

College Tuition and Income Inequality

By ZHIFENG CAI AND JONATHAN HEATHCOTE*

This paper evaluates the role of rising income inequality in explaining observed growth in college tuition. We develop a competitive model of the college market, in which college quality depends on instructional expenditure and the average ability of admitted students. An innovative feature of our model is that it allows for a continuous distribution of college quality. We find that observed increases in US income inequality can explain more than half of the observed rise in average net tuition since 1990 and that rising income inequality has also depressed college attendance.

JEL: E21, E25, I23, I24

Keywords: College Tuition, Income Inequality, Club Goods, Peer Effects

For decades, the average cost of college tuition in the United States has been rising much faster than general inflation (Figure 1), and paying for college has become a major concern for households with children. Policymakers worry that rising tuition costs may put a college education out of reach for high-ability children from low-income households. Given these concerns, it is important to understand what is driving up tuition. In this paper, we evaluate the hypothesis that rising income inequality has been a key driver of rising tuition.

This hypothesis is motivated by the fact that colleges in the United States draw their students disproportionately from relatively high-income households (see, for example, Figure 4 in Chetty et al. 2014). Rapid income growth at the top of the income distribution in recent decades has increased these households' willingness to pay for high-quality colleges. Lower-income households have experienced much weaker income growth over the same period, but this has likely not had a fully offsetting negative impact on college demand, given that few children from such households have ever attended college.

Predicting the impact of increasing college demand on college pricing requires modeling the college market. The model we develop follows the existing literature in recognizing two determining factors in the quality of a college education. The first is the amount of instructional resources devoted to each student. The second is the average ability of the student body, which could be interpreted as capturing

* Cai: Department of Economics, Rutgers University, 75 Hamilton St, New Brunswick, NJ 08901, USA., (email: zhifeng.cai@rutgers.edu). Heathcote: Research Department, Federal Reserve Bank of Minneapolis, 90 Hennepin Ave., Minneapolis, MN 55401, USA., (email: heathcote@minneapolisfed.org). We thank Richard Rogerson, Grey Gordon, Sebastian Findeisen, Lance Lochner, Betsy Caucutt, John Bailey Jones, Oksana Leukhina and Todd Schoellman for helpful comments. We thank Job Boerma for research assistance. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.




FIGURE 1. COLLEGE TUITION AND FEES (\$2016)

Note: This figure plots sticker and net tuition and fees for public four-year (in-state) and private nonprofit four-year colleges.

average IQ or college preparedness. Schools with higher average student ability might be more attractive to college applicants for two reasons: (i) they offer better prospects for learning from peers, and (ii) they offer social and professional connections to people who are likely to be successful post-graduation. To the extent that student ability is an important and a relatively inelastic input in producing college quality, increased demand for college will drive up equilibrium (quality-adjusted) tuition and not simply lead to an increase in the supply of high quality college spots.

Households in our model differ with respect to household income and the ability of the household child. Colleges can observe both income and ability (e.g., by observing test scores) and, in principle, can price discriminate in both dimensions. Households face tuition schedules for colleges of different quality levels and decide whether to send their child to college and, if so, to which quality of college.

On the supply side, the technology for producing college quality is a function of instructional expenditure per student and the average ability of enrolled students. We assume this technology exhibits constant returns to scale, and we allow for free entry at any quality level. This implies a competitive setting in which colleges have no pricing power and take equilibrium tuition schedules as given. Each potential college chooses a quality level at which to enter and, conditional on a chosen quality level, seeks to deliver that quality as cheaply as possible by optimally balancing resource spending versus the ability composition of the student body. The model delivers a continuous equilibrium distribution of college quality, a feature that is novel relative to the existing literature.

As in other “club good” models, the characterization of a competitive equilibrium is complicated by the fact that club members (students) are both consumers and inputs into production, which implies a large number of market-clearing conditions. In particular, for each college quality level, the number of students demanding college spots and the ability composition of those students must be consistent with colleges’ choices about the number and composition of students to “employ” as quality-producing inputs.

Given this complication, all the existing literature in the club good tradition assumes a very small number of different college quality types. The primary theoretical contribution of our paper is to explore a competitive setting with constant returns which allows for a continuous distribution of quality. This has both theoretical and practical advantages. The theoretical advantages are that we can prove an equilibrium exists and that the equilibrium is efficient. The main practical advantage is that we can compare the equilibrium model distribution of college characteristics with US data, which include thousands of different colleges. Relatedly, our continuous distribution of college quality can change smoothly when we change income inequality or other drivers of college demand.

While the model features a continuous distribution of college quality and thus a continuum of market-clearing conditions and prices, it is nonetheless quite tractable. As in other club good college models, colleges offer lower tuition to high-ability students, internalizing that such students contribute more to college quality. We prove that this ability discount is linear in ability, because quality depends on the average ability of enrolled students. There is no equilibrium price discrimination by income: any such discrimination would present an opportunity to profitably skim off high-income households. Equilibrium tuition increases with college quality, which implies the usual pattern of sorting: holding ability fixed, higher-income students match in a positive assortative fashion with higher-quality schools. Combining these insights, we show that it is possible to solve for equilibrium by iterating across the quality distribution: at each quality level (i) the density of college spots satisfies total demand, (ii) baseline tuition is such that colleges make zero profits, and (iii) the tuition discount per unit of ability equates the average ability of students wishing to attend with the average ability of students that colleges want to admit.

In the first part of the paper, we characterize equilibrium in closed form in a version of the model with no resource inputs in producing college education, two ability types, and a uniform distribution for household income. We use this closed-form example to gain intuition about what determines equilibrium college prices and the distribution by quality of college spots in a club good environment. We use it also to gain insight about how these objects vary with income inequality. The comparative statics are striking. In particular, changing income inequality has absolutely no impact on the equilibrium allocation of households across colleges of different qualities and changes only equilibrium tuition pricing.

This result motivates the second part of our paper, in which we calibrate a richer version of the model and use it to explore the role of rising income inequality and other factors in driving observed changes over time in college tuition, college attendance, and the distribution of college quality. The quantitative version of the model adds several salient features of the US college market. First, we introduce drop-out risk to reflect the fact that a large share of students who enroll in college do not graduate. We use evidence from the National Longitudinal Survey of Youth (NLSY) to calibrate how drop-out rates vary by ability. Second, we model a range of subsidies that impact the net cost of college and thus enrollment choices. Residence-based subsidies capture tuition discounts enjoyed by in-state students enrolled at public universities. Need-based subsidies, such as Pell Grants, reduce the cost of attendance for students from low-income families. With these features, the model generates rich variation in the net cost of tuition: at any given quality level, different students face different net costs depending on ability (institutional aid), family income (need-based aid), and residence status (in-state versus out-of-state).

We set preference parameters so that the model replicates both out-of-pocket spending on college education and enrollment rates. A key model input is the joint distribution over household income and ability. We estimate a Pareto lognormal distribution for household income, using data from the Survey of Consumer Finances. We discipline the correlation between income and ability using the NLSY (where ability is proxied by Armed Forces Qualifying Test scores). The calibrated model generates a realistic enrollment pattern by family income and by student ability. It also generates distributions of sticker and net tuition across colleges that are similar to those observed empirically.

Our key quantitative experiment is to explore the implications of the change in the US household income distribution between 1989 and 2016 for the pattern of college attendance, the distribution of college quality, and the shape of equilibrium tuition schedules. Over this period, we find evidence of both a general increase in income dispersion and a significant fattening of the right tail of the distribution. These changes can account for several key features of the data. First, rising income inequality drives up average net tuition. The combined effects of rising income inequality and rising average income can account for the entire 53 percent increase in net tuition observed in the data between 1990 and 2016,

with higher inequality playing the dominant role. Second, rising income inequality can account for a widening gap between average sticker tuition and average net tuition paid, reflecting increasingly large institutional discounts for desirable (high-ability) applicants (see Figure 1). Third, rising inequality depresses college enrollment, while higher average income boosts it.

We also explore the effects of changes to college subsidies and find that rising subsidies to students, in the form of more generous need-based aid and larger in-state discounts, have moderated growth in net tuition and have boosted enrollment.

Related Literature: There is a very large empirical literature on peer effects in education (see Epple and Romano 2010 and Sacerdote 2014 for excellent surveys) and a related literature that integrates peer effects into structural “club goods” models of the college market (see Rothschild and White 1995 and Epple and Romano 1998 for important early contributions). In these models, and in ours, students are important inputs into the production of perceived college quality. To motivate this assumption, we now briefly discuss some empirical evidence on the importance of peers.

At a broad level, there is strong evidence that in education settings, parents and students care a lot about peers. In K–12 education, parents are willing to pay high house price premia in order to enroll their children in higher scoring and whiter school districts (e.g., Black 1999; Boustan 2012). In school districts like New York and Boston there is also fierce competition to attend more academically selective magnet schools (Abdulkadiroğlu, Angrist and Pathak, 2014). At the college level, quality rankings like the ones published by US News are strongly correlated with measures of peer ability.¹

Why do parents care so much about their children’s peers? The famous Coleman (1966) report emphasized that the educational backgrounds and aspirations of children’s peers seemed a very important predictor of student performance, while variation in resource spending across schools was not large enough to be a central factor. More recent research has tried to quantify peer effects more precisely and to uncover the mechanisms at work. Sacerdote (2001) found that having high GPA roommates in college tends to raise a student’s own GPA, and Carrell, Fullerton and West (2009) find even larger effects for squadron peer groups at the US Air Force Academy. Zimmerman (2003) finds positive effects on grades from having high verbal SAT-scoring roommates.

It is not easy to distinguish causal effects of peers on performance from more mechanical links between individual outcomes and the average outcomes of the groups to which an individual belongs (see Manski 1993 and Angrist 2014). Using a regression discontinuity approach that in principle addresses these econometric concerns, Hoekstra (2009) finds that attending a flagship state university boosts

¹US News explicitly assigns a ranking weight of only 10 percent to “student selectivity” (mostly SAT and ACT scores). However, the other inputs to their rankings are effectively highly correlated with these scores. Wai, Brown and Chabris (2018) report a correlation of 0.892 between college average SAT score and US News National University Rank.

future earnings by 20 percent, while Zimmerman (2019) finds that in Chile being admitted to a highly selective college improves the chances of landing a top job. Dale and Krueger (2014) argue that the returns to attending a more selective college are small, but Chetty et al. (2020), following a similar approach, find large causal effects of colleges on future earnings, with a relatively small role for selection by unobserved student characteristics into college. Note, however, that these papers do not address to what extent better outcomes at more selective colleges reflect the presence of more desirable peers versus other factors.

Another important question is how peers matter. The typical assumption in the literature is that average peer ability is the relevant measure (Arnott and Rowse, 1987), though there is interesting recent work investigating and finding evidence that other properties of the ability distribution matter and that different types of students benefit differentially from different sorts of peers (see, e.g., Lazear, 2001 and Lavy, Paserman and Schlosser, 2012).

Many studies on the effects of peers have focused on relatively narrow measures of academic performance like grades and test scores. In principle, however, parents and children may care about peers for a much wider range of reasons, including friendship networks and the pool for potential spouses.² In addition, college is an important source of connections that are useful for future career development (see, e.g., Marmaros and Sacerdote, 2002).³ Parents and students are likely acutely aware of these sorts of peer considerations, even if they have not been the prime focus of academic research into peer effects.

On the structural modeling side, several papers build on the influential contribution of Epple and Romano (1998). Our model shares many key ingredients with this literature: college quality depends on expenditure and average student ability, tuition varies with ability in a way that reflects these peer effects, and equilibrium student stratification across colleges of different qualities reflects both differences in ability and in family income. However, Epple and Romano (1998) and most subsequent work assume a small number of colleges, which can be justified by positing large economies of scale. In contrast, we assume constant returns, which translates into an equilibrium with many small colleges, and a continuous quality distribution. Thus we avoid the equilibrium existence problems that typically arise when there are only a few college clubs, each of which is large relative to the size of the economy. These existence problems are discussed by Ellickson et al. (1999) and Scotchmer (1997) and have to do with the fact that when clubs are optimally large relative to the economy, partitioning the population into an integer number of optimally sized groups is typically not possible. Because of

²For example, according to a Facebook study, at least 28 percent of college graduates marry people they meet in college. This suggests that by attending a college whose graduates earn more on average, a student can expect higher spousal earnings and thus higher household income, independent of any impact on the student's own earnings. And if you want to marry a billionaire, you have a much better chance if you attend an elite university. Priscilla Chan met Mark Zuckerberg at Harvard, for example.

³Steve Ballmer was fortunate to be Bill Gates' dorm hall neighbor at Harvard. Sergey Brin and Larry Page met at a new student orientation at Stanford.

these problems, Epple and Romano (1998) are forced to focus on approximate equilibria. They note that one could solve the existence problem by allowing for “constant costs of schooling,” which would “lead to an infinite number of schools serving infinitely refined peer groups.” They note that while such a model “is extremely interesting, it is quite complex and not yet tractable” (p. 59). This infinitely refined peer group model is the one we solve.

A second problem with assuming only a small number of competing colleges is that in such an environment, the natural model for competition is strategic oligopoly. Each college would then choose a pricing (or quality) strategy, where each strategy specifies best responses given the strategies of its competitors. Instead, most of the existing papers in the literature assume that each college takes competitors’ prices (or students’ willingness to pay) as given when choosing its own price. We are able to sidestep the difficult task of modeling strategic interactions between colleges: price-taking is the natural assumption in our competitive setting in which colleges are all small.

Caucutt (2001, 2002) takes a different approach. In her model, there is a small number of different college types but a large number of colleges of each type. Each college club is small relative to the economy as a whole, and there are no equilibrium existence problems. Households buy probabilities of attending different colleges, building on Cole and Prescott (1997). One limitation of her approach is that she assumes only two different income levels, which implies a very small number of different school types in equilibrium. In contrast, we assume a continuous income distribution, and we do not need to introduce lotteries (which we do not observe in practice) in order to ensure type-independent equilibrium allocations.

Two important recent papers that model the college market are Epple, Romano and Sieg (2006) and Fu (2014). Both papers structurally estimate rich models. Neither paper is focused on exploring the drivers of rising college tuition. In addition, there are important differences in terms of how college supply is modeled between these papers and ours. In particular, distinctive features of our model are that we allow for entry in the college market, without imposing any restrictions on the number of colleges or the qualities at which they can enter.

