The Demographic Cliff and the Market for Higher Education: Implications for Public Finance*

Dirk Krueger† Alexander Ludwig‡ Irina Popova§

September 9, 2026

Abstract

The decline in U.S. fertility since 2008 will reduce the number of high-school graduates in the coming decades and leave the population size smaller and its composition older. We study the response of the higher education market to this "demographic cliff" and quantitatively characterize how that response shapes the aggregate and fiscal consequences of population aging. To do so, we construct a general equilibrium life-cycle model of the higher education market with heterogeneous colleges, college quality and major choice, and subsequent labor market outcomes of workers in different occupations. In the model, the cross-sectional distribution of college quality and major specialization is an equilibrium outcome. Scarcer labor raises the return to schooling, and the resulting expansion of per capita college attendance and of the high-skill tax base absorbs a substantial part of the fiscal cost of aging. The quantitative importance of this margin depends on the structure of the higher education market: when colleges are homogeneous in quality and major, the response of enrollment, and with it the implied effective labor supply, is substantially smaller. We then show that recent cuts in public education funding are particularly costly against this backdrop.

JEL Classification: D15, D31, E24, I24 **Keywords:** Equilibrium Model of College Market, Education Subsidies, Demographic Cliff, Technological Change, College Major Choice.

*We thank our discussant Andrew Glover, Jonathan Heathcote, Bob Inman, Kjetil Storesletten, Oksana Leukhina, Robeerto Pinheiro, Oliko Vardishvili and Gianluca Violante as well as participants of the 2025 SED in Copenhagen, the 2025 Omaha Macroeconomic Forum, 2026 CES Meetings Atlanta, the Richmond Macro Workshop, the Sogang Global Symposium, the Atlanta, Cleveland and Richmond FED's, Cambridge, Duke, Grinnell, GWU, HWI Halle, IDB, Loughborough, Mannheim, Oxford, Purdue, Sciences Po, UC Irvine, Regensburg, and the Wharton Labor and Public Economics Seminar for helpful comments. A first version of this paper was completed while Krueger was visiting KFG 50 (CASFI) at the University of Bonn, whose hospitality is gratefully acknowledged.

†University of Pennsylvania, NBER and CEPR

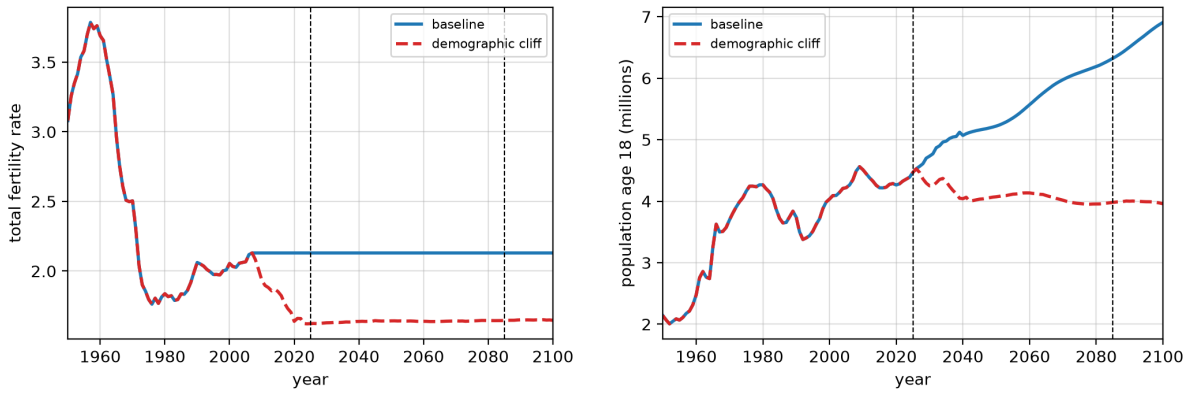
‡EUI, Florence, Goethe University Frankfurt, CEPR

§University of Bonn and ECONtribute

1 Introduction

The U.S. is undergoing a demographic cliff that is expected to have a profound impact on higher education and the labor market in the coming decades. Just between 2025 and 2032 the number of domestic high school graduates is expected to fall by 5% (NCES Digest of Education Statistics, 2023). Figure 1 shows the actual and predicted fertility rate in the US and illustrates the number of potential college students lost due to the decline of the fertility rate since 2008. The college supply side is already responding to this demographic shift. About 100 institutions closed between AY 2022/23 and AY 2023/24 and 2020 enrollments were about 10% lower than in 2010 (2023 NCES Digest of Education Statistics). Not only will lower enrollments eventually translate into smaller graduating classes, but there are concerns that shortages in specific areas such as STEM fields, will be especially prominent.[1]

Figure 1: Fertility Rate and Population Size



(a) Total Fertility Rate

(b) Population of Age 18

Notes: Total fertility rate and population of age 18, 1950–2100. Baseline (solid): fertility held at its 2007 level. Demographic cliff (dashed): the decline in fertility since 2008 carried forward. Vertical lines mark 2025 and 2085. Own calculations based on UN and HMD. The full demographic projection implied by these fertility paths, including the resulting age structure of the population, is described in Section 4.1.

The broad question we are concerned with is how private education choices and public education policy should respond to these demographic and technological trends, and how these trends alter the trade-offs involved in designing public education and redistributive

[1]BLS (2025) estimates that STEM-based occupations are expected to grow faster than non-STEM occupations. At the same time, STEM fields (both in college and in the labor market) are populated heavily by immigrants, see e.g., (NSF Board, 2024) and changes in immigration policy will (and perhaps already have) exacerbate domestic concerns about skill shortages in key labor market areas. Furthermore, there is evidence that STEM careers are getting shorter and are subject to increased risk of technological obsolescence (see, e.g., Deming and Noray, 2020).

policies? To do so, we construct a general equilibrium life cycle model of the college market with heterogeneous colleges, student college quality and major choice, and subsequent labor market outcomes of workers in different occupations. The cross-sectional distribution of college quality and college specialization is an equilibrium outcome and is shaped by the demographic structure of the economy, by public education spending and college loan policies, as well as by the endogenous relative wages that college graduates with different majors command in the labor market. We use the model to evaluate the aggregate and distributional consequences of the "demographic cliff" that will reduce the number of high-school graduates in the next decades, against the backdrop of technological change (e.g., the emergence of AI) in the labor market that changes the relative demand for STEM-intensive occupations and college majors. We then evaluate how our conclusions depend on fiscal policy, in particular on recent federal funding cuts in higher education, focusing on their impact on enrollment, the composition of college quality, tuition levels, wages, and public finances.

To conduct our analysis we embed a life cycle model with college quality and college major choice into the competitive equilibrium model of the college market in the spirit of Cai and Heathcote (2022). In our adoption of their framework, each college quality-major bundle is a different college. Households that differ in their initial resources (due to, e.g., inter-vivos transfers by their parents), their innate ability to study and their residence status (which determines whether they qualify for in-state tuition at public colleges) decide which (if any) college to attend, given their expected labor earnings processes from a given quality-major combination. After college, a household's life follows a standard life cycle with consumption-savings decisions, eventual exogenous retirement and stochastic death. On the college supply side, competitive colleges can freely enter the college market at each level of college quality, where quality is determined by a combination of average student quality as well as per-student discretionary spending. The college-quality- and student-ability-specific tuition schedule is determined in competitive equilibrium by household demand and the supply of competitive colleges. The key mechanism of the college market model operates through the ability composition of the student body. Changes in enrollment alter the share of high-ability students at each college type, which in turn affects instructional spending and the equilibrium tuition discounts offered to attract high-ability students. These pricing adjustments then feed back into households' college and major choices.

After calibrating the model to U.S. macroeconomic, demographic, and education data, we analyze three counterfactual experiments: (i) the removal of public subsidies to higher education, (ii) a demographic transition characterized by declining population growth, and (iii) the joint impact of these two shocks. Accordingly, we first use the model to characterize how the demographic cliff, modeled as a permanent decline in the U.S. fertility rate starting

in 2008 impacts the U.S. college market, the labor market and public finances in the long run. In order to quantify the importance of these demographic shifts we then consider a counterfactual (from the perspective of today) demographic scenario in which the fertility rate is at replacement level and therefore the population is constant in the long run. Finally, for both demographic scenarios we evaluate the aggregate, distributional, college market, public finance and welfare consequences of abolishing public subsidies for higher education.

The declining population due to the demographic cliff, perhaps not surprisingly, has large negative effects on macroeconomic aggregates in the long run, simply because the adult population in 2085 is about 21% smaller than in the counterfactual scenario without changes in demographics, with the working-age population about 27% and the college-age cohorts about 37% smaller. However, per-capita outcomes are less detrimental. First, the demographic cliff has modest but favorable effects on the college market. Concretely, the long-run college enrollment rate increases and the distribution of students shifts toward middle-quality colleges. However, changes in the ability composition of individuals deciding to go to college reduce the share of STEM majors in the middle segment of the quality distribution. Overall, on account of dwindling absolute numbers of high-school graduates, colleges experience a decline in absolute student numbers throughout the quality distribution. The share of high-ability students remains almost unchanged at top-quality colleges but falls substantially at low- and medium-quality colleges due to the inflow of less-prepared marginal students, which in turn increases the tuition-ability discount in the lower and middle segments of the quality distribution. These adjustments depend on the heterogeneity of the higher education market itself. Repeating the demographic experiment in a comparison economy in which colleges are homogeneous in quality and major, we find that enrollment responds substantially less to the same demographic shock, output per capita falls by more, and the labor income tax rate rises rather than remaining close to unchanged.

Given our overall focus on the interaction between public education policies and long-run demographic and technological trends (or shocks), we next evaluate a public college funding cut akin to the one recently implemented by the current administration, both under a baseline scenario with the initial steady-state population growth rate and under the demographic cliff scenario described above. We find that the long-run *aggregate* losses induced by a public funding cut are much smaller than those generated by the demographic shock. However, the subsidy cut qualitatively undoes the positive effects of the demographic shock on college enrollment, the quality distribution of students, and college tuition. When the subsidy cut takes place against the backdrop of the demographic cliff, perhaps the currently most relevant *factual* scenario, the joint impact of the two shocks on per capita output and public finances falls slightly short of the sum of their individual effects. In 2086 the labor income

tax rate is 5.91 percentage points above the baseline path, against a sum of 6.27 percentage points, and output per capita is 18.9% lower, against a sum of 19.1%. The interaction is mild, and it runs through the education margin: the demographic change does more work for education when subsidies are absent. Aging raises the return to schooling and lowers the pension replacement rate whatever the subsidy regime, but it also raises bequests per entrant, since the estates accumulated by the larger pre-cliff generations are divided among smaller cohorts. Where subsidies are in place, college is already cheap at the margin and much of that windfall is spent on other things; where they have been withdrawn, the same resources relax a constraint that now binds. The enrollment response to aging is therefore about two percentage points stronger in the economy without subsidies. This does not soften the policy reading: an economy that would have entered mid-century with attendance nine percentage points above the baseline path and a slightly lower tax rate instead enters it with attendance eighteen points below baseline and a tax rate six points higher.

1.1 Related Literature

This paper builds on and contributes to four strands of the existing literature: (i) the consequences of demographic change for higher education demand and the college landscape; (ii) equilibrium models of the college market with heterogeneous institutions, endogenous tuition, and student sorting; (iii) technological change, skill demand, and the returns to college majors; and (iv) macro public finance, education policy, and long-run fiscal sustainability.

Demographic change, higher education demand, and the changing college landscape. A large literature studies the implications of demographic change for education demand, college enrollment patterns, and institutional stress in higher education. This literature documents the decline in the number of college-aged individuals, the associated pressure on college enrollments, and the risk of institutional closures or financial distress (National Center for Education Statistics, 2023; Kelchen et al., 2024; Conesa et al., 2020). More broadly, it highlights that long-run demographic trends are likely to affect not only aggregate labor supply but also the educational composition of the workforce. Empirical work has also documented how demographic changes affect education demand and skill composition across countries and regions, with implications for labor supply and productivity, e.g. Hendricks and Schoellman (2023). In the quantitative macroeconomics literature, Ludwig et al. (2012) show that endogenous human capital accumulation is an important adjustment channel dampening the welfare consequences of demographic aging in a general equilibrium OLG model. This paper builds on this literature by moving from descriptive evidence on enrollment decline, college financial hardship, and college closures to an equilibrium anal-

ysis of how colleges adjust tuition schedules, educational quality, and program offerings in response to demographic change.

Equilibrium models of the college market and the supply side of higher education. A second strand studies higher education as an equilibrium market with heterogeneous institutions, endogenous tuition, admissions, student sorting, and financial aid. Important contributions in macroeconomics and industrial organization include Cai and Heathcote (2022), Alon et al. (2025), Capelle and Matsuda (2025), Arcidiacono (2005), Epple et al. (2006), Fu (2014), and Marto and Wittmann (2024). Related work also emphasizes quality differences across colleges, mismatch, and the consequences of differential college access for inequality and intergenerational mobility (Blair and Smetters, 2021; Hendricks et al., 2021; Hendricks, Koreshkova, and Leukhina, 2021; Hendricks et al., 2024). This paper builds directly on this literature, but embeds the college market into a broader macro–public finance environment in which demographic transition, technological change, early-life human capital accumulation, and government policy jointly shape the equilibrium structure of higher education.

Technological change, skill demand, and the composition of human capital. Another related literature studies how technological change affects the labor market returns to education, occupational structure, and the demand for specific skills. This includes work on skill-biased technological change and capital–skill complementarity (Krusell et al., 2000; Caunedo et al., 2023), as well as more recent evidence on job displacement, skill obsolescence, and changing returns to college majors and occupations (Deming and Noray, 2020; Braxton and Taska, 2023; Hendricks and Schoellman, 2023; Kogan et al., 2023). Closely related empirical work studies major choice, the returns to different fields of study, and the changing task content of jobs (Altonji et al., 2016; Altonji and Zimmerman, 2018; Altonji et al., 2012, 2014). Cross-country empirical work also highlights the role of human capital in technology adoption and diffusion (Sunde and Vischer, 2015; Cervellati et al., 2023). This paper contributes to this literature by providing a quantitative framework in which the returns to different college majors, the equilibrium composition of human capital across fields of study, and the response of the college market to technological and demographic shocks are jointly determined in general equilibrium.

Macro public finance, education policy, and long-run fiscal consequences. This paper is also closely connected to a literature in macro public finance that studies how taxes, transfers, education subsidies, public insurance, and borrowing constraints shape human capital accumulation, inequality, and long-run fiscal outcomes. Educational choices affect not only private earnings and mobility, but also productivity of future cohorts, thus the future tax base and the sustainability of public finances. In particular, this paper builds on

recent work on education finance reform and the distributional consequences of education policy (Krueger et al., 2025). Its contribution is to analyze these public-finance questions in a unified dynamic setting in which the structure of higher education is endogenous and policy reforms can be evaluated along both transition paths and long-run equilibria.

An important overall contribution of this paper is to build an integrated quantitative framework in which education and human capital investment decisions, the college market, labor-market outcomes, and public policy interact jointly over demographic and technological transitions.[2] This allows us to study not only aggregate enrollment and average skill levels, but also the composition of human capital, the allocation of talent across institutions and fields of study, and the implications of these adjustments for inequality, intergenerational mobility, and fiscal sustainability.

2 The Model

Time is indexed by t and in each period a new cohort of size $N_{0,t}$ is born. The size of newborn cohorts $N_{0,t}$ is exogenous but may vary over time. There is a continuous mass of households in each cohort. Households live from age $j = 0$ (which represents the college phase or the first period of the labor market in case of not going to college) to at most age $j = J$, so that life after college starts at age $j = 1$. The retirement age j_r is exogenous and after this age households face age-specific survival risk: the probability of surviving from age $j-1$ to age j is given by ψ_{j-1}, with $\psi_j = 1$ for $j \leq j_r - 1$ and $\psi_J = 0$. Therefore, the number of individuals of age j alive in period t is defined recursively as

$$N_{t+1,j+1} = \psi_j N_{t,j} \tag{1}$$

We will describe the model in a stationary recursive setting, therefore suppressing time subscripts t until we introduce unexpected MIT transitions. The stationary population structure is then denoted by the vector $(N_0, ...N_J)$ and determined by (1) as well as the normalization $N_0 = 1$.

2.1 Household Endowments, Preferences, Budget Sets and Decision Problems

2.1.1 Endowments and Budget Constraints

Households are endowed with innate ability e which falls into a discrete set $e \in \mathcal{E} = \{e_1, e_2, \ldots, e_E\}$, and we denote by $\mu_e(.)$ the measure of households of a given e. Following

[2] Modeling intergenerational linkages explicitly — with endogenous parental investment in the human capital of their children and intergenerational wealth transfers, in the spirit of our own earlier work (Krueger et al., 2025) — would be an interesting extension of the current framework. We abstract from intergenerational extensions and endogenous pre-college human capital in the current paper.

Cai and Heathcote (2022), we assume that in addition to ability households also differ by residence status $r \in \mathcal{R} = \{r^i, r^o\}$ which solely determines the generosity of government subsidies received by universities when these households become students (and hence the tuition these students pay).We denote by $\mu_r(.)$ the measure of households in residence r. In addition, households start with initial resources $b \in \mathcal{B}$, interpreted as initial wealth transfers from parents, which are drawn from an exogenous child ability-specific probability measure $F_e(.)$. Thus, initial household heterogeneity is summarized by the vector $(e, r, b) \in \mathcal{E} \times \mathcal{R} \times \mathcal{B} \equiv \mathcal{S}$. Probability measures over cohort age j populations will be denoted generically by Φ_j and are endogenous for ages $j \geq 1$, but the measure Φ_0 is exogenous and determined by the initial distributions (μ_e, μ_r, F_e) for the economically newborn (aged 18 in the real wold) cohort.[3]

Households can choose to attend colleges of different varieties, where a given college type is characterized by a quality $q \in \mathcal{Q} = \{0, [q_{\min}, q_{\max}]\}$, (and $q = 0$ is non-college and $q_{\min}$ is the lowest and $q_{\max}$ is the maximum college quality available on the market), as well as a program $p \in \mathcal{P}$ or specialization. We think of $\mathcal{P}$ of a small finite and fixed set, which in our application consists of three elements ($p = 0$, no-college, $p = 1$, non-STEM college and $p = 2$, college with STEM specialization). In contrast, a college of any quality in the interval $[q_{\min}, q_{\max}]$ can endogenously emerge in equilibrium. Households choose a college type $(q, p) \in \mathcal{Q} \times \mathcal{P}$. College students are subject to exogenous college completion risk: with probability $\pi_d(q)$ they do not complete college. Since we assume that dropouts are liable for the full tuition, dropout risk only affects household choices, but has no impact on the college supply side.

