3 plots the average number of applications by high-school GPA and parental income for students in the top three deciles of GPA, and the results are very similar to ones with SAT scores.

[9]Figure A.3 in the Appendix shows the corresponding fractions by high-school GPA. The share of students who applied to a selective college is slightly lower when using GPA: about 60% for the very top GPA students and 8% for the median GPA.

pected probabilities across the parental income distribution for different SAT percentiles, holding the other covariates at their means. The probability of applying to *any* college is very high and increases mildly with income. A student at the 90th percentile of the SAT whose parental income is between $35k and $55k has a 90% probability of applying to college, while a peer with income between $195k and $215k has a 97% probability.

In contrast, the probability of applying to a *selective* college conditional on applying at all varies a lot more with income and more so for for high scorers (the interaction between parental income and SAT is positive and significant, while parental income by itself is not). A student at the 90th percentile of the SAT whose parental income is between $35k and $55k has a 40% probability of applying to a selective college. The probability for a peer in the same SAT percentile with income between $195k and $215k is 16 p.p. higher, at 56%. Differences in application rates to selective colleges are therefore not accounted for by high-school GPA, gender, race, parental education, or high-school characteristics. A note of caution is nonetheless warranted. Parental income may proxy for the parental or counseling support needed to navigate the application process at selective colleges (e.g., Bailey and Dynarski (2011)).

2.4. *Who is admitted to selective colleges?*

A second hypothesis for why low-income high-ability students do not attend selective colleges is that they apply but are not admitted. Differences in admission rates by income exist, but they are modest relative to the differences in application rates documented above.

The average admission rate at selective colleges among those who apply is fairly high, at 73%, since the set of colleges includes public universities with relatively high admission rates.[10] Panel (a) of Figure 2.3 shows the fraction of applicants to a selective college who were admitted to one, by parental income. The admission rate increases with parental income. More than 80% of students whose parents earn $195k or more were admitted, against about 68% of applicants whose parents earn between $35k and $55k. Students with better SAT scores also have an easier time getting in, as panel (b) shows. More than 80% of applicants in the top two SAT deciles are admitted. Panel (c) zooms in the top three SAT deciles. Admission rates are broadly similar across the income distribution among top scorers. For students in the 9th and 10th SAT deciles, the average admission rate lies between 75% and 95% for incomes above $55k.[11]

Table A.5 in the Appendix reports logit estimates for whether a student was admitted to any four-year not-for-profit college *conditional on applying to college* (left columns) and to a

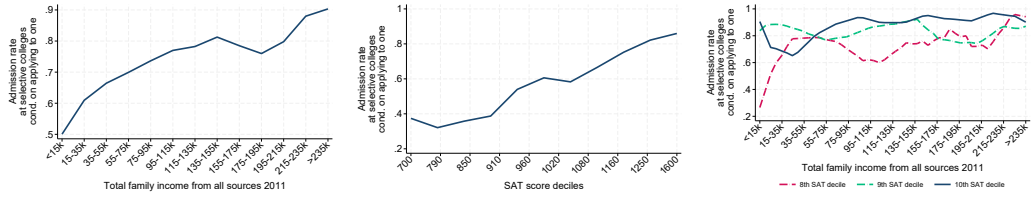
[10]The enrollment-weighted acceptance rate at selective colleges in IPEDS is lower than in the HSLS (Table A.2).

[11]Figure A.4 in the Appendix shows the corresponding rates by high-school GPA. Admission rates are flat below the median GPA and above 90% for the top GPA decile. Admissions rates for 8th and 9th GPA deciles are lower for low-income students than for the same SAT deciles and they are flat in parental income for the very top high-school GPA students.

selective college *conditional on applying to a selective college* (right columns), and Figure A.4 plots the implied probabilities. The probability of being admitted to some college conditional on applying is 98% for a student in the 90th percentile of the SAT, regardless of parental income. In contrast, the probability of being admitted to a selective college conditional on applying to one increases with income (the coefficient on parental income is positive and significant and the interaction with SAT is essentially zero, so the income gradient is similar across the SAT). For a student at the 90th percentile of the SAT, the probability of admission is 78% when parental income is between $35k and $55k, and 86% when it is between $195k and $215k.

The upshot so far is that differences in selective-college enrollment across the parental income distribution are to a large extent due to differences in who applies. Admission rates also increase with income, but the gap (about 8 p.p. for a top scorer across most of the income distribution) is half the size of the gap in application rates.

FIGURE 2.3.—Admission rates at selective colleges



(a) By parental income (b) By SAT score (c) By parental income and SAT

Note: Panel (a) shows the fraction of applicants to a selective college who were admitted to one, by parental income. Panel (b) shows the same fraction by SAT score decile. Panel (c) shows the same fraction by parental income for the top three deciles of the SAT distribution (8th decile in pink, 9th in green, 10th in blue). Figure A.4 in the Appendix replicates panel (b) and (c) using high-school GPA. *Source:* HSLS and authors' calculations.

2.5. *Who enrolls after being admitted?*

Students may be admitted to a selective college and yet prefer not to attend one partly because the cost is higher than they expected, or they received an attractive offer from a nonselective college closer to home or where friends were also admitted. These factors may disproportionately affect lower income students. To shed light on this, we focus on students who were admitted to both a selective and a nonselective college and look at which one they chose. In short, lower income students are more likely than their higher income peers to choose the nonselective college, but the gap is small once demographic characteristics are accounted for.

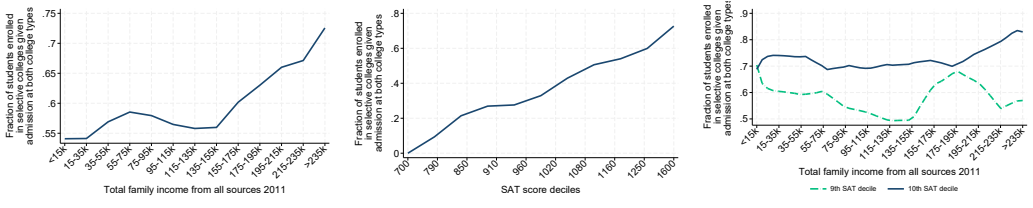
Panel (a) of Figure 2.4 shows the fraction of students admitted to both types of colleges who enrolled in a selective one, by parental income. Although most students prefer the selective college, lower income students are less likely to. Fewer than 60% of students whose parents earn less than $155k choose the selective college, against higher shares above that income level. Panel (b) shows the fraction by SAT decile. Close to 73% of students in the top SAT decile choose the selective college. Panel (c) focuses on the top two SAT deciles across parental

income. About 73% of students in the very top SAT decile choose the selective college. The share rises to about 83% for the highest-income students.[12]

Table A.6 in the Appendix reports logit estimates for whether a student enrolled in any college (left columns) or in a selective college (right columns) *conditional on being admitted to both selective and nonselective colleges*, and Figure A.5 plots the implied probabilities. Once demographic characteristics are controlled for, the probability of choosing the selective college is flat in parental income (the coefficient on parental income is small and insignificant, as is its interaction with SAT). For students at the 90th percentile of the SAT, the predicted probability of enrolling in a selective college is above 60% at all income levels.

Almost all students admitted to both types of colleges who do not enroll in a selective college enroll in a nonselective one (about 95% of such cases). The remainder defer entry for a year or forgo college. The affordability evidence in the next subsection suggests that costs weigh heavily on low-income families. Need-based aid can to some extent be anticipated from college websites, but many selective colleges also award merit-based aid that depends on an evaluation of the full application and cannot be predicted with certainty at the application stage. While colleges at the very top (e.g., Harvard, Yale) cover all expenses for low-income students, the net price a given student will face at a given selective college is uncertain until the offer arrives.

FIGURE 2.4.—Enrollment in a selective college among students admitted to both selective and non selective colleges



(a) By parental income (b) By SAT score (c) By parental income and SAT

Note: Panel (a) shows the fraction of students who enrolled in a selective college after being admitted to both selective and nonselective colleges, by parental income. Panel (b) shows the same fraction by SAT score decile. Panel (c) shows the same fraction by parental income for the top two deciles of the SAT distribution (9th decile in green, 10th in blue). Figure A.5 in the Appendix replicates panels (b) and (c) using high-school GPA. *Source:* HSLS and authors' calculations.

2.6. *Who expects to afford college?*

The evidence so far shows that the income gradient in selective-college enrollment among top students arises mostly at the application stage. A natural question is whether low-income families expect college to be unaffordable before they apply. In the first HSLS follow-up in 2012, parents and students were separately asked whether they would be able to afford college if the student were accepted. Then in the 2013 Update, students not attending college were

[12] Very few students in the 8th SAT decile applied and were admitted to both selective and nonselective colleges, so that decile is omitted. Figure A.5 shows that enrollment rates are similar when using high-school GPA.

asked whether affordability was the main reason. Figure A.6 in the Appendix shows how these answers vary with parental income.

A year before their child would enroll, more than half of parents with income below $75k expected not to be able to afford college, against less than 10% of parents earning $155k or more (panel (a)). Students' own expectations show the same gradient at low incomes: more than 30% of students whose parents earn less than $75k expected not to be able to afford college (panel (b)).[13] Finally, the share of students who report not attending college in 2013 because they could not afford it (panel (c)) also declines steeply with parental income.[14]

2.7. *Taking stock*

Among high-scoring students, the income gap in selective-college enrollment is large. It arises mainly because lower income students are much less likely to apply to selective colleges (a 16 p.p. gap between the $35–55k and $195–215k income groups at the 90th SAT percentile after controls), to a lesser extent because they are somewhat less likely to be admitted when they apply (about 8 p.p.), and only marginally because they turn down selective offers when they hold them (a difference that is not statistically distinguishable from zero after controls). Low-income families also expect college to be unaffordable at much higher rates than wealthier families, even when the student is a top scorer. The model in the next section is designed to account for these facts jointly: students apply under uncertainty about both admission and financial aid, and colleges' pricing responds to who applies.

3. Model

Motivated by the facts presented above, this section presents an equilibrium model of the college market in which students choose *how many* and *which type* of colleges to apply to under uncertainty about both admission and financial aid, and colleges set admission standards and tuition schedules taking into account who applies. The model is used to study whether expectations about financial aid and admissions can account for the observed gaps in selective-college enrollment across the income distribution, and to evaluate policies that operate through applications (fee waivers), information (the SAT), and aid (Pell grants).

Overview. The economy is populated by a unit measure of heterogeneous individuals who live for two periods, young and old. Young and old cohorts of equal measure coexist, so that generations overlap. Young individuals start life with a parental transfer y and an ability ℓ,

[13] Parents and students share similar concerns about college affordability by SAT score and high-school GPA, i.e., students with worse grades are more likely to report that they do not expect to be able to afford college.

[14] Figure A.6 in the Appendix restricts attention to students who did not apply to college. Lower income non-applicants are more likely to cite affordability as the main reason for not applying.

jointly distributed according to $\hat{\mu}(y, \ell)$, and decide whether to work or to invest in their human capital by attending college. Attending college requires applying and being admitted.[15]

There are two tiers of colleges: selective, indexed by H, and nonselective, indexed by L, $s \in \{H, L\}$. Each tier consists of N_s identical colleges that differ in endowment income, fixed costs, and productivity.[16] A student chooses an application *portfolio* $\mathbf{n} = (n_H, n_L)$, i.e., the number of selective and nonselective colleges to apply to, with $n_s \in [0, N_s]$. Applying is costly, both financially and in terms of effort, and the cost increases with the number of applications.

We interpret ℓ as the student's *own assessment* of her ability, which may differ from a college's assessment of the student's ability, σ. Think of σ as what admission committees infer from the applicant's SAT scores, high-school grades, essays, and extracurricular activities. Given the population distribution $\hat{\mu}$, the two assessments are noisy signals of each other, that is, σ is distributed around ℓ, and, by Bayes' rule, ℓ is likewise imperfectly predicted by σ.[17] In addition, each college's admission decision is subject to idiosyncratic evaluation noise, so that applying to several colleges of the same tier provides insurance against rejection. Admitted students receive college-specific tuition offers that depend on (y, σ).[18] Once admissions and offers are known, all uncertainty is resolved and students choose which offer, if any, to accept.

Colleges maximize the value they add to their students on the labor market, which is determined endogenously as a function of the average ability of their student body and their average instructional spending per student. Each college competes by choosing its admission standard, its instructional spending, and a tuition schedule for each student type (y, σ), taking as given the standards and schedules of all other colleges. Because a student holding several offers accepts at most one, the demand each college faces—and hence its optimal price—depends on the pricing and admission policies of the competing colleges. Finally, a government taxes labor income to subsidize colleges and to fund need-based grants.

Timing. The timing of events in the first period is as follows:

1. Individuals choose an application portfolio (n_H, n_L), including not applying, and pay the application costs.

[15] College means a four-year institution. The outside option pools work and two-year colleges, which are nearly all open-admission and carry little of the admission and financial-aid risk at the center of the model.

[16] Colleges within a tier are replicas of each other and, as discussed below, we focus on the equilibrium in which they behave alike. The number of colleges per tier matters only through the number of distinct admission draws a student can obtain and through the competition among identical offers at the enrollment stage.

[17] We deliberately take no stand on whether students or colleges are better informed about a latent "true" ability. Each side acts on its own assessment. What the mechanism requires is only that applications must carry information colleges do not already possess. To avoid carrying a latent ability as a third state variable, earnings load on ℓ, so the student's own assessment is the payoff-relevant index.

[18] Students' parental transfer is assumed to be fully observable to colleges. In order to receive federal grants or loans students must complete the Free Application for Federal Student Aid (FAFSA), which states students' parental income and financial assets and is observable by the colleges the student applies to. About 77% of students in the HSLS complete the FAFSA.

2. Colleges form their assessment σ of each applicant's ability, and set admission standards, tuition schedules, and instructional spending. Each college admits a student if her evaluation clears its admission standard and admitted students receive tuition offers according to the tuition schedules.
3. Students choose which offer to accept, if any, and enroll.
4. Individuals make their consumption and savings decisions.

In the second period, individuals work, college graduates repay their loans, and the government collects taxes. In equilibrium, students correctly anticipate colleges' admission standards and tuition schedules when choosing where to apply, consistent with rational expectations.

3.1. *Students*

We describe the student's problem by backward induction: the payoffs from attending college and from working, the enrollment decision given a set of offers, and the application decision.

3.1.1. *Value from attending college and from working after high school*

Individuals have preferences over consumption, c, and discount the future at rate β. Attending college s entails a non-pecuniary utility cost $\nu_s(\ell)$, which is decreasing in ability as higher-ability students may find college easier (and may turn negative for the most able) and which may differ across tiers as attending a selective college may be more demanding. This psychic cost of attendance follows Cunha et al. (2005) and Heckman et al. (2006), who find that such costs are needed to reconcile observed enrollment with the pecuniary returns to college.

While in college, the student finances her consumption, tuition, and application costs out of parental transfers, grants, and student debt. Students may borrow up to a limit $\underline{a}_s$.[19] Grants depend on the student's income and are funded either by the government, $P(y)$, or by outside sources, $Gr(y)$. Government-funded grants capture the federal Pell Grant program and are paid for with labor income taxes on workers. Outside grants are taken as exogenous. The financial cost of submitting $n = n_H + n_L$ applications is $\psi(n)$, which is weakly increasing in n.

After college, the student enters the labor market and repays her student loan at the gross interest rate R. Labor earnings are the product of the wage per efficiency unit w, net of taxes τ, and the individual's human capital. If the student graduates, which happens with probability δ_s, her human capital is $Z_s \ell^{\alpha_s}$, where $Z_s > 1$ is the value added of the college attended and

[19] This borrowing limit is motivated by the limits on federal student loans set by the Department of Education. Consistent with the federal limit, the modeled borrowing limit does not depend on the student's earnings potential.

$\alpha_s < 1$ governs the returns to ability. If she fails to graduate, her human capital is ℓ^{α_w}, with $\alpha_w < \alpha_s$, so that the return to college is larger for higher-ability individuals (for $\ell > 1$).[20]

Value of attending college. The value of attending a college of tier s at tuition $T_s(y, \sigma)$ after submitting n applications is

$$V^s(y, \ell, n, T_s(y,\sigma)) = \max_{c,c',a'} u(c) - \nu_s(\ell) + \beta u(c') \tag{3.1}$$

$$\text{s.t.} \qquad c + a' = y + Gr(y) + P(y) - T_s(y,\sigma) - \psi(n)$$

$$c' = R\,a' + (1-\tau)w\,[\delta_s Z_s \ell^{\alpha_s} + (1-\delta_s)\ell^{\alpha_w}]$$

$$a' \geqslant \underline{a}_s.$$

Value of working without a college degree. Individuals end up as workers by not applying to college, by being rejected by every college they applied to, or by declining all offers. A worker pays no tuition but has paid the application costs $\psi(n)$ if she applied. Her problem is

$$V^W(y, \ell, n) = \max_{c,c',a'} u(c) + \beta u(c') \tag{3.2}$$

$$\text{s.t.} \qquad c + a' = y + (1-\tau)w\,\ell^{\alpha_w} - \psi(n)$$

$$c' = R\,a' + (1-\tau)w\,\ell^{\alpha_w}$$

$$a' \geqslant \underline{a}_w.$$

3.1.2. *Enrollment stage*

Consider a student with parental transfer y, ability ℓ, and colleges' assessment σ who submitted $n = n_H + n_L$ applications and holds $k = (k_H, k_L)$ offers, $k_H \leqslant n_H$ and $k_L \leqslant n_L$. Since colleges within a tier are identical and price on (y, σ), all her selective offers carry tuition $T_H(y,\sigma)$ and all her nonselective offers $T_L(y,\sigma)$. At this stage, all uncertainty is resolved and she simply chooses among her $k_H + k_L$ offers and the outside option of working.

We introduce idiosyncratic preference shocks over these alternatives. Each alternative carries an independent Type I extreme value shock with scale parameter $1/\lambda_c$, $\lambda_c > 0$. The shocks capture idiosyncratic reasons to prefer one college over another (campus visits, distance, where friends enroll) that are realized after offers are in hand, and they make demand for each college continuously differentiable in its tuition, so that each college's pricing problem is smooth. The

[20] Expected human capital enters the second-period budget constraint directly, so graduation uncertainty is pooled rather than borne as consumption risk.

value of holding offers $\mathbf{k}$ is then

$$V^{O}(y,\ell,n,\mathbf{k};\mathbf{T}(y,\sigma)) = \frac{1}{\lambda_c}\log\Bigg[\underbrace{k_H\, e^{\lambda_c V^{H}(y,\ell,n,T_H(y,\sigma))}}_{\substack{\text{value of } k_H \text{ identical}\\ \text{selective offers}}} + \underbrace{k_L\, e^{\lambda_c V^{L}(y,\ell,n,T_L(y,\sigma))}}_{\substack{\text{value of } k_L \text{ identical}\\ \text{nonselective offers}}} + \underbrace{e^{\lambda_c V^{W}(y,\ell,n)}}_{\substack{\text{outside option}\\ \text{of working}}}\Bigg], \tag{3.3}$$

where $\mathbf{T}(y,\sigma) = (T_H(y,\sigma), T_L(y,\sigma))$. A student with no offers ($\mathbf{k} = (0,0)$) obtains $V^{W}(y,\ell,n)$. The probability that she accepts a *given* offer from tier s is

$$a_s(y,\ell,n,\mathbf{k};\mathbf{T}(y,\sigma)) = \frac{e^{\lambda_c V^{s}(y,\ell,n,T_s(y,\sigma))}}{k_H\, e^{\lambda_c V^{H}(y,\ell,n,T_H(y,\sigma))} + k_L\, e^{\lambda_c V^{L}(y,\ell,n,T_L(y,\sigma))} + e^{\lambda_c V^{W}(y,\ell,n)}}, \tag{3.4}$$

so that she enrolls in *some* college of tier s with probability $a_s\, k_s$. Holding more offers raises the value of the enrollment stage (the option value of applying broadly) and lowers the probability that any single offer is accepted, which is the source of competition among colleges.

3.1.3. *Application stage*

Assessments and admission. When a student applies, the colleges' assessment σ is drawn from the conditional density $g(\sigma|\ell)$ with cumulative distribution $G(\sigma|\ell)$, and is common to all colleges she applied to.[21] The conditional density satisfies the monotone likelihood ratio property in ℓ, so that high assessments are more likely to come from students who perceived themselves as high ability. College assessments are unknown to students at the time they apply.

Each college j of tier s evaluates the applicant as $\sigma + \epsilon_j$, where ϵ_j is idiosyncratic evaluation noise, independent across colleges and of (y,ℓ,σ), with continuous distribution F_ϵ and density f_ϵ centered at zero. The two sources of noise play different roles. The common assessment σ generates the information friction at the heart of the model: colleges act on σ while earnings load on ℓ, so applicants with the same payoff-relevant ability can draw different assessments and hence face different admission odds and tuition offers. The idiosyncratic noise ϵ_j generates a motive to apply to several colleges of the same tier.[22] The college admits the student if $\sigma + \epsilon_j \geqslant \underline{\sigma}_s$, where $\underline{\sigma}_s$ is its admission standard. Hence, given the assessment, an application to a college of tier s succeeds with probability

$$\pi_s(\sigma) \equiv 1 - F_\epsilon(\underline{\sigma}_s - \sigma), \tag{3.5}$$

[21] Colleges sharing a common assessment is supported empirically: in the HSLS, 97% of students admitted to a selective college were also admitted to a nonselective college, conditional on having applied to both.

[22] We assume that tuition schedules depend on (y,σ) but not on ϵ_j. That is, the idiosyncratic component affects the admission decision but does not carry information that colleges use for pricing.

which is increasing in σ and decreasing in the admission standard $\underline{\sigma}_s$. Conditional on σ, admission outcomes are independent across colleges, so the number of offers from tier s is binomial, with $k_s \sim \mathcal{B}(n_s, \pi_s(\sigma))$ and $b(k; n, \pi) = \binom{n}{k}\pi^k(1-\pi)^{n-k}$ as the binomial probability. Applying to n_s colleges raises the probability of at least one offer to $1-(1-\pi_s(\sigma))^{n_s}$. As F_ϵ becomes degenerate, admission collapses to the deterministic rule $\sigma \geqslant \underline{\sigma}_s$ and applying to more than one college per tier is redundant.

Application costs. Applying involves a financial cost, $\psi(n)$, and a psychic disutility cost, $\phi(n_H, n_L)$, meant to capture the effort of preparing applications (writing essays, filling in forms, preparing for and taking the SAT, researching which colleges to apply to). Both are weakly increasing in the number of applications and the psychic cost may also depend on the tier, as selective colleges may require more elaborate applications.

Value of an application portfolio. The expected value of applying to $\mathbf{n} = (n_H, n_L)$ colleges for a student with characteristics (y, ℓ) is

$$V^A(y,\ell,\mathbf{n};\mathbf{T},\underline{\sigma}) = \int \sum_{k_H=0}^{n_H}\sum_{k_L=0}^{n_L} \underbrace{b(k_H;n_H,\pi_H(\sigma))\,b(k_L;n_L,\pi_L(\sigma))}_{\substack{\text{probability of holding}\\ \text{offers } \mathbf{k}=(k_H,k_L)}} \underbrace{V^O(y,\ell,n,\mathbf{k};\mathbf{T}(y,\sigma))}_{\substack{\text{value of choosing}\\ \text{among offers } \mathbf{k}}} \underbrace{g(\sigma|\ell)\,d\sigma}_{\substack{\text{assessment}\\ \text{density}}} - \underbrace{\phi(n_H,n_L)}_{\substack{\text{application}\\ \text{disutility}}}, \tag{3.6}$$

where the integral is over the possible realizations of the colleges' assessment and the sum over the possible offer sets. The value of not applying is $V^A(y,\ell,\mathbf{0}) = V^W(y,\ell,0)$. Note that at the application stage the student faces two sources of risk: she may be rejected (which depends on σ and on the ϵ_j's), and she may be admitted at a price that she is unwilling or unable to pay (which depends on σ through $T_s(y,\sigma)$).