There is a set of papers that explores the potential drivers of rising college tuition. Gordon and Hedlund (2017) consider various possible factors within a variant of the model in Epple et al. (2017). They find that rising financial aid is the most important factor that is pushing up tuition. In their model, a single monopolistic college seeks to maximize the quality of education per student enrolled. When more public aid increases students’ ability to pay, this monopolist responds by increasing spending on quality-increasing inputs. In our model, in contrast, more need-based aid induces more low-income households to enroll in college, and the market responds by expanding at the low-quality end of the distribution, thereby driving *down* average tuition.

Jones and Yang (2016) argue that rising college tuition reflects service sector

disease: productivity in higher education is assumed to be constant, but the cost of college professors continues to rise, reflecting productivity growth and a rising college wage premium in the rest of the economy. Thus, in their model, rising income inequality plays a supply-side role in driving up the cost of college. We explore the role of rising supply-side costs and find they play a very small role in explaining tuition trends, relative to changes on the demand side.

There is a strong positive empirical correlation between family income and college attendance and, conditional on attendance, a strong correlation between family income and proxies for college quality (see, for example, Belley and Lochner 2007; Chetty et al. 2020, 2017; and a recent column in the *New York Times*).⁴ Our model predicts similar correlations. An important finding of Belley and Lochner (2007) is an increase in the importance of family income relative to student ability in predicting college attendance between the 1979 and 1997 waves of the NLSY. Similarly, Chetty et al. (2017) report in their table 2 that the share of enrollment at selective and highly selective private and public universities from households in the bottom 60 percent of the household income distribution declined from 2000 to 2011. When we feed in observed changes in the household income distribution and in college subsidies, we find that our model delivers changes in enrollment patterns similar to those reported by Belley and Lochner (2007) and Chetty et al. (2020). In particular, income becomes a more important predictor of college attendance, relative to ability, with large increases in model college enrollment from low-ability households in the top half of the income distribution.

I. Model

Households: The economy is populated by a continuum of measure 1 of households, each containing a parent and a college-age child. The baseline model is static, and within each household, the parent and child operate as a single decision-making unit. Households are heterogeneous with respect to parental income y , child “ability” a , and residence status r .

Ability is a summary statistic for any child characteristics that determine the child’s potential contribution to college quality. Ability also affects the probability of successfully graduating from university, conditional on enrolling. We assume ability is drawn from a finite discrete set of values A . Let μ_a denote the share of the population with ability value a , with $\sum_{a \in A} \mu_a = 1$. Conditional on ability, household income is continuously distributed, where the cumulative distribution function $F_a(y)$ is ability-type specific. The distribution function has a strictly positive compact support.

Given our focus on explaining tuition patterns we want to confront the large differential between tuition for in-state students at public universities versus tu-

⁴Gregor Aisch, Larry Buchanan, Amanda Cox, and Kevin Quealy, “Some Colleges Have More Students from the Top 1 Percent than the Bottom 60. Find Yours,” TheUpshot, *New York Times*, January 18, 2017, <https://www.nytimes.com/interactive/2017/01/18/upshot/some-colleges-have-more-students-from-the-top-1-percent-than-the-bottom-60.html>.

ition for out-of-state students or for private university students. We interpret this differential as arising because colleges receive additional public subsidies for admitting in-state students, which they pass on in the form of discounted tuition. Model residence status $r \in \{i, o\}$ determines whether a child is eligible for these discounts: i denotes in-state, and o denotes out-of-state. The empirical counterpart to model out-of-state students will be students enrolled in private universities or in public universities as out-of-state students.⁵ Heterogeneity in residence status is a simple and stylized way to capture heterogeneity in student access to subsidized public colleges.⁶ The fractions of in- and out-of-state students are μ_i and $\mu_o = 1 - \mu_i$. We assume that residence status is uncorrelated with parental income and child ability.

Each household chooses whether to enroll their child in college and, if so, in which quality of college. Expected utility is increasing in nondurable consumption c and in the quality of the college in which the child enrolls q . However, the utility value of college quality is realized only if the child successfully graduates, which happens with probability γ_a . Expected utility is given by

$$(1) \quad \mathbb{E}_{|a} [u(c, q)] = \log c + \varphi (\gamma_a \log(\kappa + q) + (1 - \gamma_a) \log \kappa),$$

where the preference parameters φ and κ are common across households. The parameter κ can be interpreted as the reservation utility level associated with a less-than-college education. If $\kappa > 0$, then utility is bounded below for children not graduating from college.

Interpreting this utility function at face value, households are willing to pay for college because it delivers a direct consumption value. Alternatively, the same utility specification can be motivated as reflecting a setting in which college is valued indirectly as an investment technology that can increase child earnings. We will further discuss this alternative interpretation in Section V.

Households take as given tuition schedules $t(q, y, a, r)$ that specify the out-of-pocket tuition charged by colleges offering quality q . The dependence of tuition on ability reflects the fact that different ability types are differentially attractive to colleges, and thus colleges will price discriminate. We allow tuition to depend on income in principle, but we will shortly show (Proposition 1) that any equilibrium can be supported by a tuition function that is independent of income.

Let \mathcal{Q} denote the set of college qualities available in equilibrium. One element of \mathcal{Q} is $q = 0$, which corresponds to the choice not to enroll in college. If the household enrolls the child in college so that $q > 0$, then in addition to tuition, it pays a fixed cost ω , corresponding to the household earnings that are forgone when the child is in college instead of at work. The household can partially defray

⁵In a previous draft of the paper, we explicitly differentiated between private and public universities.

⁶There is significant variation across US states in the size of the public four-year college sector, the range of quality offered within that sector, and the size of state subsidies. State and local support for higher education in 2018 ranged from \$94 per capita in New Hampshire to \$704 per capita in Wyoming (State Higher Education Executive Officers Association SHEF Report, 2019).

the cost of tuition through direct public financial aid, $p(y)$. Public financial aid depends on income, reflecting the fact that federal and grants are often means-tested. Tuition is not refunded if a student enrolls but fails to graduate. The household problem, for a household with idiosyncratic state (y, a, r) is therefore

$$\begin{aligned} & \max_{c \geq 0, q \in \mathcal{Q}} \mathbb{E}_{|a} [u(c, q)] \\ & \quad \text{s.t.} \\ & c + \mathbb{1}_{\{q > 0\}} [t(q, y, a, r) + \omega - p(y)] = y. \end{aligned}$$

Let $c(y, a, r)$ and $q(y, a, r)$ describe the decision rules that solve this problem.

Colleges: Colleges can enter and supply college spots at any quality level in a feasibility set $\Omega = [0, q_{max}]$.⁷ The technology for producing college spots is constant returns to scale, where quality depends on the average ability of the student body and expenditure (per student) e on quality-enhancing goods and services. Each college admits a continuous mass of students. Let η_a denote the fraction of these students who are of ability type a . Quality (per student) at a college admitting students of average ability \bar{a} and spending e on instructional inputs is

$$(2) \quad q = \bar{a}^\theta e^{1-\theta},$$

where $\bar{a} = \sum_{a \in A} \eta_a a$. The parameter θ determines the relative importance of average student ability versus goods inputs in producing quality. Note that this production technology embeds an assumption that any mix of enrolled students with the same average ability makes an identical contribution to college quality.⁸

Colleges must also pay a fixed resource cost ϕ per student enrolled, which captures administration and other costs that do not directly enhance quality. The fact that the technology for supplying college spots is constant returns to scale supports the existence of an equilibrium in which all colleges are small, in the sense that they enjoy no pricing power. The logic is that each college competes against other colleges offering identical quality.⁹

Colleges observe the income, ability, and residence status of applicants. They seek to maximize profits or, equivalently, to provide a given market value of education at the lowest possible cost.¹⁰ In addition to tuition revenue, they also

⁷The upper bound q_{max} is useful for establishing equilibrium existence. This bound can be arbitrarily large and thus will not restrict entry in equilibrium.

⁸Note that in the model, we have effectively defined the units of ability as measuring the marginal value of students to colleges: if one student has twice the ability of another, they are twice as valuable on the margin as a peer and twice as desirable to colleges. We will later show that an implication of this scaling choice for ability is that the equilibrium tuition function is linear in ability.

⁹Even if a college had a monopoly at a given quality level, given a continuous quality distribution it would still face near-identical competitors and thus enjoy no pricing power.

¹⁰Some other papers in the literature assume that colleges seek to maximize college quality. In a competitive environment in which colleges take tuition schedules as given, quality maximization would imply a degenerate college quality distribution, with all colleges operating at the highest feasible quality level.

receive per student public subsidies $s(q, a, r)$ that potentially depend on the quality they deliver and on the characteristics of students they admit. In particular, colleges will receive more public support if they admit more in-state students. Note that we assume subsidies to *colleges* are independent of student income (in contrast to subsidies to *students*).

Conditional on ability, profit-maximizing colleges will strictly prefer to admit only the students who generate the most revenue. Let $v(q, a) = \max_{y,r} t(q, y, a, r) + s(q, a, r)$ denote revenue from admitting the highest-revenue students (by income and residence status) of ability level a .

Suppose a college has decided to supply college education at quality level q . The input mix sub-problem for supplying mass one spots at quality q is

$$\begin{aligned} & \max_{\{\eta_a\} \geq 0, e \geq 0} \left\{ \sum_{a \in A} \eta_a v(q, a) - e - \phi \right\} \\ & \text{s.t.} \\ & q = \bar{a}^\theta e^{1-\theta} \text{ and } \sum_{a \in A} \eta_a = 1, \text{ where } \bar{a} = \sum_{a \in A} \eta_a a. \end{aligned}$$

Let $\{\eta_a(q)\}_{a \in A}$ and $e(q)$ denote the solution to this problem, and let $\pi(q)$ denote corresponding profit per student. Given a revenue schedule $v(q, a)$, colleges will optimally supply zero mass of college spots at qualities q where $\pi(q)$ is negative, will be indifferent about the mass of spots to supply if $\pi(q) = 0$, and will want to supply an infinite mass of spots if $\pi(q)$ is strictly positive.

Let $\chi(Q)$ denote the measure of college places in colleges of quality $q \in Q \subset \Omega$. This is a key equilibrium object. In contrast, the size distribution of colleges within any given quality level is indeterminate, given our constant returns to scale quality production function.

Definition of Equilibrium: An equilibrium in this model is a measure χ and functions $t(q, y, a, r)$, $c(y, a, r)$, $q(y, a, r)$, $\{\eta_a(q)\}_{a \in A}$, $e(q)$, and $\pi(q)$ that satisfy the following conditions:

- a. Household optimization: Given $t(q, y, a, r)$, the household choices $q(y, a, r)$ and $c(y, a, r)$ solve the household's problem for all (y, a, r) .
- b. College optimization: Given $v(q, a) = \max_{y,r} \{t(q, y, a, r) + s(q, a, r)\}$, the college input choices $\{\eta_a(q)\}_{a \in A}$ and $e(q)$ solve the college's problem for all $q > 0$, and $\pi(q)$ is the associated profit per student.
- c. Zero profits: For all $Q \subset \Omega$, $\int_Q \pi(q) d\chi(q) = 0$, and $\pi(q) \leq 0$ for all $q \in Q$.
- d. Goods market clearing:

$$\sum_r \mu_r \sum_a \mu_a \int c(y, a, r) dF_a(y) + \int_0^{q^{max}} e(q) d\chi(q) + (1 - \chi(0))(\omega + \phi) = \sum_a \mu_a \int y dF_a(y) + S,$$

where S denotes the total value of all public subsidies

e. College market clearing: For all a and $Q \subset \Omega$,

$$(4) \quad \mu_a \sum_r \mu_r \int \mathbb{1}_{\{q(y,a,r) \in Q\}} dF_a(y) = \int_Q \eta_a(q) d\chi(q),$$

where $\mathbb{1}_{\{\cdot\}}$ is an indicator function and where for all y and $r \in \{i, o\}$ and for all $q^* \in Q$,

$$(5) \quad q(y, a, r) = q^* \Rightarrow (y, r) \in \arg \max \{t(q^*, y, a, r) + s(q^*, y, a, r)\}.$$

Condition c here is the zero profit condition that follows from free entry and perfect competition. It states that profits are not strictly positive at any quality level and that average profits are identically zero over any quality values at which a positive measure of college spots are supplied. Condition d is the goods market clearing condition. In addition to the variable cost e , each student attending college also consumes a fixed resource cost $\omega + \phi$. The college market clearing conditions are described in condition e. For each ability type and for each possible set of college qualities, the number of students who wish to attend a college in that quality set must equal the corresponding number of spots supplied. And furthermore, all the students who want to attend must deliver the maximum possible revenue for colleges, conditional on ability. The fact that there are many such market-clearing conditions reflects the club good setting.

A. Equilibrium Characterization

We start by observing that any equilibrium allocation can be supported by a tuition function that does not depend on income:

PROPOSITION 1: *Consider any equilibrium with an equilibrium tuition schedule $\tilde{t}(q, y, a, r)$. The same allocation can be supported by an alternative tuition schedule that is given by*

$$(6) \quad t(q, a, r) = v(q, a) - s(q, a, r),$$

where

$$v(q, a) = \max_{y,r} \{\tilde{t}(q, y, a, r) + s(q, a, r)\}.$$

Corollary. Subsidies to in-state students pass through one-for-one into lower tuition:

$$t(q, a, o) - t(q, a, i) = s(q, a, i) - s(q, a, o) \text{ for all } q, y, a.$$

The intuition for this proposition is as follows. First, if colleges at quality q are admitting any students of ability a , they will take those delivering the highest revenue, and by definition, those students will pay $t(q, a, r)$. Second, if $\tilde{t}(q, y, a, r) < t(q, a, r)$ for some idiosyncratic state (y, a, r) , then it must be the

case that such households do not want to attend colleges of quality q (otherwise, condition 5 would be violated). But then the same allocation can be supported when $\tilde{t}(q, y, a, r)$ is replaced by $t(q, a, r)$, because raising tuition only makes attending even less attractive. Henceforth, our measure of model tuition will refer to the function $t(q, a, r)$ as defined above.

Note that while equilibrium sticker tuition is independent of income, low-income students still benefit from need-based financial aid via subsidies that go to students, $p(y)$, and thus will pay lower net tuition, all else equal.