College Age In the first period of life, in addition to the wealth transfer by their parents b, households earn an exogenous income y_0, standing in for working full-time at the non-college wage. If they attend college, which we denote by indicator $\mathbb{1}_{q>0}$, households lose a fraction ω of that income, which represents the opportunity cost of time required to study (and which is the purpose of introducing y_0 in the first place). Gross income is therefore $(1 - \mathbb{1}_{q>0}\omega)\, y_0$ on which households pay taxes according to a possibly non-linear tax function $T(\cdot)$. Net labor income is $y_0^n(q) = y_0\,(1 - \mathbb{1}_{q>0}\omega) - T\,(y_0\,(1 - \mathbb{1}_{q>0}\omega))$. Besides labor income taxation, households are also subject to proportional capital income and consumption taxes, with rates τ_k and τ_c, respectively.

If a household attends college ($\mathbb{1}_{q>0}$), she receives a potentially means tested transfer that is a fraction $\varsigma(b)$ of the ability e, residence r and quality q dependent college tuition $t(e, r, b; q, p)$ which is determined in equilibrium. Denoting by c_0 consumption and

[3]Concretely, denote by $\mathbf{B}(\mathcal{E})$, $\mathbf{B}(\mathcal{R})$, and $\mathbf{B}(\mathcal{B})$ the (Borel) σ-algebras of $(\mathcal{E}, \mathcal{R}, \mathcal{B})$ and by $\mathbf{B}(\mathcal{S})$ the σ-algebra of $\mathcal{S}$. Then for any $E \subset \mathbf{B}(\mathcal{E})$, $R \subset \mathbf{B}(\mathcal{R})$ and $B \subset \mathbf{B}(\mathcal{B})$ we have $\Phi_0(E, R, B) = \mu_e(E)\mu_r(R)\sum_{e \in E} F_e(B)$.

by a_1 asset choices of households in the first period of life, the budget set for the first period of life is defined by

$$(1+\tau_c)c_0 + a_1 = b + y_0\left(1 - \mathbb{1}_{q>0}\omega\right) - T\left(y_0\left(1 - \mathbb{1}_{q>0}\omega\right)\right) - \mathbb{1}_{q>0}(1-\varsigma\left(b\right))t(e,r,b;q,p) \tag{2}$$

$$a \geq -\underline{a}_1(q) \tag{3}$$

where $\underline{a}_1(q)$ denotes the education-specific borrowing limit: only households that choose $q > 0$ are allowed to borrow for their college education (at a level we specify in the calibration section), whereas for those with $q = 0$ we impose a tight borrowing constraint at zero on the asset choice a_1. The direct utility of going to college is given by

$$\mathbb{1}_{q>0}\left(\nu_0 + \nu_1 \ln(q) + \mathbb{1}_{p=2}\nu_2(e)\right) \tag{4}$$

where ν_0 captures the utility flow from attending any college, $\nu_1 \ln(q)$ encodes the direct preference for high-quality college and $\nu_2(e)$ measures the (dis-)utility for studying STEM, which is assumed to be student e ability-dependent.

Working Life Ages After college, for all ages $j \in \{1, \ldots, j_r\}$ households participate in the labor market; if they have chosen $q = 0$, their labor market career commences at age $j = 0$ already. Their labor income is (potentially) stochastic and depends on the quality of a college attended as well as the program studied. The deterministic age-productivity component depends only on whether a worker attended college. Denote by $s \equiv \mathbf{1}\{q > 0\} \in \{0, 1\}$ the college attendance indicator, and let $\gamma_j(s)$ be the base age-productivity profile common to all workers with schooling status s. For college graduates, productivity additionally depends on college quality q and program choice p.

Motivated by Mincer-type earnings regressions, we assume that log productivity depends linearly on a transformed measure of college quality and major. In particular, we map college quality into an index $f(q)$ that is increasing in q but exhibits diminishing returns, reflecting the idea that the marginal productivity gains from attending increasingly higher-quality colleges eventually flatten out. Productivity for college graduates is then given by

$$\gamma_j(q,p) = \exp(\gamma_j(s) + f(q)\left(\beta_1 + \beta_2 \mathbf{1}\{p=2\}\right)), \tag{5}$$

so that log-productivity and thus log-wages are given by:

$$\log(\gamma_j(q,p)) = \gamma_j(s) + f(q)\left(\beta_1 + \beta_2 \mathbf{1}\{p=2\}\right) \tag{6}$$

where $\mathbf{1}\{p = 2\}$ indicates a STEM program graduate. The parameters β_1, and β_2 determine the overall college wage premium and the STEM premium. The function $f(q)$ is increasing and concave in q, implying that productivity gains from college quality eventually saturate. In the quantitative implementation we use an exponential saturation transformation $f(q) = 1 - \exp(-\beta_3 q)$, where β_3 governs the speed at which returns to college quality flatten out, and therefore how strongly earnings differ across the observed range of college qualities[4].

Labor markets are segmented by $p \in \{0, 1, 2\}$, and the wage for labor with training p is given by $w(p)$. Realized gross labor income is given by

$$y_j(q, p) = \gamma_j(q, p) w(p)$$

and net labor income of a type (q, p) worker then is $y_j^n(q, p) = y_j(q, p) - T(y_j(q, p))$, and a typical working age budget constraint is given by

$$(1 + \tau_c)c_j + a_{j+1} = y_j^n(q, p) + (1 + r(1 - \tau_k))a_j \tag{7}$$

Retirement For retired people $j > j_r$, the budget constraint is similar to (7), but labor income $y_j(q, p)$ is replaced by retirement benefits $pen_j(q)$. Households that die prematurely leave accidental bequests that are confiscated by the government and form part of tax revenue.

2.1.2 Preferences

In each period, households derive per-period utility from a period utility function $u(c)$. We further assume that during college $j = 0$ households derive direct utility benefits from attending college (q, p) given by equation (4).

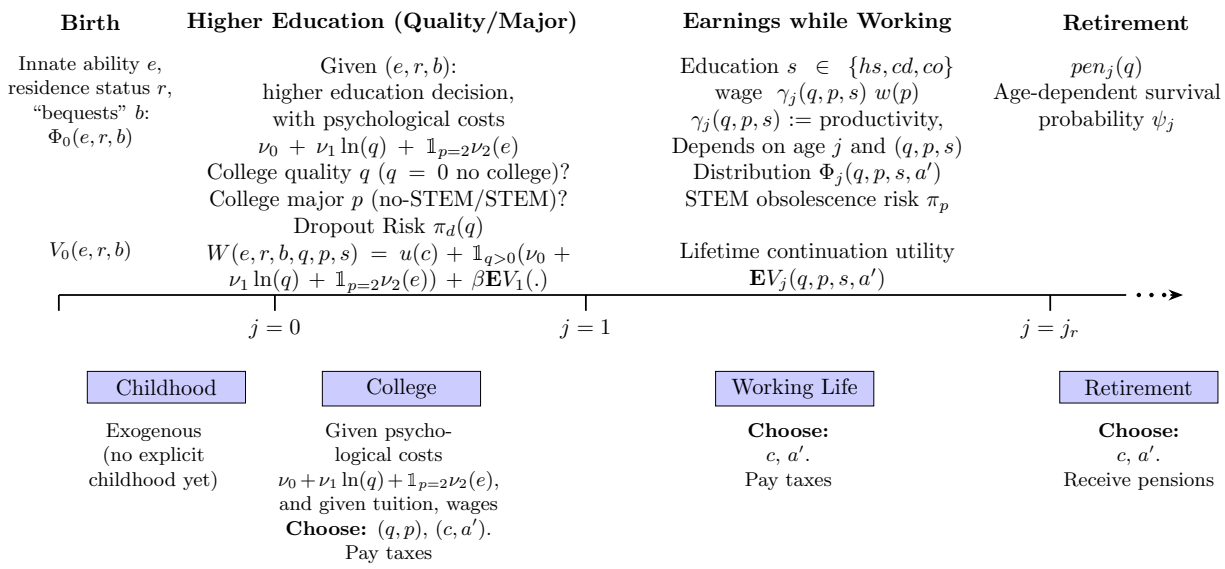
2.2 Timing of Events

At every age $j \geq 1$ households simultaneously make consumption and savings choice. At age $j = 0$, the timing of events is as follows. A household of type (e, r, b) first decides on whether to go to college, and conditional on going, what quality and program combination (q, p) to choose. Finally the household makes consumption-savings choices; at the end of the period the relevant state vector of the household is (e, q, p, a) with associated cross-sectional measure Φ_1.

[4]In the current version of the paper, labor market productivity does not depend explicitly on innate ability e, and thus e is not currently a state variable for ages $j \geq 1$. At the end of every working age $j \geq 1$ there is a chance that an individual with state $p = 2$ turns into an individual with state $p = 1$. Consequently labor productivity is deterministic, conditional on the type of a worker (e, q, p), but is rendered stochastic by the random transitions of q at the end of the college period to model college dropout risk and p (in every period of working life).

The timing of events, and consequently the model itself, is summarized succinctly in Figure 2.

Figure 2: Time Line of the Model



Notes: Time line of the model for period $j=0$ and for an arbitrary working period $j \in \{1,...,j_r-1\}$ and retirement period $j \in \{j_r,...,J\}$.

2.3 Recursive Formulation of the Household Problem

For all ages $j \geq 1$, households solve a standard consumption-savings problem with labor income being determined by the productivity process $\gamma_j(.)$ and the risk of skill obsolescence if $p=2$. To capture the risk that STEM-specific skills become obsolete over the working life, we allow the program state p to evolve stochastically: in each period, a worker with $p_j = 2$ transitions to $p_{j+1} = 1$ with probability $\pi_p \in (0,1)$. Denote the value function at the beginning of model age $j=1$ as $V_1(e,q,p,a)$. It depends on the education states (q,p), initial ability e (all of which determine labor productivity) as well as asset holdings a (which can be negative if the household borrowed to pay the tuition in college). We spell out the recursive problem(s) giving rise to $V_1(e,q,p,a)$ in Appendix A; this includes determining all age-specific value functions $\{V_j(e,q,p,a)\}$ and associated policy functions $\{(c_j(e,q,p,a)), a_{j+1}(e,q,p,a)\}$, for $j \geq 1$.

We now describe the optimization problems the household solves sequentially during age $j=0$, backward from the consumption-saving choice to the college attendance, quality and

major choice; we spell them out in the main paper as these are the more novel aspects of our model.

2.3.1 The Consumption-Saving Choice

After having made higher education choices (q, p) the relevant state vector of the household is (e, q, p, a) and the household solves the problem:

$$\begin{aligned} W(e,r,b,q;p) &= \max_{c_0,a_1} \{u(c_0) + \beta V_1(e,q,p,a_1)\} \qquad s.t. & (8)\\ (1+\tau_c)c_0 + a_1 &= b + y_0^n - \mathbb{1}_{q>0}(1-\varsigma(b))t(e,r,b;q,p) & (9)\\ a_1 &\geq -\underline{a}_1(q), c_0 \geq 0, q \in \mathcal{Q} & (10) \end{aligned}$$

where $y_0^n = y_0(1 - \mathbb{1}_{q>0}\omega) - T(y_0(1 - \mathbb{1}_{q>0}\omega))$ is net labor income during the college phase (which depends on whether the household chooses to attend college or not). The resulting policy functions are denoted by $(c_0(e,r,b,q,p), a_1(e,r,b,q,p))$. It is understood that for $q = 0$ (the household not attending college) the p-dimension of the value function is set to $p = 0$, standing in for non-college labor.

2.3.2 The College Decision

Working backwards, at the beginning of economic life a household with initial state (e, r, b) chooses what college quality and field of study (q, p) to study. As discussed before, in the theory we think of each (q, p)-combination as a separate college type.

$$V_0(e,r,b) = \max_{q\in\mathcal{Q},p\in\mathcal{P}} \{\mathbb{1}_{q>0}(\nu_0 + \nu_1 \ln(q) + \mathbb{1}_{p=2}\nu_2(e)) + W(e,r,b,q,p)\} \tag{11}$$

where $W(e, r, b, q, p)$ was defined in the previous section. We denote the associated optimal college choice functions by $q(e, r, b)$ and $p(e, r, b)$.

2.4 Cross-Sectional Distributions Implied by Household Choices

Given the college choice functions and the initial exogenous probability measure Φ_0 across characteristics (e, r, b) we can now define the cross-sectional distributions at each decision stage for households of age $j = 0$, as well as at the beginning of age $j = 1$.

In order to prepare the definition of a college market equilibrium we first specify the cross-sectional college quality distribution implied by the household desired college attendance decisions as well as the distribution of initial household characteristics. First, the mass of individuals seeking to attend a college of quality $q \in Q$ and specialization p is given by:

$$\chi(Q,p) = \sum_e \sum_r \int_b \mathbb{1}_{\{q(e,r,b)\in Q, p(e,r,b)=p\}} d\Phi_0(e,r,b) \tag{12}$$

Note that the same equation applies to $(Q,p)=(\{0\},0)$ which determines the set of households not attending college. The mass of households of ability e wanting to attend a college of type (Q,p) is given by

$$\Phi_e(Q,p)=\sum_r\int_b \mathbb{1}_{\{q(e,r,b)\in Q,p(e,r,b)=p\}}d\Phi_0(\{e\},r,b) \tag{13}$$

and thus for all Q with $\chi(Q)>0$ and for all p the conditional distribution of student ability e (i.e., the share of high-ability students) in quality set Q and major p is given by

$$\Phi(e|Q,p)=\frac{\Phi_e(Q,p)}{\chi(Q,p)}. \tag{14}$$

Of course, the set Q can be a singleton, as long as that singleton has positive mass (as will be the case once we discretize the college quality set $\mathcal{Q}$).

At the end of the college age $j=0$, or equivalently, at the beginning of age $j=1$ the remaining state variables are given by (e,q,p,a) and the cross sectional probability measure, for each $e\in\mathcal{E}$, and for all sets $(Q,P,A)\in(\mathbf{B}(\mathcal{Q}),\mathbf{B}(\mathcal{P}),\mathbf{B}(\mathcal{B}))$ with $\{0\}\neq Q$ is given by

$$\Phi_1(\{e\},Q,P,A)=\sum_r\int_b \mathbb{1}_{\{q(e,r,b)\in Q,p(e,r,b)\in P,a_1(b,e,r,q(e,r,b),p(e,r,b))\in A\}}d\Phi_0(\{e\},r,b) \tag{15}$$

For $q=0$ (households not going to college) it is understood that the college program dimension is irrelevant (and set to $p=0$):

$$\Phi_1(\{e\},\{0\},\{0\},A)=\sum_r\int_b \mathbb{1}_{\{q(e,r,b)=0,p(e,r,b)=0,a_1(b,e,r,q(e,r,b),p(e,r,b))\in A\}}d\Phi_0(\{e\},r,b) \tag{16}$$

Finally in order to appropriately aggregate the consumption and asset choice at the end of period 0, we need the cross-sectional population measure $\bar{\Phi}_0$ across characteristics (e,r,b,q,p) (since this is what the consumption-savings policies depend upon). This distribution is given by

$$\bar{\Phi}_0(\{e\},\{r\},B,Q,P)=\int_{b\in B}\mathbb{1}_{\{q(e,r,b)\in Q,p(e,r,b)\in P\}}d\Phi_0(\{e\},\{r\},b). \tag{17}$$

where again it is understood that for $q=0$ the college program dimension is irrelevant.

2.5 College Supply

Having characterized the demand of households for college of different qualities, we now build upon Cai and Heathcote (2022)'s equilibrium model of college supply. To do so, we

assume that colleges operate in a competitive environment with free entry into any quality segment $q \in \mathcal{Q} = [0, [q_{\min}, q_{\max}]]$ and program $p \in \mathcal{P}$ of the college market. That is, a college is characterized both by its quality q and the program p it offers.

Our model treats college as an investment good: college quality q and program p raise graduates' future labor market productivity, and therefore the college choice is fundamentally a human capital investment decision. This distinguishes our setting from Cai and Heathcote (2022), whose benchmark model treats college as a consumption good, with important implications for equilibrium properties and the solution algorithm; see Section 3.4 for a detailed discussion.

Suppose a college has a student ability distribution given by $\{\eta_e\}$ where η_e is the share of students of ability e; recall that $e \in \mathcal{E}$ falls in a finite set. Denote by

$$\bar{e} = \sum_e \eta_e e \tag{18}$$

the average student ability in college q. Given $\bar{e}$ colleges decide on instructional spending per student i and produce quality q according to CRS production function

$$q = A_q \left(\frac{\bar{e}}{\bar{e}^{\text{avg}}}\right)^{\theta} \left(\frac{i}{i^{\text{avg}}}\right)^{1-\theta}, \tag{19}$$

where $A_q > 0$ is a TFP parameter, and $\bar{e}^{\text{avg}}$ and i^{avg} denote the mean levels of student ability and instructional spending in the economy, respectively.

In addition to tuition revenue, colleges also receive per student public subsidies $s(e, r; q, p)$ that potentially depend on the quality and program a college delivers and on the characteristics of students the college admits. Conditional on ability e, profit-maximizing colleges strictly prefer to admit only students who generate the most revenue. Let $v(q, p, e) = \max_{b,r} t(e, r, b; q, p) + s(e, r; q, p)$ denote revenue from admitting the highest-revenue students (by resources b and residence status r) of ability level e. Producing programs p also comes at a program specific fixed cost $\kappa(p)$. For a given q, colleges maximize profits:

$$\pi(q, p) = \max_{i, \{\eta_e\}} \left\{ \sum_{e \in E} \eta_e v(q, p, e) - i - \kappa(p) \right\} \tag{20}$$

subject to (19).