Optimal application portfolio. The student chooses among the $(N_H+1)(N_L+1)$ portfolios, including not applying. She selects the portfolio $\mathbf{n} \in \{0,\dots,N_H\}\times\{0,\dots,N_L\}$ maximizing over all the application portfolios $V^A(y,\ell,\mathbf{n};\mathbf{T},\underline{\sigma}) + \epsilon^a_{\mathbf{n}}$, where $\epsilon^a_{\mathbf{n}}$ are independent Type I extreme value shocks with scale parameter $1/\lambda_a > 0$. As at the enrollment stage, these shocks smooth the discrete choice. The probability of choosing the application portfolio $\mathbf{n}$ is

$$p(\mathbf{n}|y,\ell) = \frac{e^{\lambda_a V^A(y,\ell,\mathbf{n})}}{\sum_{\mathbf{n}'} e^{\lambda_a V^A(y,\ell,\mathbf{n}')}}. \tag{3.7}$$

The optimal portfolio trades off the cost of an additional application against (i) the higher probability of receiving at least one offer, (ii) the option value of choosing among more offers, and (iii) the possibility of a favorable financial-aid draw. Since the tuition offered depends on σ, and the colleges' assessment is centered on the student's own assessment ℓ, students who assess

their own ability highly expect more favorable assessments from colleges and hence better offers. As a result, they are more willing to bear the application costs of selective colleges. This is the channel through which the applicant pool at each income level is selected on ability.

The number of colleges per tier matters only through two payoff-relevant channels: (i) it caps the number of applications a student can send to that tier (and hence the admission insurance she can buy), and (ii) it determines how many identical offers she can hold at the enrollment stage, which raises the own-price elasticity each college faces and disciplines markups.[23]

Demand for college. Consider a particular college j of tier $s \in \{H, L\}$, and denote by $-s$ the other tier. A student of type (y, ℓ) with assessment σ enrolls in college j if she applied to it, was admitted, and accepted its offer. Since colleges within a tier are identical, a student who applies to n_s colleges of tier s applies to college j with probability n_s/N_s. College j admits her with probability $\pi_s(\sigma)$. If she also holds k'_s other tier-s offers and k_{-s} offers from the other tier, the student accepts college j's offer with probability $a_s(y, \ell, n, \mathbf{k})$, where the offer set $\mathbf{k}$ comprises her $1 + k'_s$ tier-s offers and her k_{-s} other-tier offers. Summing over portfolios and offer sets, the demand for college j from a student of type (y, ℓ, σ) is

$$d_s(y,\ell,\sigma;\mathbf{T},\underline{\sigma}) = \sum_{\mathbf{n}:\, n_s \geqslant 1} \underbrace{p(\mathbf{n}|y,\ell)}_{\substack{\text{portfolio}\\ \text{choice}}} \underbrace{\frac{n_s}{N_s}}_{\substack{\text{applies to}\\ \text{college } j}} \underbrace{\pi_s(\sigma)}_{\substack{\text{admitted by}\\ \text{college } j}}$$

$$\times \sum_{k'_s=0}^{n_s-1} \sum_{k_{-s}=0}^{n_{-s}} \underbrace{b(k'_s; n_s-1, \pi_s(\sigma))\, b(k_{-s}; n_{-s}, \pi_{-s}(\sigma))}_{\substack{\text{probability of other}\\ \text{offers } (k'_s, k_{-s})}} \underbrace{a_s(y,\ell,n,\mathbf{k};\mathbf{T}(y,\sigma))}_{\substack{\text{accepts college}\\ j\text{'s offer}}}. \tag{3.8}$$

The demand a college faces thus depends on its own tuition schedule and admission standard through a_s and π_s, and on those of the competing colleges through the distribution of the student's other offers and their prices. Integrating over assessments, the probability that a student of type (y, ℓ) enrolls in some college of tier s is $q_s(y, \ell) \equiv N_s \int d_s(y, \ell, \sigma)\, g(\sigma|\ell)\, d\sigma$ (used to aggregate labor supply and government outlays) and the probability of not attending college is

$$q_w(y, \ell) = 1 - q_H(y, \ell) - q_L(y, \ell). \tag{3.9}$$

The student's problem nests the simplest explanation for the enrollment gaps of Section 2: that college quality is a normal good. Utility is concave, parental transfers relax the budget constraint, and borrowing limits bind hardest for the poorest students, so demand for the selective tier rises with income even without any information friction. Income effects alone, however, cannot explain (i) why the gap opens at the application stage rather than at enrollment, where

[23] As colleges within a tier are identical, any allocation of a student's applications is optimal (we assume they are split evenly across them as the n_s/N_s term in the college demand below shows) and tier-level outcomes are invariant to replicating colleges beyond the portfolio sizes students choose.

prices are actually paid; (ii) why it is largest at the colleges whose net prices for low-income students are the lowest in the market, or (iii) why more than half of lower-income families expect, before applying, to be unable to afford college.

3.2. *Colleges*

Each college chooses its instructional spending, admission standard, and tuition schedule taking as given the tuition schedules, admission standards, and value added of the other colleges (including those of its own tier). Colleges therefore compete à la Bertrand in tuition schedules and an equilibrium is a Nash equilibrium in $(\underline{\sigma}_s, T_s(\cdot,\cdot))$. Two informational restrictions shape this competition. First, colleges observe (y, σ) but neither the student's own assessment ℓ nor her application portfolio or other offers, so tuition schedules are functions of (y, σ) only. Second, colleges take students' application probabilities $p(\mathbf{n}|y, \ell)$ as given, that is, they internalize how their tuition and admission standard affect enrollment conditional on applications (equation (3.8)), but not how they affect the composition of the applicant pool.[24]

College value added. The value a college adds to its students is

$$Z_s = \xi_s \, Q(I_s/\kappa_s, L_s/\kappa_s), \tag{3.10}$$

where $Q(I_s/\kappa_s, L_s/\kappa_s)$ is the college's quality and ξ_s the efficiency with which quality is transformed into human capital. Quality is increasing in instructional spending per student, I_s/κ_s, and in the average ability of the student body, L_s/κ_s (the latter allows for peer effects). Instructional spending, I_s, is chosen directly by the college, while enrollment, κ_s, and aggregate ability, L_s, are the outcome of its own and the competing colleges' policies through students' demand. Let $\mu(y, \ell, \sigma) \equiv g(\sigma|\ell)\hat{\mu}(y, \ell)$ denote the joint distribution of student types and assessments. Enrollment and aggregate ability at a college of tier s are

$$\kappa_s = \int d_s(y, \ell, \sigma; \mathbf{T}, \underline{\sigma}) \, d\mu(y, \ell, \sigma), \tag{3.11}$$

$$L_s = \int \ell \, d_s(y, \ell, \sigma; \mathbf{T}, \underline{\sigma}) \, d\mu(y, \ell, \sigma). \tag{3.12}$$

Budget constraint. Colleges balance their budget:

$$I_s + C_s(\kappa_s) = E_s(\kappa_s) + Tr_s(\kappa_s) + \mathcal{T}_s. \tag{3.13}$$

Besides instructional spending, colleges face other costs, C_s (e.g., administrative, maintenance). Revenue consists of tuition, $\mathcal{T}_s$, net endowment income, E_s, and government transfers,

[24] The equilibrium applicant pool is nonetheless consistent with colleges' policies, as required by the equilibrium definition below.

Tr_s. Endowment income, transfers, and costs may depend on enrollment. Tuition revenue is

$$\mathcal{T}_s = \int T_s(y,\sigma)\, d_s(y,\ell,\sigma;\mathbf{T},\underline{\sigma})\, d\mu(y,\ell,\sigma). \tag{3.14}$$

College problem. A college of tier s solves

$$\max_{I_s,\underline{\sigma}_s,\{T_s(y,\sigma)\}_{\forall(y,\sigma)}} \xi_s\, Q(I_s/\kappa_s, L_s/\kappa_s) \tag{3.15}$$

$$\text{s.t. } (3.11), (3.12), (3.13), (3.14),\ T_s(y,\sigma) \leqslant \overline{T}_s,\ \underline{\sigma}_s \geqslant \underline{\sigma}.$$

The choice of objective function merits some discussion. Since Epple et al. (2006), the literature has formalized the college's problem as maximizing quality, a function of students' ability and instructional spending, subject to a budget constraint. The alternative of profit maximization, as in Epple and Romano (1998), delivers the same allocation. Because the budget constraint binds and spending is productive, maximizing profit subject to attaining the equilibrium quality level yields the same policies with maximized profit equal to zero. So, conditional on quality, the college prices like a profit maximizer, which is why the optimal tuition below takes the form of a markup over marginal cost. We follow the quality-maximization formulation, with two differences. First, the quality of the college a student attends directly affects her earnings. Second, ability is unobservable, so colleges can only indirectly influence the average ability of their student body, which prevents perfect assortative matching across tiers.

Two of the constraints merit comment as well. The lower bound $\underline{\sigma}$ on the admission standard is set far enough below the support of σ that, given the evaluation noise, a college at the bound admits every applicant with probability one. The tuition cap $\overline{T}_s$ is the sticker price, i.e., the maximum price any student is expected to pay so that financial aid cannot be negative. Since most institutions have tuition set or capped by state authorities, and private sticker prices are highly persistent, we treat $\overline{T}_s$ as a measured object, pinned down in Section 4 by the level at which net tuition flattens at the top of the income distribution. The cap is also what gives admission standards a role, as Proposition 3.1 makes precise. Without a cap-like friction, the model predicts admission rates close to one at every income level and enrollment is rationed by prices alone. The data reject this. Admission rates at selective colleges average 73% among applicants (Section 2.4) and selective colleges routinely reject full-pay applicants, which is evidence that colleges cannot, or do not, price discriminate without limit. The sticker-price cap is the parsimonious way to encode that constraint.[25]

[25] Deeper foundations for the cap deliver the same reduced form: regulated tuition at public institutions; the salience of sticker prices to deter applications; and reputational or charitable-status considerations that limit tuition increases. Section 5 reports the sensitivity of our policy counterfactuals to the level of the caps.

Each tier also pools public and private institutions, and the distinction does not align with selectivity as both tiers are majority public (55% of selective and 73% of nonselective enrollment, Table A.2). The two features that distinguish public colleges are in the model: state appropriations are the per-student government transfer in the budget constraint (3.13) and regulated posted tuition is the cap $\overline{T}_s$. Colleges in a tier are replicas, so the model speaks to the average selective college rather than to the variation among them.[26]

PROPOSITION 3.1—SCREENING BY STANDARDS REQUIRES LIMITS ON TUITION: Suppose that for every (y, σ) tuition is given by the interior optimum (3.16) with demand elasticity $\varepsilon_s(y, \sigma) > 1$, so the cap never binds and the markup is well defined, and suppose the domination condition of Appendix B.1 holds. Then, the optimal admission standard is at its lower bound, that is, screening operates through prices alone and applicants with sufficiently high parental transfers are never excluded.

PROOF: See Appendix B.1. *Q.E.D.*

Optimal tuition schedule. For student types (y, σ) at which the tuition cap does not bind, the optimal tuition is a markup over marginal cost,

$$T_s(y, \sigma) = \underbrace{m_s(y, \sigma)}_{\text{Markup}} \times \underbrace{MC_s(y, \sigma)}_{\text{Marginal cost}}, \tag{3.16}$$

where the markup depends on the elasticity of demand of students with observables (y, σ) with respect to the college's own price, holding rivals at their equilibrium schedules,

$$m_s(y, \sigma) = \frac{\varepsilon_s(y, \sigma)}{\varepsilon_s(y, \sigma) - 1}, \text{ where } \quad \varepsilon_s(y, \sigma) \equiv -\frac{\partial \ln \int d_s(y, \ell, \sigma; \mathbf{T}, \underline{\sigma})\, \mu(y, d\ell|\sigma)}{\partial \ln T_s(y, \sigma)},$$

and the marginal cost has two components,

$$MC_s(y, \sigma) = \underbrace{C_{s\kappa} - \frac{Q_\kappa}{Q_I} - E_{s\kappa} - Tr_{s\kappa}}_{\substack{\text{Costs net of transfers} \\ \text{per student}}} - \underbrace{\frac{Q_L}{Q_I} \widehat{\mathbb{E}}_s[\ell|y, \sigma]}_{\text{Posterior ability discount}}, \tag{3.17}$$

where $C_{s\kappa} \equiv C_s'(\kappa_s)$, $E_{s\kappa} \equiv E_s'(\kappa_s)$, and $Tr_{s\kappa} \equiv Tr_s'(\kappa_s)$ denote marginal administrative costs, endowment income, and transfers with respect to enrollment, $Q_\kappa < 0$, $Q_I, Q_L > 0$, and $\widehat{\mathbb{E}}_s[\ell|y, \sigma] \equiv \int \ell\, \partial_{T_s} d_s\, \mu(y, d\ell|\sigma) \big/ \int \partial_{T_s} d_s\, \mu(y, d\ell|\sigma)$ is the expected ability of a student with

[26] In the data, some elite colleges guarantee full-need aid to low-income admits while others charge them substantial net prices, and college objectives include considerations the value added function omits, from in-state mandates to legacy and athlete preferences (Arcidiacono et al., 2022). Modeling these variations is left to future research.

observables (y, σ), weighted by her demand response. When the cap binds, $T_s(y,\sigma) = \overline{T}_s < m_s(y,\sigma)MC_s(y,\sigma)$. The derivations are in Appendix B.2.

The first component of marginal cost is common to all students. An additional student requires more instructional spending and dilutes endowment and transfers per student. The second component accounts for the externality a student imposes on the average ability of the student body. A student whose expected ability exceeds the college's average receives a discount, and one below it pays a penalty. Since the colleges' assessment is centered on the student's own, the posterior is increasing in σ, and so is the ability discount. Crucially, the posterior also differs by income level, so it depends on the *equilibrium* composition of applicants at each (y, σ). If at some income level only students who assess their own ability highly apply, any applicant from that income level carries a higher posterior and receives a larger discount. This mechanism helps explain why low-income students receive high levels of financial aid at selective colleges: only those with the highest own assessments apply.

The markup reflects two forms of competition. Across tiers, a student who also holds a nonselective offer is more price sensitive. Within tiers, a student holding several selective offers is even more price sensitive, since the offers are close substitutes. The more selective colleges a student applies to, the lower the markup any one of them can charge her. Policies that change the number of applications therefore affect tuition through the elasticity of demand.

Optimal admission standard. Raising the standard $\underline{\sigma}_s$ lowers the admission probability $\pi_s(\sigma)$ of every applicant, most strongly for those whose assessment is close to the threshold (since $\partial\pi_s/\partial\underline{\sigma}_s = -f_\epsilon(\underline{\sigma}_s - \sigma)$). The optimal standard equates the tuition revenue lost from these marginal students to the cost they impose (Appendix B.3):

$$\int \left[\underbrace{T_s(y,\sigma)}_{\text{Tuition}} - \underbrace{\left(C_{s_\kappa} - \frac{Q_\kappa}{Q_I} - E_{s_\kappa} - Tr_{s_\kappa} \right)}_{\substack{\text{Costs net of transfers} \\ \text{per student}}} + \underbrace{\frac{Q_L}{Q_I}\ell}_{\substack{\text{Ability} \\ \text{contribution}}} \right] \underbrace{\frac{\partial d_s(y,\ell,\sigma)}{\partial \underline{\sigma}_s}}_{\substack{\text{Demand response of} \\ \text{marginal applicants}}} d\mu(y,\ell,\sigma) = 0. \tag{3.18}$$

Since $\partial d_s/\partial\underline{\sigma}_s < 0$ concentrates on students with σ near $\underline{\sigma}_s$, condition (3.18) states that tuition must cover marginal cost net of the ability discount on average over the marginal applicants. Students with low assessments carry a low posterior ability and, because of the tuition cap, cannot be charged enough to compensate the college for lowering the average ability of its student body. The college excludes them by raising its standard. As the evaluation noise vanishes, the weight $\partial d_s/\partial\underline{\sigma}_s$ concentrates on $\sigma = \underline{\sigma}_s$ and (3.18) reduces to the simpler condition that the average revenue from the lowest-assessment admitted students equals their marginal cost.

3.3. *Government*

The government taxes labor income and uses the revenue to finance Pell Grants and college subsidies. The tax base consists of all workers without a college degree (whether or not they applied) while young and old, and college-educated workers while old. Aggregate labor supplied by young and old individuals is

$$N = \underbrace{2\int \ell^{\alpha_w}\, q_w(y,\ell)\, d\hat{\mu}(y,\ell)}_{\substack{\text{Labor supplied by young and old}\\ \text{high-school workers}}} + \underbrace{\sum_{s\in\{H,L\}} \int [\delta_s Z_s \ell^{\alpha_s} + (1-\delta_s)\ell^{\alpha_w}]\, q_s(y,\ell)\, d\hat{\mu}(y,\ell)}_{\text{Labor supplied by college-educated workers}}, \quad (3.19)$$

where $q_w(y,\ell)$ is defined in (3.9). The government budget constraint is

$$\tau w N = \underbrace{\sum_{s\in\{H,L\}} \int P(y)\, q_s(y,\ell)\, d\hat{\mu}(y,\ell)}_{\substack{\text{Transfers to students}\\ \text{(Pell grants)}}} + \underbrace{\sum_{s\in\{H,L\}} N_s\, Tr_s(\kappa_s)}_{\text{Transfers to colleges}}. \quad (3.20)$$

3.4. *Equilibrium*

Definition. A stationary equilibrium in the college market consists of: (1) value functions of attending a college of tier s and of working, V^s and V^W, with the associated consumption and savings policies, for each student type and number of applications; (2) values of holding offers, V^O, and acceptance probabilities, a_s, for each student type, number of applications, and offer set $\mathbf{k}$; (3) values of application portfolios, V^A, and portfolio probabilities, $p(\mathbf{n}|y,\ell)$, for each student type; (4) college demand functions $d_s(y,\ell,\sigma)$ for each tier s; (5) an admission standard $\underline{\sigma}_s$, a tuition schedule $T_s(y,\sigma)$, instructional spending I_s, enrollment κ_s, aggregate ability L_s, and value added Z_s for each tier s; (6) a labor income tax rate τ; such that

1. Given τ, w, Z_s, $T_s(y,\sigma)$, $Gr(y)$, and $P(y)$, V^s solves (3.1) and V^W solves (3.2);
2. Given V^s, V^W, and $\mathbf{k}$, V^O and a_s are given by (3.3) and (3.4);
3. Given V^O and $\underline{\sigma}$, V^A and $p(\mathbf{n}|y,\ell)$ are given by (3.6) and (3.7);
4. Given $p(\mathbf{n}|y,\ell)$, a_s, and π_s, college demand d_s satisfies (3.8);
5. Given college demand and the policies of all other colleges, $(\underline{\sigma}_s, T_s(\cdot,\cdot), I_s)$ solve the problem (3.15) of each college of tier s, with Z_s, κ_s, and L_s given by (3.10), (3.11), and (3.12), and the solution satisfies (3.16)–(3.18);
6. The government budget constraint (3.20) holds.

3.5. *Existence and computation*

Peer effects can generate multiple equilibria. A higher Z_s raises the value of attending, drawing in students who assess their own ability highly, and the resulting rise in the student body's

average ability raises Z_s in turn. The same logic runs in reverse, so optimistic and pessimistic beliefs about a college's value added could each be self-fulfilling. We therefore establish existence analytically and locate equilibria numerically.

Existence. The taste shocks and the evaluation noise make all individual responses (portfolio choices, enrollment decisions, and admission probabilities) smooth in college policies and every equilibrium object lives in a compact convex set that the equilibrium updates respect (tuition in a closed interval containing all candidate optima below the sticker price, admission standards within the assessment grid, value added bounded by budget feasibility, and the tax rate between 0 and 1). So, the equilibrium mapping above is a continuous self-map and Brouwer's theorem applies (see Appendix B.4).

Computation. The peer-effect feedback operates entirely through the aggregate state (Z_H, Z_L). Given the college value added, every other equilibrium object is determined (applications, enrollment, tuition, admission standards, spending), and hence so is the *actual* value added. An equilibrium is therefore a fixed point of the map $\Phi(Z_H, Z_L)$.[27] We solve for it by root-finding (Appendix B.5) and search the feasible range of value added for other solutions (Appendix B.6). The search finds a single equilibrium in which both tiers operate, with $Z_H > Z_L$. The remaining configurations involve one tier operating no colleges and fail the equilibrium definition. The selection criterion of Epple et al. (2006), by which the quality ranking follows the endowment ranking, is therefore not needed at our estimates.

4. TAKING THE MODEL TO THE DATA

4.1. *Functional forms*

Preferences. Individuals have logarithmic preferences over consumption, $u(c) = \log(c)$. College students face non-pecuniary costs of completing college that are linear in ability, $\nu_s(\ell) = \nu_{0_s} - \nu_{1_s}\ell$, where $\{\nu_{0_s}, \nu_{1_s}\}$ vary across college tiers $s = \{H, L\}$.

Application costs. The application process involves a financial cost and a non-pecuniary psychic cost, both increasing in the number of applications. The financial cost is linear in the number of applications, $\psi(n) = \psi_0 + \psi_1(n-1)$ for $n \geqslant 1$, where ψ_0 is the cost of the first application and ψ_1 the fee per additional application.[28] The psychic cost has a tier-specific fixed component, incurred when the student prepares her first application to a college of that tier, and a common linear component in the number of applications: $\phi(n_H, n_L) = \phi_H \mathbb{1}\{n_H \geqslant$

[27] As instructional spending per student is limited by the tuition cap, transfers, and endowment income net of fixed costs, and average ability by the most able students in the population, each college tier's value added is bounded.

[28] CommonApp, a platform used to submit college applications, grants fee waivers to students who meet certain criteria (e.g., participation in the federal free or reduced price lunch program), and almost half of Common App colleges do not charge an application fee. We therefore set ψ_1 to the average posted fee without score-reporting and other costs, so it should be read as an upper bound on the marginal financial cost for many students.

$1\} + \phi_L \mathbb{1}\{n_L \geqslant 1\} + \phi_1 (n_H + n_L - 1)$ for $n_H + n_L \geqslant 1$. The fixed components capture the effort of researching and preparing a first application to each tier (selective colleges may require more elaborate essays), while ϕ_1 captures the marginal effort of each additional application once materials can be reused.

Assessments and admission noise. The colleges' assessment is normally distributed around the student's own assessment, with variance σ_g^2, i.e., $\sigma|\ell \sim N(\ell, \sigma_g^2)$, truncated at zero. The idiosyncratic evaluation noise is normal with mean zero and variance σ_ϵ^2, so that the admission probability at a college of tier s is $\pi_s(\sigma) = \Phi((\sigma - \underline{\sigma}_s)/\sigma_\epsilon)$.

College technology. College quality is a Cobb–Douglas function of instructional spending per student and the average ability of the student body, $Q(I/\kappa, L/\kappa) = (1/\kappa)\, I^{1-\rho_L} L^{\rho_L}$, where ρ_L is common across both college tiers.