Is this prediction of no price discrimination by income counterfactual? In practice, a portion of institutional financial aid is labeled “need-based,” which might be interpreted as income-based price discrimination. However, many colleges promise to meet the “full demonstrated financial need” of admitted students but are “need-sensitive” at the admission stage.¹¹ Among applicants who will need significant aid, these schools presumably admit only the strongest. Thus, aid that they describe as “need-based” actually has a “merit-based” component. This pattern will also arise endogenously in the equilibrium of our model. Because higher ability students pay lower tuition, they tend to choose higher quality colleges for any given level of household income. It follows that at any given quality level, the student body will tend to be a mix of relatively rich but low ability students who pay relatively high tuition and relatively poor but higher ability students who enjoy tuition discounts. In the context of our model, these tuition discounts truly reflect ability, but they will be systematically negatively correlated with household income, thanks to the equilibrium pattern of sorting.¹²

PROPOSITION 2: *Any equilibrium can be supported by a revenue schedule that is linear in ability, that is, one that takes the form*

$$v(q, a) = b(q) - d(q)(a - a_{\min}),$$

where a_{\min} is the lowest value in the set A . The quality-dependent constant $b(q)$ defines revenue from the lowest ability type, while $d(q) > 0$ denotes the revenue discount per unit of ability.

Proposition 2 greatly simplifies equilibrium characterization. A linear revenue function ensures that the revenue from admitting any set of students is linear in the average ability of the students in that set. Thus, we can rewrite the college

¹¹Only about 100 US colleges and universities are need blind at the admissions stage, and only half of those promise to meet full demonstrated need conditional on admission. And even within that small subset, some of this “aid” takes the form of student loans.

¹²See Chetty et al. (2020) Online Appendix Table XII for evidence of the same sorting pattern by income and ability within “Ivy Plus” colleges.

problem as

$$(7) \quad \begin{aligned} \max_{\bar{a}, e} \{ & b(q) - d(q)(\bar{a} - a_{\min}) - e - \phi \} \\ & s.t. \\ & q = \bar{a}^\theta e^{1-\theta} \\ & a_{\min} \leq \bar{a} \leq a_{\max}, \end{aligned}$$

where \bar{a} is the average ability of students admitted.¹³ The first-order condition is

$$(8) \quad \frac{d(q)}{1} = \frac{\theta \frac{q}{\bar{a}}}{(1-\theta) \frac{q}{e}}.$$

The left-hand side is the ratio of the price of a marginal increase in average ability relative to the price of a marginal increase in e . The right-hand side is the corresponding ratio of marginal products. If the equilibrium revenue schedule declines steeply with ability ($d(q)$ is large), then colleges will choose a high ratio of instructional inputs relative to average student ability.¹⁴

We now turn to households' choices about college attendance and quality. Absent need-based financial aid, the optimal choice for quality is increasing in household income, holding fixed ability. This is because college quality is a normal good and because equilibrium tuition, as just discussed, is independent of income. This property simplifies equilibrium computation because it means that moving up the college quality distribution, college spots for each ability type will be filled in a strictly ordered fashion by income, with the highest quality colleges taking students from the top of the income distribution.¹⁵

B. Existence and Pareto Efficiency

We now discuss equilibrium existence and efficiency. For ease of exposition, in this section we focus on a version of the model that abstracts from government intervention.

Existence has been a long-standing issue in college models with peer group effects (Epple and Romano, 1998, 2010). In our environment, with a constant

¹³Note that \bar{a} is a continuous choice variable, even though a takes a discrete set of values, because each college enrolls a continuum of students and can vary average student ability smoothly by varying the enrollment shares η_a .

¹⁴One can further show, with the zero-profit condition, that the revenue function $v(\cdot)$ must satisfy effective marginal cost (EMC) pricing as in Epple and Romano (2008). That is, the revenue function must consist of three components: a fixed cost component; a variable expenditure component; and a component reflecting the student's peer group value:

$$v(q, a) = \underbrace{\phi}_{\text{Fixed Cost}} + \underbrace{e(q)}_{\text{Variable Cost}} + \underbrace{\frac{\theta}{(1-\theta)} \frac{e(q)}{\bar{a}(q)} (a - \bar{a}(q))}_{\text{Peer Value}},$$

where $e(q)$ and $\bar{a}(q)$ are solutions to the colleges' problem. This function is unique to admitted students.

¹⁵Need-based financial aid potentially breaks this result, depending on the schedule for $p(y)$.

returns technology and a continuum of agents, we are able to prove a formal existence result, one that does not rely on introducing lotteries.¹⁶

PROPOSITION 3: *A competitive equilibrium exists.*

We next prove that the competitive equilibrium is Pareto efficient. A crucial assumption underlying this result is that clubs are competitive price takers. This implies that the peer group effect is a *local* externality within a college and is correctly priced in the competitive market: higher-ability students are charged lower tuition. Thus, as argued by Ellickson et al. (1999), a club goods economy is conceptually no different from a non-club goods economy, in the sense that (type-specific) club memberships can be treated as ordinary goods traded in a competitive market.

PROPOSITION 4: *A competitive equilibrium, absent government subsidies, is Pareto efficient.*

The proof here closely follows the standard proof of the First Welfare Theorem. The result hinges on the assumption that households care about quality directly and derive utility only from their own child's college quality, so college is a pure consumption good. High-income households will spend a lot on tuition at expensive schools – even if their children are low ability – because their marginal utility from non-college consumption is low. Thus, they have a high willingness to pay for the experience and prestige associated with high quality colleges, or for gaining access to a more attractive pool of potential spouses and professional connections. Similarly, if poor households choose not to send their children to college, that will simply reflect the fact that those households prefer to spend their limited income on consumption goods.

In Section V, we discuss an alternative interpretation of the model, in which the value of college has an investment component. In this case, allocations are *not* efficient.

II. A Closed-Form Example

Before calibrating the model described above, we first show that in one special case, equilibrium allocations can be characterized in closed form. This example is useful because it clearly illustrates that the club good nature of the college market has important implications for college pricing and for the effects of changes in income inequality on the allocation of students to colleges and the tuition they pay. It also provides a proof of existence by construction in this special case.

¹⁶The proof of existence consists of two steps. In the first step, we establish that a competitive equilibrium exists when there are finitely many different types of college. This part of the proof draws heavily on Ellickson et al. (1999). The second step extends the existence result to a continuum of club types. To do so, we adopt the methodology from Caucutt (2001) and construct a sequence of approximate economies with finitely many college types, and show that the limit of that sequence is the equilibrium for the economy with a continuum of college types.

The special case is one in which $\theta = 1$, so average student ability is the only determinant of college quality. In addition, $\omega = \phi = 0$, so there are no fixed resource costs associated with attending college. The preference parameter φ is equal to one. There are no public subsidies of any form. There are two ability types, denoted by a_l and a_h , and each half the population is of one type. The income distribution is independent of ability and uniform: $y \sim \mathcal{U}(\mu_y - \frac{\Delta_y}{2}, \mu_y + \frac{\Delta_y}{2})$, where μ_y denotes average income and Δ_y defines income dispersion. Let $\Delta_a = a_h - a_l$ denote ability dispersion and $\mu_a = \frac{a_h + a_l}{2}$ denote average ability.

Note that given $\theta = 1$, the production function implies that $q = \bar{a}$. Note also that the support of possible college qualities is $[a_l, a_h]$.

PROPOSITION 5: *Under the parameterization described above, the model has a competitive equilibrium in which the measure of college spots by quality and tuition schedules are described, respectively, by*

$$\begin{aligned} \chi(Q) &= \frac{2}{\Delta_a} \frac{2}{(4 + \pi)} \int_Q \left[\left(\frac{a_h - q}{\Delta_a} \right)^2 + \left(\frac{q - a_l}{\Delta_a} \right)^2 \right]^{-2} dq \quad \forall Q \subset (a_l, a_h), \\ \chi(a_h) &= \chi(a_l) = \frac{2}{4 + \pi}, \\ (9) \quad t(q, a_i) &= \mu_y \frac{1}{\kappa + q} \left[1 - \frac{2}{4 + \pi} \frac{\Delta_y}{\mu_y} \arctan \left(2 \frac{(\mu_a - q)}{\Delta_a} \right) \right] (q - a_i), \end{aligned}$$

where π is the mathematical constant and \arctan is the inverse tangent function.

The equilibrium tuition function satisfies what Epple and Romano (1998) term “effective marginal cost pricing”:

$$t(q, a_i) = \sum_j \eta_j \frac{\partial t(q, a_j)}{\partial q} (q - a_i).$$

This optimality condition equates the tuition revenue from admitting a student of ability a_i to the marginal cost of admitting such a student via the peer group externality. In particular, admitting a lower than current average ability student lowers college quality proportionally to $q - a_i$ and thus lowers the tuition the college can charge to all students.¹⁷

The equilibrium allocations described in equation (9) have several interesting properties. First, the quality distribution of college spots is independent of the

¹⁷See Epple and Romano (1998) (p. 40) for more discussion. To see that this condition is satisfied, note that the equilibrium tuition function (eq. 9) can be written as $t(q, a_i) = G(q) (q - a_i)$ where

$$G(q) = \mu_y \frac{1}{\kappa + q} \left[1 - \frac{2}{4 + \pi} \frac{\Delta_y}{\mu_y} \arctan \left(2 \frac{(\mu_a - q)}{\Delta_a} \right) \right].$$

It can be shown that $\sum_j \eta_j \frac{\partial t(q, a_j)}{\partial q}$ is equal to $G(q)$.

income distribution parameters μ_y and Δ_y and is also independent of the preference parameter κ . Given perfect assortative matching from income to quality, this result implies that the equilibrium college quality choice for a household depends only on the household *rank* in the income distribution and their child's ability. The quality distribution of college spots has two mass points at the lowest and highest possible quality colleges, $q = a_l$ and $q = a_h$. In between these values, the distribution is continuous, symmetric, and single peaked. The lowest quality schools are filled with low-ability students drawn from the bottom of the income distribution, while the highest quality schools are filled with high-ability students drawn from the top. Outside these income ranges, students attend mixed-ability schools.

In contrast, the income distribution parameters μ_y and Δ_y do appear in the equilibrium tuition functions. Tuition is an increasing but nonlinear function of quality.¹⁸ The tuition ability discount $d(q) = \frac{t(q, a_l) - t(q, a_h)}{\Delta_a}$ is increasing (decreasing) in income dispersion Δ_y for $q \geq \mu_a$ ($q \leq \mu_a$).

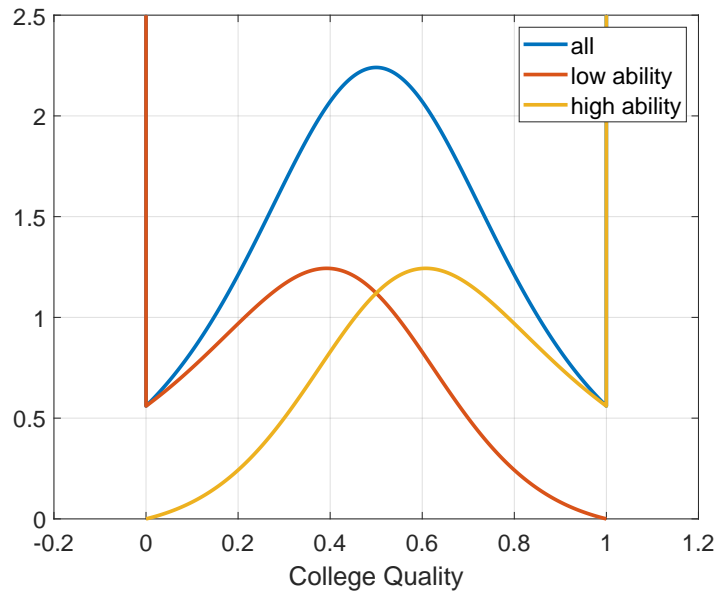


FIGURE 2. COLLEGE DISTRIBUTION $\chi(q)$

Figures 2 and 3 plot $\chi(q)$ and $t(q, a_i)$ for an example in which $a_l = 0$, $a_h = 1$, $\mu_y = 1$, and $\kappa = 5$. We consider two different values for income inequality: low inequality, with $\Delta_y = 0.1$; and high inequality, with $\Delta_y = 1.9$. The allocation of students to colleges, depicted in Figure 2, is exactly the same in both cases.

¹⁸It is easy to check that $t(q, a_l) \geq 0$ and $t(q, a_h) \leq 0$ for all $q \in [a_l, a_h]$.

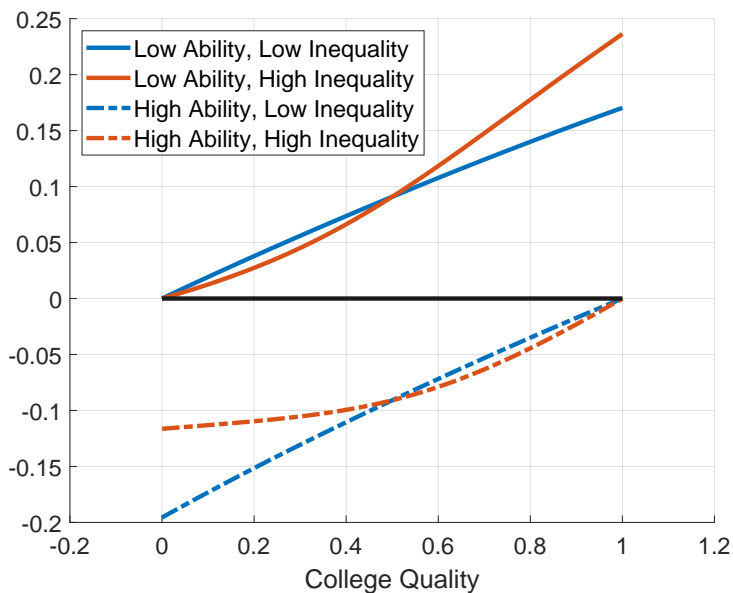


FIGURE 3. COLLEGE TUITION SCHEDULES

When income dispersion is increased, Figure 3 indicates that low-ability students pay higher tuition at high quality colleges, while high-ability students at the same colleges enjoy larger tuition subsidies (recall that average tuition, weighted by ability, is always equal to zero, thanks to free entry). One could interpret this pattern as an increase in sticker tuition, coupled with greater merit aid for high-ability students, such that the gap between sticker tuition and net tuition widens.