When ability can take only two values, with e^l and e^h denoting low and high ability, respectively, we can rewrite the above problem as follows:

$$\pi(q,p) = \max_{i,\eta(e^h)} \left\{ \eta(e^h)[v(q,p,e^h) - v(q,p,e^l)] + v(q,p,e^l) - i - \kappa(p) \right\} \tag{21}$$

subject to

$$q = A_q \left(\frac{\bar{e}}{\bar{e}^{\text{avg}}} \right)^{\theta} \left(\frac{i}{i^{\text{avg}}} \right)^{1-\theta} \tag{22}$$

$$\bar{e} = \eta(e^h)(e^h - e^l) + e^l. \tag{23}$$

Thus, as in the general case, colleges choose per student expenditures and average student ability with the latter being equivalent to choosing the share of high-ability students $\eta(e^h)$ in the case with two ability levels. The first-order conditions to the problem result in the following optimality condition:

$$d(q,p) := -\frac{v(q,p,e^h) - v(q,p,e^l)}{e^h - e^l} = \frac{\theta}{1-\theta}\frac{i}{\bar{e}} > 0 \tag{24}$$

where $d(q,p)$ is the per-unit of ability discount that high-ability students receive relative to low-ability students, on account of the fact that they raise average ability and thus quality of the college. The ability tuition discount for each college type (q,p) is the key equilibrium "pricing" function that the market clearing condition for each college sub-market has to pin down. All other equilibrium elements as well as the qualitative features of the equilibrium then follow from the properties of $d(q,p)$. To see this, first note that the optimality condition (24) can be combined with (23) to determine optimal college inputs as

$$i(q,p) = \left[\frac{1-\theta}{\theta} d(q,p) \right]^{\theta} \frac{q}{A_q} \left(\bar{e}^{\text{avg}} \right)^{\theta} \left(i^{\text{avg}} \right)^{1-\theta} \tag{25}$$

$$\bar{e}(q,p) = \left[\frac{1-\theta}{\theta} d(q,p) \right]^{\theta-1} \frac{q}{A_q} \left(\bar{e}^{\text{avg}} \right)^{\theta} \left(i^{\text{avg}} \right)^{1-\theta} \tag{26}$$

In terms of intuition, we can think of a given college sub-market (q,p) as being characterized by a low equilibrium quality discount $d(q,p)$ or equivalently, from equation (26) by a high average quality of students $\bar{e}(q,p)$, or (see equation (23)) by a high share of high-ability students,

$$\eta(e^h; q, p) = \frac{\bar{e}(q,p) - e^l}{e^h - e^l} = \frac{\left[\frac{1-\theta}{\theta} d(q,p) \right]^{\theta-1} \frac{q}{A_q} \left(\bar{e}^{\text{avg}} \right)^{\theta} \left(i^{\text{avg}} \right)^{1-\theta} - e^l}{e^h - e^l}. \tag{27}$$

To determine the implications for the tuition schedule, denote by $b(q,p) = v(q,p,e^l)$ the per-student equilibrium revenue a college obtains from a low-ability student. Revenues and tuition for all types are then given by

$$\begin{aligned} v(q,p,e^l) &= b(q,p) && (28)\\ v(q,p,e^h) &= b(q,p) - d(q,p)(e^h - e^l) && (29)\\ t(q,p,e) &= v(q,p,e) - s(q,p,e) \qquad \text{for all } e. && (30) \end{aligned}$$

Since in equilibrium each college in operation makes zero profits, it follows that

$$\begin{aligned} 0 &= \eta(e^h)[v(q,p,e^h) - v(q,p,e^l)] + v(q,p,e^l) - i - \kappa(p) && (31)\\ 0 &= \frac{\bar{e} - e^l}{(e^h - e^l)}[-d(q,p)(e^h - e^l)] + b(q,p) - i(q,p) - \kappa(p) && (32) \end{aligned}$$

which gives as the total revenue of the college from low types (the sticker tuition including government subsidies for low ability students):

$$\begin{aligned} b(q,p) &= \kappa(p) + i(q,p) + \eta(e^h;q,p)d(q,p)(e^h - e^l) && (33)\\ &= \kappa(p) - e^l d(q,p) + \left(\frac{1}{1-\theta}\right)\left[\frac{1-\theta}{\theta}d(q,p)\right]^{\theta}\frac{q}{A_q}\left(\bar{e}^{\text{avg}}\right)^{\theta}\left(i^{\text{avg}}\right)^{1-\theta} && (34) \end{aligned}$$

That is, sticker tuition (including public subsidies) have to cover program specific exogenous per student costs $\kappa(p)$, the optimally chosen discretionary per-student spending $i(q,p)$ as well as the equilibrium discount given to the share $\eta(e^h)$ of high-ability attendees. College revenues and tuition for all types follow from (28) to (30), all as a function of the equilibrium ability discount $d(q,p)$. For future reference, denote the equilibrium measure of college quality and programs by $\chi(.)$, where $\chi(Q,p)$ is the mass of colleges with quality $q \in Q$ that offer program p. The measure of students of ability e^h admitted to program p with quality in Q is given by

$$\int_{q\in Q} \eta(e^h,q,p)d\chi(q,p)$$

From the supply side, the mass of students seeking to attend a college with quality in Q and program p given by

$$\sum_r \int \mathbf{1}_{\{q(e,r,b)\in Q, p(e,r,b)=p\}} d\Phi_0(e,r,b)$$

The discount function $d(q,p)$ adjusts in equilibrium such that for all (Q,p) and for both e these two quantities coincide.

The theoretical properties of the college market equilibrium — namely, independence of equilibrium tuition from students' initial resources, linearity of equilibrium tuition in student ability, and differences in equilibrium tuition schedules and student ability composition across programs ($p=1$ and $p=2$) — are derived in Section 3.1 and Appendix B.2.

2.6 Production and the Firm's Problem

Denote by $L(p)$ the aggregate labor efficiency units trained in a given program p (e.g., non-college, college non-stem, college STEM), aggregated across all households across initial wealth b, ability e, residence status r and that attended colleges of quality q. We distinguish between labor trained in different college programs because we posit that the extent to which labor of type $p \in \{0,1,2\}$ can be substituted with capital differs across p. Aggregate college labor of type $p \in \{1,2\}$ in efficiency units is given by

$$L(p) = \sum_e \int_q \int_a \sum_{j=1}^{j_r-1} \mathbb{1}_{\{q(e,r,b)=q, p(e,r,b)=p\}} \gamma_j(q,p) \Phi_j(e, dq, p, da) N_j. \tag{35}$$

Note that this formulation assumes that labor inputs of college graduates from different quality q colleges are perfect substitutes, but that graduates from higher quality colleges are more productive. Likewise, denote aggregate non-college ($p=0$) labor by

$$L(0) = \sum_e \int_a \sum_{j=1}^{j_r-1} \mathbb{1}_{\{q(e,r,b)=0, p(e,r,b)=0\}} \gamma_j(0,0) \Phi_j(e, 0, 0, da) N_j. \tag{36}$$

where $\Phi_j(e,0,0,.)$ is the asset probability distribution of individuals of age j and ability e that have chosen not to go to college $(q,p)=(0,0)$.

Denote by $K(p)$ the capital stock used by labor of type p and $\mathbf{K}$ the vector $(K(0),...K(P))$. Total output in the economy is then produced according to the aggregate constant returns to scale production function

$$Y = F(\mathbf{K}, \mathbf{L}) = \left[\sum_{p \in \{0,\mathcal{P}\}} \lambda(p) y (\Psi K(p), \Gamma(p) L(p))^{\nu} \right]^{\frac{1}{\nu}} \tag{37}$$

$$= \sum_{p \in \{0,\mathcal{P}\}} \left[\lambda(p) \left(\alpha(p) (\Psi K(p))^{\rho(p)} + (1-\alpha(p)) (\Gamma(p) L(p))^{\rho(p)} \right)^{\frac{\nu}{\rho(p)}} \right]^{\frac{1}{\nu}}. \tag{38}$$

where $y(\Psi K(p), \Gamma(p)L(p))$ is output produced in the "sector" that uses labor from program p, the term $\Gamma(p)$ captures labor-augmenting technological progress (which may differ across p) and the factor Ψ captures capital-augmenting technological change (assumed to be uniform across all capital stocks $K(p)$). The weights $\lambda(p)$ with $\sum_p \lambda(p) = 1$ determine the relative importance (or productivity) of the p-specific sectors, and the elasticity of substitution $\frac{1}{1-\nu}$ between output produced with different labors is determined by the parameter ν. Within each p, production is a CES between capital and labor, with weights $\alpha(p)$ on capital and p-specific substitution elasticity between labor and capital $\sigma(p) = \frac{1}{1-\rho(p)}$. We will calibrate the model such that $\sigma(0) > \sigma(1) > \sigma(2)$, so that the highest ranked college program ($p = 2$, STEM-trained college graduates) is most complementary to capital, and non-college labor is most substitutable with capital. These substitution elasticities will determine the direction and the size of relative and absolute wage effects of changes in the supply of labor induced by changes in demographics and/or education policy or by changes in labor and/or capital productivity designed to mimic potentially skill-biased technological change.

The production technology (38) is operated by a representative firm that operates in perfectly competitive factor and output markets, hires labor efficiency units of the different types p for type-specific wages $w(p)$, rents capital K_t at a rental rate r_t and also decides how to allocate the capital stock across different sectors so that $\sum_p K_t(p) = K_t$. The capital stock depreciates at rate δ. Note that since the firm can freely allocate capital across all uses p, the marginal products across all $K_t(p)$ are equalized and equal to the rental rate plus depreciation, $r_t + \delta$.

Since labor is differentiated by p and outputs across "industries" are not perfect substitutes, there are 3 separate labor markets and associated wages for p-specific labor efficiency units. The firm optimality conditions (as well as the algorithm to compute equilibrium factor prices r_t and $w_t(p)$ from the capital and labor market equilibrium conditions) are provided in equations (127) and (128) in the computational appendix, Section E.3.

2.7 Government

The government provides subsidies $s(e, r; q, p)$ to colleges, proportional to sticker tuition $t(e, r; q, p)$.[5] In addition, the government finances need-based financial aid $\varsigma(b)$, which captures federal and state grants and is specified in detail in Section 4.

The government balances in each period both the general tax-and-transfer system budget and the pension system budget. Expenditures within the general tax-and-transfer system consist of an exogenous stream of non-education-related spending and an endogenous stream

[5] As in Cai and Heathcote (2022), we assume that public subsidies are provided only to in-state students.

of education-related spending, as summarized above. Revenues are generated from taxes on consumption, capital income, and labor income.

We assume that the government has to balance the following per-period budget:

$$\begin{aligned} & \sum_r \sum_e \int_b \mathbb{1}_{\{q(e,r,b)=q, p(e,r,b)=p\}} \left(s(e,r;q,p) + \varsigma(b) t(e,r;q,p) \right) \Phi_0(e,r,db) N_0 \\ = \; & T^c + T^k + T - G \end{aligned} \tag{39}$$

where the left-hand side captures education-related expenditures (public subsidies $s(e,r;q,p)$ and financial aid $\varsigma(b)$), the focus of our analysis. On the net revenue side, T^c denotes consumption tax revenue, T^k denotes capital income tax revenue (including revenue from taxing accidental bequests of deceased households), T denotes labor income tax revenue:

$$\begin{aligned} T^c &= \tau_c \sum_r \sum_e \int_b \int_q \sum_p c_0(e,r,b,q,p) \bar{\Phi}_0(e,r,db,dq,p) N_0 \\ &\quad + \tau_c \sum_{j=1}^{J} \sum_e \int_q \sum_p \int_a c_j(e,q,p,a) \Phi_j(e,dq,p,da) N_j \\ T^k &= r \tau_k \sum_{j=1}^{J} \sum_e \int_q \sum_p \int_a a_j(e,q,p,a) \Phi_j(e,dq,p,da) N_j \\ &\quad + \sum_{j_r}^{J} \sum_e \int_q \sum_p \int_a (1-\psi(j)) a_j(e,q,p,a) \Phi_j(e,dq,p,da) N_j \\ T &= \sum_r \sum_e \int_b \int_q \sum_p T(y_0(b,e,r,q,p)) \bar{\Phi}_0(e,r,db,dq,p) N_0 \\ &\quad + \sum_{j=1}^{J} \sum_e \int_q \sum_p \int_a T(y_j(e,q,p,a)) \Phi_j(e,q,p,a) N_j \end{aligned}$$

and G stands for exogenous non-education related spending G. The pension system budget constraint is also balanced period by period, with benefits paid to current retirees equaling contributions collected from current workers.

2.8 Equilibrium Definition: College Market and Government

A steady state equilibrium is a measure of colleges $\chi(Q,p)$, college tuition $t(b,e,r;q,p)$, average ability of admitted students $\bar{e}(q,p)$, per student college expenditures $i(q,p)$ and profit $\pi(q,p)$ functions of colleges, household choices $c_0(b,e,r)$ and $\{c_j(e,q,a;p)\}_{j=1}^J$, $a_1(b,e,r)$ and $\{a_{j+1}(e,q,p,a)\}_{j=1}^J$, and $q(e,r,b)$ and $p(e,r,b)$, and government policies that satisfy the following conditions:

1. Given college tuition $t(e, r, b; q, p)$ and government policies (subsidies to colleges, student financial aid and taxes), the household consumption $c_0(e, r, b, q, p)$ and $c_j(e, q, p, a)$, savings $a_1(e, r, b, q, p)$ and $a_{j+1}(e, q, p, a)$ and college quality $q(e, r, b)$ and college major $p(e, r, b)$ choices solve the household optimization problem for all (e, r, b).

2. For all $q > 0$ in the feasible college quality set Ω and all $p \in \mathcal{P}$, the maximal revenue function is determined as $v(q, p, e) = \max_{b,r}\{t(e, r, b; q, p) + s(e, r; q, p)\}$, the college input choices $\bar{e}(q, p)$ and $i(q, p)$ solve the college's profit maximization problem, and $\pi(q, p)$ is the associated profit per student.

3. Zero profits: For all college quality subsets $Q \subset \Omega \setminus \{0\}$ and all $p \in \mathcal{P}$ we have $\int_Q \pi(q, p) d\chi(q, p) = 0$, and $\pi(q, p) \leq 0$ for all $q \in Q$ and $p \in \mathcal{P}$.

4. Goods market clearing:

$$C + \delta K + G + I^{college} + FC^{college} = Y \tag{40}$$

   where $I^{college}$ denotes aggregate expenditure of colleges on students, $FC^{college}$ is the total fixed cost paid by colleges, and E denotes the aggregate financial aid provided to students by the government.[6]

5. College market clearing: for all $e \in \mathcal{E}$, all $Q \subseteq \Omega \setminus \{0\}$ and all $p \in \mathcal{P}$

$$\sum_r \int \mathbf{1}_{\{q(e,r,b) \in Q, p(e,r,b)=p\}} d\Phi_0(e, r, b) = \int_{q \in Q} \eta(e, q, p) d\chi(q, p) \tag{41}$$

   where

$$(q(e, r, b) = q^*, p(e, r, b) = p^*) \;\Rightarrow\; (b, r) \in \arg\max\{t(e, r, b; q^*, p^*) + s(e, r; q^*, p^*)\} \tag{42}$$

6. The per-period government budget constraint in equation (39) holds in every period.

[6] Written out explicitly, aggregate consumption is given by

$$C = \sum_r \sum_e \int_b \int_q \sum_p c_0(e, r, b, q, p) \bar{\Phi}_0(e, r, db, dq, p) N_0 + \sum_{j=1}^{J} \sum_e \int_q \sum_p \int_a c_j(e, q, p, a) \Phi_j(e, dq, p, da) N_j$$

3 Theory, Computation and Thought Experiments

3.1 Theoretical Results

The college market equilibrium in our model inherits two fundamental properties first established in Cai and Heathcote (2022): (i) equilibrium tuition is independent of students' initial household resources b, and (ii) equilibrium tuition is linear in student ability e. As shown in Appendix B.2, both properties carry over to our setting in which college is an investment good and colleges offer distinct programs $p \in \mathcal{P}$. Building on these properties, the tuition expressions in Section 2.5 (equations (28)–(30)) immediately imply the following results, which are specific to our setting and concern differences in equilibrium tuition schedules and student ability composition across programs ($p = 1$ and $p = 2$).

Proposition 1. *Conditional on the equilibrium quality discount function $d(q, p)$, and for a given college quality q the difference in tuition for STEM programs ($p = 2$) and non-STEM programs ($p = 1$) is given by*

$$\begin{aligned} t(q,2,e) - t(q,1,e) &= \kappa(2) - \kappa(1) - [s(q,2,e) - s(q,1,e)] - [d(q,2) - d(q,1)]\, e^l \\ &+ \left(\frac{1}{1-\theta}\right)\left[\frac{1-\theta}{\theta}\right]^{\theta} \frac{q}{A_q} \left(\bar{e}^{\text{avg}}\right)^{\theta} \left(i^{\text{avg}}\right)^{1-\theta} \left[d(q,2)^{\theta} - d(q,1)^{\theta}\right] \end{aligned}$$

1. *If $d(q, p) = d(q)$, that is, the ability discount is the same across different programs, then for all $q > 0$ and all e*

$$t(q,2,e) - t(q,1,e) = \kappa(2) - \kappa(1) - [s(q,2,e) - s(q,1,e)] \tag{43}$$

   and for schools of a fixed quality q, tuition cost differences by program fully reflect cost differences and differential subsidies, and optimal quality composition $\bar{e}$ and expenditures i are independent of p, for all q.