College budget constraint. The budget constraint differs across college tiers in terms of fixed costs, endowment, and government transfers. The college's problem and the IPEDS data identify only the *net* resource function $E_s(\kappa) - C_s(\kappa)$, so we assign all enrollment-varying components to endowment income and all fixed components to costs. Endowment income is a quadratic polynomial in tier enrollment, in the spirit of Epple et al. (2017), $E_s(\kappa) = E_{1_s}\kappa + E_{2_s}\kappa^2$, with $E_{2_s} < 0$, so that net endowment income per student declines as the college grows.[29] Fixed costs are $C_s(\kappa) = C_{0_s}$, so any constant component of endowment income is netted against the fixed cost. The government subsidizes colleges on a per-student basis, $Tr_s(\kappa) = Tr_{1_s}\kappa$.

Need-based aid. To qualify for Pell grants, students must demonstrate sufficient financial need. The grant covers the gap between students' expected family contribution (EFC) and the net cost of attendance up to a maximum set by the Department of Education. Pell grants are thus modeled as $P(y) = \max\{P_0 - y, 0\}$, and exogenous private grants as $Gr(y) = \max\{g_0 - g_1 y, \underline{g}\}$.

4.2. *Data*

Three datasets are used to estimate the model: (i) the IPEDS for college-level finances (tuition revenue, government appropriations, endowment income, and costs); (ii) the HSLS for the distribution of student characteristics and the application and enrollment moments (both described in Section 2.1); and (iii) the 2012 cohort of the Beginning Postsecondary Students Longitudinal Study (BPS), which details all financial aid received by students who began college in 2011–2012 and is used to set non-institutional grants, the responsiveness of financial aid to SAT scores and parental income, and the tuition caps (all dollar amounts are expressed

[29] The quadratic term is observationally equivalent to rising marginal operating costs. It is also what makes an interior scale optimal: a college does not expand without bound because growth dilutes per-student resources, consistent with the puzzle of elite colleges' flat enrollment documented by Blair and Smetters (2021).

in 2013 dollars). The income gradients of applications and enrollment documented in Section 2 are deliberately left untargeted and serve as validation.

4.3. *External parameters*

Table 4.1 lists the parameters set outside the estimation: the distribution of students' characteristics, $\hat{\mu}(y, \ell)$, discount factor, application smoothing parameter, returns to ability, and graduation probabilities, $\{\beta, \lambda_a, \alpha_s, \alpha_w, \delta_H, \delta_L\}$, borrowing constraints, grants, and financial application costs, $\{\underline{a}_s, \underline{a}_w, P_0, g_0, g_1, \underline{g}, \psi_0, \psi_1\}$, the number of colleges and their budget constraint, $\{N_H, N_L, C_{0_H}, C_{0_L}, E_{1_L}, E_{2_H}, E_{2_L}, Tr_{1_H}, Tr_{1_L}, \overline{T}_H, \overline{T}_L, \underline{\sigma}\}$, and aggregate prices, $\{R, w\}$.[30]

Distribution of student characteristics. The distribution of student characteristics over income and ability, $\hat{\mu}(y, \ell)$, is informed by the HSLS data. Parental income, y, corresponds to the student's Expected Family Contribution (EFC), i.e., the resources available to attend college before financial aid, and so the empirical counterpart of the model's parental transfers.[31] The student-side assessment, ℓ, corresponds to the fitted values of a regression of SAT scores on high-school GPA and rich demographic and school controls, which is the component of test performance the student can forecast from her own record.[32] The college-side assessment, σ, is the committee's reading of the full application file and our empirical proxy is the realized SAT score. The construction anchors the two assessments and disciplines their joint distribution (Appendix C.1). What the proxy requires is that meaningful assessment risk remain at the application stage, not that students ignore their scores, i.e., a student who knows her SAT still does not know how committees will weigh her full file.

Timing, preferences, and aggregate prices. The first period accounts for 4 years (time spent in college) and the second for 60 years (time spent working). Appendix C.2 maps a full life-cycle model into the two-period model.[33] The gross interest rate is set to $R = (1.0386)^4$, the 2013–2014 Federal Student Loan Program rate compounded over the four-year period, and the discount factor to $\beta = 1/R$. The wage rate, $w = 2.7$, reproduces the average wage in the Current Population Survey. The scale parameter of the application-stage taste shock is fixed at a large value, $\lambda_a = 40$, so that it acts purely as a computational smoothing device (doubling or halving λ_a leaves the equilibrium unchanged to four digits).

[30] The selective tier's endowment slope, E_{1_H}, is estimated internally to govern the selective tier's size.

[31] The EFC is determined according to rules set by the Department of Education using the FAFSA. It is observed for students who completed the FAFSA and attended college. For the remainder we compute it from reported household income using the Department of Education's formula (Appendix C.1).

[32] The regression deliberately excludes parental income, so that ℓ is an ability index rather than an income proxy; the signal law $g(\sigma|\ell)$ is income-invariant by assumption.

[33] Since the first period corresponds to four years, model units are rescaled by \$40,000. For example, $y = 1$ corresponds to an Expected Family Contribution of \$10,000 per year over four years.

Returns to ability. The parameters governing the labor market returns to ability are taken directly from Abbott et al. (2019), who find a higher ability gradient for wages among college graduates than among high school graduates. Averaging their estimates for men and women, we set $\alpha_s = 0.78$ for all college graduates and $\alpha_w = 0.55$ for workers who did not attend college.

Dropout risk. We abstract from dropout risk and set $\delta_H = \delta_L = 1$. Completion is a post-enrollment margin, so the moments that discipline the model (applications, admissions, and pricing) carry little information about it, and constant within-tier completion rates would largely be absorbed by the efficiency parameters ξ_s. The abstraction is also conservative: completion rises with ability (Bound et al., 2010), so dropout raises the value of sorting high-ability students into high-value-added colleges. Section 7 relaxes this assumption, reestimating the model with graduation probabilities that rise with ability at tier-specific gradients (completion risk strengthens the sorting forces at the center of the analysis).

Financial application cost. The average application fee in IPEDS is about \$50 (Table A.2) and the SAT registration fee in 2013 was \$51. We set the cost of the first application, inclusive of test registration, to $\psi_0 = 0.0025$ (\$100) and the fee per additional application to $\psi_1 = 0.00125$ (\$50). These costs are incurred once over the four-year period and are small relative to tuition.

Need-based grants. The maximum Pell Grant in 2013 was \$5,645 per year, so $P_0 = 0.5645$. The exogenous grant function is estimated from the BPS data, with $g_0 = 0.43$, $g_1 = 0.19$, and $\underline{g} = 0.15$ (\$1,500 per year). Appendix C.3 provides the details.

Borrowing constraints. The aggregate loan limit of the Federal Student Loan Program for dependent undergraduates is \$31,000, so the student borrowing limit is $\underline{a}_s = -0.775$, which corresponds to \$31,000 over four years. Workers who did not attend college face no risk, so their constraint does not bind and $\underline{a}_w = 0$.

Number of colleges. We set the number of colleges per tier to $N_H = N_L = 4$, so that portfolios of up to eight applications are feasible. This covers the observed support of application counts as 90% of students in the HSLS submit six or fewer applications, and students in the top SAT decile submit about five. Resolving the model at the estimated parameters with $N_s = 5$ or 6 leaves prices, sorting, admission standards, and value added essentially unchanged (value added moves by less than 1.5%), while the average number of applications rises mechanically with the number of available colleges (from 2.7 to 3.5 and 4.2).

Colleges' budget constraint. Costs and endowments are estimated using IPEDS data (Appendix C.4). The quadratic terms of the endowment function are $E_{2_H} = -5.8$ and $E_{2_L} = -0.44$, which implies that net endowment income per student falls much more steeply with scale at the selective tier, consistent with its small equilibrium size. The linear component $E_{1_L} = 0.11$ reproduces nonselective colleges' measured endowment income at observed enrollment (E_{1_H} is internally estimated to match the selective enrollment share). Fixed costs are $C_{0_H} = 0.086$

and $C_{0_L} = 0.25$. Estimating the linear relationship between government transfers and enrollment on the IPEDS sample implies that selective colleges receive more government subsidies per enrolled student than nonselective colleges (about \$10,000 vs. \$6,200 per year), so we set $Tr_{1_H} = 1.0$ and $Tr_{1_L} = 0.62$.[34] The BPS data pin down the tuition caps $\overline{T}_s$. Average tuition net of college-specific financial aid flattens at about \$25,000 per year at selective colleges and \$12,000 at nonselective ones, so we set $\overline{T}_H = 2.5$ and $\overline{T}_L = 1.2$.

Admission standard. The lower bound on the admission standard is set four standard deviations of the evaluation noise below the lowest assessment on the grid, $\underline{\sigma} = \sigma_{\min} - 4\sigma_\epsilon$, so that a college at the bound admits every applicant with probability ≈ 1. At the estimated parameters, $\underline{\sigma} = -0.906$. Because colleges admit on $\sigma + \epsilon_j$, an admission standard of 0 would still reject roughly two in five of the lowest-assessment applicants through the evaluation noise alone.

TABLE 4.1

EXTERNAL PARAMETERS

Variable	Description	Value	Source
Distribution of student characteristics			
$\mu_{\ell_0}, \sigma^2_{\ell_0}$	Log-ability distribution at y = 0, mean, variance	-0.144, 0.155	HSLS
μ_y, μ_{ℓ_1}	Mean distribution at y > 0, income, ability	-0.301, 0.135	HSLS
$\sigma^2_y, \sigma^2_{\ell_1}, \sigma_{y\ell}$	Variance and covariance at y > 0, income, ability, income–ability	2.825, 0.138, 0.174	HSLS
Preferences			
β	Discount factor (per 4-year period)	0.859	$\beta R = 1$
R	Gross interest rate (per 4-year period)	1.0386^4	USDE
λ_a	Application-stage smoothing parameter	40	Normalization
Labor market			
w	Wage rate	2.700	CPS
α_s, α_w	Returns to ability, college, high-school graduate	0.780, 0.550	Abbott et al. (2019)
δ_H, δ_L	Graduation prob., selective, nonselective	1.000, 1.000	Normalization
Application costs, grants, and borrowing constraints			
ψ_0, ψ_1	Financial application cost, first app., add. app.	0.0025, 0.00125	IPEDS, College Board
P_0	Pell Grant maximum	0.5645	USDE
$g_0, g_1, \underline{g}$	Exogenous grant, intercept, slope, minimum	0.430, 0.190, 0.150	BPS
$\underline{a}_s, \underline{a}_w$	Borrowing limit, student, worker	-0.775, 0.000	USDE, Normalization
College market structure, costs, and revenues			
N_H, N_L	Number of colleges, selective, nonselective	4, 4	HSLS
C_{0_H}, C_{0_L}	Fixed costs, selective, nonselective	0.086, 0.250	IPEDS
$E_{2_H}, E_{1_L}, E_{2_L}$	Endowment, selective (quadratic), nonselective (per student, quadratic)	-5.800, 0.110, -0.440	IPEDS
Tr_{1_H}, Tr_{1_L}	Gov. transfers per student, selective, nonselective	1.000, 0.620	IPEDS
$\overline{T}_H, \overline{T}_L$	Tuition sticker price, selective, nonselective	2.500, 1.200	BPS
$\underline{\sigma}$	Admission standard lower bound (open admission)	-0.906	Normalization

4.4. *Estimated parameters*

The remaining parameters relate to the disutility of applying to colleges, $\{\phi_H, \phi_L, \phi_1\}$, the disutility of attending college, $\{\nu_{0_H}, \nu_{0_L}, \nu_{1_H}, \nu_{1_L}\}$, the extreme value scale parameter for choosing among admission offers, λ_ϵ, the variance of the colleges' assessment and the standard deviation of the evaluation noise, $\{\sigma^2_g, \sigma_\epsilon\}$, the college-tier efficiency and the quality elasticity, $\{\xi_H, \xi_L, \rho_L\}$, and the selective tier's endowment slope, E_{1_H}. These parameters are

[34] Selective colleges receive more appropriations per student even though a smaller share of their enrollment is in public institutions (55% vs. 73%) because the public colleges in the selective tier are disproportionately large state flagships with high per-student appropriations.

jointly estimated by minimizing the distance between data and model moments

$$\min_{\Theta} \sum_{j=1}^{J} \left[\frac{M_j^{\text{data}} - M_j(\Theta)}{M_j^{\text{data}}} \right]^2,$$

where M_j^{data} denotes the jth empirical moment and $M_j(\Theta)$ the corresponding model moment at the vector of estimated parameters Θ. Because all choice probabilities are available in closed form, the objective function is deterministic and smooth in Θ and the model moments are computed exactly. Table 4.2 presents the estimated parameters with standard errors and the moments that most identify them.[35]

Application disutility costs. The fixed components of the disutility of applying, ϕ_H and ϕ_L, are identified by the fractions of students applying to each college tier only and to both tiers. In the HSLS, 2.2% of high-school graduates apply only to selective colleges, 45.0% apply only to nonselective colleges, and 16.9% apply to both.[36] The marginal cost of an additional application is identified by the level and ability profile of application counts (2.4 among all students, and about 5 among students in the top SAT decile), with ϕ_1 estimated at zero with tight standard errors, so the data attribute the disutility of applying entirely to the per-tier fixed components.[37]

Attendance disutility costs. The disutility cost of attending college varies by college tier and depends linearly on ability. Enrollment decisions conditional on being accepted help identify the intercepts (ν_{0_H} and ν_{0_L}), while the relative application rates across SAT quintiles help identify the slopes (ν_{1_H} and ν_{1_L}). The estimates imply that attending a selective college carries the lower non-pecuniary cost for nearly all students. The corresponding consumption equivalents (Table C.4 in Appendix C.6) are smaller than those estimated in the literature (e.g., Abbott et al. (2019)), because the model accounts for the fact that students must apply and be admitted in order to attend, i.e., when admission itself screens out many students, a relatively small psychic cost suffices to match non-enrollment.

Enrollment taste shock. The scale parameter of the enrollment shock, λ_c, guides the option value of having several offers to choose from. When it is low, a large fraction of students who applied only to nonselective colleges decline their offers because their applications were motivated by option value rather than by the value of attending. Thus, enrollment conditional on applying only to nonselective colleges identifies λ_c.

Assessment variance and evaluation noise. The variance of the colleges' assessment, σ_g^2, is chosen to match the average responsiveness of tuition (minus institutional grants) with respect

[35] As the moments are drawn from three datasets whose cross-moment covariances we cannot estimate, we weight proportional deviations equally (see Appendix C.5 for details).

[36] These shares are computed from recorded portfolios, which the HSLS censors at three listed colleges.

[37] Table C.4 in Appendix C.6 reports the implied consumption equivalent of removing the disutility of applying.

to SAT scores, controlling for EFC and college fixed effects, using BPS data.[38] The evaluation noise, σ_ϵ, governs how much admission risk remains for a given assessment and hence both the return to submitting several applications to the same tier and the steepness of admission rates in test scores. Given σ_g^2 and ϕ_1, it is identified by the admission rate at selective colleges of applicants in the top SAT decile relative to the average admission rate among all applicants. The noisier the evaluation, the flatter the admission profile in test scores and the more applications high scorers submit. Although several of these moments vary with the SAT, they identify distinct parameters through distinct margins. The level of application counts pins down ϕ_1, the slope of admission rates in SAT scores identifies σ_ϵ, and the tier composition of applications across scores determines ν_{1_S}, holding the others fixed.

TABLE 4.2

ESTIMATED PARAMETERS AND TARGETED MOMENTS

Variable	Description	Estimate (s.e.)	Targeted moment	Data	Model	Source
Preferences						
ϕ_H, ϕ_L	Disutility of applying, selective, nonselective	0.189, 0.154	% applying to selective colleges only	2.199	2.269	HSLS
		(0.026), (0.024)	% applying to nonselective colleges only	45.036	41.139	HSLS
			% applying to both tiers	16.944	14.586	HSLS
ϕ_1	Disutility per additional application	0.000	Average number of applications	2.400	2.727	HSLS
		(0.019)	Average number of applications, top SAT decile	5.000	5.656	HSLS
ν_{0_H}, ν_{0_L}	Disutility of attending (intercept), selective, nonselective	1.737, 3.035	Selective enrol. rate (admitted to both tiers)	0.588	0.686	HSLS
		(0.311), (0.430)	Nonselect. enrol. rate (applied to both, reject. at select.)	0.931	0.925	HSLS
ν_{1_H}, ν_{1_L}	Disutility of attending (slope), selective, nonselective	0.324, 0.857	Selective rel. application rate (SAT quintile 5/4)	3.360	2.884	HSLS
		(0.268), (0.396)	Nonselect. rel. application rate (SAT quintile 4/2)	1.839	1.772	HSLS
λ_c	Extreme value scale when choosing offers	4.494	Nonselect. enrol. rate (applied to nonselect. only)	0.886	0.789	HSLS
		(0.267)				
Technology and costs						
σ_g^2	Variance of colleges' assessment	0.129	Tuition responsiveness to SAT ($10k per sd)	-0.102	-0.105	BPS
		(0.013)				
σ_ϵ	Std. dev. of evaluation noise	0.239	Selective admission rate, all applicants	0.734	0.742	HSLS
		(0.378)	Selective admission rate, top SAT decile/all	1.171	1.202	HSLS
ρ_L	Quality elasticity of ability	0.807	Avg net tuition, selective ($10k)	1.410	1.588	BPS
		(0.026)	Avg net tuition, nonselective ($10k)	0.500	0.614	BPS
			Instr. spending per student, selective ($10k)	2.075	1.650	IPEDS
			Instr. spending per student, nonselective ($10k)	0.860	0.537	IPEDS
ξ_H, ξ_L	Value-added efficiency, selective, nonselective	1.139, 1.657	College wage premium	0.600	0.600	Abbott et al. (2019)
			Cross-tier value-added gap (logs)	0.113	0.113	Chetty et al. (2020)
E_{1_H}	Net endowment per student, selective	0.535	% enrolled in selective colleges	10.030	9.178	HSLS
		(0.131)				

Note: All parameters are estimated jointly. Each is grouped with the moments that mostly identify it per sensitivity analysis. Standard errors in parentheses, from the minimum-distance sandwich formula with the exact moment Jacobian and data-moment variances built from the surveys' cell sizes (Appendix C.5).

College quality elasticity. The elasticity of college quality with respect to students' ability, ρ_L, is identified by the admission-rate profile together with the levels of net tuition and instructional spending. So, a higher ρ_L deepens the ability discount, which shapes who is admitted and enrolls at the selective tier and lowers tuition revenue, and hence spending, per student. The estimated ρ_L implies that colleges value ability heavily (roughly four-fifths of college quality and the remaining fifth is instructional spending), which is what the steep tuition–SAT gradient and the dollar moments jointly demand.

[38] At the estimated σ_g^2, the correlation between a student's own assessment ℓ and the colleges' assessment σ is 0.76. Its empirical counterpart, that is, the correlation between realized SAT scores (the counterpart of σ) and their fitted values from the regression on GPA and controls (the counterpart of ℓ) is 0.74 in the HSLS (Appendix C.1). Since σ_g^2 is identified from a price moment, this serves as an untargeted check on the implied informativeness of the two assessments' joint distribution.

College efficiency. Colleges' efficiency parameters, ξ_H and ξ_L, are backed out from two conditions. First, the enrollment-weighted average college wage premium equals 0.6, in line with Abbott et al. (2019), which can be written as $0.16 \log Z_H + 0.84 \log Z_L = 0.6$, where 0.16 is the selective tier's share of college enrollment. Second, the cross-tier value-added gap satisfies $\log Z_H - \log Z_L = 0.113$: the observed gap in the tiers' median earnings (Table A.2) net of the difference in their students' ability. The adjustment removes the part of the earnings gap that reflects who enrolls rather than what colleges add, in the spirit of Chetty et al. (2020). Given the equilibrium $(I_s/\kappa_s, L_s/\kappa_s)$, these two equations determine (ξ_H, ξ_L). Finally, the linear endowment term of selective colleges, E_{1_H}, is estimated to match their enrollment share. A higher endowment per student allows the selective tier to sustain a larger student body at given quality.

4.5. *Model fit and validation*

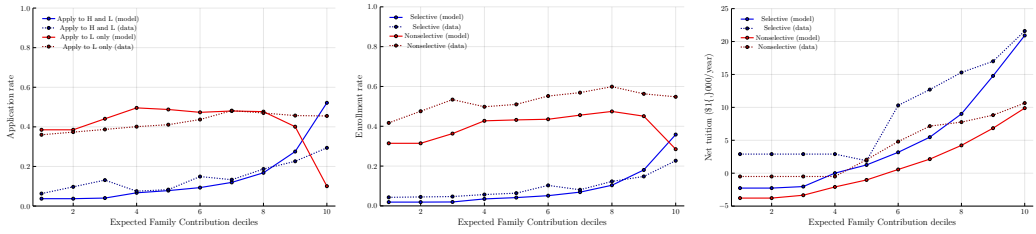
Student application and enrollment. The model accounts for the low application and enrollment rates of low-income students despite the presence of substantial financial aid. Panels (a) and (b) in Figure 4.1 shows model-predicted average application and enrollment rates alongside the HSLS data by EFC decile. The model reproduces the rising application rate to the selective tier and the rising enrollment at both tiers across the EFC distribution. Although the means of application and selective-enrollment rates are targeted in the estimation, their variation by EFC as well as the enrollment rates at nonselective colleges are not. The visible exception is the top EFC decile, where the model sends half of students to both tiers (against 29% in the data) and understates nonselective-only portfolios (10% against 46% in the data). At the estimated zero marginal application cost, adding a selective application with weakly positive option value is free for these students. And because reported portfolios in the HSLS are truncated at three listed colleges, the data overstate the nonselective-only share and understate the both-tier share.[39]

The intensive margin of applications lines up. Students in the bottom EFC deciles submit 1.8 applications on average against 4.6 in the top decile (about half, as in the data) and the gap narrows sharply conditional on applying (4.1 against 6.4), mirroring the pattern documented in Section 2.3. The one margin the model misses is the income gradient of admission rates. In the model low-income applicants are strongly selected on ability so they are more likely to be admitted to selective colleges than in the data, although the aggregate admission rate is matched well (74% against 73%).[40]

[39] Also, nonpecuniary considerations outside the model (e.g., geography, fit) keep many high-income students in single-tier portfolios in the data.

[40] Matching the unconditional gradient would require income to enter admissions directly rather than only through the applicant pool (e.g., need-aware admission standards or applications whose strength rises with income).

FIGURE 4.1.—College application, enrollment rates, and average net tuition



(a) Application rates (b) Enrollment rates (c) Average net tuition

Note: Panel (a) displays the average application rates by Expected Family Contribution (EFC) decile to nonselective colleges and to both tiers. The fraction of students only applying to selective colleges is small both in the model and the data. Panel (b) displays the average enrollment rates to each college tier. The data moments are from the HSLS. Panel (c) shows the average net tuition (tuition minus all grants and $T(y,\sigma) - P(y) - Gr(y)$ in the model) by Expected Family Contribution decile, in the model and in the BPS data.

College finances, enrollment, and student composition. Table 4.3 compares college-level statistics produced by the model to the ones calculated in the data. The selective tier's enrollment share is matched closely (9.2% vs. 10.0%) and the nonselective share, whose level is untargeted, comes in at 39.5% (against 52.3% in the data). The model's outside option of working attracts more of the marginal nonselective students than the data suggests.[41] The model reproduces the key cross-tier patterns. Average student EFC at selective colleges is about three times that at nonselective ones, instructional spending per student is three times as high, and spending per student at the selective tier exceeds net tuition revenue.