Why is the distribution for college quality χ completely insensitive to income inequality? A partial intuition is that the support of the college quality distribution cannot expand in our environment: the bounds are always $q = a_l$ and $q = a_h$. Clearly, in order for households at the extremes of the income distribution to choose feasible values for quality, tuition schedules must move when income inequality is increased. The fact that the richest low-ability households are now richer increases the relative demand from low-ability households for high quality colleges. In equilibrium, a rise in low-ability tuition at high quality colleges induces these richer households to leave their quality choices unchanged. Similarly, rising income inequality leaves the poorest high-ability students poorer. A rise in high-ability tuition at low quality colleges induces these households not to downgrade college quality.

Note that the implications of widening inequality in this club good setting are starkly opposite to those in a conventional non-club goods model. In a conven-

tional model, increasing income dispersion would lead the rich to increase consumption of college quality, while the poor would reduce quality, and the price per unit of quality would remain unchanged. In the club good model, in contrast, all the effects of widening income inequality show up in changes in equilibrium tuition, with no impact on the allocation of college quality. Thus, this simple example offers a powerful motivation for exploring the potential role of widening income inequality on college tuition in a richer calibrated model, to which we now turn.

III. Quantitative Application

Our baseline calibration is for 2016.¹⁹ We assume that the distribution for household income (conditional on child ability) is Pareto lognormal, a parametric functional form that closely approximates the actual distribution of income in the United States (see Heathcote and Tsujiyama 2019). Thus, log household income is given by $\log y = x_1 + x_2$, where x_1 and x_2 are independent random variables, x_1 is normally distributed with mean μ_y and variance σ^2 , and x_2 is exponentially distributed with exponential (Pareto index) parameter α . This distribution transitions smoothly from an approximately lognormal distribution over most of the income distribution toward a Pareto distribution in the right tail. We estimate the parameters μ_y , σ^2 , and α using microdata on log total household income from the 2016 Survey of Consumer Finances. One important strength of this survey is that households at the top of the income distribution are not underrepresented, which is important for being able to estimate the Pareto parameter α . Because we are interested in income for households that are making decisions about college, we restrict our sample to households between ages 40 and 59. The maximum likelihood estimates for σ^2 and α are 0.55 and 1.67, implying a variance of log income equal to $0.55 + 1.67^{-2} = 0.91$.²⁰

As in the closed-form example, our baseline calibration assumes two ability types.²¹ We assume that the income distributions conditional on ability are both Pareto lognormal, with the same estimated values for σ^2 and α . To allow for correlation between household income and child ability, we index the level parameter by ability, setting $\mu_{a_h} = \mu_y + \delta$ and $\mu_{a_l} = \mu_y - \delta$. To estimate δ , we turn to the 1997 NLSY and use AFQT scores as a proxy for child ability. We rank children by these scores and set δ to replicate the ratio of average family income for households with children in the top versus the bottom half of the AFQT score

¹⁹See the Data Appendix for more details on the construction of the statistics used in calibration and model-data comparisons.

²⁰These estimates are similar to those from other studies (e.g., Piketty and Saez, 2003). Measured income dispersion may actually understate true income dispersion, thanks to non-classical measurement error in survey data (Gottschalk and Huynh, 2010).

²¹We develop a specialized computational algorithm for the two-ability type case, which is much faster than our general algorithm that handles more ability types. In the online computational appendix, we show that introducing more ability types has a negligible impact on the distribution of college enrollment and tuition.

distribution, which is 1.59.²²

Hendricks, Koreshkova and Leukhina (2018) report graduation rates (i.e., one minus dropout rates) by AFQT score, conditional on enrollment in four year colleges, again using the 1997 NLSY. These rates are 78 percent for students in the top half of the AFQT distribution and 52 percent for students in the bottom half. Thus, we set the graduation rate probabilities to $\gamma_{ah} = 0.78$ and $\gamma_{al} = 0.52$.

We set the reservation utility parameter, κ , to target a model college enrollment rate of 50.7 percent. Given the graduation probabilities described above, this is the enrollment rate consistent with the observed mix of above and below median ability students graduating from four year colleges, as well as with an aggregate graduation rate of 36.1 percent, which is the share of individuals aged 25–29 who report at least a bachelor’s degree in the CPS.

The remaining model parameters are set to replicate a set of first moments involving tuition, financial aid, subsidies and costs for the universe of nonprofit four-year private and public colleges.²³ Data on tuition, aid, and costs are from the College Board and the National Center for Education Statistics (NCES). Many statistics are reported separately for public and private four-year colleges. Recall that we do not separately model the two sectors. Rather we focus on the distinction between public students who are in-state and all other students, taking the view that private students and public out-of-state students are effectively active in the same market. We assume that universities receive additional government subsidies in proportion to the number of in-state students that they admit and therefore offer those students discounted tuition. In 2017, 70 percent of students in four-year colleges were in public universities, and of those, 78 percent were in-state students. We therefore set μ_i , the share of in-state students in the population, to target a share of in-state students in the model of $0.7 \times 0.78 = 54.6$ percent. The implied μ_i , is 0.529. This is slightly lower than the target enrollment share because subsidized in-state students are more likely to choose to enroll.

The preference parameter φ is set to replicate average college tuition paid. We focus on net tuition, defined as sticker tuition minus government and institutional aid. This reflects the average amount students pay out of pocket and is therefore a good gauge of the strength of the preference for college. Average net tuition for the 2016–17 academic year was \$3,770 for public in-state students, \$14,190 for private university students, and \$19,050 for public out-of-state students.²⁴

²²We set μ_y to target a value for average income:

$$\bar{y} = 0.5 \left[\exp \left(\mu_y + \delta + \frac{\sigma^2}{2} \right) \frac{1}{\alpha} + \exp \left(\mu_y - \delta + \frac{\sigma^2}{2} \right) \frac{1}{\alpha} \right].$$

²³Around one-tenth of individuals with a bachelor’s degree are graduates from for-profit colleges, which could be conceptualized as low-quality private schools in the context of our model. However, because of data limitations, we exclude the for-profit sector when computing the statistics targeted in our calibration.

²⁴We observe average sticker tuition for public out-of-state students but not average net tuition. We estimate average net tuition for this group by assuming that the sticker versus net tuition differential for out-of-state students is the same as for public in-state students.

Given the respective enrollment shares for the three types of students, average net tuition was \$9,250.

We next translate this dollar value into an estimate of net tuition at four year colleges as a share of total consumption spending, which is an object we can readily compute in both model and data. In 2016, there were 8.34 million students enrolled in four-year schools out of a total US population of 323.4 million, implying an aggregate attendance rate of 2.58 percent.²⁵ In 2016, private consumption per capita was \$39,417. Thus, aggregate net tuition spending on four-year colleges was $2.58 \times \frac{9,249}{39,417} = 0.61$ percent of total consumption.²⁶ We set φ so that model average net tuition times the target enrollment rate (0.507) is equal to 0.61 percent of average model income.²⁷

Forgone earnings ω is an important additional component of the cost of attending college. We set ω equal to 20 weeks of median weekly earnings for full-time workers aged 16 to 24, which was \$10,020 in 2016 (Current Population Survey).²⁸ We exclude room and board from our measure of college prices, on the grounds that similar living expenses apply irrespective of college attendance.²⁹

We next describe the model for student financial aid, $p(y)$. We assume that this aid has a general component p_0 and a need-based component p_1 that is available to students with family income below a threshold y^* :

$$p(y) = p_0 + \mathbb{1}_{\{y \leq y^*\}} p_1.$$

The main form of federal need-based aid is Pell Grants. In 2016–17, 32 percent of undergraduate students received a Pell Grant, and the average grant size was \$3,800. A large share of state grant aid is also need-based, translating into \$982 per student. We assume the students receiving need-based state grants are the same as those receiving Pell Grants, implying $p_1 = 3,800 + 982/0.32 = \$6,870$. We then set the threshold y^* such that in equilibrium, 32.0 percent of enrolled students come from families with income below y^* , which requires y^* equal to

²⁵Given the observed 50.7 percent enrollment rate, and assuming the U.S. economy is in steady state, this translates into each college enrollee spending $0.0258/0.507 = 5.1$ percent of their lifetime in college, or 4.1 years, assuming an 80 year lifetime.

²⁶The share of college tuition and fees in the Consumer Price Index is larger at 1.8 percent. However, the CPI college category includes two-year schools and graduate and professional schools. Also, the CPI tuition measures are closer to capturing sticker price than actual price paid (e.g., Schwartz and Scafidi 2004).

²⁷Average model income is the value for \bar{y} that solves $\frac{0.507 \times 9,250}{\bar{y}} = 0.0258 \times \frac{9,250}{39,417}$, which implies $\bar{y} = \$774,590$. The reason this is a large number is that in our static model, the ratio of tuition to income should be thought of as average lifetime tuition to average lifetime income: a student with average household income attending college and paying average net tuition will spend $9,250/774,590 = 1.2\%$ of household income on tuition.

²⁸We chose 20 weeks, thinking that students can work full-time when school is out (typically 16 weeks a year), and part-time during the school year. The College Board (2016, p. 8) suggests using living expenses as a proxy for the cost of forgone earnings. The average cost of room and board at 4-year colleges in 2016–17 was \$10,875, which is similar to our earnings-based estimate.

²⁹Increased top tail income inequality will translate into greater demand for high quality accommodation and food and thus can potentially rationalize observed growth over time in the cost of room and board.

71.4% of mean model income.³⁰

We pin down the general component of the subsidy p_0 residually, exploiting the accounting identity that the sum of average government aid plus institutional aid is equal to the difference between average sticker tuition and average net tuition. Average sticker tuition in 2016–17 was \$19,152 (\$9,650 for public in-state students, \$33,480 for private students, and \$24,930 for public out-of-state students). Our estimate for average institutional aid per student is \$5,808 (see below). Thus, $p_0 = \$19,152 - \$9,250 - \$5,808 - 0.32 \times \$6,870 = \$1,896$.

We turn next to college subsidies. Our model for the form of these subsidies is

$$\begin{aligned} s(q, a, o) &= \bar{s} \\ s(q, a, i) &= \bar{s} + \max\{(1 - \lambda)t(q, a, o), 0\}. \end{aligned}$$

Thus, colleges receive a fixed per student subsidy \bar{s} for every student they enroll. In addition to various forms of government support, this captures any other source of non-tuition income, such as endowment income and private gifts. Colleges that admit in-state students receive an additional per student subsidy, which is a fraction $(1 - \lambda)$ of the equilibrium tuition that a similar out-of-state student would pay at a school of the same quality. Note that given eq. 6, this model implies that in equilibrium, in-state students pay a constant fraction λ of what identical out-of-state students pay: $t(q, a, i)/t(q, a, o) = \lambda$.³¹

Average sticker tuition for public out-of-state students in 2016–17 was \$24,930, while the corresponding number for in-state students was \$9,650, suggesting an average extra subsidy per in-state student of \$15,280 (or \$8,343 per student overall). We set λ to replicate this average per student subsidy across all in-state students, which implies $\lambda = 0.49$.

Next, we turn to the ratio between ability levels, a_h/a_l , and the share parameter, θ , which defines the relative importance of average student ability versus instructional spending in determining college quality. These parameters jointly determine the importance of the club good feature of the model. Increasing a_h/a_l or increasing θ both make high ability students more attractive for colleges, which translates into more institutional aid being used to recruit such students. Average institutional aid per student in 2015–16 was \$2,274 at public four-year institutions and \$14,055 at private institutions, for an overall average of \$5,808. We normalize $a_h = 1$ and set a_l so that the model replicates this value, which implies $a_l = 0.375$.³² Note that in the model, all low ability students pay full sticker tu-

³⁰Our model for need-based aid is simpler than the actual Pell grant program. The size of Pell grant a student can expect depends on the student's Expected Family Contribution (which depends on family income) and also on the cost of college attendance. However, there is an upper bound on the maximum possible grant, which was \$5,920 in 2016.

³¹From the College Scorecard data, we find that the empirical ratio of sticker tuition for in-state versus out-of-state students does not systematically vary with the *level* of tuition across public universities, consistent with this subsidy model.

³²The units of ability and income in the model are both arbitrary, and thus we are free to choose the units of each. Given parameter values and a particular choice of units, one can characterize a competitive

ition, while all institutional aid goes to high ability students. Because a_h/a_l and θ are not sharply separately identified, following much of the existing literature we simply set $\theta = 0.5$, implying that peer effects and expenditure are equally important in delivering quality (see, e.g., Caucutt 2002 and Epple, Romano and Sieg 2006).³³ We will also experiment with $\theta = 0.25$ and $\theta = 0.75$ in Section IV.D.

The only remaining parameters are the fixed cost for admitting students ϕ and the fixed per student subsidy \bar{s} . From the college problem, it is clear that colleges care only about the net fixed cost of enrolling students, $\phi - \bar{s}$. In equilibrium, colleges make zero profits. We can therefore estimate the net fixed cost using the following equation equating college sector income to expenses:

$$\underbrace{E[t(q, a, r)]}_{\$19,152-\$5,808} + \underbrace{E[s(q, a, r) - \bar{s}]}_{\$8,343} = \underbrace{E[e(q)]}_{\$17,077} + \underbrace{(\phi - \bar{s})}_{\$4,610}.$$

The first term on the left-hand side is average tuition income (average sticker tuition minus average institutional aid). We identify variable resource spending e with the sum of the NCES expenditure components “instruction” and “student services.” In 2016, instruction and student services spending per student averaged \$22,749 at private four-year schools and \$14,646 at public schools, for an overall average of \$17,077 per student. From the equation above, the implied net fixed cost is \$4,610. The calibration is summarized in Table 1.

The model is computed on a discretized grid of college quality. The upper bound q_{max} is set large enough that it does not restrict entry of colleges at the top. We use 400 grid points for quality, and we have checked that making the grid finer does not affect equilibrium prices or allocations. We set the maximum level of income to be 20 times the average, which corresponds to the 99.8th percentile of the SCF income distribution. We develop an iterative algorithm that exploits the assortative pattern of sorting by income to college quality, conditional on ability, residence status, and need-based aid eligibility. A detailed description of the algorithm is in the online appendix.