2. *Now suppose that for fixed quality q, the quality discount is strictly decreasing in p. Then, for a given q STEM programs have a better student quality composition $\bar{e}$ and lower discretionary spending i (but might still have higher cost $\kappa(p) + i$). Tuition differences between programs might be higher or lower than cost differences, and tuition for STEM programs ($p = 2$) might actually be lower than for non-STEM program of the same quality q for some e (possibly even for both e).*

Proposition 2. *For a given college quality q average student quality $\bar{e}$ is strictly increasing in p, that is, strictly higher in STEM colleges than in non-STEM colleges.*

Proof. We want to show that in equilibrium $\bar{e}(q,2) > \bar{e}(q,1)$. From equation (26), this is equivalent to showing that the equilibrium tuition discount satisfies $d(q,2) < d(q,1)$. Suppose not, that is, suppose that $d(q,2) = d(q,1) = d(q)$; if the inequality is reversed, then the argument goes through unchanged. If $d(.)$ is independent of p for fixed q, then from (25) and (26), both programs set identical optimal inputs: $i(q,2) = i(q,1)$ and $\bar{e}(q,2) = \bar{e}(q,1)$. By case 1 of Proposition 1, tuition differences reduce to $t(q,2,e) - t(q,1,e) = \kappa(2) - \kappa(1) - [s(q,2,e) - s(q,1,e)]$, which is independent of e. Since $\nu_2(e)$ is strictly increasing in ability e, high-ability students have a strictly stronger preference for STEM relative to non-STEM than low-ability students. The market clearing condition (41) then requires STEM programs to enroll a strictly higher fraction of high-ability students, i.e., $\bar{e}(q,2) > \bar{e}(q,1)$ – a contradiction. We therefore conclude $d(q,2) < d(q,1)$, and the result follows from (26). □

3.2 Sources of Inefficiency of the Equilibrium and the Scope for Higher Education Subsidies

In our framework, several sources of inefficiency could potentially call for government intervention, and higher education subsidies specifically. First, the presence of borrowing constraints prevent households from smoothing consumption and investment optimally across periods. Although we permit households to borrow to attend college, the potential risk of dropping out of college as well as the risk of post-college low productivity might push households into the borrowing constraint later in life, in turn inducing suboptimal (relative to a complete financial markets benchmark) higher education choice (both at the extensive margin – whether to go to college – and along the intensive margin – what quality college to go to). This friction is especially relevant for low-resource households.

Second, households face uninsurable idiosyncratic risks - most notably, the risk of failing to complete college and the risk of skill obsolescence later in life (e.g., in STEM occupations). In the absence of private insurance markets for such shocks, households must bear them individually, potentially leading to precautionary behavior and inefficient underinvestment in education. Reducing the private costs of attending college through public subsidies mitigates this risk (but of course also induces a moral hazard issue: some individuals might attend college even though private returns are low since the government pays for part of the cost).

Third, government fiscal policies interact with these frictions. In our incomplete markets model, attending a (high-quality) college and/or studying a high potential wage major turns an individual into a high tax payer, a fiscal externality that the individual does not take into account when making higher education decisions. More broadly, subsidies to colleges, financial aid for students, and the progressive labor income tax code might positively impact the distribution of lifetime utilities across individuals. At the same time, these interventions

may also distort individual incentives, for instance by affecting equilibrium tuition in the college market or altering the returns to education through the tax-and-transfer system.

3.3 College Market: Mechanisms and Intuition

This subsection summarizes the mechanisms governing tuition schedules and student composition in the college market. To isolate the role of ability composition, which in equilibrium is jointly determined by student demand and college supply, we consider comparative statics within a given college cell (q, p), holding college quality fixed. Thus, the results in Section 5 should be interpreted as partial-equilibrium responses of tuition schedules to changes in the share of high-ability students, rather than as comparative statics of the full equilibrium in which the distribution of college quality itself adjusts.

College market equilibrium. The college-side optimality conditions imply that the equilibrium ability discount satisfies

$$d(q,p) = \frac{\theta}{1-\theta}\frac{i(q,p)}{\bar{e}(q,p)}, \tag{44}$$

where $\bar{e}(q,p)$ denotes average student ability at colleges of type (q,p). With two ability types e^l and e^h, average ability is

$$\bar{e}(q,p) = e^l + \eta(q,p)\Delta e, \tag{45}$$

where $\Delta e = e^h - e^l > 0$ and $\eta(q,p)$ denotes the share of high-ability students attending colleges of type (q,p).

Let $b(q,p)$ denote total revenue obtained from a low-ability student:

$$b(q,p) = \kappa(p) + i(q,p) + \eta(q,p)d(q,p)(e^h - e^l). \tag{46}$$

Revenues by ability are therefore

$$v(q,p,e^l) = b(q,p), \tag{47}$$

$$v(q,p,e^h) = b(q,p) - d(q,p)(e^h - e^l). \tag{48}$$

Substituting into net tuition $t(q,p,e) = v(q,p,e) - s(q,p,e)$ yields

$$\begin{aligned} t(q,p,e^l) &= \kappa(p) - s(q,p,e^l) + i(q,p) \\ &\quad + \eta(q,p)\, d(q,p)\, \Delta e, \end{aligned} \tag{49}$$

$$\begin{aligned} t(q,p,e^h) &= \kappa(p) - s(q,p,e^h) + i(q,p) \\ &\quad - (1-\eta(q,p))\, d(q,p)\, \Delta e. \end{aligned} \tag{50}$$

Comparative statics with respect to ability composition. An increase in the share of high-ability students raises average ability ($\partial \bar{e}/\partial \eta = \Delta e > 0$), which allows colleges to produce a given quality with less instructional spending ($\partial i/\partial \eta < 0$), which in turn reduces the equilibrium ability discount ($\partial d/\partial \eta < 0$). Thus, a larger share of high-ability students increases average ability but lowers both instructional spending and the equilibrium ability discount.

Implications for tuition schedules. Differentiating tuition with respect to $\eta(q,p)$ gives

$$\frac{\partial t(q,p,e^l)}{\partial \eta} = \frac{\partial i(q,p)}{\partial \eta} + \Delta e \left(d(q,p) + \eta(q,p) \frac{\partial d(q,p)}{\partial \eta} \right), \tag{51}$$

$$\frac{\partial t(q,p,e^h)}{\partial \eta} = \frac{\partial i(q,p)}{\partial \eta} + \Delta e \left(d(q,p) - (1-\eta(q,p)) \frac{\partial d(q,p)}{\partial \eta} \right). \tag{52}$$

The effect of ability composition on tuition operates through its impact on average ability. When $\eta(q,p)$ increases, average ability rises, instructional spending falls, and the equilibrium ability discount shrinks.

For low-ability students both channels reduce tuition, implying $\partial t(q,p,e^l)/\partial \eta < 0$. For high-ability students the two effects work in opposite directions: lower instructional spending reduces tuition, while the reduction in the ability discount increases it. The net effect is therefore ambiguous.

However, the difference between the tuition responses satisfies

$$\frac{\partial t(q,p,e^h)}{\partial \eta} - \frac{\partial t(q,p,e^l)}{\partial \eta} = -\Delta e \frac{\partial d(q,p)}{\partial \eta} > 0. \tag{53}$$

Hence, changes in ability composition shift low-ability tuition downward relative to high-ability tuition.

Summary. Three key comparative statics follow from the model:

1. A higher share of high-ability students raises average ability and lowers both instructional spending and the ability discount.

2. When ability composition improves, tuition for low-ability students declines unambiguously, while the response of high-ability tuition is ambiguous.

3. Holding subsidies fixed, an increase in the share of high-ability students reduces the tuition gap between ability types.

These comparative statics provide the intuition for the equilibrium responses of the college market in the quantitative experiments in Section 5.

3.4 Computational Solution of the College Market Equilibrium: College as Investment, Major Heterogeneity, and Credit Frictions

Our treatment of college as an investment good has important implications for both the equilibrium properties of the college market and the computational solution strategy, and it marks a key departure from Cai and Heathcote (2022), whose benchmark model treats college as a consumption good.

Cai and Heathcote (2022) establish two polar results for the investment-good version of their model. First, with a *frictionless* credit, the competitive equilibrium features *no* sorting by income: only ability determines college enrollment (their Proposition 7). Second, when credit frictions are *severe*, the investment model collapses to the consumption model (their Proposition 8), which delivers perfect sorting by income conditional on ability. In the consumption-good (or severe-friction) case, the equilibrium can be solved by a simple and efficient algorithm that exploits the existence of resource thresholds $b^*(q, e)$: for each college quality q and ability level e, there is a cutoff initial wealth above which households enroll, and the college market clears by establishing a mapping from quality q to these wealth thresholds.

Our model occupies the intermediate case: credit market frictions are present but not severe, and college is an investment good. In this intermediate regime neither polar result applies: income sorting is partial, the monotone threshold structure breaks down, and the equilibrium cannot be solved by a simple pass from quality q to income cutoffs. To our knowledge, no existing paper has numerically solved a college market model in this intermediate credit friction case of the investment-good model. Our computational approach is general and efficient enough to handle rich models of the education market that allow for empirically plausible credit frictions, and developing it constitutes an important methodological contribution of the paper. The solution algorithms are described in detail in Appendix E, which contains two algorithms: one for a discrete college quality grid (accommodating a large number of quality levels) and one for a continuous college quality distribution.

3.5 Thought Experiments

To study how demographic forces and education policy interact with the endogenous college market, we conduct three thought experiments. The first experiment isolates the effect of changes in public funding for higher education. The second examines the consequences of demographic transition. The third combines both forces.

Experiment I: Cut of Public Subsidy to Colleges. In the first experiment we remove public subsidies to colleges. This reform raises the net tuition faced by students and allows us to study how the college market adjusts when higher education becomes more expensive. In equilibrium, tuition schedules, enrollment, and the composition of students across college quality and programs adjust endogenously. These first-order changes in college demand and pricing interact with general equilibrium wages and the interest rate and affect public finances through changes in the tax base and the labor income tax rate required to balance the government budget.

Experiment II: Demographic Cliff. In the second experiment we study the consequences of a demographic transition characterized by a persistent decline in fertility. We model this as a transition from an initial steady state with positive population growth to a final steady state with zero population growth. The resulting demographic shift leads to a smaller and older population and therefore alters equilibrium wages and the interest rate through changes in labor supply and capital accumulation. In addition, demographic aging puts pressure on the public pension system. The Social Security budget constraint must be balanced in every period, and in our baseline specification the adjustment occurs through a decline in the pension replacement rate. These general equilibrium adjustments - through both factor prices and pension benefits - affect incentives to invest in higher education and thereby change both college enrollment and the composition of students across college quality and programs. As a result, tuition schedules and ability discounts adjust endogenously in the college market, and the demographic transition also affects the government budget through its impact on the tax base and public finances.

Experiment III: Demographic Cliff with Subsidy Cut. Finally, we combine the two forces by analyzing a demographic transition in an environment in which public subsidies to higher education are removed. This experiment allows us to study how demographic change and fiscal pressures jointly affect equilibrium wages, tuition schedules, and the allocation of students across college quality and majors.

Details on the calibration of the demographic transition and the construction of the population paths used in these experiments are provided in Section 4.1.

4 Calibration

Each period in the model corresponds to four years. Households enter economic life at model age $j_a = 0$ (biological age 18). College education lasts one model period, so all households enter the labor market full time at model age $j_c = 1$ (biological age 22). The statutory retirement age is $j_r = 12$ (biological age 66), and the maximum lifespan is $J = 20$ (biological age ≈101). We calibrate the initial steady state of the model to broadly approximate the U.S. economy in the 2010s.

4.1 Demographic Dynamics

We calibrate the survival probabilities $\{\psi_j\}_{j=0}^{J}$ as the weighted average of female and male survival rates of the population of age 18 and older in year 2025 taken from the Human Mortality Database (HMD). To consistently model the demographic distribution along the transition of the economy to the final steady state, we introduce the notion of age-specific fertility rates into the demographic model, which we calibrate from the year 2025 age-specific fertility rate distribution taken from the United Nations World Population Prospects Database.

To link the population growth rate to the empirical distributions of fertility and survival rates, let the age specific fertility rates in our unisex model be $\bar{\zeta} \cdot \zeta_j$, where we normalize $\sum_j \zeta_j = 1$ so that $\bar{\zeta}$ represents the total fertility rate (TFR). Also note that $\zeta_j > 0$ for $j \in \{\underline{j}, \ldots, \bar{j}\}$ and $\zeta_j = 0$, otherwise. The newborn population in $t+1$ is then

$$N_{t+1,0} = \bar{\zeta} \sum_{j=0}^{J} \zeta_j N_{t,j}. \tag{54}$$

For given $\{\zeta_j, \psi_j\}_{j=0}^{J}$, we get from (1) and (54) the following relationship between the TFR and the population growth rate n follows, see Appendix C.1:

$$\bar{\zeta} = \frac{1}{\sum_{j=0}^{J} \zeta_j \frac{1}{(1+n)^{1+j}} \prod_{i=0}^{j-1} \psi_i}. \tag{55}$$

Observe that for $n = 0$ and $\psi_j = 1, \forall j$, $\bar{\zeta} = 1$. For $\psi_j < 1$ for at least one $j \in \{\underline{j}, \ldots, \bar{j}\}$ we however require $\bar{\zeta} > 1$.

We set the initial steady state population growth rate to the average US adult population growth rate of pre-covid years 2000 to 2018 of $n_{iss}\% = 1\%$, which from equation (55) and given our calibration of $\{\zeta_j, \psi_j\}_{j=0}^{J}$ requires a TFR of $\zeta_{iss} = 1.1427$. The final steady state total fertility rate consistent with a population growth rate of $n_{fss}\% = 0\%$ is $\zeta_{fss} = 1.0107$. We impose this lower fertility rate from year 2008 onwards, which induces the transitional dynamics in the low population growth scenario of Figure 3.

Table 1: Calibration Parameters: First Stage

Symbol	Description	Value	Target/Source
	Demographics		
j_a, j_c, j_r, J	Entry, college exit, retirement, max. age	0, 1, 12, 20	Ages 18, 22, 66, 101
$\{\psi_j\}$	Survival probabilities	—	HMD life tables
	Ability and programs		
E, P	No. of ability types; no. of programs	2; 2	[low, high]; [non-STEM, STEM]
e^h, e^l	High and low ability levels	1.0; 0.375	Cai and Heathcote (2022)
$\mu(e^h)$	Population share, high ability	0.5	Cai and Heathcote (2022)
	College quality production		
A_q	TFP, quality production function	1.77	Mean instructional expenditure
θ	Ability elasticity	0.5	Normalization, Cai and Heathcote (2022)
$\bar{\kappa}$	Avg. college fixed cost	≈\$9,000	Altonji and Zimmerman (2018)
$\kappa(2)/\kappa(1)$	STEM/non-STEM fixed cost ratio	1.71	Altonji and Zimmerman (2018)
	Household preferences, endowments and productivity		
σ, φ	Risk aversion; Frisch elasticity	1; 0.6	—
ℓ, ℓ^{stud}	Avg. and student hours worked	1/3; $\ell/4$	—; NCES
ω	Earnings opportunity cost of college (fraction of y_0 foregone)	20 wks	Cai and Heathcote (2022)
$\bar{b}$	Mean initial resources	\$61,200	Avg. IVT transfer, PSID 2013
ϱ	Resource ratio, high/low ability	1.35	Cai and Heathcote (2022)
σ^b	Variance, initial resource dist.	0.55	Cai and Heathcote (2022)
$\mu(r^i)$	Share of in-state students	0.529	Cai and Heathcote (2022)
π_p	STEM obsolescence prob. (per period)	0.25	Deming and Noray (2020)
$\{\gamma_j(s{=}0)\}$	Age-productivity profile, non-college	—	PSID 1967–2011
$\{\gamma_j(s{=}1)\}$	Age-productivity profile, college	—	PSID 1967–2011
$\underline{a}_1$	College borrowing limit (1 model period = 4 yrs)	\$45,000	Krueger and Ludwig (2016) (\$11,397/yr)
	Aggregate production		
ν	Outer CES elasticity	1	Perfect substitutes
α, δ	Capital share; depreciation rate	33.3%; 5%	—
$\rho(p{=}1)$, $\rho(p{=}2)$	Capital-labor subst. elas., by program	−0.11; −0.43	Caunedo et al. (2023)
$\rho(p{=}0)$	Capital-labor subst. elas., non-college	0.33	Caunedo et al. (2023)
	Government		
$\bar{s}$	Avg. per-student college subsidy rate (in-state)	49%	Cai and Heathcote (2022)
ς_0	Financial aid fraction of sticker tuition	0.21	Cai and Heathcote (2022)
ς_1	Financial aid cap (avg. Pell Grant + state aid)	\$6,870	Cai and Heathcote (2022)
τ_c, τ_k	Consumption and capital income tax rates	5%; 36%	Legislation; Trabandt and Uhlig (2011)
ξ	Labor income tax progressivity	0.18	Heathcote et al. (2017)
τ^p	Social Security payroll tax	12.4%	Legislation
G/Y	Non-education govt. spending/GDP	13.8%	Current value

Notes: First-stage (exogenously set) parameters. HMD = Human Mortality Database; IVT = inter vivos transfers.

Table 2: Calibration Parameters: Second Stage

Symbol	Description	Value	Target/Source
	Household preferences and endowments		
β	Discount factor	0.9882	Interest rate
ν_0	Psych. cost of college attendance	-2.64	College enrollment rate
ν_1	Taste for college quality (warm-glow motive)	0.25	Avg. annual net tuition per student (NCES 2000–2019)
$\nu_2(e)$	Ability-specific STEM disutility	$[-1.65, -1.20]$	STEM shares by ability
ς_d	Slope of graduation rate in quality; $1-\pi_d(q) = \varsigma_d \cdot q$	0.293	Average college dropout rate
β_1, β_2, β_3	Parameters of $\gamma(q,p)$, see (5)	0.55; 0.30; 1.5	Avg. college wage premium; avg. STEM wage premium; earnings gap, top/bottom q quartile
	Aggregate production		
Γ	Wage normalization	1.34	$w(q{=}0) = 1$
	Government		
$\underline{b}_1$	Financial aid eligibility threshold	$\approx 0.7\,\bar{y}$	Share of enrolled students receiving aid
τ_ℓ	Labor income tax level parameter	0.226	Marginal labor income tax rate at avg. earnings

Notes: Second-stage (endogenously calibrated) parameters.