TABLE 4.3

COLLEGE FINANCES, ENROLLMENT, AND STUDENT COMPOSITION

		Selective colleges		**Nonselective colleges**	
		Model	*Data*	*Model*	*Data*
κ_S	% enrolled	9.2	10.0	39.5	52.3
$\mathcal{T}_S/\kappa_S$	Average net tuition ($1k/year)	15.9	14.1	6.1	5.0
I_S/κ_S	Instr. spending per student ($1k/year)	16.5	20.8	5.4	8.6
Y_S/κ_S	Average student EFC ($1k/year)	27.4	30.4	9.8	15.5
y_S	*Parental income Distribution (%)*				
	Q1	4.9	9.6	21.7	23.4
	Q2	12.6	12.1	32.3	16.4
	Q3	16.8	23.1	23.5	28.1
	Q4	65.6	55.2	22.4	32.2

Notes: Tuition is net of institutional aid but not of federal or private grants. EFC levels and income distributions are untargeted.

Table 4.3 also shows the distribution of parental transfers (EFC in the data) within each college. The model generates a substantial low-income presence at the selective tier without any assumed preference for diversity (about 4.9% of selective enrollment from the bottom EFC quartile against 9.6% in the data, and a majority from the top quartile, 65.6% against 55.2%). In this class of models colleges have a strong incentive to enroll top-income students, who pay higher tuition and are on average higher ability. Epple et al. (2017) overpredict the high income share and appeal to unmodeled social objectives such as affirmative action to close the gap. Our application-and-admissions system closes most of it endogenously. Because low-income

[41] The gap is concentrated among below-median-ability students for whom nonselective college and work are substitutes.

students who apply to the selective tier are strongly self-selected on ability, colleges find it profitable to admit them and to offer them large discounts.

Two forces generate this economic diversity endogenously. First, the technology weights ability heavily ($\rho_L = 0.81$), so colleges willingly trade tuition revenue for high-ability students and many of whom are lower income. Second, selection works through applications. Since low-income students are less likely to apply to the selective tier, those who do apply are very high ability in equilibrium, so the college is confident that enrolling them raises its quality and offers them large tuition discounts.

The same selection rationalizes a well-known puzzle. Dale and Krueger (2002) find little earnings return to attending a more selective college among students admitted to the same set of colleges. In the model, students admitted to both tiers who choose the nonselective college are of *higher* average ability than those who accept the selective offer. Two selection forces drive this. First, composition: a low-income student reaches the both-admitted pool only if her ability is exceptional, so enrollment then hinges on the aid offer rather than on her ability. When the offer falls short, the most able students in the pool prefers to enroll in another college. Second, as offers condition on the colleges' assessment, students whose assessment underrates their own (i.e., they have a high ℓ given σ) receive worse financial offers than their valuation of attending and thus prefer decline. The decliners' ability advantage then offsets the selective tier's higher value added in earnings when old. So, comparing earnings across their choices understates the return to selectivity, as Dale and Krueger document.

Panel (c) Figure 4.1 shows the average net tuition in the model and in the data by EFC decile within each college tier.[42] The model reproduces the steep income gradient of net tuition at both tiers and matches its level at the top of the EFC distribution, where the tuition caps bind. In the lower deciles the model offers more financial aid than the data. Net prices turn negative because enrolled low-income students are disproportionately students with high assessments and colleges are willing to offer them larger ability discounts.[43]

5. Policy Counterfactuals

This section uses the estimated model to evaluate three families of policies that operate on different margins of the enrollment process: (i) lowering the cost of applying (fee waivers and streamlined applications); (ii) expanding need-based aid (Pell grants), and (iii) relaxing borrowing constraints. Throughout, colleges reoptimize their tuition schedules, admission standards,

[42] In the data, net tuition is defined as the sticker-price tuition minus all grants available to the student. In the model, net tuition for a student with parental transfers y and assessment σ corresponds to $T(y, \sigma) - P(y) - Gr(y)$.

[43] About −$2,300 per year at selective colleges and −$3,800 at nonselective ones in the bottom three deciles in the model, against roughly zero in the data.

and instructional spending, value added adjusts, and the government budget balances through the labor income tax.

One property of the equilibrium feedback shapes every experiment below. The two college tiers respond asymmetrically to the loop between expected and delivered value added. At the selective tier, a direct improvement is reinforced in equilibrium (i.e., better students attract spending, which attracts further better students), while at the nonselective tier, the re-sorting of students and spending following a policy can have the opposite effect (the Jacobian of $\Phi(Z_H, Z_L)$ in Appendix B.6 quantifies both responses). So subsidies that expand students' resources can reduce enrollment and welfare once colleges respond.

The experiments that follow are summarized along two dimensions: (i) welfare through consumption equivalent variation, and (ii) the allocation of talent through a sorting efficiency index, which measures the share of the aggregate earnings gap between randomly assigned seats and seats perfectly sorted by students' own assessments that the equilibrium allocation closes, holding each counterfactual's enrollments and per student spending fixed.[44] The baseline closes 94% of that gap. The last row of Table 5.1 reports the index for every experiment.

5.1. *Lowering application costs*

Policies that encourage applications are among the cheapest and most widely advocated interventions in this market, and the model is well suited to evaluate them, since the number and composition of applications is an equilibrium object. We consider two experiments. The first eliminates the financial cost of applying altogether ($\psi_0 = \psi_1 = 0$), mimicking universal fee waivers. The second halves the fixed psychic costs ϕ_H and ϕ_L, capturing a streamlined application technology. Table 5.1 reports the results.

Universal fee waivers have negligible effects. The average number of applications rises from 2.73 to 2.74 and enrollment is essentially unchanged. The financial cost of applying is small relative to the other costs and returns at stake, so removing it does not change behavior by much. The operative friction is therefore not the fee itself but the effort cost of engaging with a college tier. Halving the psychic costs of applying, by contrast, raises the average number of applications from 2.7 to 3.3. However, average enrolled ability at selective colleges falls (1.79 to 1.69) and with it selective value added (2.00 to 1.90), while top-decile-ability students (who already submit 6.3 applications in the baseline) expand their portfolios further (6.7), insuring themselves against the noisier competition. This is the model's mechanism at work. As the

[44] Formally, for an equilibrium with allocation m and value added Z, let $Y(m, Z)$ be aggregate lifetime earnings. The index is $\big(Y(m, Z) - Y(m^R, Z^R)\big)/\big(Y(m^S, Z^S) - Y(m^R, Z^R)\big)$, where m^S assigns the tiers' seats to the highest own-assessment students in order, m^R assigns them randomly, and Z^S, Z^R recompute each tier's value added at its reassigned student body through the quality technology, holding spending fixed.

marginal applicants come disproportionately from below the top of the ability distribution, the informativeness of the applicant pool gets diluted at each income level with more applications.

Economic diversity at the selective tier collapses. The bottom-income-quartile share of selective enrollment falls from 4.9% to 1.8%. When every low-income student applies, applying no longer distinguishes the highest-ability among them and colleges withdraw the confident discounts that made their low-income admits attractive. The average welfare change is −0.15% (Table D.2 in Appendix D). The lesson from general equilibrium is that policies that expand applications without targeting ability can worsen both the allocation of talent and the economic diversity they aim to promote, even as they increase access at the application stage.

TABLE 5.1

COUNTERFACTUAL EQUILIBRIA

		Baseline	*Fee waiver*	$\phi_S/2$	*Pell 7.4k*	*Pell 25k*	$\underline{a}_S \downarrow$	$\text{corr}(\ell,\sigma)=0.68$	*AI*	$\text{corr}(\ell,\sigma)=1$
Panel A: Selective colleges										
κ_H	% enrolled	9.18	9.20	8.68	9.44	9.69	9.59	9.18	8.17	7.61
Z_H	Value added	2.00	2.00	1.90	2.03	2.14	2.14	1.95	1.79	1.80
I_H/κ_H	Instr. spending per student ($1k/yr)	16.50	16.55	16.04	17.13	22.38	21.85	16.59	16.50	13.47
L_H/κ_H	Average student ability	1.79	1.79	1.69	1.80	1.80	1.81	1.73	1.56	1.64
$\mathfrak{T}_H/\kappa_H$	Avg. tuition collected per student ($1k/yr)	15.88	15.93	15.68	16.41	21.56	21.06	15.97	16.46	13.88
	Avg. tuition net of all grants ($1k/yr)	13.26	13.27	13.61	13.19	5.90	17.59	13.37	14.81	11.77
$\underline{\sigma}_H$	Admission standard[a]	1.54	1.54	1.46	1.63	1.99	1.93	1.49	1.16	1.02
	Bottom-quartile-income share (%)	4.94	5.38	1.83	6.29	11.85	9.58	4.67	0.00	2.79
	Top-quartile-income share (%)	65.61	65.30	76.59	61.69	37.86	44.66	66.59	88.90	81.94
	Top-decile-ability share (%)	56.95	56.95	46.14	58.02	55.91	57.30	50.17	33.83	43.81
Panel B: Nonselective colleges										
κ_L	% enrolled	39.50	39.30	38.42	37.39	24.79	26.51	38.77	38.25	38.93
Z_L	Value added	1.79	1.78	1.72	1.73	1.46	1.58	1.76	1.70	1.83
I_L/κ_L	Instr. spending per student ($1k/yr)	5.37	5.28	3.97	4.44	1.54	2.20	4.67	3.29	4.76
L_L/κ_L	Average student ability	1.28	1.28	1.31	1.28	1.33	1.36	1.29	1.34	1.35
$\mathfrak{T}_L/\kappa_L$	Avg. tuition collected per student ($1k/yr)	6.14	6.06	4.87	5.47	5.41	5.49	5.52	4.21	5.59
	Avg. tuition net of all grants ($1k/yr)	1.07	0.99	-0.34	-0.83	-17.27	0.23	0.39	-1.05	0.72
$\underline{\sigma}_L$	Admission standard[a]	$\underline{\sigma}$	$\underline{\sigma}$	$\underline{\sigma}$	$\underline{\sigma}$	$\underline{\sigma}$	$\underline{\sigma}$	$\underline{\sigma}$	$\underline{\sigma}$	-0.67
	Bottom-quartile-income share (%)	21.69	21.68	22.49	23.35	28.28	22.53	22.22	22.73	20.03
	Top-quartile-income share (%)	22.44	22.32	20.13	20.37	11.29	20.01	21.44	18.60	24.07
	Top-decile-ability share (%)	6.44	6.45	9.88	6.12	9.25	8.56	8.20	13.16	11.50
Panel C: Economy-wide										
τ	Tax rate (%)	1.73	1.73	1.74	1.91	3.95	1.46	1.74	1.74	1.64
n	Number of applications, all students	2.73	2.74	3.33	2.69	2.55	2.45	2.79	3.21	2.14
	Number of applications, top-decile ability	6.25	6.33	6.74	6.49	7.21	7.14	6.72	6.24	4.57
	Average welfare change (CEV, %)[b]	0.00	-0.04	-0.15	-0.65	-3.65	-1.16	-0.46	-0.68	0.72
	Sorting efficiency (% of feasible gap)[c]	93.97	93.85	89.94	92.42	86.10	93.10	92.87	87.84	95.30

Notes: Each column is a counterfactual general equilibrium at the estimated parameters. [a]An entry $\underline{\sigma}$ means the standard rests at the open-admission lower bound $\underline{\sigma}$ =-0.906, i.e., the college admits every applicant. [b]Consumption equivalent variation relative to the baseline, population-weighted. The welfare impact by income and ability is in Appendix D. [c]Share of the aggregate-earnings gap between random assignment of seats and perfect sorting by assessments that the equilibrium allocation closes, at each counterfactual's seats and spending.

5.2. *Expanding need-based aid*

The maximum federal Pell grant was $5,645 per year for the cohort we study. We evaluate an increase to $7,395 (the 2023–24 statutory maximum) and, as an upper bound, an increase to $25k, which equals the tuition cap at selective colleges and effectively makes selective college free for zero-EFC students conditional on admission. Both experiments are financed with the labor income tax. Table 5.1 reports college-level outcomes. Table D.2 in Appendix D reports the welfare effect by income and ability.

The increase to $7,395 has modest aggregate effects. Enrollment barely moves (9.4% and 37.4% at selective and nonselective colleges, against 9.2% and 39.5% in the baseline), the selective admission standard goes up (1.54 to 1.63), and the labor tax rises from 1.7% to 1.9%.

Colleges capture part of the additional aid as net tuition at selective colleges increases as aid expands. This policy is a transfer from taxpayers to the colleges and the few students whose enrollment it enables, and the average welfare change is -0.7%.

The \$25,000 expansion triggers a large response of the college market. Demand from newly eligible students rises sharply and the selective tier responds by raising its admission standard (from 1.54 to 1.99) and extracting much of the aid through higher prices. Average tuition collected per student rises by more than a third at selective colleges (from \$15,900 to \$21,600 per year), while instructional spending per student rises from \$16,500 to \$22,400. Tuition net of grants falls (from \$13,300 to \$5,900), with taxpayers financing the difference. The surge in college demand is converted into more selectivity and quality rather than additional seats.

The price response is a shift of the tuition schedules and not a composition effect. Evaluating each tier's schedule at the *baseline* student mix, selective colleges raise the tuition charged to Pell-eligible students by 80 cents for every dollar of additional aid.[45] The capture rate is a direct implication of student-specific pricing. A college that observes a student's aid reprices that student one for one up to the tuition cap. Because the caps are held at their baseline levels, the capture reflects smaller institutional discounts (the margin documented by Turner (2017)). To a first order, the pass-through rate of aid to tuition is thus one minus the share of enrollment priced at the cap, about 0.72 at the baseline. The measured 0.80 exceeds that benchmark because the expansion also raises admission standards and improves the admitted pool.

Total enrollment falls from 49% to 34%, driven entirely by the nonselective tier (39.5% to 24.8%). Students holding \$25k of aid direct their demand at selective colleges, driving the selective admission standard up, and nonselective colleges compete on price to retain students. Their net tuition turns negative, per-student spending collapses from \$5,400 to \$1,500, and value added erodes from 1.79 to 1.46. The composition at selective colleges shifts markedly toward lower-income students, with the bottom-quartile share rising from 5% to 12% and the top-quartile share falling from 66% to 38%. Top-decile-ability students insure themselves against the tougher admissions with broader portfolios (6.3 to 7.2 applications), while the economy-wide average falls slightly as discouraged marginal applicants exit. The tax rate needed to finance the program more than doubles, from 1.7% to 4.0%.

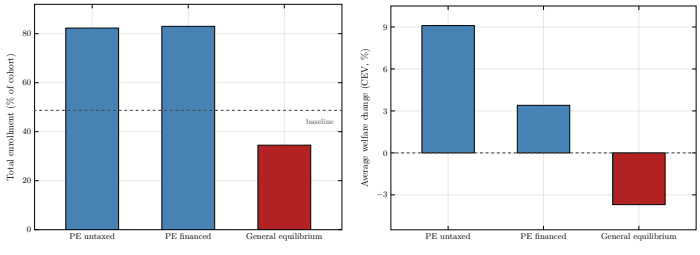
The winners are top-decile-ability students outside the top income quartile (CEVs of $+3.7\%$ in the bottom quartile, declining to $+0.4\%$ in the third) and the losers are everyone else (most heavily upper-middle-ability students in the top income quartile, who are squeezed by higher standards, higher prices, and the higher tax, loosing -12% in CEV). The average welfare

[45] The 0.80 is the enrollment-weighted increase in the selective tier's tuition schedule $T(y, \sigma)$ between the baseline and the counterfactual, holding the student mix at its baseline composition and restricting to (y, σ) cells whose Pell grant increases, divided by the enrollment-weighted increase in the grant. Empirically, Lucca et al. (2018) and Turner (2017) document precisely this pass-through of federal aid into posted and net prices.

change is −3.7%. The allocation of talent deteriorates along with welfare. Sorting efficiency falls from 94% to 86%, the largest decline of any experiment we consider.

The role of the college market's response. To isolate the role of the admissions system and college pricing, we recompute the $25,000 experiment holding colleges' tuition schedules, admission standards, and value added fixed at their baseline values, resolving only the students' application, enrollment, and consumption decisions. With the tax rate also held at baseline, total enrollment rises from 49% to 82% and the average welfare gain is 9.1%. Allowing the tax rate to adjust while still holding college behavior fixed requires a labor tax of 7.6%, which is substantially higher than in general equilibrium and still leaves an average welfare gain of 3.4%. In general equilibrium, the response of tuition and standards turns the 33 p.p. enrollment increase into a 15 p.p. decline and the average welfare change is −3.7%. Government financing thus accounts for a little under half of the gap between the naive and the equilibrium evaluation, while the remainder is the college market's response through tuition, admission standards, and value added. Figure 5.1 summarizes the decomposition.

FIGURE 5.1.—The $25,000 Pell expansion: partial vs. general equilibrium



(a) Total enrollment (b) Average welfare change

Note: "PE untaxed" holds colleges' tuition schedules, admission standards, and value added at baseline and leaves the program unfinanced; "PE financed" additionally balances the government budget through the labor income tax; "General equilibrium" lets colleges and the tax adjust. The dashed line in the left panel is baseline enrollment (48.7%).

Sensitivity to the tuition caps. Because sticker prices might respond to a large aid expansion, Table D.1 in the Appendix recomputes the baseline equilibrium under caps 10% below and 10% above their measured values. The conclusions are qualitatively similar under the new caps.

5.3. *Relaxing borrowing constraints*

Finally, we raise the federal loan limit from $31,000 to $80,000 over four years ($\underline{a}_s = -2.0$), holding all other parameters at baseline.[46] As Table 5.1 shows, the additional borrowing capacity of previously constrained students is largely captured by colleges. Average net tuition rises from $13,300 to $17,600 per year at selective colleges, the selective admission standard jumps

[46]See Lochner and Monge-Naranjo (2012) for the evidence on credit constraints in education and Lochner and Monge-Naranjo (2016) for the design of student loan programs.

from 1.54 to 1.93, top-decile-ability students expand their portfolios from 6.3 to 7.1 applications, total enrollment falls from 49% to 36% (almost entirely out of the nonselective tier), and the selective tier's composition shifts sharply toward lower-income students (the top-income-quartile share falls from 66% to 45%). The impact on welfare mirrors the Pell expansion. Top-decile-ability students gain modestly (CEVs up to +1.3%), everyone else loses, with the average welfare falling 1.2%. Unlike the Pell expansion, the policy reduces the tax rate (from 1.7% to 1.5%). As enrollment falls, so do per-student government transfers, and the fiscal saving is rebated through lower taxes. The pass-through of borrowing capacity into tuition echoes the empirical evidence in Lucca et al. (2018).

6. The Role of Application Informativeness

This section studies the general equilibrium effects of changing the informativeness of applications. Motivated by the widespread suspension of standardized test requirements, the first experiment makes colleges' assessments less informative by increasing its variance by 50%, equivalent to adding roughly 45 SAT points of noise to colleges' reading of each applicant, or to lowering the correlation between the assessments of colleges and students from 0.76 to 0.68. The second experiment aligns the two sides' information, so that colleges and students share the same ability assessment and the only admission risk is the colleges' idiosyncratic evaluation noise. Table 5.1 in Section 5 reports college-level outcomes for all counterfactuals and Table D.3 in Appendix D shows the welfare impact by income and ability.

6.1. *The informational content of the applicant pool*

If only the highest ability among low-income students apply, the assessments colleges form about them are highly informative and colleges offer them generous aid with confidence. To illustrate this mechanism, Table 6.1 shows what happens to the selective tier when all low-income students apply in the same way as their higher-income peers.[47] With the substantial increase in low-ability, low-income applicants, any given assessment of a low-income applicant is now much more likely to come from a low-ability student, and so the selective tier loses the capacity to sort students that sustains its higher value added. The selective admission standard collapses to its lower bound, average enrolled ability falls, and selective value added drops below the nonselective tier's (1.76 against 1.95, from 2.00 and 1.79 in the baseline). High-

[47] Specifically, students in the bottom half of the parental transfer distribution are assigned the application portfolio probabilities of the highest-income students. The equilibrium is then recomputed holding these application pools fixed, while tuition schedules, admission standards, value added, and enrollment decisions adjust. This exercise isolates the effect of the composition of application portfolios on the allocation of students.

ability low-income students do enroll at selective colleges at a higher rate (0.41 to 0.49, Table 6.1), but those colleges no longer deliver a quality premium.

The composition of the applicant pool thus governs how informative colleges' assessments are, and hence tuition. Through the lens of the model, low-income students receive high financial aid because they are less likely to apply, resolving the "puzzle" of low application rates coexisting with generous aid. In this exercise, high-ability low-income students are not crowded out of selective colleges (their selective enrollment rate rises from 0.41 to 0.49), but the colleges they enter no longer deliver a quality premium. When applications respond endogenously, the bottom-quartile income share of selective enrollment falls from 4.9% to 1.8%, reversing the economic diversity the policy aims to promote.

TABLE 6.1

UNIFORM APPLICATIONS BY INCOME

		Baseline	*Uniform applications*
κ_H, κ_L	% enrolled, selective, nonselective	9.18, 39.50	7.77, 39.75
Z_H, Z_L	Value added, selective, nonselective	2.0, 1.79	1.76, 1.95
$L_H/\kappa_H, L_L/\kappa_L$	Average enrolled ability, selective, nonselective	1.79, 1.28	1.72, 1.32
σ_H, σ_L	Admission standard, selective, nonselective	1.54, $\underline{\sigma}$	$\underline{\sigma}$, $\underline{\sigma}$
	Enrollment rate, top-decile-ability students (bottom half of transfers), selective, nonselective	0.41, 0.59	0.49, 0.50

Note: Students in the bottom half of the parental transfer distribution are assigned the application portfolios of the highest-income students. The equilibrium is recomputed with only the application portfolios fixed. An entry $\underline{\sigma}$ denotes that the college admits every applicant, which happens for both college tiers in the counterfactual.

6.2. *Noisier versus perfectly informative assessments*

Noisier assessments. Consider the use of standardized tests as part of the technology that makes colleges' assessments informative, so that a counterpart of removing them is an increase in the assessment variance σ_g^2.[48] We consider a 50% increase in the variance at both tiers, which lowers the correlation between students' and colleges' ability assessments from 0.76 in the baseline to 0.68.[49] Noisier ability assessments change the market on three margins. First, sorting deteriorates. The share of top-decile-ability students at selective colleges falls from 57% to 50%, average enrolled ability falls (1.79 to 1.73), and the value-added gap between the tiers narrows (Z_H falls from 2.00 to 1.95 while Z_L slips to 1.76). Second, the selective admission standard eases slightly (1.54 to 1.49) as a marginal applicant's assessment is now less indicative of low ability. Third, and new relative to models without an application margin, applications rise and they rise most at the top. Top-decile-ability students expand their portfolios from 6.3 to 6.7 applications, insuring against noisier admissions. As ability-based discounts shrink and the value added of college falls, top-decile-ability students lose 2.2 to 2.5% of lifetime consumption at every income level, while students of middling ability are nearly unaffected.

[48] Bleemer (2024) and Borghesan (2025) also study the impact of admission regimes that de-emphasize test scores.

[49] In practice, test optional adoption was concentrated among selective colleges. In our quantitative analysis, the assessment's empirical counterpart is the realized SAT, which all colleges observe. Cross-college differences in screening are instead absorbed by the college-specific evaluation noise.