A. Model Predictions

To assess whether ours is a reasonable model of the college sector, we examine what it implies for some key non-targeted moments. In particular, does the model successfully replicate who enrolls in college, and how much tuition they pay?

equilibrium. It is straightforward to see how the same equilibrium must be rescaled under alternative choices for units. For example, if we multiply the units for the output / consumption good by a factor m , then the same equilibrium is preserved as long as the preference parameter κ is scaled by a factor $m^{1-\theta}$. Similarly, if we multiply the units for ability by a factor n , the same equilibrium is preserved as long as κ is scaled by a factor n^θ .

³³The standard deviation of test scores on standardized tests cannot be used to identify the ratio a_h/a_l since test score dispersion is a choice of the test designer, rather than an informative empirical moment that can be targeted.

TABLE 1—CALIBRATION

Income distribution		Aid	
μ_y	13.77	Residence-based	
δ	0.232	μ_i, μ_o	0.529, 0.471
σ^2	0.548	λ	0.490
α	1.670	Need-based	
Preferences		y^*/\bar{y}	0.714
φ	0.049	p_1	\$6,870
κ	36.89	General	
Technology		p_0	\$1,896
θ	0.50	$\phi - \bar{s}$	\$4,610
a_h, a_l	1, 0.375		
$\gamma_{a_h}, \gamma_{a_l}$	0.781, 0.520		
ω	\$10,020		

In terms of enrollment, we focus on the income and ability composition of students enrolling in college (we match the in-state versus out-of-state mix by construction). Chetty et al. (2020) report average parental family income at the college level, constructing these estimates from federal income tax records.³⁴ Average parental income for children in four-year universities is 56.0 percent larger than the average for all children. The corresponding statistic for the model is very similar at 56.7 percent, implying that the model delivers realistic sorting into college by family income. Using NLSY 97 data we computed the fraction of college graduates for each quintile of the parental income distribution (see Table 2). Our model is able to match the empirical pattern quite well, except at the very bottom, where the model predicts zero enrollment.³⁵

Again using the NLSY 97, we estimate that 74.9 percent of above median ability children enroll in college, compared with 26.5 percent of below median ability children. The corresponding enrollment rates in the model are 80.2 percent and 21.2 percent, indicating that the model does a good job replicating the pattern of college enrollment by ability.³⁶

Next, we compare the model's implications for the cross-sectional distributions of both sticker and net tuition. Our calibration procedure ensures that the model

³⁴Parental income is averaged over 1996–2000, a five-year period in which the child (and potential college attendee) is aged 15–19.

³⁵Note that the model predicts that more children enroll and graduate from college at the 2nd family income quintile than at the 3rd. This is due to our Pell grant specification, which implies that low ability students right below the income eligibility threshold attend college while those just above it do not. See figure 7, panel C.

³⁶Recall that our calibration targets the aggregate enrollment rate, and that graduation rates are higher for high ability students. Because the model slightly over-predicts (under-predicts) high ability (low ability) enrollment, it delivers a slightly higher graduation rate than the data: 36.9 percent versus 36.1 percent,

matches the empirical values for average sticker and net tuition, but the pattern of dispersion in tuition across schools of different qualities is endogenous and untargeted. Data on tuition at the college level are available from the College Scorecard, which is a tool provided by the US Department of Education to facilitate comparison shopping across colleges.³⁷

Figure 4 plots the distributions of net and sticker tuition, model against data. Overall, the distributions look broadly similar. One discrepancy is at the very top of the distribution, where published tuition and fees at elite private universities top out at around \$55,000 in the data, while the model predicts a right tail of even more expensive colleges.³⁸ Table 2 indicates that the coefficients of variation for net and sticker tuition across all four-year colleges in the United States are 0.99 and 0.77. The analogous statistics from the model are 1.31 and 0.80.

The model also generates a fairly realistic joint distribution across colleges, in terms of other observable college characteristics: average household income of students attending, and the share of high-ability students enrolled.³⁹ These moments are summarized in Table 2. Note, in particular, that the model replicates the positive correlations observed in the data between sticker tuition, net tuition, family income and ability: higher tuition colleges tend to enroll students from more affluent backgrounds and higher-ability students. The model slightly overpredicts dispersion in average family income.⁴⁰

Overall, our relatively simple calibrated model successfully replicates some important features of the US college market: the number of people who attend, attendees' average ability and family income, and observed variation in college tuition. We now explore and attempt to better interpret the equilibrium patterns of college enrollment and the nature of equilibrium college pricing.

Figure 5 plots the cumulative distribution function (CDF) for student enrollment under the baseline calibration, conditional on enrolling. The different colors within the distribution reflect different types of students, where each type faces


³⁷The College Scorecard reports sticker tuition and fees and the average net price of attendance for undergraduates receiving Title IV aid, which is the full cost of attendance less federal, state, and institutional grant aid. This measure includes living expenses. To construct a net tuition measure that excludes living expenses, we estimate living expenses at the college level as the difference between the average annual full cost of attendance and tuition and fees.

³⁸Note, however, that wealthy families at elite private colleges are informally expected to make voluntary donations on top of paying published tuition. Thus, published tuition arguably understates the true cost of attendance for these families. Applications from the wealthiest families are typically known as "development cases" and are handled separately from regular admissions, with enrollment offers more or less explicitly tied to actual or expected donations (see Golden 2006). For such students, sticker price may vastly understate the true cost of admission.

³⁹We have no information at the college level about AFQT scores, which we used to measure the share of high ability students enrolled in college. The College Scorecard does report various percentiles of SAT and ACT scores for admitted students. We assume these scores are normally distributed with college-specific means and variances, estimate the fraction of students whose score lies in the top half of the national test score distribution, and identify this fraction as the share of the college student body that is high ability. Another measurement issue is that in the College Scorecard data, 26% colleges do not report information on test scores.

⁴⁰The model also underpredicts dispersion in the fraction of high ability students. However, as noted above in footnote 39, our empirical measure for the share of high ability students is likely quite noisy, which may account for this discrepancy.

(a) Sticker Tuition Distribution (2016, \$1,000)



(b) Net Tuition Distribution (2016, \$1,000)




FIGURE 4. STICKER AND NET TUITION DISTRIBUTIONS

TABLE 2—NON-TARGETED MOMENTS COMPARISON

	Data	Model
<i>Enrollment Patterns</i>		
Family income enrolled / Mean	1.560	1.567
Share of high ability enrolled	0.749	0.802
Share of low ability enrolled	0.265	0.212
Graduation rate	0.361	0.369
Graduation rate by parental income quintile		
5th (top)	0.637	0.673
4nd	0.471	0.459
3rd	0.341	0.496
2th	0.214	0.216
1st (bottom)	0.141	0.000
<i>College-level Moments</i>		
Standard Deviation / Mean		
Net tuition	0.99	1.31
Sticker tuition	0.77	0.80
Avg. family income	0.51	0.92
Fraction of high ability	0.26	0.10
Correlation		
Sticker tuition vs. Net tuition	0.83	0.98
Net tuition vs. Family income	0.60	0.97
Net tuition vs. Fraction of high ability	0.22	0.71
Family income vs. Fraction of high ability	0.59	0.77

a type-specific net tuition schedule. There are eight different student types in the model, corresponding to each possible combination of residence status (determining eligibility for in-state tuition discounts), family income relative to the threshold y^* (determining eligibility for Pell Grants and other need-based aid), and ability (determining eligibility for institutional aid).

Several features of the plot are worth noting. First, no students attend very low-quality colleges: demand for such colleges is weak, given the lost-earnings cost ω of attending college in combination with a positive reservation quality κ when not attending. Second, need-based aid eligible students are clustered in relatively low quality colleges (the corresponding conditional CDF's flatten out around normalized quality equal to 3). This suggests strong sorting by income in college enrollment, in line with Belley and Lochner (2007) and Chetty et al. (2020).

To better understand quality choices by student type, Figure 6 plots net tuition by type as a function of quality, while Figure 7 plots the corresponding college quality decisions. Tuition is increasing in quality, in a non-linear fashion. It is lowest for high ability in-state students (who receive both in-state subsidies and

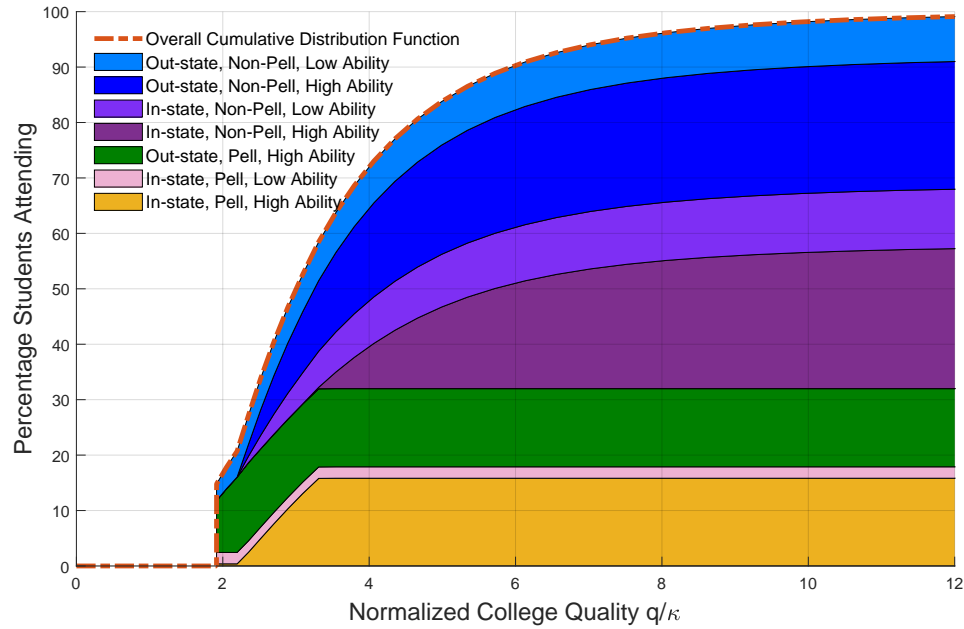


FIGURE 5. CUMULATIVE QUALITY DISTRIBUTION AND DECOMPOSITION

merit-based institutional aid) and highest for low-ability out-of-state students (who benefit from neither form of tuition discount). Net tuition for those eligible for need-based aid is simply reduced by the amount of that aid, $p_1 = \$6,869$. Note also that tuition for high-ability students rises more slowly with quality than tuition for low-ability students.⁴¹

Each panel of Figure 7 plots college enrollment choices for households of a different combination of ability and residence status. All high-ability students with family income above the need based aid eligibility threshold y^* enroll in college, as do some poorer households. For low ability students, who face higher tuition, quality choices are non-zero only at higher income values, indicating that a smaller fraction of this group attends college. Note that some low ability in-state students with family income just below y^* enroll (the disconnected blue line segment), even though similar students from families with income just above y^* do not.

Conditional on enrolling, quality choices are generally increasing in income for

⁴¹Note that net tuition is negative for some household types at low-quality colleges, reflecting the fact that the combination of general, need-based, and merit-based aid can exceed sticker tuition. However, the total cost of attending college, inclusive of the \$10,020 opportunity cost of work, is positive for all types.

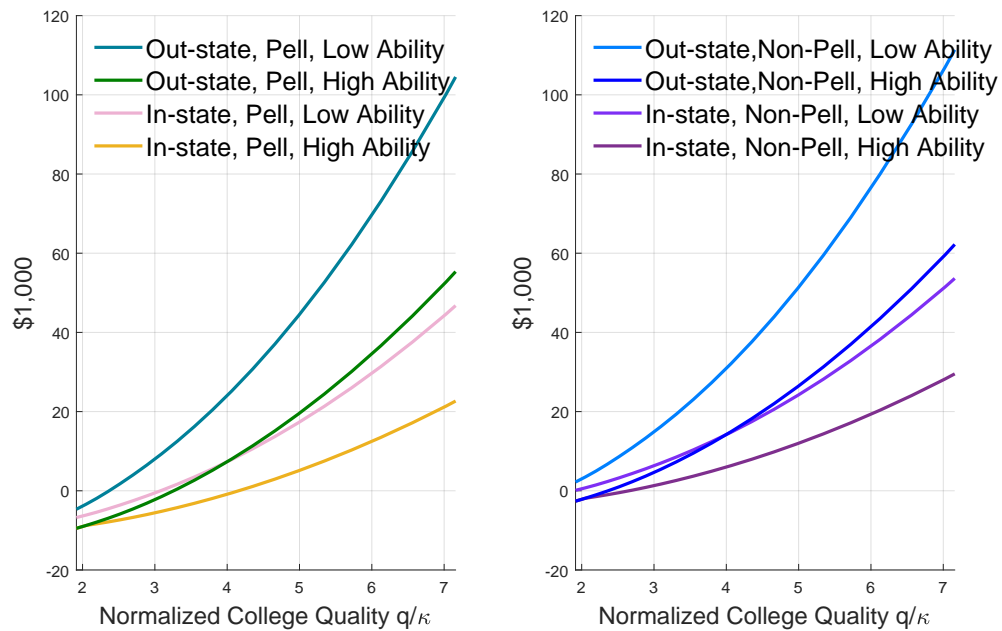


FIGURE 6. NET TUITION SCHEDULES FOR DIFFERENT TYPES

a given household type.⁴² In addition, because high-ability students face flatter tuition schedules, they choose higher quality colleges than otherwise identical low-ability students.

Figure 7 highlights that the effects on college enrollment and quality choices of different forms of tuition discount are quite different. Need-based aid incentivizes some low-income students to enroll in college, especially those with high ability. But because the amount of need-based aid is fixed, this form of subsidy does not incentivize choosing a higher quality college, conditional on enrolling: the quality choices of students just below and just above y^* in Figure 7 are similar. Thus, in Figure 5 students receiving government need-based aid tend to be clustered in relatively low quality colleges.⁴³

In contrast, in-state-resident subsidies in the model do little to encourage enrollment, because at low quality colleges, tuition is already low and proportional in-state discounts are therefore small. Thus, the share of in-state students en-

⁴²Note that a student from a household with income slightly below y^* who is attending college is effectively slightly richer than a household slightly above y^* , since the former receives need-based aid p_1 . As a result, the former chooses a slightly higher quality college.