Table 3: Targeted Moments: Empirical Values and Data Sources

Moment	Emp. Value	Data Source
Interest rate (annual)	3.5%	—
College enrollment rate	51%	PSID 2011–2017
Avg. annual net tuition per student (2010 USD, all college types)	\$15,000	NCES 2000–2019
Avg. STEM share among college graduates	24%	NSF 2024 report
STEM share by ability group	—	No direct target; plausible ability gradient
Avg. college wage premium	80%	PSID 2011–2017
College dropout wage premium (rel. to HS graduates)	20%	PSID 2011–2017
Average college dropout rate	50%	NLSY97
Earnings difference between top and bottom quartile of q distribution	98%	Leukhina (2023)
Avg. STEM wage premium	57%	PIAAC 2011 & 2012
Share of enrolled students receiving financial aid	32%	Cai and Heathcote (2022)
Marginal labor income tax rate at avg. earnings	29%	—

Notes: Targeted moments, empirical values, and data sources.

4.1.1 Demographics Dynamics under Demographic Cliff

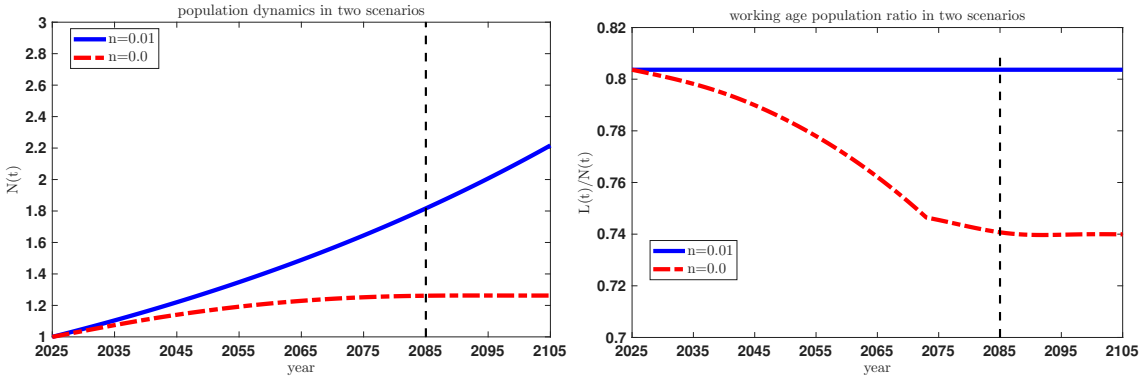
We now describe how the demographic transition used in the thought experiments II and III is constructed.

We define the demographic cliff as the drop of the fertility rate in 2008 shown in Figure 1. We focus on the adult population. Those born in year 2007, immediately before the drop of fertility, are of age 18 in year 2025 and thus year 2025 constitutes the base year of our analysis. The year 2025 "benchmark" scenario features a population growth rate of $n_{iss}\% = 1\%$ annually, which we calibrate on the basis of past data, see Section 4.1. We then assume a demographic transition such that in the long run the population converges to a final steady state population distribution with a zero population growth rate, $n_{fss} = 0$, see Figure 1, which shows that the population of age 18 stabilizes in the long-run around 2080 (and so does the total population). We model this transition through a drop in fertility while keeping mortality rates constant at their year 2025 values.

Panel (a) of Figure 3 provides a graphical illustration of the aggregate population transitions, where the blue line represents the dynamics of the adult population in the initial steady state high population growth scenario, in which we normalize the size of the adult population to 1 in year 2025. The red dashed line represents the population in the demographic cliff (low population growth) scenario. In our main experiment we look at a horizon of 60 years and therefore evaluate the drop of the population between the high and the low population growth scenario in year 2085, when the demographic transition is close to a steady state, see Panel (b) of the figure, which shows that the working age to adult population ratio has almost stabilized in 2085. With these assumptions, the size of the population in the low growth population scenario (red dashed line) in 2085 is about 69.44% of the population size achieved in 2085 in the high population growth scenario. Together with the drop of the population growth rate from 1% to 0% these are the three calibration parameters for our demographic cliff experiment of our quantitative steady state comparison, in which we additionally approximate the year 2085 population dynamics in the low population growth scenario by assuming a zero population growth steady state. As the red dashed line of Panel (b) further illustrates, by 2085 the working age to adult population ratio will have dropped from 80.36% to 74.07%, a reduction by 6%p, which shows how the reduction of fertility induces a relative shortage of labor in the economy.

The transition experiment of Section 5.7 requires the full population path rather than the two stylized steady states just described. Figure 4 shows that input, for the same two scenarios and over the entire horizon of the transition. The baseline holds fertility at its pre-decline level, while the cliff scenario carries the observed decline in fertility since 2008 forward; the two paths share the same history and the same realized mortality, and long-

Figure 3: Population Dynamics



(a) Population Size

(b) Working Age Population Ratio

run net migration is held at a constant rate that is common across scenarios. History is identical through 2026: the first cohorts born after the fertility decline reach age 18 in that year, and the adult population paths first differ in 2027. Since the model is solved on a four-year grid, 2030 is the first model period at which the two series differ. By the end of the horizon the adult population reaches about six times its 1950 level under the baseline against about three times under the cliff, so the long-run difference in population size is close to a factor of two. By our assumption of zero population growth in the long run, which we impose in all scenarios, the two scenarios converge to the same final age structure (the working-age population ratio settles at 0.660 in both). Along the way, however, the two age structures differ substantially. In 2085 the working-age ratio is 0.670 under the cliff against 0.724 under the baseline, old-age dependency 0.493 against 0.381, and the young-adult share 0.102 against 0.127. These levels differ from the ratios reported above for Figure 3 because the two figures are different objects: Figure 3 shows the stylized scenario used for the steady-state calibration, normalized to 2025, whereas Figure 4 shows the full population path initialized in 1950 and carried forward to the terminal age structure. The gaps peak in the early twenty-second century, at 5.6 and 12.0 percentage points for the working-age ratio and old-age dependency, and close only as both scenarios approach the common plateau. The two scenarios are those underlying Figures 1 and 3; here we show the forecast over a longer horizon, up to the terminal condition imposed on the projection.

We read the transition results at the focal year 2086, some sixty years after the first cohorts born after the fertility decline reach adulthood. By that date the difference between the two age structures has essentially reached its full size. The working-age ratio is 0.672 under the cliff against 0.726 under the baseline, a gap of 5.4 percentage points, and old-age

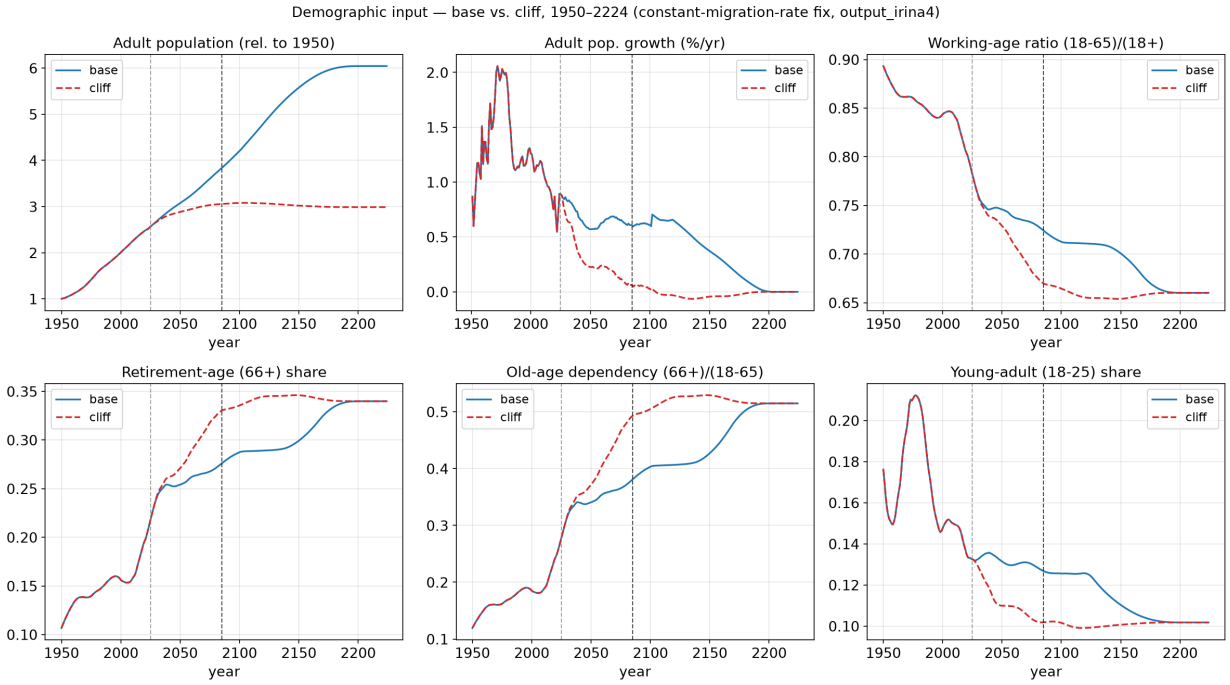


Figure 4: Demographic input, 1950–2224. Baseline (solid) versus demographic cliff (dashed). Panels: adult population relative to 1950; the annual growth rate of the adult population; the working-age population ratio (18–65)/(18+); the retirement-age (66+) share of adults; old-age dependency (66+)/(18–65); and the young-adult (18–25) share of adults. Dashed vertical lines mark 2025 and 2085. The two scenarios share the historical period and differ only in fertility after 2008; long-run net migration is held at a common constant rate, so both converge to the same terminal age structure.

dependency is 0.489 against 0.378, a gap of 11.1 percentage points; both are close to the largest values these gaps attain at any date along the projection, and the working-age gap widens by less than half a percentage point thereafter. The focal year is thus the earliest date at which the demographic burden that the college market has to absorb is realized in full, which is why we use it throughout when comparing a reform path with the baseline path.

4.2 Households

Preferences. Recall that the psychological cost of attending college depends on child ability and is specified as $\nu_0 + \nu_1 \ln(q) + \mathbb{1}_{p=2}\nu_2(e)$. The parameter ν_0 is calibrated to match the average college enrollment rate, while $\nu_2(e)$ governs the ability-specific shares of students in STEM majors. Finally, ν_1 controls the taste for college quality and is calibrated to match the average net tuition per student (net of grants and scholarships), yielding $\nu_1 = 0.25$. Qualitatively, this parameter captures an additional, non-pecuniary direct utility from college

quality beyond the labor market returns of attending a better (and more expensive) college. Households discount utility at rate β, which is calibrated such that the implied steady-state annual interest rate in general equilibrium equals 3.5%.

Endowments, Constraints and Costs. Our calibration of many elements of the endowment calibration closely follows Cai and Heathcote (2022). Accordingly, the share of in-state students in the population is set to 0.529. We assume two ability types, $e \in \{e^l, e^h\}$. The high-ability type is normalized to $e^h = 1$ and the low-ability type is set to $e^l = 0.375$. The population shares of the two ability types are set to $\mu(e^h) = \mu(e^l) = 0.5$.

The initial resource distribution $F_e(b)$ is log-normal with ability-specific means. The mean initial resources are set to $\bar{b} = \$61{,}200$, corresponding to the average inter-vivos transfer in the PSID (2013). The ratio of mean resources for high- to low-ability households is $\varrho \equiv \bar{b}(e^h)/\bar{b}(e^l) = 1.35$, calibrated to match the ratio of average family income for households in the top versus bottom half of the AFQT score distribution. The standard deviation of log initial resources within each ability group is $\sigma^b = 0.55$.

An alternative, richer calibration strategy - which we plan to implement in future work - constructs the ability distribution directly from SAT test score data. Specifically, denoting SAT composite scores by $h \in \mathcal{H} = \{h_1, \ldots, h_H\}$, we build the probability distribution function $\phi(h)$ by assigning the midpoints of SAT composite score bins to the corresponding percentile ranks, see https://blog.prepscholar.com/sat-percentiles-and-score-rankings. The lowest 1 percent of students is assigned an SAT score of $h_1 = 480$ (the midpoint of the 400–660 range), the next 1 percent gets $h_2 = 675$, and so forth, up to $h_H = 1585$ for the top percentile. We then map this fine score distribution into the ability grid $e \in \mathcal{E} = \{e_1, \ldots, e_E\}$, where $E < H$, as follows. First, we normalize the bounds of the support, setting $e_1 = h_1/h_H = 480/1585 \approx 0.3$ and $e_E = 1$. Second, the probability mass at the bounds, $\phi(h_1)$ and $\phi(h_H)$, is assigned to $\phi(e_1)$ and $\phi(e_E)$, respectively. Third, for each interior score $h_j \in \{h_2, \ldots, h_{H-1}\}$, we locate its two neighboring gridpoints e_{i-1} and e_i and distribute the mass $\phi(h_j)$ by linear interpolation: $\Delta\phi(e_{i-1}) = \left(1 - \frac{h_j - e_{i-1}}{e_i - e_{i-1}}\right) \cdot \phi(h_j)$ is added to e_{i-1} and the remainder to e_i.

The graduation rate $1 - \pi_d(q)$ is assumed to be linear in college quality: $1 - \pi_d(q) = \varsigma_d \cdot q$. The slope parameter $\varsigma_d = 0.293$ is calibrated to match an average college dropout rate of 50%, as implied by NLSY97. As an additional validation, the model-implied graduation gradient across college quality tiers is close to the empirical gradient by freshmen SAT scores in IPEDS (2015). After graduation, a worker with STEM degree ($p = 2$) transitions to $p = 1$ with probability $\pi_p = 0.25$ per model period, capturing the risk of STEM skill obsolescence over the working life. While in college, students cannot work full time and, in addition to incurring the psychological cost described above, must also forgo earnings. We assume that

students forgo 20 weeks of median weekly earnings of full-time workers aged 16 to 24 while enrolled.

Earnings differences by college quality and major/program of study. Following Hendricks et al. (2021), we classify colleges into four quality types. The lowest quality type consists of two-year community colleges offering transferable associate degrees. The other three types are defined based on the average SAT scores of their freshmen. Type 2 includes four-year institutions with freshmen's average SAT scores in the first tercile of the SAT distribution (least selective public and private colleges). Type 3 includes four-year institutions with scores in the second tercile (many flagship universities and directional schools). Finally, Type 4 includes the most selective four-year institutions with scores in the third tercile (Ivy League, selective private schools, most flagship universities, and many other selective public universities).

As described in Leukhina (2023), after controlling for family income and AFQT scores, graduates from Type 2, Type 3, and Type 4 colleges earn on average 65%, 75%, and 98% more, respectively, than graduates from Type 1 colleges.

Regarding earnings differences by college major, in calibrating $\gamma_j(q,p)$ we target that average earnings of STEM workers are about 57% higher than those of non-STEM workers, based on the Programme for the International Assessment of Adult Competencies (PIAAC, 2011–2012).

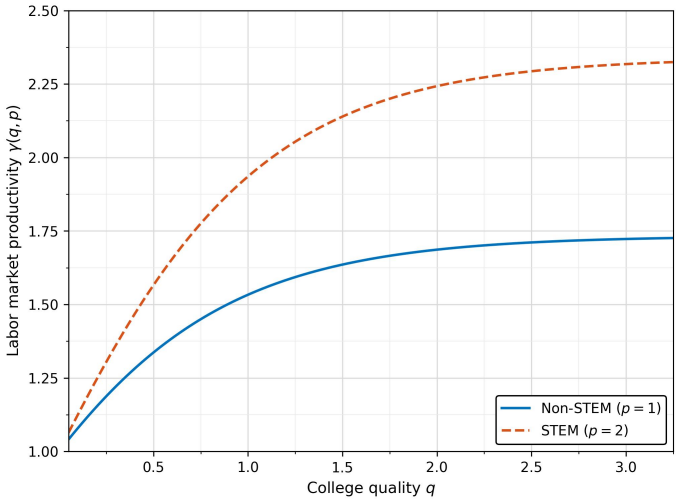
The productivity function $\gamma_j(q,p)$ is specified in equation (5). The three parameters β_1, β_2, and β_3 are endogenously calibrated to match the average college wage premium of 80% (PSID, 2011–2017), the average STEM wage premium, and the wage gap between colleges in the bottom and top quartiles of the q distribution, respectively. We obtain $\beta_1 = 0.55$, $\beta_2 = 0.30$, and $\beta_3 = 1.5$, where β_3 enters the saturation function as $f(q) = 1-\exp(-\beta_3 q)$. Figure 5 illustrates the resulting $\gamma(q,p)$ productivity schedule. Additionally, the age-productivity profiles $\{\gamma_j(s)\}$ are calibrated to PSID earnings data (1967–2011) and are consistent with a wage premium for college dropouts of approximately 20% relative to high-school graduates (PSID, 2011–2017).

4.3 College Supply

The quality production function (19) involves three parameters: the TFP parameter A_q, the ability elasticity θ, and the program-specific fixed costs $\kappa(p)$. We calibrate them as follows.

The TFP parameter A_q is calibrated to match the empirical mean of instructional expenditure per student, $\bar{i}^{\text{avg}}$. Following Cai and Heathcote (2022), the ability elasticity θ is currently set to 0.5 as a normalization. This parameter governs the strength of the club-good mechanism: it directly determines the per-unit ability discount $d(q,p) = \frac{\theta}{1-\theta}\frac{i(q,p)}{\bar{e}(q,p)}$ –

Figure 5: Labor Market Productivity



(a) labor productivity by quality q, conditional on major p

Notes: The figure plots labor market productivity $\gamma_j(q,p)$ as a function of college quality q for non-STEM ($p = 1$, solid line) and STEM graduates ($p = 2$, dashed line), as specified in equation (5). Productivity is normalized to 1 for non-college workers ($q = 0$). The saturation function $f(q) = 1 - \exp(-\beta_3 q)$ implies diminishing returns to college quality. Calibrated parameters: $\beta_1 = 0.55$, $\beta_2 = 0.30$, $\beta_3 = 1.5$.

the tuition advantage that accrues to higher-ability students. In future work we plan to calibrate θ to match the ratio of sticker tuition charged by colleges in the top versus the bottom decile of the quality distribution, and to identify ability dispersion $e^h - e^l$ separately from the average generosity of institutional merit-based financial aid observed in the data.