The AI scenario. The rapid adoption of AI in writing is expected to lower the fixed application costs and by making application materials easier to produce and more uniform, it also plausibly drives the two sides' ability assessments apart as a college reading machine-polished materials learns less about the student behind them. We compute an experiment that activates both margins at once: fixed application costs are halved and the assessment variance rises by 50%, lowering the correlation between students' and colleges' assessments from 0.76 to 0.68 (column *AI* of Table 5.1). The two margins compound. Selective value added falls below the values of cheaper applications alone and of noisier assessments alone, the selective admission standard eases from 1.54 to 1.16, applications rise from 2.7 to 3.2 per student, the top-decile-ability share at selective colleges falls from 57% to 34%, and sorting efficiency falls to 87.8 (a larger loss than the two experiments' losses combined). The average welfare cost, 0.68% of lifetime consumption, likewise exceeds the sum of its components' and is borne by high-ability students, who lose 3.3 to 4.4% at every income level (Table D.3 in Appendix D). The composition effect is strong as the bottom quartile income share of selective enrollment falls to zero. With the pool diluted and the assessments degraded at the same time, applying no longer marks a low-income student as exceptional. The two margins of AI adoption are complements in destroying the information on which the selective tier and its low-income students rest.

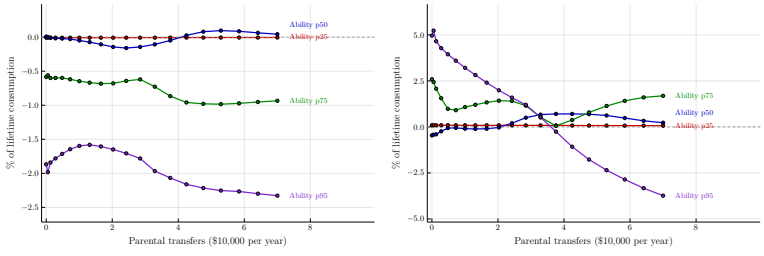
Perfectly informative assessments. At the opposite extreme, we now raise the correlation between the two ability assessments from 0.76 to 1, so that colleges and students agree on a student's ability (i.e., $\ell = \sigma$). Students now know ex ante the price they will face at every college, and the only remaining admission risk is the idiosyncratic evaluation noise. Instead of better sorting, perfect information dissolves much of what makes the selective tier special. Its informational advantage is now available to every college and the value added of the two tiers converges. The selective admissions standard falls sharply (1.54 to 1.02), net tuition and instructional spending fall, and selective enrollment shrinks from 9.2% to 7.6%. With little quality difference left to chase and prices known in advance, students stop paying for selectivity and stop hedging as the average application portfolio shrinks from 2.7 to 2.1, and top-decile-ability students cut from 6.3 to 4.6 applications.

Table D.3 in Appendix D and Figure 6.1 show the welfare consequences by parental income and ability. The average gain is +0.7% at zero fiscal cost (the tax rate falls slightly). The top-decile-ability students in the bottom income quartile gain 5.3% of lifetime consumption, with the gain declining monotonically in income (4.7%, 3.7%, 0.1%). Their highest-income peers, who no longer benefit from being identifiable in a noisy pool, gain essentially nothing.

Together with the previous experiment, these results imply that the noise in the admissions system operates like a regressive tax on high ability. Noisier assessments cost top-ability students 2 to 2.5% regardless of income, while removing the noise delivers its largest gains to

high-ability lower-income students. Contrast this with the $25k Pell expansion of Section 5.2, which delivered a comparable gain to the same group at an average welfare *cost* of 3.7%. Better information delivers what the Pell grant promises at zero fiscal cost, although the model does not price the information improvement itself.

FIGURE 6.1.—Welfare effect of changing the informativeness of assessments



(a) A 50% higher assessment variance (b) Perfectly informative assessments

Note: Consumption-equivalent variation by parental transfers, for students at the 25th, 50th, 75th, and 95th percentiles of the ability distribution. Positive values are welfare gains.

The output cost of admission noise. How much talent misallocation does the admission system's noise cause in output terms? The question is the college market counterpart of the economy-wide allocation of talent accounting in Hsieh et al. (2019). Holding fixed the number of students each tier enrolls and its per student instructional spending at their baseline values, perfect sorting raises aggregate lifetime earnings by 1.2%, of which about 0.5 p.p. is the direct effect of matching stronger students to the selective tier and the remainder 0.7 p.p. the peer-effect driving the rise in the selective value added. Random assignment, by contrast, lowers aggregate earnings by 18%.

The last row of Table 5.1 reports the sorting efficiency index. Perfect information is the only counterfactual that raises it (to 95.3%), while each policy of the previous section lowers it. This means that information and subsidies move the market in opposite directions. Making assessments informative reallocates the existing seats toward the students who benefit from them most, while expanding aid increases students' demand, raising tuition and admission standards. The welfare effect by income and ability makes the contrast concrete. Both interventions deliver large gains to top-ability students in the poorest quartile (5.3% under perfect information and 3.7% under the Pell expansion, Table D.3), but information achieves it by improving the allocation of talent (0.7% average welfare gain) while the subsidy achieves it at everyone else's expense (average welfare falls by 3.7% and everyone outside the top ability decile loses).

7. GRADUATION RISK

The analysis so far treats enrollment as the end of the risky part of a college career: everyone who enrolls graduates. In the data, barely three in five entrants at four-year colleges finish

within six years, and completion rises steeply with both test scores and college selectivity. This section extends the model with graduation probabilities that vary by tier and ability, reestimates it, and asks what completion risk changes. Three results emerge. First, the extended model fits the baseline moments slightly better than the benchmark while matching eight new completion moments. Second, completion risk overturns the benchmark's prediction for perfectly informative assessments, that is, instead of converging in terms of value added, the two tiers diverge. Third, completion risk sharply limits the scale of aid policies the market can absorb. In the benchmark, the demand an aid expansion creates comes from marginal students, so colleges cannot convert the aid into quality: nonselective value added falls in every aid experiment of Section 5. Under graduation risk, those students expect not to finish, so aid does not draw them in. The money raises tuition and spending instead. This section leaves the benchmark analysis of Sections 5 and 6 as is and its lesson is that completion risk strengthens the sorting forces at the heart of the paper and changes who gains from information policy.

7.1. *Graduation probabilities and estimation*

A student at tier s now graduates with probability $\delta_s(\ell) = \Lambda(a_s + b_s\,\ell)$, where Λ is the logistic function, so that completion rises with the student's assessment ℓ at tier-specific gradients. Everything else in the model is unchanged. The parameters (a_s, b_s) are estimated jointly with all others by adding eight moments: six-year bachelor's completion rates by tier and admissions-test quartile from the 2012 BPS, the cohort that also disciplines net tuition. Because the BPS records realized test scores (the counterpart of colleges' assessments σ), the model moments condition on σ rather than on ℓ. The information structure handles the gap between the observed test score and the student's assessment that governs completion.[50] The level of completion within a tier pins down a_s, while its slope across quartiles pins down b_s.[51]

Table C.5 in Appendix C.7 reports the fit. The model matches the selective tier's completion profile closely (0.62 to 0.89 across quartiles against 0.62 to 0.85 in the data) and draws a monotone profile through the noisier nonselective data. The estimated gradients are $(a_H, b_H) = (-4.85,\ 3.95)$ and $(a_L, b_L) = (-4.38,\ 3.26)$, so completion rises steeply and significantly with the student's assessment at both tiers. The fit of the original moments improves slightly as completion risk relieves pressure on the non-pecuniary value of attending college.

[50] For tier s and quartile q, the moment is $\mathbb{E}[\delta_s(\ell) \mid \text{enrolled in } s,\ \sigma \in Q_q]$, integrating over the model's joint distribution of (ℓ, σ) among enrollees. Because the two assessments are imperfectly correlated, completion rises less steeply across test-score quartiles than across ℓ, and the estimated slopes b_s correct for that attenuation.

[51] The two conditions backing out college (ξ_H, ξ_L) are restated on expected earnings.

7.2. *Information policy under graduation risk*

Completion risk reverses the benchmark's answer to the perfect-information experiment of Section 6. In the benchmark, aligning colleges' assessments with students' own compresses the market, i.e., the nonselective tier retains the students with strong assessments it can now identify and the tiers' value added converges. With graduation risk the same experiment makes the tiers diverge. Selective value added rises while nonselective value added falls (Z_H from 2.21 to 2.25 and Z_L from 1.65 to 1.48, a widening of the gap by 37%), and nonselective per-student spending falls by 42%. The mechanism runs through who is worth competing for. Under graduation risk, a student with a strong assessment is doubly valuable to the nonselective tier, for her quality contribution and because she would almost surely finish. When assessments become perfectly informative, the selective tier identifies and takes precisely those students, and what remains of the nonselective pool no longer sustains its value added or its tuition revenue.

The market as a whole sorts better: the sorting-efficiency index of Section 5 rises from 90.6% to 93.2%, a larger improvement than in the benchmark (94.0% to 95.3%). The welfare consequences nonetheless invert (Table D.4 in Appendix D). In the benchmark, perfect information raised average welfare by 0.7% and delivered its largest gains to top-decile-ability students from the poorest families. Under graduation risk the average effect turns negative (−0.9%) and the incidence flips: top-decile-ability students still gain, but their gains now rise with income (from 2.5% in the bottom quartile to 5.6% in the top), while students in the 60–90% ability range lose 2.7 to 4.5% of lifetime consumption.[52]

7.3. *Graduation risk and the market's capacity to absorb aid*

In the benchmark, the demand an aid expansion creates at the nonselective tier comes from students with weak assessments, so the tier cannot convert the aid into quality. In Table 5.1, nonselective value added falls in every aid experiment, and with it the willingness to pay that colleges capture through tuition. Graduation risk removes that brake. Because students discount college by their own completion probability, the marginal student with a weak assessment (and therefore poor completion odds) does not respond to aid at all. The aid fuels instead higher tuition and per-student spending on the students who remain. Each of these experiments has a well-defined absorption limit, above which a fixed point in college value added ceases to exist (see Appendix B.7).

[52] A 50% increase in admission-noise variance reallocates students with strong assessments toward the nonselective tier, raising its spending per student and average ability assessments. As quality rises, students' willingness to pay for a nonselective spot increases and colleges convert that willingness to pay into tuition and further spending. Past a modest increase in admission noise, this fixed point no longer closes and an equilibrium fails to exist.

8. Conclusion

This paper traces the gap in selective-college attendance between high- and low-income students to the application stage. Using restricted-use data that follow the same students through application, admission, and enrollment allows us to decompose the role of each stage. For a top scorer, parental income substitutes for about 130 SAT points in the decision to apply, while income gaps at the later stages are modest. An equilibrium model of the college market explains why: low-income students apply only when their ability is exceptional. This self-selection makes colleges confident enough to fund the low-income applicants they do see.

We use the estimated model to evaluate expansions of financial aid, cheaper applications, and changes in the correlation between students' and colleges' assessments of a student's ability. The experiments deliver a common lesson: every intervention depends on the response of college prices and admission standards. Encouraging applications without targeting ability dilutes the information in the applicant pool and crowds out the students the policy aims to help. In the benchmark, making assessments informative instead improves welfare and sorting, with the largest gains going to high-ability low-income students. Under graduation risk, sorting improves even more strongly, but the welfare gains shift up the income distribution.

The model developed here can be a laboratory for answering important questions on higher education policy. Richer student-side margins and a public–private distinction on the supply side would sharpen the quantitative lessons. We leave these extensions to future research.

References

B. Abbott, G. Gallipoli, C. Meghir, and G. L. Violante. Education policy and intergenerational transfers in equilibrium. *Journal of Political Economy*, 127(6):2569–2624, 2019. [4, 28, 29, 30, 31, 32, A22]

T. Alon, D. Capelle, and K. Matsuda. University research and the market for higher education. *American Economic Journal: Macroeconomics*, 2026. Forthcoming. [4]

R. J. Andrews, S. A. Imberman, and M. F. Lovenheim. Recruiting and supporting low-income, high-achieving students at flagship universities. *Economics of Education Review*, 74:101923, 2020. [4]

P. Arcidiacono, J. Kinsler, and T. Ransom. Legacy and Athlete Preferences at Harvard. *Journal of Labor Economics*, 40(1):133–156, 2022. [22]

M. J. Bailey and S. M. Dynarski. *Inequality in Postsecondary Education*, pages 117–132. Russell Sage Foundation, 2011. [1, 10]

P. Belley and L. Lochner. The changing role of family income and ability in determining educational achievement. *Journal of Human Capital*, 1(1):37–89, 2007. [4, 8]

E. P. Bettinger, B. T. Long, P. Oreopoulos, and L. Sanbonmatsu. The Role of Application Assistance and Information in College Decisions: Results from the H&R Block Fafsa Experiment. *The Quarterly Journal of Economics*, 127 (3):1205–1242, 07 2012. [2]

P. Q. Blair and K. Smetters. Why Don't Elite Colleges Expand Supply? NBER Working Paper 29309, 2021. [26]

Z. Bleemer. Top percent policies and the return to postsecondary selectivity. Working paper, Princeton University, 2024. [4, 40]

A. L. Bodoh-Creed and B. R. Hickman. College Assignment as a Large Contest. *Journal of Economic Theory*, 175: 88–126, 2018. [4]

E. Borghesan. The heterogeneous effects of changing SAT requirements in admissions: An equilibrium evaluation. Working paper, University of Michigan, 2025. [40]

J. Bound, M. F. Lovenheim, and S. Turner. Why have college completion rates declined? an analysis of changing student preparation and collegiate resources. *American Economic Journal: Applied Economics*, 2(3):129–157, 2010. [28]

Z. Cai and J. Heathcote. College tuition and income inequality. *American Economic Review*, 112(1):81–121, 2022. [4]

D. Capelle and K. Matsuda. The great gatsby goes to college: Tuition, inequality and intergenerational mobility. Working paper, International Monetary Fund, 2026. [4]

S. E. Carrell, R. L. Fullerton, and J. E. West. Does your cohort matter? measuring peer effects in college achievement. *Journal of Labor Economics*, 27(3):439–464, 2009. [4]

E. Caucutt and L. Lochner. Early and late human capital investments, borrowing constraints, and the family. *Working Paper*, 2017. [4]

R. Chetty, J. N. Friedman, E. Saez, N. Turner, and D. Yagan. Income Segregation and Intergenerational Mobility Across Colleges in the United States*. *The Quarterly Journal of Economics*, 02 2020. [1, 4, 6, 31, 32]

M. Colas, S. Findeisen, and D. Sachs. Optimal need-based financial aid. *Journal of Political Economy*, 129(2): 492–533, 2021. [4]

F. Cunha, J. Heckman, and S. Navarro. Separating uncertainty from heterogeneity in life cycle earnings. *oxford Economic papers*, 57(2):191–261, 2005. [15]

S. B. Dale and A. B. Krueger. Estimating the Payoff to Attending a More Selective College: An Application of Selection on Observables and Unobservables*. *The Quarterly Journal of Economics*, 117(4):1491–1527, 11 2002. [4, 34]

S. B. Dale and A. B. Krueger. Estimating the effects of college characteristics over the career using administrative earnings data. *Journal of human resources*, 49(2):323–358, 2014. [4]

E. Dillon and J. Smith. Determinants of the match between student ability and college quality. *Journal of Labor Economics*, 35:45–66, 2017. [4]

K. Donovan and C. M. Herrington. Factors affecting college attainment and student ability in the U.S. since 1900. *Review of Economic Dynamics*, 31, 2019. [4]

S. Dynarski, C. Libassi, K. Michelmore, and S. Owen. Closing the gap: The effect of reducing complexity and uncertainty in college pricing on the choices of low-income students. *American Economic Review*, 111(6):1721–1756, 2021. [2]

D. Epple and R. E. Romano. Competition between private and public schools, vouchers, and peer-group effects. *American Economic Review*, pages 33–62, 1998. [21]

D. Epple, R. Romano, and H. Sieg. Admission, tuition, and financialaid policies in the market for higher education. *Econometrica*, 74:885–928, 2006. [4, 21, 25, A20]

D. Epple, R. Romano, S. Sarpca, and H. Sieg. A general equilibrium analysis of state and private colleges and access to higher education in the u.s. *Journal of Public Economics*, 155:164–178, 2017. [4, 26, 33]

M. Folch and L. Mazzone. Go big or buy a home: The impact of student debt on career and housing choices. *Review of Economic Dynamics*, 59, 2026. [4]

C. Fu. Equilibrium tuition, applications, admissions and enrollment in the college market. *Journal of Political Economy*, 122(2):225–281, 2014. [4, A22]

G. Gordon and A. Hedlund. Accounting for the rise in college tuition. Working Paper 21967, National Bureau of Economic Research, February 2016. [4]

G. Gordon and A. Hedlund. Accounting for the rise in college tuition. Working paper, Federal Reserve Bank of Richmond, 2021. [4]

J. J. Heckman, L. J. Lochner, and P. E. Todd. Earnings functions, rates of return and treatment effects: The mincer equation and beyond. *Handbook of the Economics of Education*, 1:307–458, 2006. [15]

L. Hendricks, C. Herrington, and T. Schoellman. College quality and attendance patterns: A long-run view. *American Economic Journal: Macroeconomics*, 13(1):184–215, 2021. [4]

L. Hendricks, T. Koreshkova, and O. Leukhina. College access and intergenerational mobility. Working paper 2024-030, Federal Reserve Bank of St. Louis, 2024. [4]

C. M. Herrington. Public education financing, earnings inequality, and intergenerational mobility. *Review of Economic Dynamics*, 18(4):822–842, 2015. [3]

M. Hoekstra. The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach. *The Review of Economics and Statistics*, 91(4):717–724, 11 2009. [4]

C. M. Hoxby and C. Avery. The missing "one-offs": The hidden supply of high-achieving, low income students. *Brookings Papers on Economic Activity*, 2014. [1, 4, 6]

C.-T. Hsieh, E. Hurst, C. I. Jones, and P. J. Klenow. The allocation of talent and U.S. economic growth. *Econometrica*, 87(5):1439–1474, 2019. [1, 4, 42]

F. Ionescu. The federal student loan program: Quantitative implications for college enrollment and default rates. *Review of Economic Dynamics*, 12(1):205–231, 2009. [3]

F. Ionescu and N. Simpson. Default risk and private student loans: Implications for higher education policies. *Journal of Economic Dynamics and Control*, 64(1):119–147, 2016. [3]

A. Kapor. Distributional Effects of Race-Blind Affirmative Action. Princeton University Economics Department Working Paper 2020-51, 2020. [4]

T. Kotera and A. Seshadri. Educational policy and intergenerational policy. *Review of Economic Dynamics*, 25: 187–207, 2017. [3]

D. Krueger and A. Ludwig. On the optimal provision of social insurance: Progressive taxation versus education subsidies in general equilibrium. *Journal of Monetary Economics*, 77(C):72–98, 2016. [3]

D. Krueger, A. Ludwig, and I. Popova. Shaping inequality and intergenerational persistence of poverty: Free college or better schools. Nber working paper 32467, National Bureau of Economic Research, 2024. [4]

O. Leukhina. The changing role of family income in college selection and beyond. *Federal Reserve Bank of St. Louis Review*, pages 198–222, 2023. [4, 8]

L. Lochner and A. Monge-Naranjo. Credit constraints in education. *Annual Review of Economics*, 4:225–256, 2012. [38]

L. Lochner and A. Monge-Naranjo. Student loans and repayment: Theory, evidence, and policy. In *Handbook of the Economics of Education*, volume 5, pages 397–478. Elsevier, 2016. [38]

L. J. Lochner and A. Monge-Naranjo. The nature of credit constraints and human capital. *American Economic Review*, 101(6):2487–2529, October 2011. [3]

D. O. Lucca, T. Nadauld, and K. Shen. Credit Supply and the Rise in College Tuition: Evidence from the Expansion in Federal Student Aid Programs. *The Review of Financial Studies*, 32(2):423–466, 2018. [4, 37, 39]

M. Morazzoni. Student debt and entrepreneurship in the US. Working paper, University College London, 2023. [4]

E. G. Moschini and G. Raveendranathan. College financial aid application frictions. *International Economic Review*, 2026. Forthcoming. [4]

J. Mountjoy and B. Hickman. The returns to college(s): Estimating value-added and match effects in higher education. Working paper, University of Chicago, Booth School of Business, 2020. [4]

D. Restuccia and R. Rogerson. Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics*, 11(4):707–720, 2008. [1, 4]

B. Sacerdote. Peer Effects with Random Assignment: Results for Dartmouth Roommates*. *The Quarterly Journal of Economics*, 116(2):681–704, 05 2001. [4]

L. J. Turner. The economic incidence of federal student grant aid. Working paper, University of Maryland, 2017. [4, 37]

O. Vardishvili. The macroeconomic cost of college dropouts. Working paper, University of California, Irvine, 2023. [4]

J. Wright. Quality, inequality, and mobility: American vs. Canadian systems of higher education. Working paper, University of Minnesota, 2025. [4]

J. Wright and A. Zheng. From preschool to college: The impact of education policies over the lifecycle. Iza discussion paper 17301, IZA Institute of Labor Economics, 2024. [4]

D. J. Zimmerman. Peer effects in academic outcomes: Evidence from a natural experiment. *Review of Economics and statistics*, 85(1):9–23, 2003. [4]

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

ONLINE APPENDIX

APPENDIX A: EMPIRICAL ANALYSIS

A.1. *Student characteristics*

A license for the restricted-use version of the High School Longitudinal Study (HSLS) of 2009 can be requested from the National Center for Education Statistics (NCES) *here*. We drop from the original sample students with missing parental income or SAT score, students attending a foreign institution in 2013, and students who did not answer the question about whether they were attending a postsecondary institution as of November 2013. The resulting sample includes 18,913 of the 25,206 interviewed students. Table A.1 contrasts the demographic characteristics of our sample with those of the full HSLS. Our sample closely matches the demographic composition of the original set of interviewed students.

TABLE A.1

HSLS SUMMARY STATISTICS

Demographic characteristics (%)	Full data	Sample
White	51.7	52.5
Asian	3.6	3.5
Black/African-American	13.7	13.4
Hispanic	22.0	21.9
Other race	9.1	8.7
Male	50.4	49.7
Female	49.6	50.3
Parent has less than high school	10.5	9.9
Parent has HS diploma or GED	44.2	44.5
Parent has Associate's	16.0	15.9
Parent has Bachelor's	19.7	20.0
Parent has Master's or higher	9.7	9.6
Lives with 2 parents	68.2	68.4
Lives with 1 parent or none	31.8	31.6
At least 1 parent is working	88.9	89.1
Average number of household members	4.3	4.3
High school is public	93.2	93.1
High school is private	6.8	6.9
High school is in a city	30.9	30.6
High school is in a suburb	28.4	28.5
High school is in a town	12.1	12.1
High school is in a rural area	28.7	28.7
High school is in the Northeast	17.7	17.8
High school is in the Midwest	21.9	22.2
High school is in the South	37.7	37.6
High school is in the West	22.6	22.4
Observations	25,206	18,913

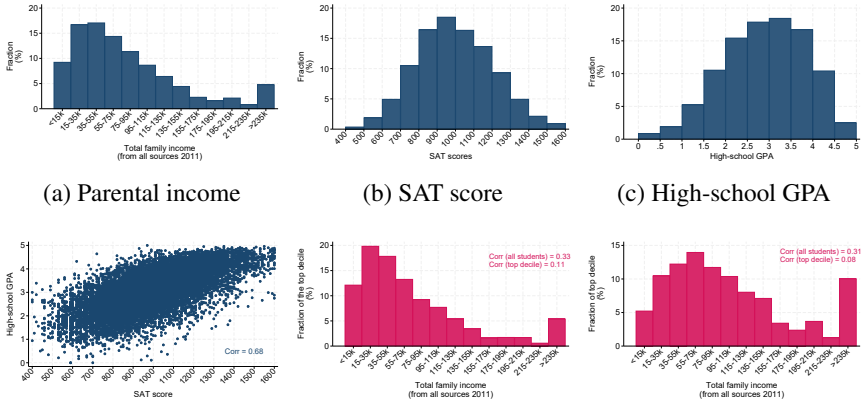
Panel (a) of Figure A.1 plots the parental income distribution in our sample. Parental income is a categorical variable capped at $235k. The median income is in the $55k–75k range and the interquartile range spans $15k–35k to $95k–115k.