⁴³This is also a qualitative feature of the data: 64.2 percent of students at for-profit colleges are Pell Grant recipients (Kelly, Holian and Archer, 2019), compared with around 15 percent at Ivy League universities (see <https://www.romanhighered.com/data-viz/2018/1/2/pell-grant-recipients-in-the-ivy-league>).

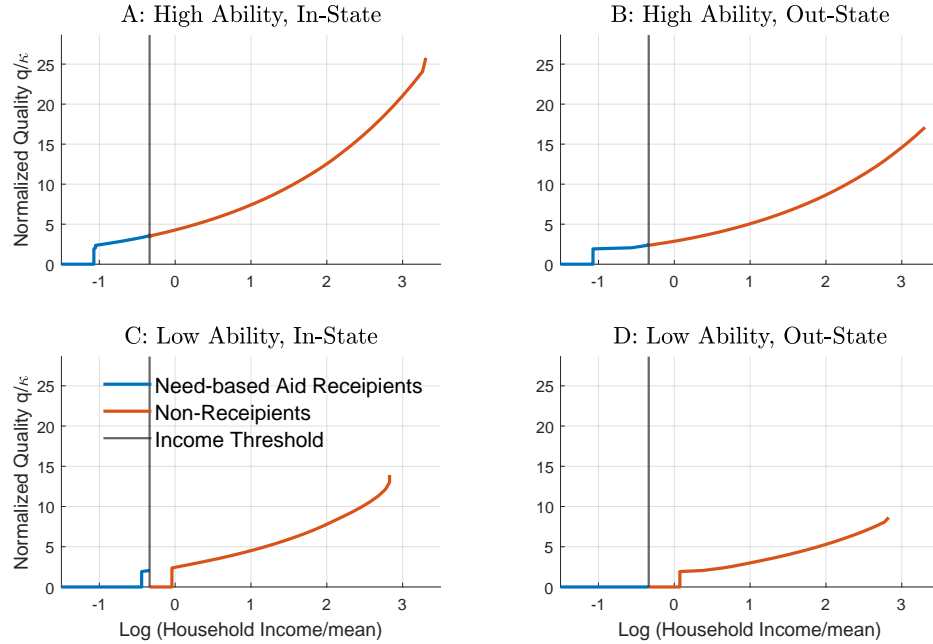


FIGURE 7. STUDENT ALLOCATION ACROSS COLLEGES

rolling in college is similar to their population share: 54.8 percent versus 53.2 percent. However, because in-state discounts are proportional to tuition, the marginal cost of attending higher quality colleges is lower for in-state students, which incentivizes them to choose higher quality colleges (see Figure 7).

As noted above, the endogenous tuition discounts that high ability students enjoy (institutional aid) incentivize both higher enrollment and higher quality choices, because high ability students face both lower and flatter tuition schedules.⁴⁴

We now turn to the supply side of the market. Figure 8 shows how college inputs vary by quality. Panels A and B plot the level of expenditure per student by quality and the fraction of high-ability students enrolled. Higher quality schools spend more per student and also generally admit a student pool with higher average ability. Panel C indicates that the tuition discount per unit of ability $d(q)$ (equivalently, the equilibrium differential between tuition for high- and low-ability students) is increasing in college quality. Thus, as quality increases, the effective price of the ability input rises, inducing colleges to increase the ratio of expenditure to average student ability (Panel D; see also eq. 8). Because the

⁴⁴Higher enrollment also reflects the fact that high-ability students tend to come from more affluent families.

ratio e/\bar{a} increases in quality, the ratio of expenditure to quality is also increasing in quality (from eq. 2, this ratio is $(e/\bar{a})^{1-\theta}$). This convexity in the equilibrium expenditure level (Panel A) translates into convexity in the equilibrium tuition schedules in Figure 6. Intuitively, because high ability students command ever larger tuition discounts moving up the quality distribution, colleges choose instead to spend ever larger amounts on instructional inputs and pass those expenses on in the form of higher tuition.⁴⁵

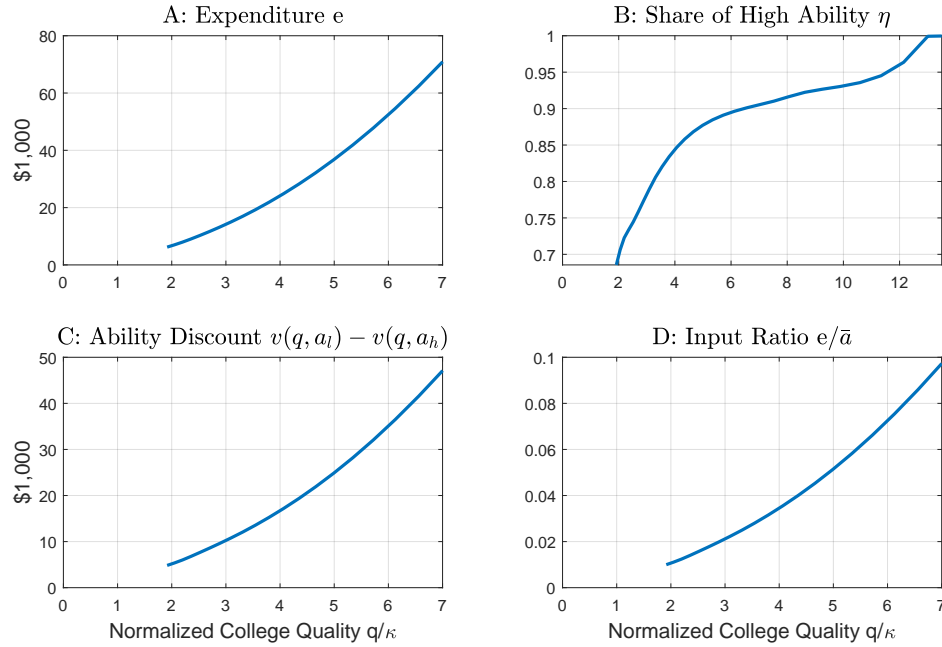


FIGURE 8. COLLEGE INPUTS

IV. Understanding Changes in College Tuition

We estimate average net tuition in 1990 to be \$6,034, compared with \$9,250 in 2016. The college graduation rate, computed as the share 25-29 year-olds with a college degree, increased from 23.2% in 1990 to 36.1% in 2016.⁴⁶

⁴⁵Note that while there is a unique optimal value for average student ability at each quality level, colleges are indifferent about admitting any mix of students by residence status. Thus, some colleges at a given quality level might admit only out-of-state students (those could be interpreted as private colleges), while the rest (public colleges) could admit a mix.

⁴⁶An alternative way to measure the graduation rate would be as the fraction of high school graduates attaining a college degree. This rose from 27.4% in 1990 to 39.1% in 2016.

We now explore possible drivers of this increase in college tuition. Our focus is on the role of widening income inequality. We re-estimate the income distribution parameters σ^2 and α using household income data from the 1989 wave of the Survey of Consumer Finances. We find that $\sigma^2 = 0.48$ and $\alpha = 2.4$. Thus, the implied variance of log income in 1989 is $0.48 + 2.4^{-2} = 0.65$, compared with 0.91 in 2016. Note that most of the increase in the variance over this period is attributable to a heavier right tail in the income distribution.⁴⁷ Figure 9 plots the estimated exponentially modified Gaussian (EMG) distribution for log household income in the two years. Note that while mean *log* income is very similar in both years, mean *level* income is 24.7 percent larger in 2016.

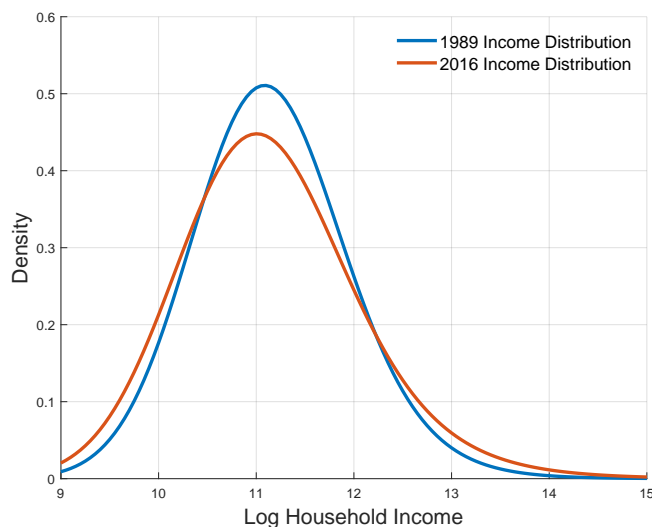


FIGURE 9. ESTIMATED EMG DISTRIBUTION OF LOG HOUSEHOLD INCOME (SCF)

Another set of changes we can measure and explore are those connected to changes in public support for higher education. Assuming no change in the fixed cost of college education per student ϕ , we can measure the change in total public support per student between 1990 and 2016 as the change in average educational expenditure per student minus the change in average net tuition.⁴⁸ We estimate

⁴⁷Data from the online appendix to Piketty and Saez (2003) indicate a similar increase in the Pareto parameter. For example, one can estimate the Pareto parameter as average income conditional on being above the x^{th} percentile of the income distribution, relative to this average minus income at the x^{th} percentile. Applying this formula to the Piketty-Saez data at the 90th percentile of the income distribution implies Pareto coefficients of 2.20 in 1989 and 1.83 in 2014.

⁴⁸While the difference between college spending and net tuition paid primarily reflects various forms of government subsidy, it also captures other sources of non-tuition revenue such as endowment income or alumni giving. Thus, our “public support” label includes these sources of income, in addition to government subsidies.

TABLE 3—SUMMARY CHANGES IN COLLEGE MARKET, 2016 VERSUS 1990

	2016 Data	1990 Data	% Growth
Net tuition	\$9,250	\$6,034	53.3
Expenditure per student e	\$17,077	\$10,503	62.6
Total subsidies per student net of ϕ	\$7,828	\$4,469	75.2
Need-based aid	\$2,198	\$1,377	59.6
In-state subsidies	\$8,343	\$5,413	54.1
General subsidies to colleges net of ϕ	-\$4,609	-\$2,396	-16.9
General subsidies to students	\$1,896	\$76	
Enrollment	0.507	0.327	+18.0pp
Share in-state	0.546	0.581	-3.5pp
Share Pell	0.32	0.30	+2.0pp
Graduation	0.361	0.233	+12.8pp

average educational spending in 1990 to be \$10,503, compared with \$17,077 in 2016.

The fact that spending has increased much more than out-of-pocket tuition indicates growth of \$3,359 in per student public subsidies. In Table 3, we decompose this \$3,359 growth in public support into (1) growth in public need-based aid, (2) growth in public subsidies to in-state students, and (3) growth in general per student public subsidies. We estimate the values of these three components, on a per student basis, to be \$821, \$2,930, and -\$393, respectively. Thus, need-based aid has risen, the value of in-state subsidies has risen quite markedly, and the value of general subsidies has declined slightly. The composition of these general subsidies has also changed, according to our estimates. In particular, general subsidies to colleges, \bar{s} , have *declined* by \$2,213, while general subsidies to students, p_0 , have *risen* by \$1,820. In the context of our competitive model, whether colleges or students receive per student subsidies is irrelevant for the equilibrium allocation and for net tuition, but a shift in subsidies towards students implies a commensurate increase in sticker tuition.⁴⁹

A. *The Effect of Rising Income Inequality*

We first explore the implications of changes in income inequality for college enrollment and tuition. Column (2) of Table 4 reports how this change in inequality changes some key predictions of the model. In order to isolate the impact of changing income dispersion, the mean parameter μ is re-scaled so that average household income \bar{y} for the economy in column (2) is identical to the one in column (1).

⁴⁹We do not experiment with changing the opportunity cost of attending college ω , since inflation-adjusted median earnings for workers between the ages of 16 and 24 barely changed between 1990 and 2016.

TABLE 4—EFFECTS OF CHANGES IN INCOME INEQUALITY AND SUBSIDIES

	(1)	(2)	(3)	(4)	(5)
	2016 Model	1989 Ineq.	1989 Mean	1989 Dist.	Dist.+ Subs.
<i>Parameters changed</i>	–	σ^2, α	\bar{y}	$\bar{y}, \sigma^2, \alpha$	all in Table 5
Net tuition	\$9,250	\$7,359	\$7,746	\$5,921	\$7,476
Expenditure	\$17,077	\$13,904	\$13,982	\$10,944	\$12,389
Enrollment	0.507	0.562	0.428	0.477	0.447
Family inc. enrolled / mean	1.567	1.407	1.747	1.543	1.625
Share high ability enrolled	0.802	0.889	0.714	0.807	0.746
Share low ability enrolled	0.212	0.235	0.143	0.147	0.149
Quality / κ	3.724	3.604	3.569	3.448	3.594

TABLE 5—CHANGES IN INCOME DISTRIBUTION AND SUBSIDY PARAMETERS

2016 Income distribution		1989 Income distribution	
\bar{y}	\$774,590	\bar{y}	\$621,221
σ^2	0.548	σ^2	0.478
α	1.67	α	2.40
2016 Subsidies		1990 Subsidies	
μ_i	0.529	μ_i	0.578
λ	0.490	λ	0.536
y^*/\bar{y}	0.714	y^*/\bar{y}	0.884
p_1	\$6,870	p_1	\$4,590
p_0	\$1,896	p_0	\$76
$\phi - \bar{s}$	\$4,609	$\phi - \bar{s}$	\$2,396

Note: \bar{y} denotes average lifetime income. See footnotes 25 and 27 for details.

The table indicates that the observed change in income inequality over this period predicts large changes in net tuition and enrollment. Given income inequality in 1989, average net tuition in the model is \$7,359, suggesting that changing income inequality alone can account for over half of the observed rise in net tuition over this period (\$1,891 out of \$3,216, or 58.8%). Average net tuition rises because households at the top of the income distribution became much richer and thus more willing to pay for expensive high quality colleges. At the same time, the counterpart to fast income growth at the top of the distribution was declining relative income for households in the middle of the distribution. These households were close to indifferent about attending college in 1989, and these income losses therefore drive down college enrollment. Thus, the model predicts that rising income inequality (all else held equal) reduced the enrollment rate by 5.5 percentage points.