Based on Altonji and Zimmerman (2018), the average fixed cost that colleges bear independently of quality is approximately \$9,000 per student, and the cost of producing graduates in STEM programs is approximately 1.71 times higher than in non-STEM programs, so that $\kappa(2)/\kappa(1) = 1.71$.

4.4 Aggregate Production

We parameterize the aggregate production function specified in equation (38) as follows. In the current quantitative implementation we assume perfect substitutability across labor types at the outer level, i.e. $\nu = 1$; the full derivation for the general CES case is provided in the appendix. The parameters $\rho(p = 0)$, $\rho(p = 1)$ and $\rho(p = 2)$ governing the skill- and major-specific substitution elasticities between capital and labor are based on Caunedo et al. (2023) estimates and are set to 0.33, -0.11 and -0.43, respectively.

Non-college wages are normalized to $w(0) = 1$. We achieve this by scaling Γ using (129) for $p = 0$ to get

$$\Gamma = \frac{\alpha(0)}{1-\alpha(0)} \cdot \frac{k_t(0)^{\rho(0)-1}}{r_t + \delta}$$

4.5 Government

College subsidies, student financial aid and student loans. Following Cai and Heathcote (2022), the government subsidy $s(e, r; q, p)$ is provided only to in-state students ($r = r^i$) and is set proportional to sticker tuition at rate $\bar{s} = 49\%$, so that $s(e, r^i; q, p) = \bar{s} \cdot t(e, r^i; q, p)$ and $s(e, r^o; q, p) = 0$. We calibrate the student financial aid function $\varsigma(b)$ to capture the two main components of U.S. aid programs: (i) federal need-based grants, primarily Pell Grants, which do not scale with sticker tuition, and (ii) state-level aid programs, which generally do depend on tuition. To this end, we assume that financial aid is provided as a fixed fraction of sticker tuition, but subject to a dollar cap and a resource eligibility threshold:

$$\varsigma(b) = \mathbb{1}\{b \leq \underline{b}_1\} \cdot \min(\varsigma_1,\ \varsigma_0 \cdot t(e, r; q, p)), \tag{56}$$

where $t(e, r; q, p)$ denotes the sticker tuition faced by a student of ability e, residence status r, college quality q, and program p; ς_0 is the tuition fraction; and ς_1 is the dollar cap. Currently, we assume a single resource threshold $\underline{b}_1$. All households with family resources below this threshold qualify for financial aid, up to a cap of $\varsigma_1 = \$6{,}870$, which corresponds to the average Pell Grant and state-level financial aid reported in Cai and Heathcote (2022). Households with resources above $\underline{b}_1$ do not qualify for aid. The threshold $\underline{b}_1$ is endogenously calibrated such that, in equilibrium, 32% of enrolled students receive financial aid.

Following Krueger and Ludwig (2016), the annual limit on publicly provided student loans is \$11,397, corresponding to a model-period (4-year) borrowing limit of $\underline{a}_1 = \$45{,}000$.

Taxes and (non-education) transfers. The consumption tax rate is set at 5% (see Mendoza et al. (1994)), and the capital income tax rate is fixed at 36% following Trabandt and Uhlig (2011). The progressive labor income tax code is approximated using a two-parameter tax function as in Heathcote et al. (2017),

$$T(y) = y - \tau_\ell y^{1-\xi},$$

where the progressivity parameter ξ is set to 0.18 - an average estimate across demographic groups in Heathcote et al. (2017) - and the level parameter τ_ℓ is endogenously calibrated so

that the implied marginal labor income tax rate at average earnings in the model matches its empirical counterpart.

5 Results

This section studies how demographic forces and public education policy reshape the college market and interact with fiscal policy in general equilibrium. Within the college market, an important equilibrium mechanism operates through the ability composition of the student body. Changes in enrollment alter the share of high-ability students at each college type, which in turn affects instructional spending and the equilibrium ability discount that colleges offer. Because tuition schedules depend directly on this discount, changes in student composition translate into endogenous adjustments in tuition and therefore into further shifts in college demand. Demographic changes affect the college market primarily through their impact on student composition and the resulting adjustment of tuition schedules, in addition to their effects on equilibrium wages and interest rates. Reductions in public college subsidies, in contrast, also operate through a direct price channel, since lower subsidies immediately increase the net tuition that colleges charge students.

In this section we present our basic quantitative results. Section 5.1 displays the college market equilibrium in the year 2025 on the baseline demographic path, along which the population grows at a constant rate of 1%, the long-run U.S. population growth rate (which factors in immigration) prior to 2008. Section 5.2 then subjects the baseline economy to an unexpected (MIT shock) reduction in public subsidies for higher education; we model a (possibly too) stark reform in which all public subsidies are reduced to zero. In Section 5.3 we turn to the impact of the demographic cliff on the macro economy and the college market, and Section 5.4 discusses the interaction between demographic shifts and public education funding reform. In all scenarios the change in demographics and/or public education subsidies is modeled as unanticipated. For the demographic cliff, what is unanticipated is the decline in fertility itself; once households have learned of it, the implied future decline in the college-age and adult population is perfectly foreseen. Sections 5.2 through 5.6 report these experiments along the transition rather than as a comparison of steady states: every college-market figure and every table compares the year 2086 on the reform path with the year 2086 on the baseline path, and reports alongside it how each economy moves between 2026 and 2086.[7] Section 5.6 asks how much of the economy's response to demographic change is

[7]Comparing the two paths at a common date is the transition analogue of a policy counterfactual: it nets out the growth, the ageing and the secular rise in attainment that both economies share, and it isolates what the reform does. Because the baseline path is itself far from stationary — attendance rises from 49.0 percent in 2026 to 55.4 percent in 2086 — differences measured against the initial steady state and differences measured against the contemporaneous baseline are not the same object, and the latter is the one that answers the question we pose.

carried by the heterogeneity of the college market, by repeating the demographic experiment in a comparison economy with homogeneous colleges and comparing the two responses at the same date. We organize the results around a focal year, 2086. Sections 5.2 to 5.6 report the medium- and long-run effects of each experiment at that date: they hold the year fixed and ask how the college market, the aggregates and the budget differ between an economy that received the shock and one that did not. Section 5.7 then follows the paths themselves through time, where the different timing of the education response and of the demographic burden becomes important, and 2086 is one point on those paths. It characterizes the transition induced by the demographic cliff (learned in 2010) and by its combination with an unanticipated subsidy cut in 2026; the transition for the subsidy cut alone is reported in Appendix D.

The timing is itself of economic consequence, and not merely a matter of presentation. The macroeconomic drag of aging and the response of higher education decisions to it arrive at different dates, and it is this differential timing that determines the magnitude of the effects in any given year, including the focal year. Studying the question we pose therefore requires both a transition analysis and a model in which college quality and major choice are heterogeneous: the first governs when the education response arrives relative to the demographic burden, the second how large that response can be.

5.1 Baseline Demographics

In Figure 6 we display the equilibrium allocation in the college market in the year 2025, read off the baseline transition solution at the model period closest to that year. On the baseline demographic path the population grows at a rate of 1% (with age-specific mortality rates determined directly from the data). Panel (a) shows the distribution of students across college quality levels. The histogram includes the roughly 50% of high school graduates who do not attend college and therefore select $q = 0$, which appears as the spike at the left end of the distribution.

Panel (b) displays the share of high ability ($e = e_h$) students by college quality. It shows that about 75% of high-school students not attending college come from the low-ability group. The remaining share of about 25% consists of high-ability individuals who nevertheless choose not to attend college. This outcome arises because some high-ability individuals start their life with a low b draw (which is a low-probability event but can happen), and are also subject to idiosyncratic taste shocks. Given the tuition schedule, the borrowing limits, the foregone wages while in college and the utility cost of studying, they find it optimal not to attend college despite their high innate ability. The higher the quality of the college, the larger is the ability tuition discount and the share of high-ability students

attending colleges of that quality. Given the college quality production function in which per-capita expenditures and average quality enter multiplicatively, the value of admitting a high-ability student is greater at colleges with high per-student expenditures, a property that is reflected in the equilibrium tuition schedule.

Panel (c) of Figure 6 displays the equilibrium tuition schedule. It shows two prominent (but fully expected) features: first, tuition is monotonic in school quality since per-student discretionary spending is one of the two determinants of school quality and colleges have to break even. Second, high ability students receive a tuition discount that compensates them for raising the average quality of the student body, and thus the quality of the college itself, given that average student quality is the second determinant of college quality.[8]

Panel (d) focuses on the college program (p) or major choice. Recall that studying a STEM field (which in many universities now includes economics) incurs a larger disutility, the more so the less innately able a student ($e = e_l$) is, but carries a higher wage premium in the labor market (and this effect is the stronger the better the attended college, given the log-linear (and thus multiplicative in levels) structure of labor productivity and thus wages in equation (5)). Consequently, the share of STEM students is higher at high-quality colleges, and larger among high ability students (which in turn sort more strongly into high-quality colleges).

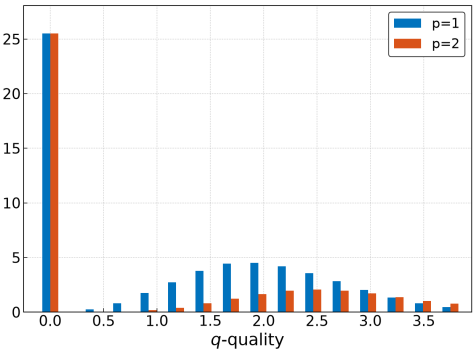
Program choice is also influenced by financial constraints. Since the model features borrowing limits, students with limited resources face tighter budget constraints when attending high-quality colleges. For low-ability students in particular, the combination of higher tuition costs and the larger disutility of STEM study makes STEM majors relatively less attractive at the top of the quality distribution. This mechanism generates the slight decline in the STEM share among low-ability students at the highest-quality colleges. Wealth effects may also play a role because households can save in the model, but the non-monotonicity in the figure is primarily driven by borrowing constraints interacting with the higher disutility of STEM fields.

5.2 Education Subsidy Cuts in General Equilibrium

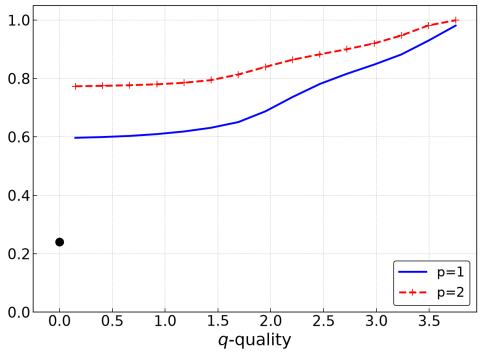
In this section we quantify how a reduction in public higher education subsidies impacts the college market equilibrium. Public subsidies are cut to zero, unanticipated, in 2026, and the economy is otherwise on its baseline demographic path; we compare it with that baseline path in 2086, sixty years after the cut. This exercise is meant to prepare the discussion for the same thought experiment when concurrently the economy is subject to a demographic

[8]The resulting tuition-ability discount in our model can be interpreted as capturing merit-based institutional financial aid in the U.S.

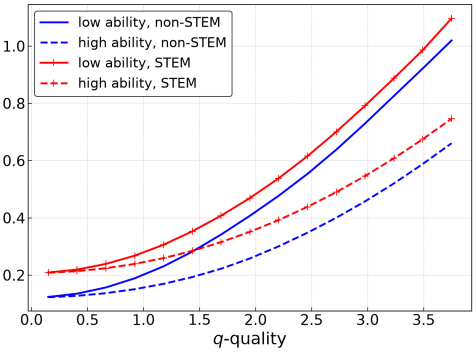
Figure 6: Baseline: The College Market in 2025



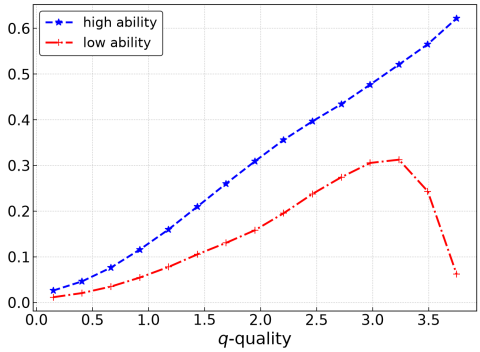
(a) Equilibrium q Histogram



(b) share of high ability students



(c) Equil. gross tuition schedule

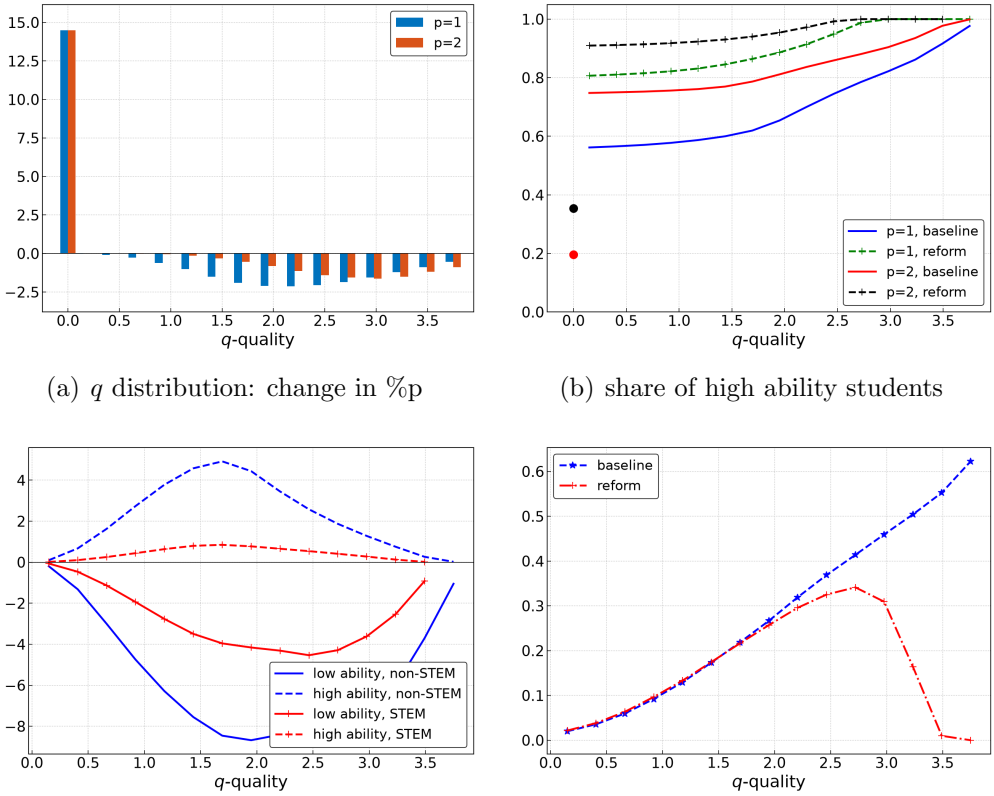


(d) share of students in STEM fields, by student ability

Notes: College market equilibrium in the year 2025 on the baseline demographic path (population growth of 1%), read off the baseline transition solution. Panel (a): distribution of the youth cohort across college qualities and majors; for $q = 0$ (non-enrollment) the bar is split into two equal subbars ($p = 1$ and $p = 2$) for visual proportionality only. Panel (b): share of high-ability students by college quality, with the marker at $q = 0$ giving that share among non-attendees. Panel (c): equilibrium gross tuition (before public subsidies), by student ability and major. Panel (d): share of students choosing STEM, by student ability.

cliff in Section 5.4. The direct effect of lowering $s(e, r; q)$ uniformly to zero is to increase net tuition for all households seeking to go to college. Of course, in equilibrium the pre-subsidy tuition schedule adjusts, and with it the equilibrium allocation of students across colleges.

Figure 7: College Funding Cut: Eliminating the 25% Subsidy



(a) q distribution: change in %p

(b) share of high ability students

(c) %-Δ in equil. gross tuition schedule

(d) share of students in STEM

Notes: Public higher education subsidies are cut from an average of 25% per student to zero (reflecting a 49% subsidy rate for in-state students, who constitute roughly half the student body), unanticipated, in 2026. Panels (a), (b) and (d) compare the college market on the path of the subsidy cut with the college market on the baseline path, both in the year 2086. Panel (a): change in the distribution of students across college quality, in percentage points of the youth cohort; for $q = 0$ (non-enrollment) the bar is split into two equal subbars ($p = 1$ and $p = 2$) for visual proportionality only, as there is no major choice for non-college individuals. Panel (b): share of high-ability students, with the two markers at $q = 0$ giving that share among non-attendees. Panel (d): share of students choosing STEM. All panels plot outcomes as a function of college quality q. Panel (c): percentage change in the equilibrium gross tuition schedule (before public subsidies), by student ability and major, comparing the two paths in the year 2086.

Panel (a) of Figure 7 shows the change in the distribution of college-age individuals across college quality levels q and fields of study p following the subsidy cut. In 2086 the share of households not attending college stands at 73.6 percent on the cut path against 44.6 percent on the baseline path, a gap of 29 percentage points. Enrollment moves away from middle-

and high-quality colleges, where tuition levels are highest, so that the surviving student body is concentrated in the lower half of the quality distribution: mean attended quality is 1.90 against 2.18 on the baseline path. Panel (b) displays the resulting change in the ability composition of students. The share of high-ability students among enrollees is 90.6 percent in 2086 on the cut path against 74.4 percent on the baseline path. As tuition rises, the college population shifts toward higher-ability students. In the model, colleges offer an ability-based tuition discount, so low-ability students face higher tuition for a given college type. At the same time, low-ability students incur higher utility costs of studying, particularly in STEM fields. In addition, our calibration implies that initial financial resources $b(e)$ are positively correlated with ability, so low-ability students are more likely to be borrowing constrained. Taken together, these forces make college attendance less attractive for low-ability students when net tuition increases, leading to a stronger sorting of high-ability students into college.