Panel (b) shows the distribution of SAT scores in our sample. The weighted mean SAT score is 974 and the median is 960. The interquartile range is 830 to 1,120. The top 1%, 5%, and 10% of students have SAT scores of at least 1,460, 1,320, and 1,250, respectively. Panel (c) depicts the distribution of high-school GPA. The weighted mean GPA is 2.71 and the median is 2.74. The interquartile range is 2.02 to 3.46. The top 1%, 5%, and 10% of students have GPAs of

at least 4.61, 4.22, and 3.97, respectively. Panel (d) displays a scatterplot of SAT scores and high-school GPA. The correlation between the two is 0.68.

Panels (e) and (f) focus on the top decile of SAT scores and of high-school GPA and show how parental income is distributed among top students. Among the top 10% of SAT test takers parental income is widely spread out. The correlation between parental income and SAT is 0.11 among top students, compared with 0.33 among all test takers (panel (e)). The correlation between parental income and high-school GPA is 0.31 among all students but only 0.08 among students in the top GPA decile (panel (f)).

FIGURE A.1.—Sample distributions: parental income, SAT scores, and high-school GPA



(a) Parental income (b) SAT score (c) High-school GPA

(d) SAT vs. GPA (e) Income, top SAT decile (f) Income, top GPA decile

Note: Panel (a): distribution of parental income (categorical, capped at $235k). Panels (b)–(d): distributions of SAT scores and high-school GPA and their joint distribution. Panels (e)–(f): distribution of parental income among students in the top decile of SAT scores and of high-school GPA. Source: HSLS.

A.2. *College characteristics*

The Integrated Postsecondary Education Data System (IPEDS) is made publicly available by the NCES. The data cover every college, university, and technical and vocational institution that participates in the federal financial aid programs (such as Pell grants and federal student loans) and include information about student enrollment, program completion, graduation rates, faculty and staff, finances, tuition and fees, and student financial aid.

We access the custom data files via the NCES website *here*. In the first tab, we select institutions by groups and restrict the IPEDS sample to 2013 as the data collection period, which is the relevant period for the HSLS cohort. We make the following additional choices:

1. We select "U.S. only" and "Title IV participating" institutions (7,375 institutions).
2. Under "degree-granting status", we choose only "degree-granting" institutions (4,861 institutions).

3. Under "sector", we keep all four-year public and private not-for-profit colleges (2,281 institutions).

We then select our variables of interest and export the dataset. We further refine the sample by focusing on colleges that offer at least a bachelor's degree according to the 2010 Carnegie Classification, and drop theological/faith-related institutions and colleges without any enrolled students in 2013. These choices restrict the sample to 1,763 colleges.

We use Barron's *Profile of American Colleges* (2015) to classify institutions into selective (tiers 1 and 2) and nonselective. Table A.2 compares colleges in each group.

TABLE A.2

COMPARISON OF COLLEGES IN EACH SELECTIVITY TIER

College characteristics	Selective	Non selective
Number of colleges	186	1,577
Fraction of undergraduate enrollment	16%	84%
Fraction public	55%	73%
Average acceptance rate	44%	68%
Average rejection rate	*56%*	*32%*
Average SAT score	1,292	1,058
Median earnings 10 years after entry	$57,803	$42,228
Application fees	$59	$40
Tuition and fees	$24,576	$12,662
Room and board	$11,474	$9,372
Net tuition		
Income <$30k	-$1,848	$645
Income $30k-$48k	-$19	$1,888
Income $48k-$75k	$4,218	$4,847
Income $75k-$110k	$9,280	$7,429
Income >$110k	$15,924	$8,626
Instructional expenditures per student	$20,746	$8,597
Endowment assets per student	$131,440	$13,271

Note: All averages are weighted by total undergraduate enrollment at each college. Data are from the 2013-2014 academic year, except the acceptance rate which was calculated by averaging over the 2012 and 2013 admissions cycles. This sample consists of all four-year, non-profit Bachelor's degree granting colleges available from IPEDS (excluding religious colleges). *Source:* IPEDS, College Scorecard.

A.3. *Additional results on college enrollment*

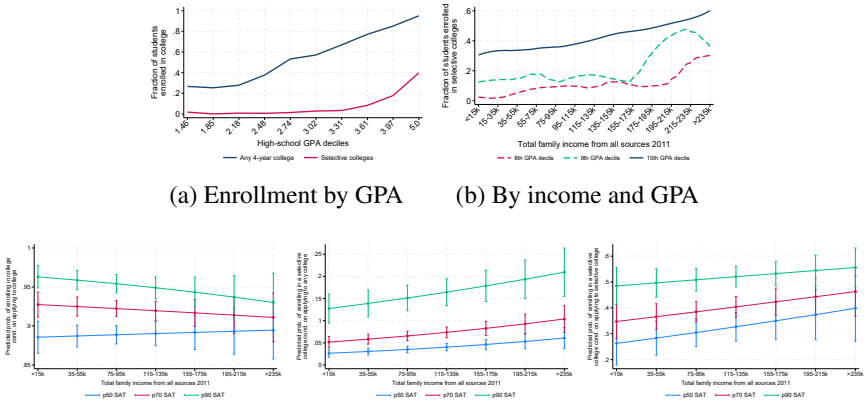
The HSLS variables used in Section 2.2 are as follows:

- The unconditional fraction of students enrolled in any four-year college and in a selective college relies on the variable `S3CLGID`. Colleges are classified as selective if they are in Barron's tiers 1 and 2. The remaining colleges are classified as nonselective.
- When conditioning on students who applied to any four-year college or to a selective college, we use the three college applications available through variables `S3CLGID` (college currently enrolled), `S3CLGAPPID1` (first other college applied to), and `S3CLGAPPID2` (second other college applied to).
- The student weight used to compute averages is `W3W2STUTR`.

Panel (a) of Figure A.2 shows the fraction of students who enrolled in any four-year college and in a selective college by high-school GPA. The average enrollment rate increases with GPA. Panel (b) depicts the fraction of students who enrolled in a selective college across parental income for the top three deciles of high-school GPA. Among these students, higher-income

students are more likely to be enrolled in a selective college than lower income students with a similar GPA.

FIGURE A.2.—College enrollment: GPA profiles and predicted probabilities conditional on applying



(a) Enrollment by GPA (b) By income and GPA

(c) Any 4-year college (d) Selective (applied to any) (e) Selective (applied to selective)

Note: Panel (a): fraction of students enrolled in any four-year not-for-profit college (blue) and in a selective college (pink) by high-school GPA decile. Panel (b): fraction enrolled in a selective college by parental income for the top three GPA deciles (8th in pink, 9th in green, 10th in blue). Panels (c)–(e): predicted probabilities of enrolling, conditional on applying, across the parental income distribution for students at the 50th (blue), 70th (pink), and 90th (green) percentiles of the SAT distribution; controls as in Table A.3, which presents the estimates.

Table A.3 presents logit estimates of whether a student enrolled in any four-year college or in a selective college, conditional on applying, on parental income, SAT score, their interaction, and demographic and high-school characteristics. Parental income enters as the ordinal income category (see Figure A.1), and SAT and GPA are standardized. Panel (c) of Figure A.2 plots the predicted probability of enrolling in any four-year college by parental income for students at the 50th, 70th, and 90th percentiles of the SAT, holding all other covariates at their means. The probability of enrolling in college conditional on applying does not vary much with parental income or SAT scores. Panels (d) and (e) show the predicted probabilities of enrolling in a selective college conditional on applying to any college (panel (d)) and conditional on applying to a selective college (panel (e)). Conditional on applying to any college, the probability of enrolling in a selective college is significantly increasing in income for all SAT percentiles. Conditional on applying to a selective college, the point estimates still imply an increasing profile of about 7 percentage points between the lowest and highest income categories for a student at the 90th percentile of the SAT, but the income coefficient is not statistically significant at conventional levels (Table A.3, right columns).

TABLE A.3

LOGIT ESTIMATION RESULTS FOR WHETHER STUDENTS ARE ATTENDING COLLEGE (CONDITIONAL ON APPLYING TO COLLEGE)

	Enrolled at any college		Enrolled at selective college (app. to any college)		Enrolled at selective college (app. to selective college)	
	Coefficient	SE	Coefficient	SE	Coefficient	SE
Parental income	0.005	(0.023)	0.071***	(0.026)	0.050	(0.034)
SAT (standardized)	0.915***	(0.150)	1.221***	(0.161)	0.715***	(0.198)
Parental income*SAT	-0.046*	(0.023)	-0.016	(0.020)	-0.020	(0.024)
GPA (standardized)	0.624***	(0.095)	1.264***	(0.126)	0.862***	(0.149)
Demographic controls	Yes		Yes		Yes	
High-school controls	Yes		Yes		Yes	
N	8,414		8,414		2,712	
Pseudo R^2	0.159		0.328		0.161	

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The dependent variable is equal to 1 if the student enrolled in any four-year, non-profit college conditional on applying to any college (left columns), or enrolled in any selective college conditional on applying to any college (middle columns), or enrolled in any selective college conditional on applying to a selective college (right columns), and 0 otherwise. Estimates are weighted by the HSLS student weight (`W3W2STUTR`). Standard errors are heteroskedasticity-robust. Demographic controls: race, gender, parental education, living situation and employment, and household size. High-school controls: private-school indicator, free-lunch share, urbanicity, and region.

A.4. *Additional results on college applications*

The HSLS variables used in Section 2.3 are as follows:

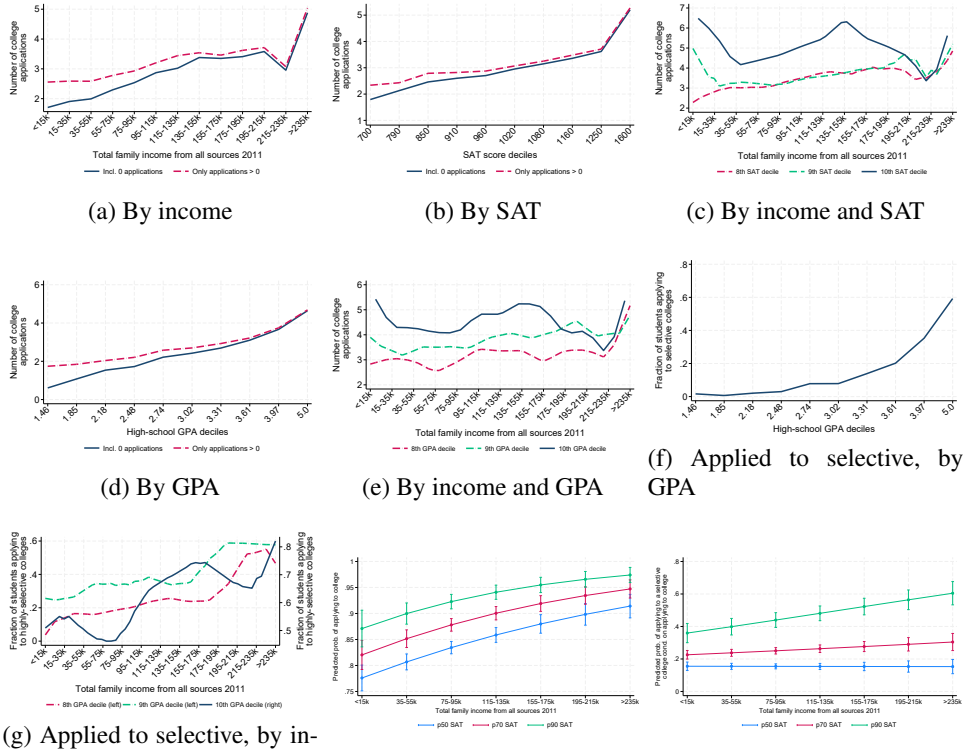
- The number of applications is measured by the variable `S3CLGAPPNUM` (the following responses are treated as missing: (i) missing; (ii) unit non-response; (iii) component not applicable; and (iv) item not administered).
- Applications to selective colleges are those submitted to an institution in Barron's tiers 1 and 2, identified using variables `S3CLGID`, `S3CLGAPPID1`, and `S3CLGAPPID2`.
- The student weight used to compute averages is `W3W2STUTR`.

Figure A.3 shows the average number of applications to four-year colleges by parental income (panel (a)), by SAT decile (panel (b)), and by parental income for the top three SAT deciles (panel (c)). Panel (a) also reports the average conditional on submitting at least one application. The conditional average is considerably higher for lower income students, indicating that many of them do not apply at all. Students in the top SAT decile submit about five applications, and the median student fewer than three (these numbers change little when conditioning on applying). Application counts are U-shaped in income among top scorers: students in the top SAT decile submit between 4 and 6 applications in nearly every income bin, while counts for the 8th decile rise from about 2 at the lowest incomes to almost 5 at the highest. Conditioning on applicants slightly raises the average number of applications for top students whose parents earn less than $15k.

Figure A.3 also shows the average number of applications to four-year colleges by high-school GPA decile (panel (d)) and by parental income for the top three GPA deciles (panel (e)). The median student submitted 2.2 applications on average and students in the top GPA decile 4.6 (about the same as when SAT scores are used). For the 8th and 9th GPA deciles, the average number of applications is increasing in parental income.

Panel (f) shows the fraction of students who applied to a selective college by high-school GPA decile and panel (g) by parental income for the top three GPA deciles. About 60% of the very top GPA students applied to a selective college, against 8% of students with the median GPA. Among students in the top GPA decile, more than 71% of those whose parents earn between $175k and $195k applied to a selective college, compared with 50% of those who earn between $35k and $55k. These shares are slightly lower than when SAT scores are used.

FIGURE A.3.—College applications: counts, selective-college applications, and predicted probabilities



(a) By income (b) By SAT (c) By income and SAT

(d) By GPA (e) By income and GPA (f) Applied to selective, by GPA

(g) Applied to selective, by income and GPA (h) Predicted prob., any college (i) Predicted prob., selective

Note: Panels (a)–(c): average number of applications by parental income (for all students and for applicants), by SAT decile, and by income for the top three SAT deciles (8th in pink, 9th in green, 10th in blue). Panels (d)–(e): the GPA counterparts. Panels (f)–(g): fraction who applied to a selective college by GPA decile, and by income for the top three GPA deciles. Panels (h)–(i): predicted probabilities of applying to any college and to a selective college (conditional on applying), for students at the 50th (blue), 70th (pink), and 90th (green) SAT percentiles; controls as in Table A.4, which presents the estimates.

Table A.4 presents logit estimates of whether a student applied to any four-year college or to a selective college on parental income, SAT score, their interaction, and demographic and high-school characteristics. Panel (h) of Figure A.3 plots the predicted probability of applying to college by parental income for students at the 50th, 70th, and 90th percentiles of the SAT, holding all other covariates at their means. Panel (i) shows the predicted probability of applying to a selective college conditional on applying to college. The probability of applying to college is increasing in income at all SAT percentiles. The probability of applying to a selective college differs by income only for top students (although the parental income coefficient is insignifi-

cant, its interaction with SAT is positive and significant) and the difference between high- and low-income top students is large.

TABLE A.4

LOGIT ESTIMATION RESULTS FOR WHETHER STUDENTS APPLIED TO COLLEGE

	Applied to any college		Applied to selective college	
	Coefficient	SE	Coefficient	SE
Parental income	0.096^{***}	(0.017)	0.020	(0.019)
SAT (standardized)	0.443^{***}	(0.112)	0.834^{***}	(0.135)
Parental income*SAT	0.035^{*}	(0.021)	0.051^{***}	(0.019)
GPA (standardized)	1.091^{***}	(0.068)	0.956^{***}	(0.081)
Demographic controls	Yes		Yes	
High-school controls	Yes		Yes	
N	11,399		11,399	
Pseudo R^2	0.247		0.322	

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The dependent variable is equal to 1 if the student applied to any four-year, not-for-profit college (left column) or any selective college (right column), and 0 otherwise. The sample is restricted to all students who applied to any four-year, not-for-profit college in the right column. Estimates are weighted by the HSLS student weight (`W3W2STUTR`). Standard errors are heteroskedasticity-robust. Demographic controls: race, gender, parental education, living situation and employment, and household size. High-school controls: private-school indicator, free-lunch share, urbanicity, and region.

A.5. *Additional results on college admissions*

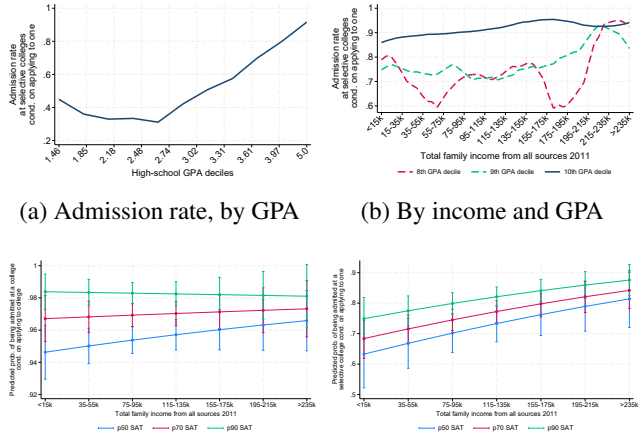
The HSLS variables used in Section 2.4 are as follows:

- We focus on the set of students who applied to any college or to a selective college, constructed as in the previous subsection from variables `S3CLGID`, `S3CLGAPPID1`, and `S3CLGAPPID2`. The student is treated as admitted to the institution she is currently enrolled in; the admission status at the other two listed colleges is reported in variables `S3APPSTATUS1` and `S3APPSTATUS2`.
- The student weight used to compute averages is `W3W2STUTR`.

Panel (a) of Figure A.4 shows the fraction of applicants to a selective college who were admitted to one, by high-school GPA. The average admission rate is flat at less than 40% below the median GPA and increasing in GPA above the median. Panel (b) reports admission rates by parental income for the top three GPA deciles. Admission rates are flat in parental income for the very top students, averaging 74%. They are lower for the 8th and 9th deciles, but do not differ significantly between low- and high-income students.

Table A.5 presents logit estimates of whether a student was admitted to any college or to a selective college, conditional on having applied to any college or to a selective college respectively, on parental income, SAT score, their interaction, and demographic and high-school characteristics. Panel (c) of Figure A.4 plots the predicted probability of being admitted to any four-year college conditional on applying, across the parental income distribution for different SAT percentiles, holding the other covariates at their means. Panel (d) shows the predicted probability of being admitted to a selective college conditional on applying to one. For any college, the probability is flat in parental income. For selective colleges, the probability of ad-

FIGURE A.4.—Admission rates: GPA profiles and predicted probabilities



(a) Admission rate, by GPA (b) By income and GPA

(c) Predicted prob., any college (d) Predicted prob., selective

Note: Panels (a)–(b): fraction of applicants to a selective college admitted to one, by high-school GPA decile and by parental income for the top three GPA deciles (8th in pink, 9th in green, 10th in blue). Panels (c)–(d): predicted probabilities of admission conditional on applying, for students at the 50th (blue), 70th (pink), and 90th (green) SAT percentiles; controls as in Table A.5, which presents the estimates.

mission is increasing in parental income. Since the interaction between income and SAT is essentially zero, the gradient is similar across the SAT distribution.

TABLE A.5

LOGIT ESTIMATION RESULTS FOR WHETHER STUDENTS WERE ADMITTED INTO COLLEGE (CONDITIONAL ON APPLYING)

	Admitted at any college		**Admitted at selective college**	
	Coefficient	SE	Coefficient	SE
Parental income	0.037	(0.035)	0.077**	(0.038)
SAT (standardized)	0.926***	(0.210)	0.398	(0.260)
Parental income*SAT	-0.035	(0.037)	-0.004	(0.033)
GPA (standardized)	1.071***	(0.126)	1.387***	(0.182)
Demographic controls	Yes		Yes	
High-school controls	Yes		Yes	
N	8,953		2,773	
Pseudo R^2	0.240		0.216	

Note: textsuperscript***$p < 0.01$, **$p < 0.05$, *$p < 0.1$. The dependent variable is equal to 1 if the student was admitted at any four-year, not-for-profit college (left column) or any selective college (right column), and 0 otherwise. The sample is restricted to all students who applied to any four-year, not-for-profit college in the left column and to all students who applied to a selective college in the right column. Estimates are weighted by the HSLS student weight (`W3W2STUTR`). Standard errors are heteroskedasticity-robust. Demographic controls: race, gender, parental education, living situation and employment, and household size. High-school controls: private-school indicator, free-lunch share, urbanicity, and region.

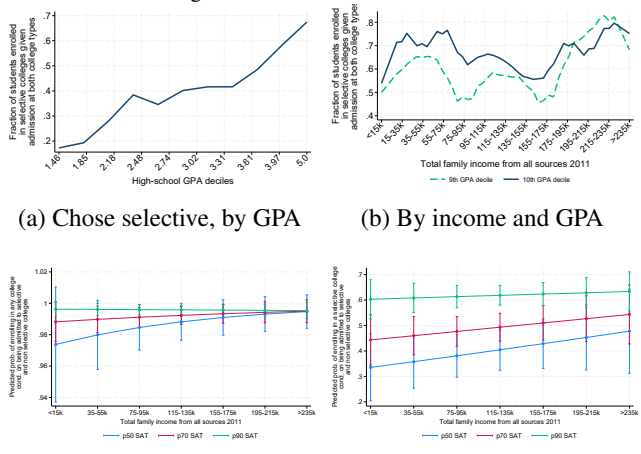
A.6. *Additional results on college enrollment conditional on being admitted*

The HSLS variables used in Section 2.5 are as follows:

- We identify the college a student is enrolled in using variable `S3CLGID` and condition on the set of students who received offers from both selective and nonselective colleges, constructed as in the previous subsection (based on `S3APPSTATUS1` and `S3APPSTATUS2`).
- The student weight used to compute averages is `W3W2STUTR`.

Panel (a) of Figure A.5 shows the fraction of students who chose a selective college after being admitted to both selective and nonselective colleges, by high-school GPA. The fraction increases sharply with GPA, with more than 67% of students in the top GPA decile choosing the selective college. Panel (b) shows the same fraction by parental income for the top two GPA deciles. Most students choose the selective college when they hold competing offers. Higher-income students are nonetheless somewhat more likely to do so than lower income students.

FIGURE A.5.—Enrollment among students admitted to both selective and nonselective colleges



(a) Chose selective, by GPA (b) By income and GPA

(c) Predicted prob., any college (d) Predicted prob., selective

Note: Panels (a)–(b): fraction of students who chose a selective college after being admitted to both tiers, by high-school GPA decile and by parental income for the top two GPA deciles (9th in green, 10th in blue). Panels (c)–(d): predicted probabilities of enrolling conditional on holding both offers, for students at the 50th (blue), 70th (pink), and 90th (green) SAT percentiles; controls as in Table A.6, which presents the estimates.