Figure 10 illustrates how changing income inequality affects the equilibrium quality distribution (Panel A), the share of high ability students by quality (Panel

C), and equilibrium tuition schedule (Panels B and D). Increasing income inequality reduces the equilibrium supply of low quality schools while increasing slightly the number of spots at very high quality schools. This helps to explain why average net tuition rises. Across most of the quality distribution, the share of high-ability students in college declines, indicating that colleges rely more on instructional spending and less on peer effects to maintain quality.

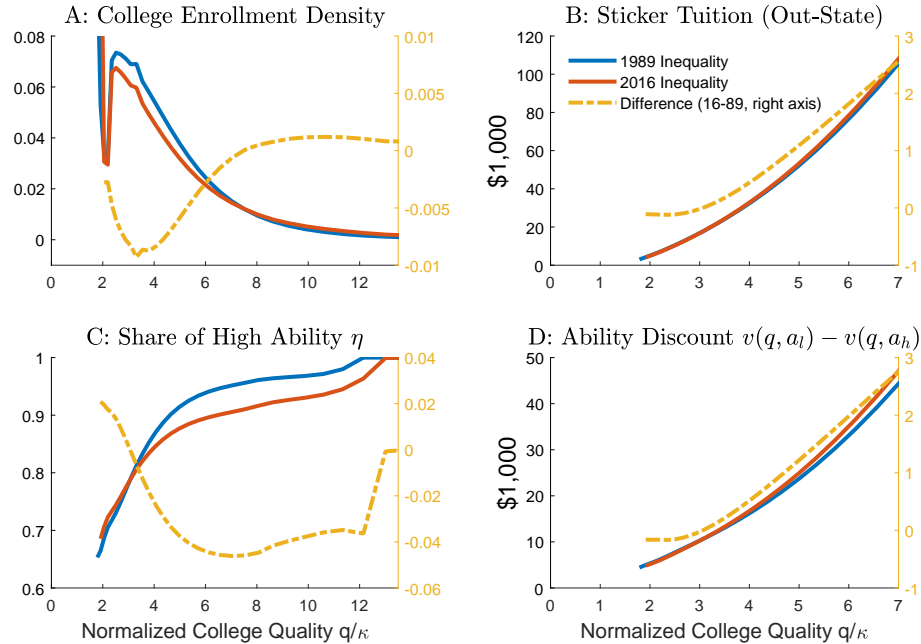


FIGURE 10. EFFECT OF CHANGING INCOME INEQUALITY

One might suspect that rising income inequality would lead to college attendance's being driven more by income and less by ability. Indeed, when inequality is increased, the share of high ability students enrolling in college falls by 8.7 percentage points, as poorer high ability students are effectively priced out. At the same time, however, increasing inequality also increases the *demand* for high-ability students. In particular, because the quality production technology features decreasing returns to expenditure, colleges seeking to satisfy increased demand for high quality want to increase both expenditures and average ability, all else equal. Thus, increased demand for high quality colleges indirectly increases the relative demand for high ability students, which translates into larger institutional tuition discounts for high-ability students (see Panel D of Figure 10). This suggests that rising income inequality has played a role in generating the observed growth in institutional financial aid and the associated rising gap between sticker and net

tuition (see Figure 1). Note that the increasing price of attracting high ability students (Panel D of Figure 10) explains why colleges at each quality level rely more on instructional expenditure and less on a high ability student body when inequality rises (Panel C).

To recap, when inequality rises both net tuition and expenditure rise as colleges respond to increased demand from households at the top of the income distribution for high quality. Moving from column (2) to column (1), net tuition grows by 25.7 percent and instructional expenditure per student rises 22.8 percent, but this extra spending leads to only a 3.3 percent increase in average quality. The increase in average quality is small for two reasons.

The first is the presence of peer effects: expenditure is only one quality-enhancing input, and the other — average peer ability — does not increase in response to rising income inequality. If quality were equal to expenditure, as in a conventional non-club good model, then average quality would obviously rise by exactly as much as expenditure per student.

The second reason the increase in average quality is small is that the increase in expenditure is concentrated in colleges at the top of the quality distribution, where expenditure is already high and the marginal product of additional spending is low. Given equal factor shares ($\theta = 0.5$), our production technology implies that average college quality is given by

$$E[q] = E[\bar{a}(q)^{0.5}e(q)^{0.5}] = \sqrt{E[\bar{a}] - \text{var}(\sqrt{\bar{a}})}\sqrt{E[e] - \text{var}(\sqrt{e})} + \text{cov}(\sqrt{\bar{a}}, \sqrt{e}).$$

When income inequality is increased, the boost to average quality from higher average expenditure is partially offset by a rise in the variance of (the square root of) expenditure, coupled with a fall in the correlation between average ability and expenditure. If we were to observe the same changes in average ability and expenditure in a model with only one type of college (so that $E[q] = \sqrt{E[\bar{a}]} \sqrt{E[e]}$), then average quality would increase by 10.8 percent.

We conclude that in order to properly quantify the rise in average college quality associated with a rise in income inequality, it is important both to model peer effects and to allow for heterogeneity in college quality: abstracting from either feature leads one to overstate the increase in average quality and thus to understate growth in quality-adjusted tuition.

B. The Roles of Other Drivers of Tuition Growth

Column (3) of Table 4 shows the effect of reducing average income (via a reduction in the parameter μ) to its value in 1989.⁵⁰ Reducing average income reduces households' willingness to pay for college and thus lowers both the enrollment

⁵⁰In this experiment, we also reduce the need-based aid threshold y^* proportionately with mean income. No other parameters are changed relative to the economy in column (1).

rate and net tuition. Thus, the model points to economic growth as one factor driving up college tuition.

Column (4) shows the combined effect of reducing both income inequality and mean income to their values in 1989. In this case, the model implies average net tuition of \$5,921, which is very similar to the actual value in 1990. The predicted enrollment rate given the income distribution in 1989 is lower than in 2016, indicating that the positive effect on enrollment of growth in average income outweighs the negative effect of widening inequality.

Column (5) of Table 4 shows how key model predictions change when, in addition to moving back to the income distribution in 1989, we also roll back subsidies to their values in 1990. In particular, we change the six subsidy parameters to the values described in Table 5. The values in 1990 for these parameters are chosen to match the four subsidy values in 1990 in Table 3 as well as the in-state and Pell-eligible enrollment shares for 1990 listed there.⁵¹

Lower subsidies per student in 1990 translate into higher tuition and a lower enrollment rate relative to the economy with subsidies at the 2016 values (compare column 5 with column 4). Thus, we find that growth in public subsidies (and other non-tuition revenue) has played a role in supporting enrollment and restraining growth in tuition.

We have also experimented with changing one aspect of the technology for producing college quality—namely the price, relative to consumption, of the variable instructional input, e . By allowing this price to rise over time, we can assess whether rising faculty salaries play a significant role in accounting for rising tuition (see, e.g., Jones and Yang 2016).⁵² We estimate that this price rose by 9.4 percent in real terms over the 1990–2016 period. We find that reducing the price of the instructional input to its 1990 value reduces average net tuition by only \$348. The change in average tuition is small because households respond to changes in instructional costs primarily by adjusting their chosen college quality rather than by changing the fraction of income they devote to college: reducing the price of instruction to its 1990 value raises enrollment by 1.5 percentage points and increases average college quality by 3.61%.

Overall, these experiments suggest that rising income inequality plays the most important role in accounting for observed growth in college tuition, while rising average income is also an important factor. Rising subsidies (and non-tuition revenue more generally) have a modest offsetting impact in the opposite direction.

⁵¹Need-based aid in 1990 is straightforward to estimate, given 1990 data on average Pell Grants and need-based state grants. Conditional on receiving aid, need-based aid has risen from \$4,589 to \$6,870, or from \$1,376 to \$2,198 on an unconditional per student basis.

Growth in subsidies to in-state students is more difficult to estimate. In 2016, we measured the in-state subsidy parameter λ , using the ratio of in-state to out-of-state sticker tuition for public colleges. We were unable to find direct estimates of out-of-state public college tuition in 1990. We therefore impute a value for out-of-state public sticker tuition in 1990 by multiplying private sticker tuition in 1990 by the 2016 ratio of public out-of-state sticker tuition to private sticker tuition.

⁵²Formally, we assume that z_t units of the consumption good y are required to produce one unit of the instructional input e . Given a competitive market for e , the equilibrium price of e at date t is then $p_t = z_t$. We normalize $z_{2016} = 1$.

Rising costs of instructional inputs play a marginal role.

C. *Changes in Enrollment Patterns*

Comparing across all the different experiments, the only parameter change that has a notable impact on enrollment of low ability students is growth in average income: the enrollment share of low-ability students is 6.9 percentage points larger in the baseline calibration for 2016 compared with when average income is reduced to its value in 1989 (column 3). For high-ability students, in contrast, all the different factors considered affect enrollment. Growth in average income raises high-ability enrollment by 8.8 percentage points (column 1 versus 3), increasing inequality reduces high-ability enrollment by 8.7 to 9.3 percentage points (column 1 versus 2, or 3 versus 4), and increasing subsidies increase enrollment by 6.1 percentage points (column 4 versus 5). Thus, while rising income inequality has reduced enrollment by high-ability students from poor families, the effect has been partially offset by growth in need-based financial aid.

Overall, the combined effects of changes in the income distribution and in subsidies are to increase the share of high-ability students enrolling in college by 5.6 percentage points and to increase the share of low-ability students enrolling by 6.3 percentage points (comparing columns 1 and 5). All of these extra low-ability students are drawn from the top half of the model family income distribution.

Belley and Lochner (2007) describe patterns of college attendance by ability (AFQT score) and family income for college-age individuals in the 1979 and 1997 NLSY.⁵³ They find that almost half of the increase in college attendance between the two waves was from students in the bottom half of the AFQT score distribution. Furthermore, these additional low-ability college students were drawn mainly from the top half of the income distribution. Thus, both model and data are consistent with the message that income has become a more important driver of college attendance, relative to student ability.

D. *The Importance of Peer Effects*

How important are peer effects for our results? To address this question we have experimented with alternative values for θ , which governs the relative importance of peers versus expenditure in producing college quality. We experimented with setting $\theta = 0.25$ and $\theta = 0.75$. In each case, we recalibrated all model parameters, following our baseline calibration strategy.⁵⁴

When peer effects are less important ($\theta = 0.25$), we find that income becomes a more important driver of college enrollment: the average family income of those enrolling in college is now 2.02 times the average, compared with 1.57 in

⁵³Note that our quantitative exercise focuses on a comparison between 1990 and 2016, while the two NLSY waves offer a comparison of the early 1980s with the early 2000s.

⁵⁴We hold the values for a_h and a_l fixed at their baseline values across these experiments, because changing those values also influences the importance of peer effects.

TABLE 6—THE ROLE OF PEER EFFECTS

	$\theta = 0.25$	$\theta = 0.5$	$\theta = 0.75$
Enrollment Pattern			
Family income enrolled / mean	2.020	1.567	1.498
Share of high ability enrolled	0.731	0.802	0.876
Share of low ability enrolled	0.283	0.212	0.139
Impact of Rising Inequality (1989 to 2016)			
Enrollment rate (change, percentage points)	- 6.30	- 5.46	- 3.74
Net tuition (change, \$)	+ 453	+1,891	+2,210

the baseline model (and 1.56 in the data). The intuition is simple: when quality depends mostly on resources, high-ability students enjoy smaller tuition discounts, and income becomes more important, relative to ability, in driving the enrollment pattern. When peer effects are more important ($\theta = 0.75$), in contrast, the model predicts that ability becomes more important in driving college enrollment: 88 percent of high ability students now enroll in college. The fact that the model generates a pattern of college enrollment by income and ability similar to the one observed in the data when $\theta = 0.5$ offers indirect support for our baseline parameter choice.

We also revisited our key experiment of changing income inequality for these alternative parameterizations. What we find is that with a larger value for θ , the model delivers a larger increase in average net tuition in response to an increase in income inequality, and a smaller decline in aggregate enrollment. These results are summarized in Table 6. The finding that as peer effects become more important, changes in inequality have a smaller impact on equilibrium quantities and a larger impact on prices is qualitatively consistent with our earlier analytical result from the simpler model described in Section II, where we showed that in the limiting case in which quality is a pure peer effect, changing inequality has zero impact on the allocation of quality and affects only equilibrium tuition.

V. Education as Investment

In this section, we briefly discuss an alternative setting in which education is treated as an investment good instead of a consumption good. We will contrast the nature of equilibrium allocations in these two settings and the role of access to a credit market in the investment model.

In the investment model, parents do not directly care about their childrens' education, but instead they care about their future income and consumption, which depends, in part, on education. Parents consume, borrow or lend, and potentially invest in college in an initial period. Children generate income and consume and repay loans in a second period. For expositional purposes, we assume away government policies such as in-state subsidies and need-based aid and also

abstract from drop-out risk.⁵⁵

Credit is in zero net supply, and households borrow and lend at an equilibrium interest rate R . Loans must be repaid, but to start, there are no additional restrictions on credit. The tuition function is given by $t(q, a)$. Taking the market interest rate and the tuition function as given, the household's problem is

$$\begin{aligned} & \max_{c_1, c_2, q, b} \{ \log(c_1) + \beta \log(c_2) \} \\ & \quad \text{s.t.} \\ & c_1 + t(q, a) \leq y + b - \mathbb{1}_{\{q > 0\}} \omega \\ & c_2 = y_2(q, a) - Rb, \end{aligned}$$

where the second period income function is

$$(10) \quad y_2(q, a) = (q + \tau)^\zeta a^\lambda, \text{ where } \tau, \zeta, \lambda > 0.$$

The interpretation is that the child's income depends on the quality of her education q and on her ability a , with corresponding elasticities defined by ζ and λ . The solution to this problem is a consumption rule $c_1(y, a)$, $c_2(y, a)$, an education choice $q(y, a)$, and a borrowing decision $b(y, a)$.

The college aspect of the model remains unchanged, with college decision rules $e(q)$ and $\eta_a(q)$ and an endogenous distribution of college quality $\chi(q)$. A competitive equilibrium can be defined similarly to the baseline college-as-consumption model, with one extra market-clearing condition for the credit market.