Panel (c) plots the percentage change in the equilibrium gross tuition schedule — the sticker price before public subsidies — as a function of college quality and student ability. Because the subsidy is proportional, the change in the tuition students actually pay adds a mechanical component to any movement of the gross schedule: had gross tuition remained unchanged, net tuition would have increased by

$$\frac{1}{1-0.25} - 1 = 33.3\%.$$

As shown in Section 3.3, the improvement in the ability composition of students following the subsidy cut affects tuition through two channels. First, higher average ability reduces the instructional spending required to produce a given quality, lowering tuition for all students. Second, the improvement in student composition reduces the equilibrium ability discount $d(q, p)$, which raises tuition for high-ability students and lowers it for low-ability students. For low-ability students both channels operate in the same direction: their gross tuition falls, by up to 9% in the middle of the quality distribution, so their net tuition increases by less than the mechanical 33.3%. For high-ability students the two effects work in opposite directions, but quantitatively the reduction in the ability discount dominates: their gross tuition rises, by up to 5%, so their net tuition increases by slightly more than 33.3%.

Panel (d) shows the share of students choosing STEM majors; in 2086 the aggregate STEM share is 23.4 percent of enrollees against 31.4 percent on the baseline path. Despite the stronger positive selection of students on ability following the subsidy cut, the STEM share declines across the entire college quality distribution and does so particularly strongly at high-quality colleges. The reason is that the increase in net tuition raises the financial cost of attending college. While higher ability would, in isolation, reduce the psychological

cost of studying STEM, this effect is dominated by the price channel. As tuition increases, students shift away from the more demanding STEM majors toward less costly alternatives. Because tuition levels are highest at the top of the quality distribution, this mechanism generates the largest decline in the STEM share at high-quality colleges. As a result, both the overall share and the absolute number of STEM students fall substantially following the subsidy cut. To the extent that future technological change increases demand for STEM skills, reductions in higher-education funding may therefore exacerbate STEM shortages in the context of this model.

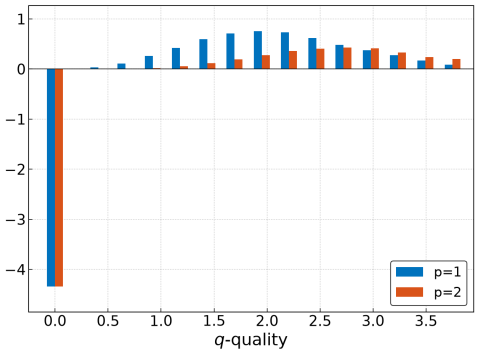
5.3 The Demographic Cliff and the College Market

We now consider the demographic cliff itself. Households learn in 2010 that fertility falls permanently, and the economy travels from that date along the population path implied by our demographic projections (see Section 4.1); the counterfactual baseline path is the one in which fertility does not fall. By 2086 the two paths differ substantially in size: the age 18–21 population on the cliff path is 63.3% of its size on the baseline path, and the adult population 79%. This matters for the comparison of not only per capita variables, but aggregate levels of economic activity, student body, and fiscal variables across the two demographic scenarios. [9]

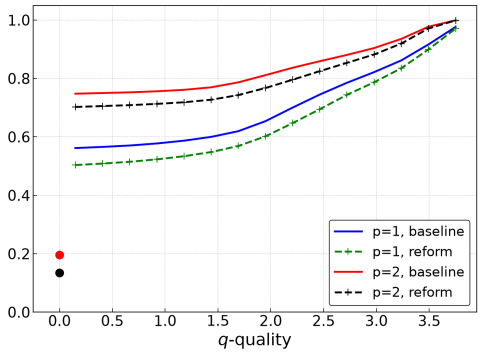
Figure 8 illustrates the general equilibrium effects of the demographic shock on the college market and households' higher-education choices. Panel (a) reports the change in the distribution of students across college quality (with $q = 0$ denoting non-attendance) and across fields p. The aging shock increases the capital–labor ratio, which lowers the return on physical capital, see Table 4 below. Since capital is more complementary to college-educated than to non-college labor, this change raises the relative return to education and increases equilibrium college enrollment. In 2086 the non-attendance share stands at 35.9 percent on the cliff path against 44.6 percent on the baseline path, a reduction of about nine percentage points, and enrollment shifts toward middle- and high-quality colleges. In addition, demographic aging puts pressure on the public pension system. In our baseline specification the Social Security budget constraint must be balanced in every period, and the adjustment occurs through a decline in the pension replacement rate. The resulting reduction in expected retirement income further strengthens incentives to invest in higher education. Consistent with these mechanisms, the relative shares of students attending middle-quality colleges also increase.

[9] If we were only interested in per-capita variables, it would not be necessary to compute how the total population changes between the two steady states. However, because we also study aggregate economic activity, the size of the student body, and fiscal variables, the change in the total population stock must be explicitly modeled.

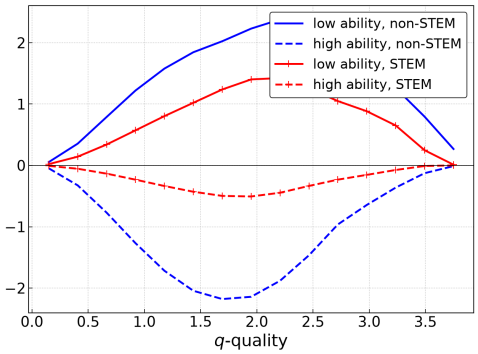
Figure 8: Demographic Cliff



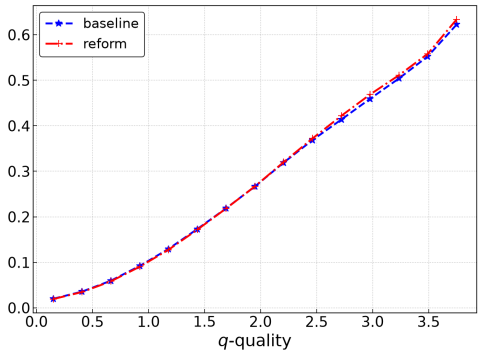
(a) q distribution: change in %p



(b) share of high ability students



(c) %-Δ in equil. gross tuition schedule



(d) share of students in STEM

Notes: The demographic cliff is learned in 2010 and the population growth rate falls from 1% to 0%. Panels (a), (b) and (d) compare the college market on the path of the demographic cliff with the college market on the baseline path, both in the year 2086. Panel (a): change in the distribution of students across college quality, in percentage points of the youth cohort; for $q = 0$ (non-enrollment) the bar is split into two equal subbars ($p = 1$ and $p = 2$) for visual proportionality only, as there is no major choice for non-college individuals. Panel (b): share of high-ability students, with the two markers at $q = 0$ giving that share among non-attendees. Panel (d): share of students choosing STEM. All panels plot outcomes as a function of college quality q. Panel (c): percentage change in the equilibrium gross tuition schedule (before public subsidies), by student ability and major, comparing the two paths in the year 2086.

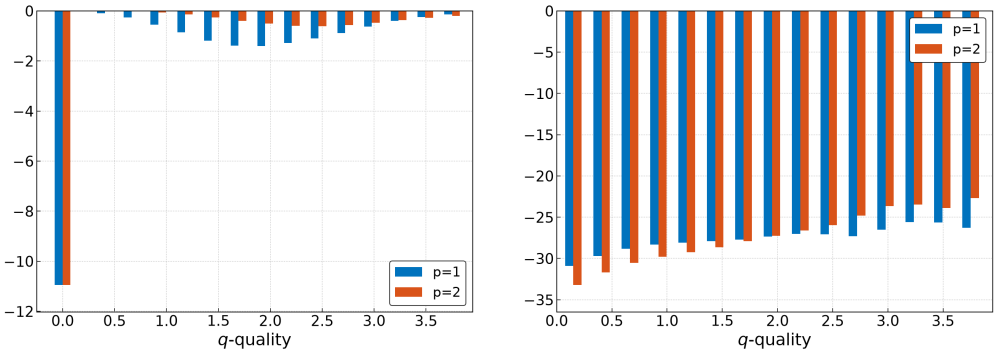
Panel (b) shows that the share of high-ability students declines at low-q colleges; averaged over enrollees it is 70.4 percent on the cliff path in 2086 against 74.4 percent on the baseline path. This occurs because the marginal students who switch from non-attendance in the initial steady state to college attendance in the final steady state have, on average, lower ability than the inframarginal students already attending college. At the same time, some high-ability inframarginal students re-sort from low-q to higher-quality colleges, further reducing the high-ability share at the lower end of the college quality distribution. The deterioration in average student ability at middle-quality colleges increases the equilibrium ability discount offered to high-ability students, consistent with the comparative statics derived in Section 3.3, which show that a decline in average ability raises the equilibrium tuition discount for high-ability students.

As shown in panel (d), the STEM share declines at middle-quality colleges following the aging shock. These colleges experience the largest inflow of students, either because previously non-attending marginal students now enroll or because some inframarginal students reallocate from lower-quality colleges. On average, these additional students have lower ability than the students already attending middle-q colleges in the baseline equilibrium. Since STEM fields are more demanding and involve higher disutility of study, the larger presence of lower-ability students reduces the share of STEM majors at those colleges. This compositional effect also lowers the aggregate STEM share. Because enrollment shifts toward middle-quality colleges, where STEM shares decline, the overall proportion of students choosing STEM falls despite the increase in total college attendance. This result illustrates how shifts in the ability composition of students across college qualities translate into changes in major choice.

Finally, note that the changes in ability composition across the q-distribution lead to noticeable adjustments in equilibrium gross tuition (see panel (c)). The largest changes occur at middle-quality colleges, where enrollment increases the most. These colleges receive an inflow of marginal students who have, on average, lower ability than the inframarginal students attending those colleges in the baseline equilibrium. As a result, average student ability declines at middle-q colleges, which raises the instructional spending required to produce a given level of quality. At the same time, the associated adjustment in the equilibrium ability discount further changes the tuition schedule across ability types. Together, these effects lead to an increase in gross tuition for low-ability students—up to about 2.4% for non-STEM majors and around 1.4% for STEM—and a decline in gross tuition for high-ability students of up to about 2.2% at middle-quality colleges.

The results thus far might suggest that the impact of the demographic cliff on the college market (and, by implication, the labor market and the macroeconomy more broadly) is

Figure 9: Demographic Cliff: Student Enrollment in 2086 Relative to the Baseline Path



(a) absolute change in age 18–21 population (baseline = 100)

(b) % change in student population relative to the baseline path

Notes: Number of college-age individuals at each point of the college quality distribution in the year 2086, on the path of the demographic cliff against the baseline path, with the college-age population on the baseline path normalized to 100. Panel (a): absolute change, so that the entries sum to the total change in the size of the youth cohort; for $q = 0$ (non-enrollment) the bar is split into two equal subbars ($p = 1$ and $p = 2$) for visual proportionality only, as there is no major choice for non-college individuals. Panel (b): the same change at each quality, divided by the number of students at that *same* quality on the baseline path, so that each bar is the proportional change in the size of a college of that quality rather than its contribution to the total. The two panels therefore weight the quality distribution differently: panel (a) is largest where baseline enrollment is largest, panel (b) where enrollment falls by the largest fraction of its own base. In 2086 the age 18–21 population on the cliff path is 63.3% of its size on the baseline path, so the entries of panel (a) sum to approximately −37.

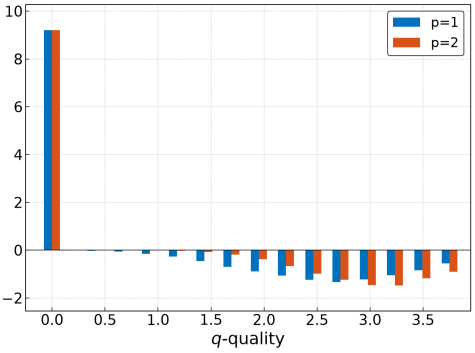
relatively benign. In percentage terms, the demographic shock raises college attendance and shifts enrollment toward middle-quality colleges as labor scarcity increases the incentives to invest in higher education. However, demographic change also implies a substantial reduction in the absolute number of individuals in the economy: by 2086 the age 18–21 population on the cliff path is 63.3% of its size on the baseline path. Figure 9 reports the absolute changes in the number of college-age individuals across the college quality distribution. While the number of individuals not attending college declines sharply, the overall contraction in population leads to much smaller reductions in the absolute number of students. In levels these losses are concentrated in the middle of the quality distribution, at q between about 1.7 and 2.2, simply because that is where baseline enrollment is largest. Expressed relative to the enrollment of the corresponding college on the baseline path, however, the losses fall steadily as quality rises, from about 31 percent at the bottom of the distribution to about 24 percent at the top: the cliff shifts students toward better colleges at the same time as it shrinks their number, so the high-quality segment is the part of the college market least affected in proportional terms, and the low-quality segment the most. Averaged over the distribution the two panels agree, since the student body as a whole is 26.8 percent smaller than on the baseline path in 2086.

5.4 Public Subsidy Cuts when Population is Aging

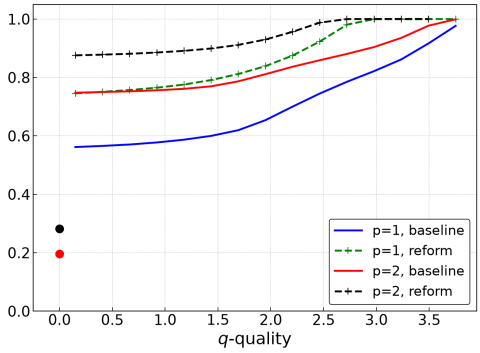
Figure 10 shows the college market response when, in addition to the demographic shock just discussed, public funding of colleges is fully withdrawn: the cliff is learned in 2010 and the subsidy rate is reduced from its baseline level to zero in 2026, as the demographic burden is building. Although reality is of course more complex than that, one could interpret this scenario as the environment the U.S. is currently best described by.

Qualitatively, the responses closely resemble those induced by the pure subsidy cut discussed in Section 5.2. In other words, the favorable effects of the aging shock on the college market in per capita terms (higher enrollment and a shift away from low-quality colleges) are largely undone by the withdrawal of public funding. The magnitudes of the responses are smaller than in Figure 7: non-attendance in 2086 is 63.0 percent under the two shocks together against 73.6 percent under the cut alone, and the enrollment gap relative to the contemporaneous baseline is 18.4 rather than 29.0 percentage points. The demographic shock, by making labor scarcer and dearer, partially offsets the disincentive effects of the subsidy cut. The offset is also visible in the ability composition: with the cut alone the student body is 90.6 percent high-ability in 2086, with both shocks 87.1 percent, since the scarcity of labor pulls marginal students back into college.

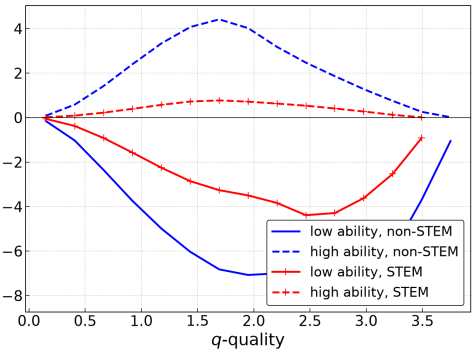
Figure 10: Demographic Cliff & Subsidy Cut



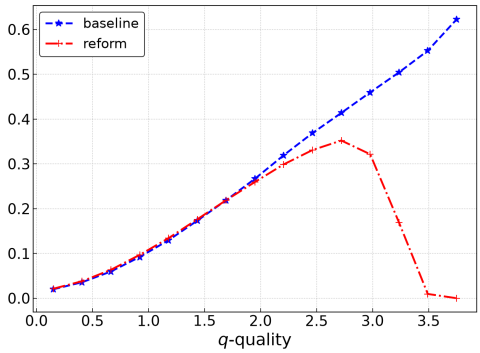
(a) q distribution: change in %p



(b) share of high ability students



(c) %-Δ in equil. gross tuition schedule

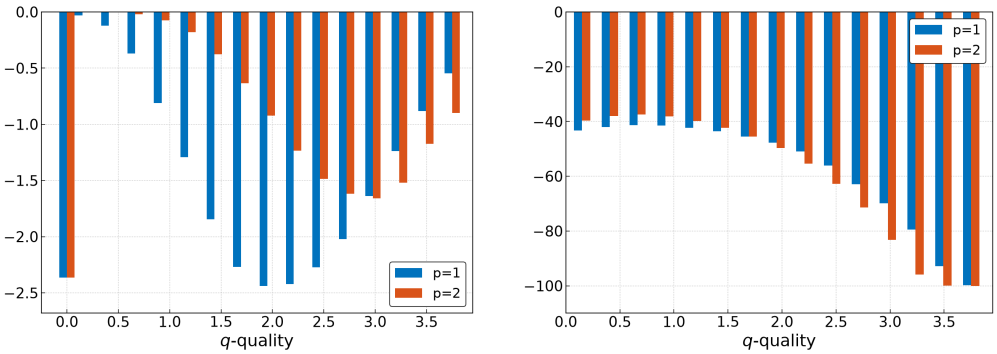


(d) share of students in STEM

Notes: The demographic cliff is learned in 2010 and public higher education subsidies are cut to zero, unanticipated, in 2026. Panels (a), (b) and (d) compare the college market on the path of the two shocks combined with the college market on the baseline path, both in the year 2086. Panel (a): change in the distribution of students across college quality, in percentage points of the youth cohort; for $q = 0$ (non-enrollment) the bar is split into two equal subbars ($p = 1$ and $p = 2$) for visual proportionality only, as there is no major choice for non-college individuals. Panel (b): share of high-ability students, with the two markers at $q = 0$ giving that share among non-attendees. Panel (d): share of students choosing STEM. All panels plot outcomes as a function of college quality q. Panel (c): percentage change in the equilibrium gross tuition schedule (before public subsidies), by student ability and major, comparing the two paths in the year 2086.

Anticipating the discussion of aggregate and fiscal effects in Section 5.5, it is worth noting that the joint experiment combines two adverse forces: a smaller population and a decline in college enrollment relative to the baseline. The result is a severe contraction of macroeconomic aggregates and of public finances, reflecting the loss of the enrollment-driven fiscal buffer that the demographic shock creates on its own. Measured along the transition the damage is close to, and slightly below, the sum of the two shocks taken separately; Section 5.5 takes up why the two do not simply add.