Table A.6 presents logit estimates of whether a student enrolled in any college or in a selective college, conditional on being admitted to both selective and nonselective colleges, on parental income, SAT score, their interaction, and demographic and high-school characteristics. Panel (c) of Figure A.5 plots the predicted probability of enrolling in any college by parental income for students at the 50th, 70th, and 90th percentiles of the SAT, holding all other covariates at their means. Panel (d) shows the predicted probability of enrolling in a selective college. For students at the 90th percentile of the SAT, the probability of enrolling in some college is close to 100% and flat in parental income; the probability of enrolling in a selective college is above 60% and also flat in parental income.

TABLE A.6

LOGIT ESTIMATION RESULTS FOR WHETHER STUDENTS ARE ATTENDING COLLEGE (CONDITIONAL ON BEING ADMITTED AT SELECTIVE AND NON SELECTIVE COLLEGES)

	Enrolled at any college		Enrolled at selective college	
	Coefficient	SE	Coefficient	SE
Parental income	0.129	(0.119)	0.048	(0.044)
SAT (standardized)	1.518**	(0.706)	0.817***	(0.242)
Parental income*SAT	-0.110	(0.088)	-0.028	(0.030)
GPA (standardized)	0.848**	(0.370)	0.516***	(0.174)
Demographic controls	Yes		Yes	
High-school controls	Yes		Yes	
N	1,797		1,797	
Pseudo R^2	0.229		0.130	

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The dependent variable is equal to 1 if the student enrolled in any four-year, non-profit colleges (in the left column) or any highly-selective colleges (in the right column), and 0 otherwise. The sample is restricted to all students who were admitted at any four-year, non-profit college (both selective and non selective). Estimates are weighted by the HSLS student weight (`W3W2STUTR`). Standard errors are heteroskedasticity-robust. Demographic controls: race, gender, parental education, living situation and employment, and household size. High-school controls: private-school indicator, free-lunch share, urbanicity, and region.

A.7. *Additional results on college affordability*

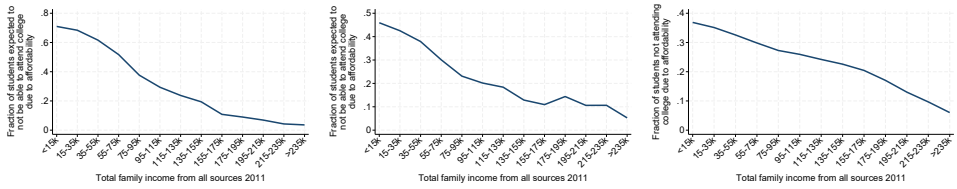
The HSLS variables used in Section 2.6 are as follows:

- Parents' expectations about college affordability are computed using variable `P2CANTAFFORD`. Parents are classified as not able to afford college if they responded "Agree" or "Strongly Agree" when asked whether they would not be able to afford to send their child to college.
- Students' expectations about college affordability are computed using variable `S2CANTAFFORD`. Students are classified as not able to afford college if they responded "Agree" or "Strongly Agree" when asked whether their family would not be able to afford to send them to college.
- The fraction of students not attending college due to affordability relies on variable `S3CANTAFFORD`, computed among students who responded either "Yes" or "No" to the question about college affordability.
- The student weight used to compute averages is `W3W2STUTR`.

Figure A.6 shows, by parental income, the fraction of students whose parents expected not to be able to afford college if the student were accepted (panel (a)), the fraction of students who themselves expected not to be able to afford college (panel (b)), and the fraction of students who report not attending college in 2013 because they could not afford it (panel (c)). Section 2.6 discusses these patterns.

Figure A.6 replicates Figure A.6 for students who did not apply to college. Most lower income non-applicants have parents who expected not to be able to afford college (panel (d)), and the majority of low-income non-applicants themselves did not expect to be able to afford college (panel (e)). More than 10% of non-applicants whose parents earn less than $75k report not applying to college because they could not afford it (panel (f)).

FIGURE A.6.—College unaffordability by parental income: all students and non-applicants



(a) Parents' expectation (b) Students' expectation (c) Not attending

(d) Parents' expectation (non-applicants) (e) Students' expectation (non-applicants) (f) Not attending (non-applicants)

Note: Panels (a)–(c): fraction of students whose parents expected (in 2012) not to be able to afford college, who themselves expected not to be able to afford it, and who report (in 2013) not attending college because they could not afford it, by parental income. Panels (d)–(f): the same objects among students who did not apply to college. Source: HSLS.

APPENDIX B: THEORY

B.1. *Proof of Proposition 3.1*

At an interior tuition optimum (3.16) with $\varepsilon_s(y,\sigma) > 1$, the price charged to each admitted type strictly exceeds her effective marginal cost and thus includes a markup. Define $T_s(y,\sigma) - MC_s(y,\sigma) = -D_s(y,\sigma)/\big(\partial D_s/\partial T_s\big) > 0$, where $D_s(y,\sigma) \equiv \int d_s(y,\ell,\sigma)\,\mu(y,d\ell,\sigma)$ is the college's demand aggregated to the cell (y,σ). Substituting this margin into the admission standard first-order condition (3.18), its left-hand side (the marginal value to the college of raising its standard) becomes

$$\int \Big[\underbrace{T_s(y,\sigma) - MC_s(y,\sigma)}_{>0} + \frac{Q_L}{Q_I}\big(\ell - \widehat{\mathbb{E}}_s[\ell\,|\,y,\sigma]\big)\Big]\,\frac{\partial d_s}{\partial \underline{\sigma}_s}\,d\mu(y,\ell,\sigma), \quad \text{with } \frac{\partial d_s}{\partial \underline{\sigma}_s} < 0.$$

Each marginal applicant contributes the markup earned on her, which is strictly positive, plus her ability relative to her cell's demand-weighted average. Within each cell, the ability term aggregates to the covariance between the demand response and ability, and this covariance vanishes as the evaluation noise concentrates the marginal applicants at $\sigma = \underline{\sigma}_s$. The proposition's domination condition is exactly that the markup term outweighs this covariance term.

Under the domination condition the bracketed term is strictly positive and $\partial d_s/\partial \underline{\sigma}_s < 0$, so the left-hand side of (3.18) is strictly negative at every $\underline{\sigma}_s$. So, whatever the admission standard, the marginal applicants more than pay for themselves and the college gains by admitting more of them. As no interior admission standard is optimal, it stays at its lower bound $\underline{\sigma}$, where every

applicant's admission probability is maximal. Screening then operates through prices alone. An applicant fails to enroll only if no price below her maximal payment, $y - \underline{a}_s + Gr(y) + P(y) - \psi(n)$, covers her effective marginal cost. The maximal payment rises with y, so applicants with sufficiently high parental transfers are never excluded. □

We verify the conclusion directly at the estimated parameters. Without the tuition caps, the value of the marginal applicant pool is strictly positive at every admission standard in the support of σ for both tiers, so the derivative of the college's objective function with respect to the admission standard is strictly negative throughout and the standard rests at its floor.

B.2. *Tuition schedule (equation (3.16))*

Consider the problem (3.15) of a college of tier s. Let λ_κ, λ_L, and λ_I be the Lagrange multipliers on the enrollment identity (3.11), the ability identity (3.12), and the budget constraint (3.13), and $\lambda_T(y,\sigma) \geqslant 0$ the multiplier on the tuition cap. Throughout, $\mu(y, d\ell, \sigma)$ denotes integration over the student's own ability assessment ℓ with (y,σ) held fixed. The first-order conditions with respect to κ, I, and L are

$$[\kappa_s] \qquad \xi_s Q_\kappa = \lambda_\kappa + \lambda_I \left[C_s'(\kappa_s) - E_s'(\kappa_s) - Tr_s'(\kappa_s)\right],$$

$$[I_s] \qquad \xi_s Q_I = \lambda_I,$$

$$[L_s] \qquad \xi_s Q_L = \lambda_L,$$

where Q_κ, Q_I, Q_L denote the partial derivatives of $Q(I_s/\kappa_s, L_s/\kappa_s)$ with respect to κ_s, I_s, and L_s. Since college demand at (y,σ) depends on $T_s(y,\sigma)$ only through the acceptance probability of students with those observables, the first-order condition with respect to $T_s(y,\sigma)$ is

$$\begin{aligned}[T_s(y,\sigma)] \qquad & \lambda_I \left[\int d_s(y,\ell,\sigma)\,\mu(y,d\ell,\sigma) + T_s(y,\sigma)\int \frac{\partial d_s(y,\ell,\sigma)}{\partial T_s(y,\sigma)}\,\mu(y,d\ell,\sigma)\right] \\ & + \lambda_\kappa \int \frac{\partial d_s(y,\ell,\sigma)}{\partial T_s(y,\sigma)}\,\mu(y,d\ell,\sigma) + \lambda_L \int \ell\,\frac{\partial d_s(y,\ell,\sigma)}{\partial T_s(y,\sigma)}\,\mu(y,d\ell,\sigma) \\ & - \lambda_T(y,\sigma) = 0.\end{aligned}$$

Substituting the multipliers, dividing by $\lambda_I \int [\partial d_s(y,\ell,\sigma)/\partial T_s(y,\sigma)]\,\mu(y,d\ell,\sigma)$, and using the definitions of $\varepsilon_s(y,\sigma)$ and $\widehat{\mathbb{E}}_s[\ell|y,\sigma]$ in the text gives, when the solution is interior (i.e., $\lambda_T(y,\sigma) = 0$),

$$T_s(y,\sigma)\left(1 - \frac{1}{\varepsilon_s(y,\sigma)}\right) = C_s'(\kappa_s) - \frac{Q_\kappa}{Q_I} - E_s'(\kappa_s) - Tr_s'(\kappa_s) - \frac{Q_L}{Q_I}\,\widehat{\mathbb{E}}_s[\ell|y,\sigma],$$

which is (3.16) with the markup $m = \varepsilon/(\varepsilon - 1)$ and marginal cost (3.17). When the cap binds, $\lambda_T(y,\sigma) > 0$ and $T_s(y,\sigma) = \overline{T}_s < m_s(y,\sigma)MC_s(y,\sigma)$.

B.3. *Admissions standard (equation (3.18))*

The college's own admission standard $\underline{\sigma}_s$ enters its college demand $d_s(y,\ell,\sigma)$ through its own admission probability $\pi_s(\sigma) = 1 - F_\epsilon(\underline{\sigma}_s - \sigma)$, with $\partial\pi_s/\partial\underline{\sigma}_s = -f_\epsilon(\underline{\sigma}_s - \sigma)$. The other colleges' admission standards are taken as given. Let $\lambda_\sigma \geqslant 0$ be the multiplier on the college's lower bound $\underline{\sigma}_s \geqslant \underline{\sigma}$. The first-order condition with respect to $\underline{\sigma}_s$ is

$$\lambda_I \int T_s(y,\sigma)\,\frac{\partial d_s(y,\ell,\sigma)}{\partial \underline{\sigma}_s}\,d\mu(y,\ell,\sigma) + \lambda_\kappa \int \frac{\partial d_s(y,\ell,\sigma)}{\partial \underline{\sigma}_s}\,d\mu(y,\ell,\sigma) + \lambda_L \int \ell\,\frac{\partial d_s(y,\ell,\sigma)}{\partial \underline{\sigma}_s}\,d\mu(y,\ell,\sigma) + \lambda_\sigma = 0.$$

Substituting the multipliers from the previous subsection and dividing by λ_I yields (3.18) when $\underline{\sigma}_s > \underline{\sigma}$ (and so $\lambda_\sigma = 0$). When $\underline{\sigma}_s = \underline{\sigma}$, the left-hand side of (3.18) is weakly negative. Since $\partial d_s/\partial\underline{\sigma}_s < 0$, the marginal applicants more than cover their marginal cost and the college would like to lower its standard further. In the limit of vanishing evaluation noise, $\partial d_s/\partial\underline{\sigma}_s$ concentrates on $\sigma = \underline{\sigma}_s$ and (3.18) becomes

$$\underbrace{\frac{\int T_s(y,\underline{\sigma}_s)\,d_s(y,\ell,\underline{\sigma}_s)\,\mu(dy,d\ell,\underline{\sigma}_s)}{\int d_s(y,\ell,\underline{\sigma}_s)\,\mu(dy,d\ell,\underline{\sigma}_s)}}_{\text{Average revenue from lowest-assessment students}} = \underbrace{C_{s_\kappa} - \frac{Q_\kappa}{Q_I} - E_{s_\kappa} - Tr_{s_\kappa}}_{\text{Costs net of transfers per student}} - \underbrace{\frac{Q_L}{Q_I}\frac{\int \ell\,d_s(y,\ell,\underline{\sigma}_s)\,\mu(dy,d\ell,\underline{\sigma}_s)}{\int d_s(y,\ell,\underline{\sigma}_s)\,\mu(dy,d\ell,\underline{\sigma}_s)}}_{\text{Average ability of lowest-assessment students}},$$

i.e., the average tuition revenue from students at the admission threshold must cover the marginal cost of schooling them net of the ability discount.

B.4. *Existence of equilibrium*

We verify that the equilibrium satisfies the hypotheses of Brouwer's fixed-point theorem.

- *Finite dimensionality.* The economy we estimate and solve has finitely many values of (y,ℓ,σ), so tuition schedules, standards, value added, and the tax rate form a finite-dimensional vector.
- *Compactness and convexity.* Each tuition entry lies in a closed interval whose upper end is the sticker price and whose lower end sits below every candidate optimum: effective marginal costs are bounded on the finite (y,σ) grid, so the interior optima (3.16) are bounded. Admission standards lie between the open-admission bound $\underline{\sigma}$ and the top of the assessment grid; with the normal evaluation noise, every applicant retains a positive admission probability at any standard in this range, so enrollment remains positive on the domain. Value added is bounded by budget feasibility, and the tax rate lies in $[0,\bar{\tau}]$ with $\bar{\tau} < 1$ chosen large enough to finance the maximal feasible outlays. The domain is a product of closed intervals, hence nonempty, compact, and convex.
- *Continuity and single-valuedness.* The Type I extreme value shocks and the evaluation noise make demand smooth in all policies. An option whose implied consumption is non-

positive is assigned value $-\infty$. Because utility diverges continuously as consumption falls to zero, that option's choice probability vanishes continuously and the option of working remains feasible everywhere on the domain, so choice probabilities stay well defined. The interior tuition optimum (3.16) is unique cell by cell, so policy updates are single-valued and continuous.

- *Self-map.* Each update respects the same bounds, that is, tuition updates are truncated to the interval above, admission standards remain in the grid's range, value added computed from feasible allocations satisfies the budget constraint, and the implied tax rate finances the realized outlays.

A fixed point of the updated map therefore exists by Brouwer's theorem. At the fixed point, the tuition entries solve the colleges' pricing condition (3.16) cell by cell and the admission standard satisfies (3.18). □

B.5. *Equilibrium algorithm*

1. Set the parameters and discretize student types (y, ℓ) and assessments σ on grids.
2. Guess the value added pair (Z_H, Z_L) (updated in step 7), the college-level aggregates and policies, $\{\kappa_s, I_s, L_s, \ T_s(y,\sigma), \underline{\sigma}_s\}$ for $s \in \{H, L\}$, the portfolio probabilities $\{p(\mathbf{n}|y,\ell)\}$ for each student type, and the tax rate τ. Steps 3–6 hold (Z_H, Z_L) fixed.
3. Given these guesses, compute V^s and V^W for each (y, ℓ, n) from (3.1) and (3.2), which have closed-form solutions for each number of applications $n \in \{0, \ldots, N_H + N_L\}$, the admission probabilities $\pi_s(\sigma)$ from (3.5), the acceptance probabilities a_s from (3.4) for each offer set $\mathbf{k}$, and the demand functions d_s from (3.8) for each (y, ℓ, σ), holding $p(\mathbf{n}|y,\ell)$ fixed.
4. Given demand and the aggregates, solve for the tuition schedules and admission standards of both tiers as a fixed point:
   (a) For each tier s and each (y, σ), update the tuition $\hat{T}_s(y,\sigma)$ using (3.16), truncated at $\overline{T}_s$, recomputing a_s and d_s at the updated price.
   (b) For each tier s, update the admission standard $\hat{\underline{\sigma}}_s$ as the root of (3.18) given $\hat{T}_s$ (or $\hat{\underline{\sigma}}_s = \underline{\sigma}$, the open-admission bound, if the left-hand side is negative there).
   (c) If $\sup|T_s(y,\sigma) - \hat{T}_s(y,\sigma)| > 10^{-5}$ or $|\underline{\sigma}_s - \hat{\underline{\sigma}}_s| > 10^{-5}$, set $T_s = \hat{T}_s$, $\underline{\sigma}_s = \hat{\underline{\sigma}}_s$ and repeat (a).
5. Given demand, tuition schedules, and admission standards, update the aggregates $\{\hat{\kappa}_s, \hat{I}_s, \hat{L}_s\}$ using (3.11), (3.12), and (3.13); the portfolio probabilities $\{\hat{p}(\mathbf{n}|y,\ell)\}$ using (3.6) and (3.7); and the tax rate $\hat{\tau}$ using (3.20).
6. Let X collect the aggregates, portfolio probabilities, and tax rate. If $\sup|X - \hat{X}| > 10^{-5}$, set X to a dampened update of $\hat{X}$ and return to step 3.

7. At the converged solution of steps 3–6, compute the *delivered* value added $\Phi(Z_H, Z_L) \equiv (\hat{Z}_H, \hat{Z}_L)$ from (3.10). If $\sup_s |\hat{Z}_s - Z_s| > 10^{-5}$, update (Z_H, Z_L) by a Newton step toward the zero of $\Phi(Z) - Z$ and return to step 3. Otherwise, stop.

The algorithm nests three loops: (i) the tuition fixed point of step 4; (ii) the dampened loop of steps 3–6 on aggregates, portfolio probabilities, and the tax rate; and (iii) the outer loop of step 7 on value added. The tuition fixed point of step 4 converges reliably because, given the aggregates, each college's problem is a smooth pricing problem in (υ, σ). The dampening in step 6 stabilizes the joint update of aggregates and portfolio probabilities.

B.6. *Computational search for equilibria*

All numerical verifications in this subsection are computed at the estimated parameters of Section 4.4 (Table 4.2).

Local properties of Φ. The Jacobian of $\Phi(Z_H, Z_L)$ at the equilibrium has diagonal entries of 0.42 (selective) and 1.78 (nonselective), and eigenvalues 0.39 and 1.81. These are the asymmetric responses behind the policy discussion of Section 5. The fixed point is regular ($\det(I - J) \approx -0.49$), so Z^* is locally isolated. The negative determinant is consistent with the fixed-point index, whose complementary fixed points are the shutdown configurations located by the scan below, which fail the equilibrium definition. Because the model is static and agents hold rational expectations, no belief-adjustment process is part of the equilibrium concept, and the super-unit eigenvalue does not bear on whether Z^* is an equilibrium. It means only that the baseline and every counterfactual are computed by root-finding on $\Phi(Z) - Z$ (step 7 of Appendix B.5) rather than by iterating Φ.

Global scan. A 15×15 grid covering $[1.20, 2.40] \times [1.10, 2.45]$, which spans the feasibility bounds on each tier's value added implied by the budget constraint and the ability distribution ($Z_H \leqslant 2.38$, $Z_L \leqslant 2.37$), is solved by continuation. Both components of $\Phi(Z) - Z$ change sign only across the cell $[1.98, 2.03] \times [1.77, 1.81]$ that brackets Z^*. Everywhere else the economy admits no stable inner solution: it collapses to a shutdown configuration, in which one tier operates no colleges, and which fails the equilibrium definition.

B.7. *Nonexistence under graduation risk*

This subsection documents the computations behind the absorption limits of Section 7.3. At the extended model's estimate, the equilibrium map Φ is a local contraction at the baseline (Jacobian eigenvalues 0.13 and 0.43), so the baseline equilibrium is computed as in the benchmark. Each policy is scaled by $t \in [0, 1]$ between the baseline and the experiment of Sections 5–6, and the equilibrium is tracked along the path in steps, each solved cold from the

previous step's solution and accepted only if strictly converged. The limits past which no equilibrium exists are: a maximum Pell grant of \$6,614; an application-cost reduction of 9.4%; 47% of the completion-parity gap; a borrowing limit of \$33,450; and an admission-noise increase of 1.23×. Fee waivers, ±10% capacity caps, and the perfect-information experiment solve at full scale.

APPENDIX C: CALIBRATION AND ESTIMATION

C.1. *Distribution of student characteristics*

Expected Family Contribution. We map the model's notion of parental income, y, to a student's expected family contribution (EFC). The EFC corresponds to the amount of resources a student has in order to attend college (before financial aid). The EFC is available for students who completed the FAFSA and attended college. For students who did not fill out the FAFSA, we calculate their EFC directly using the 2013-2014 EFC formula with data from the HSLS survey. We first compute the Adjusted Available Income (AAI), which corresponds to household income net of allowances (which depend on household size) and household assets (excluding the family's home). Since the HSLS does not report assets, we assume that they are 0. The family contribution out of the AAI is calculated from a (progressive) non-linear function of AAI according to Table C.1 below. Following the rules set by the Department of Education, a student's EFC is then the family contribution divided by the number of children that are enrolled in college. The HSLS asks if students have a sibling in college at the same time. If the answer is yes, we considered the number of children enrolled in college to be two. Panels (a) and (b) of Figure C.1 plot the resulting distribution of the EFC and its joint distribution with parental income.

TABLE C.1

FAMILY CONTRIBUTION OUT OF THEIR ADJUSTED AVAILABLE INCOME

Family's AAI	Family contribution
< -\$3,409	-\$750
-\$3,409 to \$15,300	22% of AAI
\$15,301 to \$19,200	\$3,366 + 25% of AAI over \$15,300
\$19,201 to \$23,100	\$4,341 + 29% of AAI over \$19,200
\$23,101 to \$27,000	\$5,472 + 34% of AAI over \$23,100
\$27,001 to \$30,900	\$6,798 + 40% of AAI over \$27,000
\$30,901 or more	\$8,358 + 47% of AAI over \$30,900

Source: Department of Education.

The two assessments of ability. The college-side assessment, σ, corresponds to the student's SAT score, which is what admission committees observe. The student-side assessment, ℓ, is computed as the predicted values from regressing SAT scores on high-school GPA, race dummies, gender, parental education, parental living situation, parental employment situation, number of household members, private high school dummy, the fraction of high-school students receiving free meals, and geographical dummies related to the high school attended (i.e., the

same variables as used in Section 2.1). It is the component of test performance the student can forecast from her own record and circumstances before taking the test.

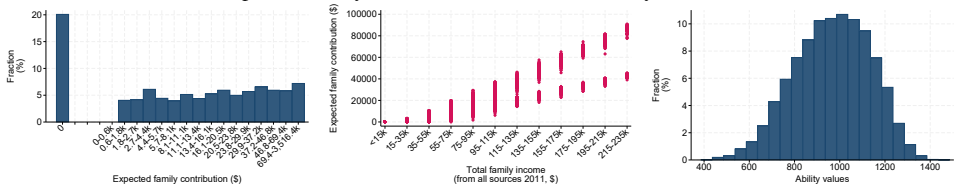
Panel (c) of Figure C.1 plots the distribution of the student-side assessment, ℓ, on a scale similar to that of the SAT. The distribution is more left skewed relative to the SAT score distribution presented in Figure A.1. The mean is 967 and the median is 974. The interquartile range corresponds to an SAT score between 849 and 1,093, so the distribution of ℓ is more concentrated around its mean than that of the SAT. The top 1%, 5%, and 10% of students have scores of 1,303, 1,226, and 1,182, respectively. Panel (d) contrasts in a scatterplot the student-side assessment with students' SAT scores. The correlation coefficient is 0.74.