PROPOSITION 6: *The competitive equilibrium in the investment model with frictionless credit is Pareto efficient.*

PROPOSITION 7: *The competitive equilibrium of the investment model with frictionless credit features no sorting by income in college enrollment and positive sorting by ability.*

The intuition is that given the ability-education complementarity (equation 10), it is efficient to admit high ability students to college, regardless of their income. Since the competitive equilibrium must be efficient (Proposition 6), the equilibrium enrollment pattern features sorting only by ability. Note that the lack of sorting by income in the frictionless investment model is at odds with the empirical patterns documented in the data by Belley and Lochner (2007) and others, especially for recent decades.

We now introduce frictions into the credit market and show that as frictions increase, this model converges to our college-as-consumption model, which features positive sorting by income.

⁵⁵Incorporating drop-out risk in an investment setup would require a discussion of whether this risk can be insured.

A. Frictional Credit Market

We now introduce an exogenous wedge ρ between the borrowing rate and the lending rate, which captures frictions in the credit market. The lending rate is R and the borrowing rate is $R + \rho$. As $\rho \rightarrow 0$, the model collapses to the frictionless case just described.⁵⁶ The households' problem is identical to the one just described, except that the second period budget constraint is now

$$c_2 = y_2(q, a) - \{R + \mathbb{1}_{\{b>0\}}\rho\} b.$$

With a positive wedge $\rho > 0$, the equilibrium is not efficient. Intuitively, poor families with high ability children may under-invest in education when credit is costly.

When the credit friction is sufficiently severe (ρ is sufficiently large) the frictional investment model collapses to the consumption model we used for our quantitative analysis. The intuition is that when the credit friction is sufficiently severe, there is no credit in equilibrium: $b(y, a) = 0$ for any pair (y, a) . Substituting the implied expression for c_2 into the objective function, the household's problem is then isomorphic to the one in the college-as-consumption model:

PROPOSITION 8: *There exists a threshold $\bar{\rho}$ such that for any $\rho > \bar{\rho}$, the investment model is isomorphic to a consumption model in which $\varphi = \beta\zeta$ and $\kappa = \tau$.*

Thus, the consumption model considered in the main text can be reinterpreted as an investment model in which the financial friction is severe. In both these models, income plays an important role in college attendance. The frictionless investment model instead implies a quite different pattern, in which enrollment is independent of income. That implication is inconsistent with the data, while our baseline college-as-consumption model delivers a realistic pattern of college enrollment across the income-ability distribution. For that reason, we have focused on the college-as-consumption model in our quantitative application.

VI. Conclusion

A satisfactory model of the college market is essential for understanding what sorts of potential students go to college, what sorts of colleges they attend, and how much they pay. It is also important for understanding how these features of the college landscape have changed over time, and for exploring the impact of possible policy interventions.

We have developed a competitive model of the college market in which college quality depends on the average ability of attending students. A novel feature of

⁵⁶This is a highly stylized model. Embedding our model of the college market into a life-cycle framework with a quantitative model of the student loan market is a possible avenue for future research (see, for example, Abbott et al. 2019).

the model is a continuous distribution of college quality. When we use a calibrated version of the model to predict the impact of the rise in top-tail income inequality since 1989, we find that greater income inequality can explain over half of the observed increase in average net tuition. By itself, greater income inequality would depress enrollment rates, but we find that growth in average income and more generous college subsidies are countervailing forces that have pushed more people into college.

Our analysis could be extended in many directions. First, while a perfectly competitive model offers a reasonable positive theory of observed outcomes, competition is likely not perfect in practice. Fillmore (2016) argues that colleges exploit information on FAFSA forms to price discriminate by income. He shows that higher family income translates into smaller tuition discount offers, even after controlling for observable proxies for ability (ACT scores and high school GPAs). This pattern emerges primarily at the most selective schools, suggesting that such colleges have some pricing power (for example, see Epple et al. 2017, 2019). One could allow colleges to earn rents (and thereby consider alternative objectives to profit maximization) by endowing colleges with an idiosyncratic non-reproducible attribute, such as location or brand name, and endowing households with heterogeneous preferences over this attribute.

A second possible extension would be to model heterogeneity in college information about student ability. At the time of admission, students would then be unable to perfectly anticipate the terms of admission offers and would therefore want to apply to multiple schools (see, e.g., Fu 2014).

A third avenue for further research is an exploration of the nature of optimal college subsidies, tying subsidies to student attributes (income or ability), the quality of the college the student attends, or both (see Findeisen, Sachs and Colas 2018). The nature of the optimal intervention will depend on the planner's social welfare function, in addition to whether the laissez-faire allocation is efficient (as when college is a pure consumption good) or inefficient.

Fourth, one could consider broader sets of attributes that may affect students' attractiveness as peers, in addition to measures of performance on standardized tests. For example, one way to introduce income-dependent pricing in a competitive framework would be to posit that students perceive value to having peers from diverse backgrounds, which would make minority or low-income students attractive to colleges that draw predominantly from rich white families.⁵⁷

A final possible application would be to develop a multi-generational extension of the model outlined in Section V to explore the propagation of inequality across generations. Consider an increase in the financial return to college quality. This will lead to an increase in investment in quality by higher-income households, which will amplify the effect on income inequality in the next generation. In

⁵⁷For example, one could define peer desirability (ability) as $a = s^\kappa y^{1-\kappa}$, with s being a test score measure and y being family income. Given $\kappa > 1$, our competitive framework would then imply larger tuition discounts for low family income students, all else equal.

turn, a fatter right tail in the income distribution for that generation will further amplify inequality in college investment. Over successive generations, a small increase in the return to college quality can potentially generate both a large increase in income inequality and a decline in intergenerational mobility.

REFERENCES

- Abbott, Brant, Giovanni Gallipoli, Costas Meghir, and Giovanni L. Violante.** 2019. "Education Policy and Intergenerational Transfers in Equilibrium." *Journal of Political Economy*, 127(6): 2569–2624.
- Abdulkadiroğlu, Atila, Joshua Angrist, and Parag Pathak.** 2014. "The Elite Illusion: Achievement Effects at Boston and New York Exam Schools." *Econometrica*, 82(1): 137–196.
- Angrist, Joshua D.** 2014. "The Perils of Peer Effects." *Labour Economics*, 30: 98–108.
- Arnott, Richard, and John Rowse.** 1987. "Peer Group Effects and Educational Attainment." *Journal of Public Economics*, 32(3): 287–305.
- Belley, Philippe, and Lance Lochner.** 2007. "The Changing Role of Family Income and Ability in Determining Educational Achievement." *Journal of Human Capital*, 1(1): 37–89.
- Black, Sandra E.** 1999. "Do Better Schools Matter? Parental Valuation of Elementary Education." *The Quarterly Journal of Economics*, 114(2): 577–599.
- Boustan, Leah Platt.** 2012. "School Desegregation and Urban Change: Evidence from City Boundaries." *American Economic Journal: Applied Economics*, 4(1): 85–108.
- Carrell, Scott E., Richard L. Fullerton, and James E. West.** 2009. "Does Your Cohort Matter? Measuring Peer Effects in College Achievement." *Journal of Labor Economics*, 27(3): 439–464.
- Caucutt, Elizabeth M.** 2001. "Peer Group Effects in Applied General Equilibrium." *Economic Theory*, 17(1): 25–51.
- Caucutt, Elizabeth M.** 2002. "Educational Vouchers When There Are Peer Group Effects—Size Matters." *International Economic Review*, 43(1): 195–222.
- Chetty, Raj, John N. Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan.** 2017. "Mobility Report Cards: The Role of Colleges in Intergenerational Mobility." National Bureau of Economic Research Working Paper 23618.

- Chetty, Raj, John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan.** 2020. "Income Segregation and Intergenerational Mobility Across Colleges in the United States." *The Quarterly Journal of Economics*, 135(3): 1567–1633.
- Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez.** 2014. "Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States." *Quarterly Journal of Economics*, 129(4): 1553–1623.
- Cole, Harold L., and Edward C. Prescott.** 1997. "Valuation Equilibrium with Clubs." *Journal of Economic Theory*, 74(1): 19–39.
- Coleman, James S.** 1966. "Equality of Educational Opportunity (COLEMAN) Study (EEOS)." *Inter-University Consortium for Political and Social Research*.
- College Board.** 2016. *Trends in College Pricing 2016*. Trends in Higher Education Series. Available at <https://trends.collegeboard.org>.
- Dale, Stacy B., and Alan B. Krueger.** 2014. "Estimating the Effects of College Characteristics over the Career Using Administrative Earnings Data." *Journal of Human Resources*, 49(2): 323–358.
- Ellickson, Bryan, Birgit Grodal, Suzanne Scotchmer, and William R. Zame.** 1999. "Clubs and the Market." *Econometrica*, 67(5): 1185–1217.
- Epple, Dennis, and Richard E Romano.** 1998. "Competition between Private and Public Schools, Vouchers, and Peer-Group Effects." *American Economic Review*, 88(1): 33–62.
- Epple, Dennis, and Richard Romano.** 2008. "Educational Vouchers and Cream Skimming." *International Economic Review*, 49(4): 1395–1435.
- Epple, Dennis, and Richard Romano.** 2010. "Peer Effects in Education: a Survey of the Theory and Evidence." in *Jess Benhabib, Alberto Bisin, and Matthew O. Jackson (eds.) Handbook of Social Economics*, vol. 1.
- Epple, Dennis, Richard Romano, and Holger Sieg.** 2006. "Admission, Tuition, and Financial Aid Policies in the Market for Higher Education." *Econometrica*, 74(4): 885–928.
- Epple, Dennis, Richard Romano, Sinan Sarpca, and Holger Sieg.** 2017. "A General Equilibrium Analysis of State and Private Colleges and Access to Higher Education in the U.S." *Journal of Public Economics*, 155(C): 164–178.
- Epple, Dennis, Richard Romano, Sinan Sarpça, Holger Sieg, and Melanie Zaber.** 2019. "Market Power and Price Discrimination in the US Market for Higher Education." *The RAND Journal of Economics*, 50(1): 201–225.

- Fillmore, Ian.** 2016. "Price Discrimination and Public Policy in the U.S. College Market." Washington University Working Paper.
- Findeisen, Sebastian, Dominik Sachs, and Mark Colas.** 2018. "Optimal Need-Based Financial Aid." Working Paper 14, Opportunity and Inclusive Growth Institute, Federal Reserve Bank of Minneapolis.
- Fu, Chao.** 2014. "Equilibrium Tuition, Applications, Admissions, and Enrollment in the College Market." *Journal of Political Economy*, 122(2): 225–281.
- Golden, Daniel.** 2006. "The Price of Admission: How America's Ruling Class Buys Its Way into Elite Colleges—and Who Gets Left Outside the Gates." *New York: Three Rivers Press*.
- Gordon, Grey, and Aaron Hedlund.** 2017. "Accounting for the Rise in College Tuition." in *Charles R. Hultan and Valerie Ramey (eds.), Education, Skills, and Technical Change: Implications for Future U.S. GDP Growth*, 357–394 NBER Books.
- Gottschalk, Peter, and Minh Huynh.** 2010. "Are Earnings Inequality and Mobility Overstated? The Impact of Nonclassical Measurement Error." *The Review of Economics and Statistics*, 92(2): 302–315.
- Heathcote, Jonathan, and Hitoshi Tsujiyama.** 2019. "Optimal Income Taxation: Mirrlees Meets Ramsey." Federal Reserve Bank of Minneapolis Working Paper.
- Hendricks, Lutz, Tatyana Koreshkova, and Oksana Leukhina.** 2018. "Sorting of Students into Colleges: Inefficiencies and Policy Implications." Working Paper.
- Hoekstra, Mark.** 2009. "The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach." *The Review of Economics and Statistics*, 91(4): 717–724.
- Jones, John Bailey, and Fang Yang.** 2016. "Skill-Biased Technical Change and the Cost of Higher Education." *Journal of Labor Economics*, 34(3): 621–662.
- Kelly, Emily, Laura Holian, and Casey Archer.** 2019. "Trends in Pell Grant Receipt and the Characteristics of Pell Grant Recipients: Selected Years, 2003–04 to 2015–16." U.S. Department of Education.
- Lavy, Victor, M. Daniele Paserman, and Analia Schlosser.** 2012. "Inside the Black Box of Ability Peer Effects: Evidence from Variation in the Proportion of Low Achievers in the Classroom." *The Economic Journal*, 122(559): 208–237.

- Lazear, Edward P.** 2001. "Educational Production." *The Quarterly Journal of Economics*, 116(3): 777–803.
- Manski, Charles F.** 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies*, 60(3): 531–542.
- Marmaros, David, and Bruce Sacerdote.** 2002. "Peer and Social Networks in Job Search." *European Economic Review*, 46(4): 870 – 879.
- Piketty, Thomas, and Emmanuel Saez.** 2003. "Income Inequality in the United States, 1913–1998." *The Quarterly Journal of Economics*, 118(1): 1–41.
- Rothschild, Michael, and Lawrence J. White.** 1995. "The Analytics of the Pricing of Higher Education and Other Services in Which the Customers Are Inputs." *Journal of Political Economy*, 103(3): 573–586.
- Sacerdote, Bruce.** 2001. "Peer Effects with Random Assignment: Results for Dartmouth Roommates." *The Quarterly Journal of Economics*, 116(2): 681–704.
- Sacerdote, Bruce.** 2014. "Experimental and Quasi-Experimental Analysis of Peer Effects: Two Steps Forward?" *Annual Review of Economics*, 6(1): 253–272.
- Schwartz, Amy Ellen, and Benjamin Scafidi.** 2004. "What's Happened to the Price of College? Quality-Adjusted Net Price Indexes for Four-Year Colleges." *Journal of Human Resources*, 39(3): 723–745.
- Scotchmer, Suzanne.** 1997. "On Price-Taking Equilibria in Club Economies with Nonanonymous Crowding." *Journal of Public Economics*, 65(1): 75–88.
- Wai, Jonathan, Matt Brown, and Christopher Chabris.** 2018. "Using Standardized Test Scores to Include General Cognitive Ability in Education Research and Policy." *Journal of Intelligence*, 6(3): 37.
- Zimmerman, David J.** 2003. "Peer Effects in Academic Outcomes: Evidence from a Natural Experiment." *The Review of Economics and Statistics*, 85(1): 9–23.
- Zimmerman, Seth.** 2019. "Elite Colleges and Upward Mobility to Top Jobs and Top Incomes." *American Economic Review*, 109(1): 1–47.