Distribution. The distribution of student characteristics, $\hat{\mu}(y,\ell)$, is a mixture of a mass point at $y=0$ and a continuous distribution for $y>0$. About 20% of students in our sample have an EFC of 0 (panel (a) of Figure C.1), so we assign a mass point of 20% to $y=0$. We assume that ability among these students follows a log-normal distribution. For the remaining 80% of students with $y>0$, income and ability follow a bivariate log-normal distribution. Panels (e) and (f) plot the ability distribution separately for students with zero and positive EFC. To summarize,

$$\hat{\mu}(y,\ell)=\begin{cases}\text{LogN}(\mu_{\ell_0},\sigma^2_{\ell_0}) & \text{if } y=0\\ \text{LogN}\left(\begin{bmatrix}\mu_y\\ \mu_{\ell_1}\end{bmatrix},\begin{bmatrix}\sigma^2_y & \sigma_{y\ell}\\ \sigma_{y\ell} & \sigma^2_{\ell_1}\end{bmatrix}\right) & \text{if } y>0.\end{cases}$$

Using the HSLS, we estimate $\mu_{\ell_0}=-0.144$ and $\sigma^2_{\ell_0}=0.155$ for the students with $y=0$. For the students with $y>0$, we estimate $\mu_y=-0.301$, $\mu_{\ell_1}=0.135$, $\sigma^2_y=2.825$, $\sigma^2_{\ell_1}=0.138$, and a covariance between income and ability of $\sigma_{y\ell}=0.174$.

FIGURE C.1.—Expected family contribution and the two ability assessments in the HSLS



(a) EFC (b) EFC vs. parental income (c) Student-side assessment, ℓ

(d) ℓ vs. SAT score (e) Ability, zero EFC (f) Ability, positive EFC

Note: Panels (a)–(b): distribution of the expected family contribution and its joint distribution with parental income (the top income category, >$235k, is left out of panel (b) for readability). Panels (c)–(d): distribution of the student-side assessment ℓ (fitted values of the SAT regression) and its joint distribution with the realized SAT score. Panels (e)–(f): distribution of ability for students with zero and positive EFC.

C.2. *Lifecycle Model*

Consider an individual who lives for $T+1$ periods, where a period lasts four years and the first period is spent in college:

$$\begin{aligned}
\max_{\{c_j, a_{j+1}\}_{j=0,\ldots,T}} \quad & u(c_0) + \sum_{j=1}^{T} \beta^j u(c_j) \\
\text{s.t.} \quad & c_0 + a_1 + \text{Tuition} = y \\
& c_j + a_{j+1} = a_j R + w(1-\tau) Z \ell^{\alpha}, \quad j = 1, \ldots, T \\
& a_1 \geqslant \underline{a}_s.
\end{aligned}$$

Assume that individuals do not save or leave bequests in the terminal period. Consolidating the working-period budget constraints into a single present-value constraint and using the Euler equation $c_{1+j} = (\beta R)^{j/\sigma_c} c_1$ (with CRRA preferences with intertemporal elasticity of substitution σ_c) collapses the working phase into one period, so the problem becomes the two-period model of Section 3:

$$\begin{aligned}
\max_{c_0, c_1, a_1} \quad & u(c_0) + \tilde{\beta} u(c_1) \\
\text{s.t.} \quad & c_0 + a_1 + \text{Tuition} = y \\
& c_1 = a_1 \tilde{R} + \tilde{w}(1-\tau) Z \ell^{\alpha} \\
& a_1 \geqslant \underline{a}_s,
\end{aligned}$$

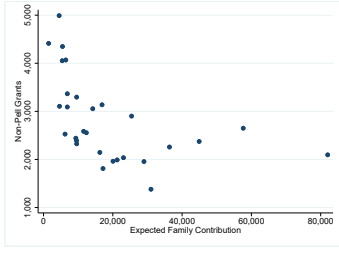
where we replaced $\tilde{\beta}u(c_1) = \sum_{j=1}^{T} \beta^j u(c_j)$ and defined

$$\tilde{\beta} = \sum_{j=1}^{T} \beta^j (\beta R)^{\frac{(j-1)(1-\sigma_c)}{\sigma_c}}, \qquad \tilde{R} = \frac{R}{\sum_{j=1}^{T} \frac{(\beta R)^{(j-1)/\sigma_c}}{R^{j-1}}}, \qquad \tilde{w} = \frac{w \sum_{j=1}^{T} \frac{1}{R^{j-1}}}{\sum_{j=1}^{T} \frac{(\beta R)^{(j-1)/\sigma_c}}{R^{j-1}}}.$$

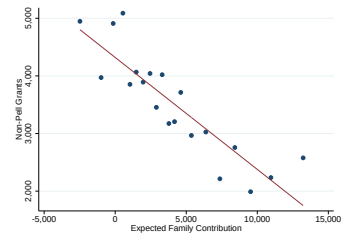
C.3. *External grants*

We use the BPS dataset to identify the external grants received by students, $Gr(y)$, that are not college-specific (need-based or merit-based) grants and federal Pell Grants. We do so by adding up all grants and aid excluding loans, Pell Grants, and institutional aid. The binned scatterplots in Figure C.2 show how these grants vary with parental income (here, Expected Family Contribution). These external grants are decreasing in EFC and tend to level off at around $15,000 (panel a). For an EFC below $15,000, we find a downward sloping relationship between EFC and non-Pell Grants using OLS and controlling for SAT scores (panel b).

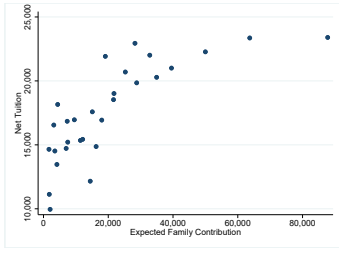
FIGURE C.2.—External grants and net tuition in the BPS



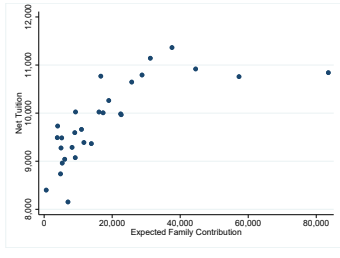
(a) External grants, all students



(b) External grants, low-income



(c) Net tuition, selective



(d) Net tuition, nonselective

Note: Panels (a)–(b): binned scatterplots of external grants (all grants excluding Pell Grants and institutional aid) against the expected family contribution, for all students and for low-income students. Panels (c)–(d): binned scatterplots of net tuition (sticker price minus college-specific grants) against the expected family contribution, by college tier. Source: BPS.

C.4. *Colleges' budget constraint*

Tuition caps. We use the BPS data to assess how net tuition (defined as sticker-price tuition minus college-specific need and merit-based grants) varies with student's parental income (mea-

sured as the EFC). Panels (c) and (d) of Figure C.2 provide binned-scatter plots of that relationship for selective and nonselective colleges. We see a leveling off of net tuition for higher-income students. This helps us set the tuition cap across the two college types, $\overline{T}_H$ and $\overline{T}_L$, which we set at \$25,000 for highly-selective colleges and at \$12,000 for less-selective colleges.

Costs and endowment. We use IPEDS data to estimate the relationship between non-instructional costs and enrollment for each college selectivity tier. There are two methods to measure costs: one is by directly adding up all non-instructional expenditure including academic support, student services, and institutional support in IPEDS. The other is by using the budget constraint in Equation (3.13), where we calculate cost by adding up all tuition revenue, net grant revenue, government appropriations, unrestricted revenue from private sources, and subtracting off total instructional expenditure. We posit the following relationship between non-instructional costs and enrollment for each college tier s

$$\text{Op Costs}_{i_s} = O_{0_s} + O_{2_s} \times \text{Enrollment}_{i_s}^2 + \varepsilon_{i_s},$$

where the costs in the left hand side are the maximum over both methods (this helps us deal with cases where the second procedure produces small or negative numbers). Table C.2 presents the results of the estimation. We follow Epple et al. (2006) to transform these estimates into our model parameters. Let $K_s = N_s \kappa_s$ denote total enrollment in colleges tier s, where it is assumed that colleges within each tier are the same. Then summing up individual cost functions yields

$$\text{Op Costs}(K_s) = \sum_{i=1}^{N_s} \left[O_{0_s} + O_{2_s} \times \frac{K_s^2}{N_s^2} \right] = N_s O_{0_s} + \frac{O_{2_s}}{N_s} K_s^2.$$

IPEDS data is also used to calculate the total level of private endowment income received by each college by adding up all private revenue received over the 2013-2016 sample period, which corresponds to unrestricted revenue from gifts, investment return from their endowment, and contributions from affiliates. The total endowment income of selective colleges is twice as large as that of nonselective colleges. In model units, this corresponds to $E_H(\kappa) = 0.12$ and $E_L(\kappa) = 0.06$. All coefficients are expressed at the tier level. The equivalent per-college normalization divides enrollment and the fixed cost by N_s and multiplies the quadratic endowment coefficient by N_s; marginal endowment income, $E_{1_s} + 2E_{2_s}\kappa_s$, and the per-student subsidy are invariant to this renormalization, so the two conventions deliver the same marginal conditions and the same equilibrium. The same invariance keeps tier-level outcomes unchanged when N_s varies, beyond the portfolio and offer-competition channels. To get the fixed costs in the model, C_{0_H} and C_{0_L}, we take the estimated constant from the regression, O_{0_H} and O_{0_L}, and multiply it by the number of colleges of each tier in our sample, as reported in Table A.2, net of the constant

term in endowment income, i.e.,

$$C_{0_s} = N_s O_{0_s} - E_{0_s},$$

where $E_{0_H} = 0.12 - E_{1_H}\kappa_H$ and we assume $E_{0_L} = 0$ (the constant component of endowment income is netted against the fixed cost, so that the model's endowment function has no constant term). Note that E_{1_H} is internally estimated to match the enrollment share of selective colleges, so E_{0_H}, and with it the net fixed cost, is recomputed at every parameter vector.

To proceed we compute the enrollment shares in each college selectivity tier observed in the HSLS data. These are $\kappa_H = 0.100$ and $\kappa_L = 0.523$. At the estimated $E_{1_H} = 0.54$, this procedure translates into $C_{0_H} = 0.086$ and $C_{0_L} = 0.25$. The linear term in the endowment function for nonselective colleges is given by $E_{1_L} = 0.06/\kappa_L = 0.11$. Without loss of generality, we place the quadratic enrollment term in the (net) endowment function. Taking the estimated coefficient on squared enrollment from the cost regression, dividing by the number of colleges, and flipping the sign yields $E_{2_H} = -5.80$ and $E_{2_L} = -0.44$.

TABLE C.2

OLS ESTIMATION RESULTS FOR COSTS

	Selective colleges	**Nonselective colleges**
Enrollment^2	1057.6***	624.2***
	(120.4)	(14.85)
Private college dummy		-0.0000919***
		(0.0000153)
Constant	0.000841***	0.000222***
	(0.000111)	(0.0000126)
N	182	1445
Adj. R^2	0.296	0.604

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Standard errors in parentheses. Variables have been scaled down by the total student population. Costs have also been scaled down by \$40,000 to fit the model units.

Government transfers. We estimate the government transfer functions for selective and nonselective colleges using the IPEDS data. Government transfers correspond to state government appropriations and federal funds colleges received. Table C.3 presents the results of the estimation of the linear relationship between government transfers and student enrollment, which help identify Tr_{1_H} and Tr_{1_L}.

TABLE C.3

OLS ESTIMATION RESULTS FOR GOVERNMENT TRANSFERS

	Selective colleges	**Nonselective colleges**
Enrollment	1.044***	0.620***
	(0.0516)	(0.0112)
N	182	1445
Adj. R^2	0.691	0.679

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Standard errors in parentheses.

C.5. *Standard errors*

The estimator minimizes the unweighted sum of squared residuals, $g_j(\Theta) = \left(M_j^{data} - M_j(\Theta)\right)/M_j^{data}$, so the weighting matrix is the identity and the asymptotic variance takes the form

$$\widehat{\mathrm{Var}}(\hat{\Theta}) = (G'G)^{-1}\, G'\, \widetilde{\Omega}\, G\, (G'G)^{-1},$$

where $G = \partial g/\partial \Theta'$ is the Jacobian of the residuals, computed by central finite differences of the exactly computed model moments and $\widetilde{\Omega}$ is the covariance matrix of the residuals of the data moments, which we approximate as diagonal.

The variances of the data moments are set moment by moment. An HSLS share moment is a sample proportion, so its variance is binomial, $p(1-p)/N$ (where N is the size of the cell in which the share is computed: 18,913 students in the full analytic sample, 11,399 who applied to any college, 2,773 who applied to a selective college, and 1,797 admitted to both tiers). A moment defined as a ratio of two application rates combines the two cells' binomial variances by the delta method. For the BPS tuition-responsiveness coefficient and the IPEDS/BPS dollar-level moments, the underlying standard errors cannot be recovered, so we assign them coefficients of variation of 10 and 5%. These are generous relative to the sampling precision of samples of this size, so the resulting parameter standard errors are conservative.

C.6. *Application and attendance disutility costs*

Table C.4 reports the consumption-equivalent magnitudes of the estimated application and attendance disutility costs.

TABLE C.4

CONSUMPTION EQUIVALENTS OF THE ESTIMATED DISUTILITY COSTS

Panel A: Application disutility

	Overall average	Enrollee average (Bottom 50% income)	Fu (2014)
Applying to selective colleges only	0.90%	0.41%	0.30%
Applying to nonselective colleges only	1.01%	0.66%	0.34%
Applying to both	2.19%	0.95%	0.45%

Panel B: Psychic costs of attending college

	Model $\nu_S(\ell)$	Abbott et al. (2019)
Selective colleges	−1.80%	10.7%
Nonselective colleges	3.35%	
Among enrollees	2.47%	4.6%

Note: Consumption equivalents as a percentage of average life-cycle consumption (about $2.5 million in the model). Panel A: disutility incurred at the chosen application portfolio; Fu (2014) values are computed assuming two applications to selective colleges, three to nonselective colleges, and six to both. Panel B: enrollment-weighted (a negative value is a net psychic benefit). Abbott et al. (2019) values are provided for reference.

C.7. *Completion moments for the graduation-risk extension*

Table C.5 reports the eight completion moments targeted by the graduation-risk extension of Section 7 and the model's fit.

TABLE C.5

SIX-YEAR BACHELOR'S COMPLETION BY TIER AND TEST-SCORE QUARTILE

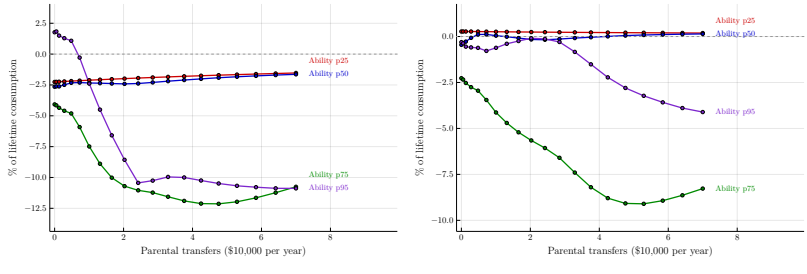
	Selective		Nonselective	
Test-score quartile	Data	Model	Data	Model
Q1	0.620 (0.026)	0.616	0.290 (0.030)	0.351
Q2	0.669 (0.017)	0.661	0.541 (0.041)	0.416
Q3	0.785 (0.012)	0.746	0.504 (0.050)	0.494
Q4	0.852 (0.012)	0.892	0.646 (0.070)	0.580

Note: Data are weighted six-year bachelor's completion rates of first-time students who entered four-year public and private not-for-profit institutions directly from high school in 2011–12, from BPS:12/17. Model completion rates are $\mathbb{E}[\delta_s(\ell) \mid \text{enrolled in } s,\ \sigma \in \Omega_q]$, with quartile cutoffs taken from the BPS. All eight moments are targeted.

APPENDIX D: WELFARE EFFECTS OF THE COUNTERFACTUALS

This appendix reports the consumption-equivalent variation by income quartile and ability group for each counterfactual in Sections 5 and 6, defined as the percentage of consumption in both periods that a student of each type would be willing to give up (negative: would need to receive) to be born into the counterfactual equilibrium rather than the baseline. Averages are population-weighted. Table D.1 reports the sensitivity of the baseline to the level of the tuition caps, discussed in Section 5.2. Table D.2 reports the welfare effect of the five policy experiments of Section 5 (the fee waiver, halved fixed application costs, the two Pell expansions, and the relaxed borrowing limit). Table D.3 reports the welfare effect of the three information experiments of Section 6 (noisier assessments, the AI scenario, and perfectly informative assessments). Figure D.1 plots the welfare effect across parental transfers for the aid and borrowing experiments (the information experiments are in Figure 6.1). Table D.4 reports the welfare effect of perfectly informative assessments in the graduation-risk extension of Section 7.

FIGURE D.1.—Welfare effect of the aid and borrowing experiments



(a) $25,000 Pell maximum (b) Relaxed borrowing limit

Note: Consumption-equivalent variation by parental transfers, for students at the 25th, 50th, 75th, and 95th percentiles of the ability distribution. Positive values are welfare gains from the policy.

TABLE D.1

SENSITIVITY TO THE TUITION CAPS

		Selective colleges			Nonselective colleges		
		Baseline	$0.9\,\overline{T}$	$1.1\,\overline{T}$	*Baseline*	$0.9\,\overline{T}$	$1.1\,\overline{T}$
κ_S	% Enrolled	9.2	8.8	9.2	39.5	39.3	39.2
	Avg. tuition net of all grants ($1k/year)	13.3	11.9	14.4	1.1	0.3	1.0
I/κ_S	Instr. spending per student ($1k/year)	16.5	15.3	17.6	5.4	4.5	5.3
L/κ_S	Average student ability	1.79	1.67	1.79	1.28	1.31	1.28
Z_S	Value added	2.003	1.868	2.033	1.789	1.772	1.785
$\underline{\sigma}_S$	Admission standard	1.54	1.35	1.53	-0.91	-0.91	-0.91
τ	Tax rate (%)	1.73	1.75	1.72	1.73	1.75	1.72
	Mean number of applications (all students)	2.73	2.64	2.71	2.73	2.64	2.71
	Mean number of applications (top-decile ability)	6.25	5.55	6.27	6.25	5.55	6.27
y	*Income distribution (%)*						
	Q1 Income	4.9	9.0	4.9	21.7	21.1	21.8
	Q2 Income	12.6	14.0	12.8	32.3	31.9	32.4
	Q3 Income	16.8	5.7	16.8	23.5	25.8	23.6
	Q4 Income	65.6	71.3	65.5	22.4	21.2	22.2
ℓ	*Ability distribution (%)*						
	0-30% Ability	0.0	0.0	0.0	0.1	0.2	0.1
	30-60% Ability	4.8	9.5	4.8	43.3	41.2	42.8
	60-90% Ability	38.3	46.5	37.5	50.1	48.7	50.7
	Top 10% Ability	56.9	44.0	57.6	6.4	10.0	6.3

TABLE D.2

WELFARE EFFECT OF THE POLICY EXPERIMENTS

	0–30% ability	30–60%	60–90%	Top 10%
Panel A: Fee waiver (average -0.04%)				
Q1 income	0.00	0.00	-0.08	-0.11
Q2 income	0.00	-0.01	-0.08	-0.09
Q3 income	0.00	-0.03	-0.10	-0.08
Q4 income	0.00	-0.07	-0.12	-0.03
Panel B: Halved fixed application costs ($\phi_S/2$) (average -0.15%)				
Q1 income	0.12	0.25	-0.70	-1.31
Q2 income	0.11	0.10	-0.95	-1.13
Q3 income	0.10	-0.05	-0.98	-0.43
Q4 income	0.10	0.15	-0.32	-0.47
Panel C: Pell maximum $7,395 (average -0.65%)				
Q1 income	-0.18	-0.02	-0.53	-0.58
Q2 income	-0.17	-0.19	-0.73	-0.66
Q3 income	-0.17	-0.83	-2.07	-1.76
Q4 income	-0.15	-1.17	-2.02	-0.73
Panel D: Pell maximum $25,000 (average -3.65%)				
Q1 income	-2.26	-2.99	-3.10	3.66
Q2 income	-2.21	-3.09	-3.58	3.47
Q3 income	-2.11	-3.70	-6.43	0.36
Q4 income	-1.85	-4.15	-12.01	-6.42
Panel E: Relaxed borrowing limit ($\downarrow \underline{a}_S$) (average -1.16%)				
Q1 income	0.28	-0.95	-2.25	1.06
Q2 income	0.27	-0.93	-2.74	1.11
Q3 income	0.26	-1.28	-3.54	1.28
Q4 income	0.23	-2.01	-5.99	1.01

Note: Consumption-equivalent variation, in percent of consumption in both periods (positive values are gains from the policy). Rows: parental-transfer quartiles; columns: ability groups. Panel averages are population-weighted; the discrete ability grid puts the four ability groups' population masses at 36%, 32%, 24%, and 8%.

TABLE D.3

WELFARE EFFECT OF THE INFORMATION EXPERIMENTS

	0–30% ability	30–60%	60–90%	Top 10%
Panel A: Noisier assessments ($\text{corr}(\ell, \sigma) = 0.68$) (average -0.46%)				
Q1 income	-0.01	-0.08	-0.86	-2.44
Q2 income	-0.01	-0.12	-0.87	-2.30
Q3 income	-0.01	-0.19	-0.85	-2.19
Q4 income	-0.01	-0.27	-1.03	-2.49
Panel B: AI scenario ($\phi_S/2$ **and** $\text{corr}(\ell, \sigma) = 0.68$) (average -0.68%)				
Q1 income	0.11	0.08	-1.54	-4.39
Q2 income	0.10	-0.08	-1.84	-4.25
Q3 income	0.09	-0.28	-2.03	-3.27
Q4 income	0.10	0.16	-1.53	-3.70
Panel C: Perfectly informative assessments (average 0.72%)				
Q1 income	0.09	0.03	3.36	5.28
Q2 income	0.09	-0.22	2.41	4.70
Q3 income	0.09	-0.20	1.89	3.69
Q4 income	0.08	0.73	0.84	0.10

Note: Consumption-equivalent variation, in percent of consumption in both periods (positive values are gains from the policy). Rows: parental-transfer quartiles; columns: ability groups. Panel averages are population-weighted; the discrete ability grid puts the four ability groups' population masses at 36%, 32%, 24%, and 8%.

TABLE D.4

WELFARE EFFECTS OF PERFECTLY INFORMATIVE ASSESSMENTS UNDER GRADUATION RISK

	0-30% ability	30-60%	60-90%	Top 10%
Q1 income	0.08	-0.42	-2.65	2.53
Q2 income	0.08	-0.86	-3.63	1.85
Q3 income	0.07	-1.16	-4.49	3.66
Q4 income	0.06	-1.38	-3.57	5.55
Average		-0.85		

Note: Consumption-equivalent variation, in percent of consumption in both periods; positive values are gains from the policy. Rows: parental-transfer quartiles; columns: ability groups. Average is population-weighted; the discrete ability grid puts the four ability groups' population masses at 36%, 32%, 24%, and 8%.