Figure 11: Demographic Cliff and Subsidy Cut: Student Enrollment in 2086 Relative to the Baseline Path



(a) absolute change in age 18–21 population (baseline = 100)

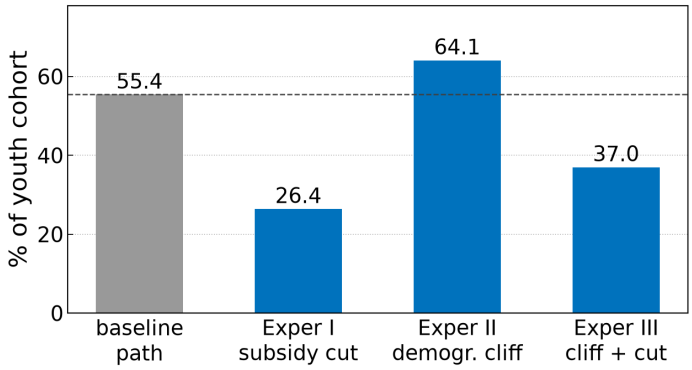
(b) % change in student population relative to the baseline path

Notes: Number of college-age individuals at each point of the college quality distribution in the year 2086, on the path of the two shocks combined against the baseline path, with the college-age population on the baseline path normalized to 100. Panel (a): absolute change, so that the entries sum to the total change in the size of the youth cohort; for $q = 0$ (non-enrollment) the bar is split into two equal subbars ($p = 1$ and $p = 2$) for visual proportionality only, as there is no major choice for non-college individuals. Panel (b): the same change at each quality, divided by the number of students at that *same* quality on the baseline path, so that each bar is the proportional change in the size of a college of that quality rather than its contribution to the total. The two panels therefore weight the quality distribution differently: panel (a) is largest where baseline enrollment is largest, panel (b) where enrollment falls by the largest fraction of its own base. In 2086 the age 18–21 population on the joint path is 63.3% of its size on the baseline path, so the entries of panel (a) sum to approximately −37.

5.5 Aggregate and Fiscal Effects of Demographic Shocks and Subsidy Cuts

The previous subsections examined how the different reforms reshape the college market equilibrium. In this section we summarize the aggregate and fiscal policy implications of the thought experiments discussed thus far. Table 4 contains the impact on interest rates and relative wages, whereas Table 5 and 6 contain the reaction of macroeconomic aggregates and fiscal variables.

Figure 12: College Enrollment in 2086 Across the Three Experiments



Notes: Share of the youth cohort attending college ($q > 0$, completers plus dropouts) in the year 2086, on the baseline path and on the path of each of the three experiments. The three experiments share the same baseline path, drawn as the dashed line.

Table 4: Wages, Interest Rate and Tax Adjustments Along the Transition, 2086 (% and %p Differences to the Baseline Path)

Reform	Interest Rate	Non-Coll. Wage	Coll. Non-STEM Wage	Coll. STEM Wage	Lab. Inc. Tax	Eff. Lab. Inc. Tax
Exper I	-0.46%p	1.73%	2.52%	2.86%	7.42%p	3.88%p
Exper II	-0.25%p	0.92%	1.36%	1.54%	-1.15%p	0.21%p
Exper III	-0.63%p	2.39%	3.46%	3.91%	5.91%p	3.83%p
sum of I and II	*-0.72%p*	*2.66%*	*3.88%*	*4.40%*	*6.27%p*	*4.09%p*

Notes: Exper I: elimination of public college subsidies (average subsidy rate cut from 25% per student to zero, reflecting a 49% subsidy rate for in-state students who constitute roughly half the student body), unanticipated in 2026. Exper II: demographic cliff, learned in 2010. Exper III: both. All entries compare the reform path with the baseline path in the year 2086: each is the difference between the two economies at that common date, so that the growth, the ageing and the secular rise in attainment common to both are netted out. The three experiments share the same baseline path, so the final row, the arithmetic sum of the first two, is a well-defined benchmark for whether the two shocks add up. Interest rate and labor income tax changes are in percentage points (%p); the interest rate is annualized; wage changes are in percent (%); the labor income tax rate is $1 - \lambda$ of the HSV schedule, i.e. the statutory average tax rate at the schedule's fixed reference earnings level, which moves only when the schedule itself moves. The final column reports the aggregate measure of the labor tax burden: total labor income tax revenue divided by total labor income, which in addition reflects the equilibrium change in the composition of earnings; it stands at 14.0% on the baseline path in 2086, against 17.9% for the statutory measure. Interest rates, wages and the tax rate are read from the converged outer-loop files of each pass; the aggregate measure from the fiscal accounts of each pass.

From Table 4 we observe that both shocks lower the equilibrium interest rate and raise wages across all skill groups. This reflects the increase in the capital–labor ratio: the demographic cliff shrinks the workforce, while the subsidy cut reduces the efficiency units that workforce supplies, so that in both cases capital becomes abundant relative to labor, the return to capital falls and the marginal product of labor rises.

The magnitude of the wage response varies across skill groups, with larger gains for college-educated workers reflecting their stronger complementarity with capital. In the de-

mographic experiment (Experiment II), wages in 2086 stand about 0.9% above their level on the baseline path for non-college workers, 1.4% for college non-STEM workers and 1.5% for STEM graduates. Under the subsidy cut (Experiment I) the same ordering holds at larger magnitudes — 1.7%, 2.5% and 2.9% — since the cut withdraws college labor in particular. When the two shocks occur together (Experiment III), wages are 2.4%, 3.5% and 3.9% above the baseline path.

The fiscal consequences of the two shocks differ in sign. The demographic cliff on its own *lowers* the labor income tax rate, by 1.15 percentage points relative to the baseline path in 2086. Aging shrinks the workforce, but it also raises college attendance — by 8.7 percentage points at that date — and the resulting expansion of the high-skill earnings base more than offsets the revenue lost to a smaller population, so that the budget balances at a lower rate. Two further margins push in the same direction: students select into higher-quality institutions and into fields with higher returns, so the earnings base widens by more than the change in the attendance rate alone would suggest, and the pension replacement rate falls by 1.7 percentage points, which shifts the financing of old-age consumption toward own lifetime earnings and reinforces the incentive to acquire skill.

How the tax burden is measured matters for reading this result, which is why Tables 4 and 6 report two tax objects side by side. The first is the statutory measure $1 - \lambda$: the average tax rate of the HSV schedule at a fixed reference earnings level, which moves only when the schedule itself moves. The second is the aggregate ratio of labor income tax revenue to labor income, which also reflects the equilibrium composition of earnings. Under the cliff the two measures split cleanly: the statutory rate falls by 1.15 percentage points, while the aggregate ratio is essentially unchanged (+0.21 percentage points). Holding the composition of earnings fixed, the lower schedule alone would reduce the ratio by about 1.2 percentage points; but the workforce shifts toward college graduates — mean labor earnings per worker are 9.3% higher on the cliff path in 2086 — and under a progressive schedule a better-paid workforce faces a higher average rate, which adds the offsetting 1.4 percentage points back. Aging with an education response thus buys a lower tax *rate*, not a lower tax *burden*: the burden relative to labor income is unchanged, but it is carried by a better-educated workforce facing a more lenient schedule. The composition margin works in reverse for the subsidy cut: the statutory rate must rise by 7.42 percentage points, but only about half of that (3.88 percentage points) shows up in the aggregate ratio, because the schedule increase applies to an earnings distribution that the cut itself has made poorer.

The subsidy cut does the opposite. Enrollment in 2086 is 29.0 percentage points below the baseline path, output per capita is 21.4% lower, and the labor income tax rate has to be 7.42 percentage points higher. Cutting public support for higher education therefore does

not save the government resources: the erosion of the tax base dominates the mechanical saving on subsidies by a wide margin, and it does so at every date after the cut, not only in the long run.

When both shocks occur together, the labor income tax rate is 5.91 percentage points above the baseline path. This is slightly *less* than the sum of the two effects taken separately, $7.42 - 1.15 = 6.27$ percentage points. The joint effect is marginally smaller than the sum on every other margin as well: enrollment is 18.4 percentage points below baseline against a sum of 20.3, output per capita 18.9% below against a sum of 19.1%, and the wage gains likewise fall short of their sum, 2.4% against 2.7% for non-college workers and 3.9% against 4.4% for STEM graduates. This does not depend on how the tax burden is measured: on the aggregate ratio the joint effect falls short of the sum by 0.26 percentage points, against 0.36 on the statutory measure. The two shocks are thus not independent, and the direction of the interaction is that the demographic change is somewhat *more* favorable to education, and hence to the budget, when it arrives in an economy without public college subsidies than when it arrives in one with them. We return to why below.

Table 5: Aggregate Variables Along the Transition, 2086 (% Differences to the Baseline Path)

Reform	Output	Capital	Aggregate Labor	Output (pc)
Exper I	-21.42%	-22.31%	-21.37%	-21.42%
Exper II	-18.90%	-9.90%	-20.35%	2.33%
Exper III	-35.73%	-29.00%	-36.81%	-18.90%
sum of I and II	*-40.32%*	*-32.22%*	*-41.72%*	*-19.09%*

Notes: Experiments and comparison as defined in the notes to Table 4. Output (pc): output per capita, using the adult population over the model's age bins from the demographic input files; under Exper I the two paths share the same population, so the aggregate and the per-capita entries coincide. Capital is net private capital (household assets less government debt). All entries are percentage changes (in levels) between the reform path and the baseline path in 2086.

Table 6: Enrollment and Fiscal Variables Along the Transition, 2086 (%p Differences to the Baseline Path)

Reform	Enrollment	Pension Repl. Rate	Lab. Tax Rate Δ (GE)	Eff. Lab. Tax Δ (GE)
Exper I	-29.00%p	-1.99%p	7.42%p	3.88%p
Exper II	8.69%p	-1.73%p	-1.15%p	0.21%p
Exper III	-18.42%p	-3.39%p	5.91%p	3.83%p
sum of I and II	*-20.31%p*	*-3.72%p*	*6.27%p*	*4.09%p*

Notes: Experiments and comparison as defined in the notes to Table 4. Enrollment: share of the youth cohort attending college (completers plus dropouts), which stands at 55.4% on the baseline path in 2086. Pension Repl. Rate: replacement rate of the pay-as-you-go pension system, which balances period by period and stands at 0.113 on the baseline path in 2086. Lab. Tax Rate Δ (GE): difference in the labor income tax rate $(1-\lambda)$, which stands at 17.9% on the baseline path in 2086, after all behavioral and price adjustments, required to balance the per-period government budget. Eff. Lab. Tax Δ (GE): difference in the aggregate measure, total labor income tax revenue over total labor income, which stands at 14.0% on the baseline path in 2086; see the notes to Table 4. A partial-equilibrium decomposition — the mechanical saving on education spending and the change in net revenues at benchmark prices — has no counterpart along the transition, where prices and the tax rate move in every period, and is therefore not reported.

Table 5 shows that both shocks contract aggregate labor, capital and output. Under the demographic cliff the contraction is a matter of the number of bodies: aggregate labor is 20.4% below the baseline path in 2086 and output 18.9% below it, yet output *per capita* is 2.3% *higher*, because the smaller cohorts are better educated and work with more capital each. Under the subsidy cut the contraction is a matter of skill rather than of numbers — the two paths share the same population — so the 21.4% decline in output is also a 21.4% decline per capita. Combining the two removes the per-capita gain entirely: output per capita is 18.9% below baseline in 2086.

The enrollment response behind these numbers is an endogenous substitution of human capital for labor that has become scarce. With $\rho(0) = 0.33$, $\rho(1) = -0.11$ and $\rho(2) = -0.43$, physical capital substitutes for non-college labor (an elasticity of substitution of about 1.5) and complements college labor, in particular in STEM occupations (about 0.7). Households respond by acquiring more human capital, so that part of the decline in raw labor input is made up along a margin that the aggregate production function itself rewards — and Experiment III shows that this margin is fragile to the funding of higher education. The pension system reinforces the same incentive: aging is absorbed through a lower replacement rate, which shifts the financing of old-age consumption toward own lifetime earnings and therefore raises the private return to acquiring skill.

It is useful to read these fiscal results as the interaction of two intergenerational transfer systems. Public subsidies to higher education are a transfer from working-age taxpayers to the young; the pension system is a transfer from the same taxpayers to the old. In the model the two have separate budgets — the pension system is pay-as-you-go and balanced period

by period through the replacement rate, while education subsidies are part of the general budget balanced by the labor income tax — but both are financed out of the labor earnings of the same working-age population, and education subsidies shape that base: how many of tomorrow's workers attend college, and at what quality, determines the efficiency units of labor from which both pension contributions and taxes are raised. Population aging tightens the pension system directly: with more retirees per contributor, the replacement rate must fall. The same shock raises college attendance, and the resulting expansion of the high-skill earnings base more than offsets the revenue lost to a smaller workforce. In this sense public spending on higher education purchases a fiscal buffer against aging — an outlay on the young that widens the base from which transfers to the old are financed.

The joint experiment shows how fragile that buffer is. Withdrawing public education funding removes the enrollment response through which aging would otherwise have broadened the base, so that the economy faces a workforce that is at once smaller and less skilled. Instead of a labor income tax rate 1.2 percentage points *below* the baseline path, it requires one 5.9 percentage points above it, and a pension replacement rate 3.4 rather than 1.7 percentage points lower. The order of magnitude of the fiscal cost is set by the education margin, not by the demographic shock itself.

Why, then, does the joint effect fall slightly short of the sum of the two separate effects? Because the demographic change does more work for education when subsidies are absent. Aging raises the return to schooling and lowers the pension replacement rate, and both strengthen the incentive to acquire skill whatever the subsidy regime. But the demographic change also carries a resource arm: the estates bequeathed at mid-century were accumulated by the larger pre-cliff generations and are divided among smaller cohorts, so bequests per entrant rise substantially. What differs across regimes is not so much the size of that windfall as what it is used for. Where subsidies are in place college is cheap at the margin, and a larger share of the additional resources is consumed rather than spent on attendance; where subsidies have been withdrawn, the same resources relax a constraint that now binds, and more of them go into college. The demographic shock is therefore worth more, in enrollment terms, in the economy without subsidies, which is why the damage the two shocks do together is a little less than the damage they do apart. Section 5.7.2 takes this up along the transition path, where the size of the inherited windfall can be read off directly.

The interaction is mild, and it does not soften the policy reading. Withdrawing public funding for higher education is a particularly costly policy against the backdrop of population aging, since it removes the margin at precisely the point where it would have done the most work: an economy that would have entered mid-century with attendance nine percentage

points above baseline and a slightly lower tax rate instead enters it with attendance eighteen points below baseline and a tax rate six points higher.

The timing of the education response is therefore central, and a comparison at a single date cannot speak to it. Section 5.7 shows that higher education decisions respond as soon as households learn about the demographic change, partly in anticipation of the future scarcity of labor and of the accompanying decline in the pension replacement rate, whereas the macroeconomic drag of aging materializes only as the large cohorts retire. There is consequently a window in which the education margin is already at work while the demographic burden is still building, and in which education policy is correspondingly potent. This suggests that a cut in education funding falling inside that window would be especially damaging, since it removes the response exactly when it would otherwise have been largest, whose transition path we report in Section D. How much work that margin can do is itself a property of the structure of the higher education market: Section 5.6 shows that when colleges are homogeneous in quality and major, the enrollment response to aging is substantially smaller, by between a third and a half depending on the calibration, and the fiscal buffer correspondingly weaker. That configuration is closer to the more uniform public university systems found in much of continental Europe, with a compressed quality distribution and a lower college wage premium.

5.6 Does Heterogeneity in College Quality and Majors Matter for the Macroeconomic Effects of Population Aging?

The experiments so far take the heterogeneity of the college market as given. We now ask how much of the economy's response to demographic change is carried by that heterogeneity, by repeating the demographic experiment in a comparison economy in which colleges are homogeneous in quality.

The comparison economy. We restrict the range of admissible college qualities to a narrow interval: in the comparison economy the quality grid has the same number of nodes as in the baseline but is confined to $[0.99\bar{q},\ 1.01\bar{q}]$, where $\bar{q}$ is the enrollment-weighted mean quality in the baseline economy. Colleges therefore remain formally differentiated — a continuum of distinct qualities with free entry and zero profits — so that the equilibrium construction of Section 2.5 carries over unchanged, while in economic terms the college sector is homogeneous. We implement the restriction in this way rather than by imposing literally a single quality, because with one admissible quality the college's problem is no longer a choice over a continuum and the free-entry argument would have to be restated. This is a restriction on the technology available to colleges rather than a computational approximation: the baseline grid discretizes the idea that any quality may be produced, whereas here the permitted

range is itself narrowed. We impose the analogous restriction on the second dimension of college heterogeneity and remove the distinction between programs, so that students no longer choose between fields that differ in their earnings profiles and in their psychic costs. College quality and major choice are therefore shut down jointly, and the comparison measures their combined contribution rather than that of either dimension on its own. Separating the two is feasible within the same framework, and we leave it to future work.

Tuition continues to condition on student ability, so the ability-discount schedule and the college's optimality conditions are unchanged; what disappears is the sorting of students of different ability across colleges of different quality and across fields of study. We additionally impose that tuition may not condition on family income. In the baseline model this independence is an equilibrium property rather than an assumption: with a continuum of differentiated colleges, a college offering an income-contingent schedule can be undercut, since a competitor can always construct a college whose quality and tuition attract the preferred type. That argument relies on product differentiation and hence does not survive the restriction to a single quality, so we impose income-independent tuition directly, in the same way that we impose the restriction on quality. Keeping the pricing technology identical across the two economies is what makes the difference in their response to demographic change attributable to quality dispersion. The comparison economy is thus an equilibrium under restrictions, and a different model rather than a special case of the baseline; we use it as a benchmark, not as a policy counterfactual.

Collapsing the quality distribution raises college attendance at unchanged preference parameters: the restriction replaces the bottom of the quality ladder with a mid-quality option, which draws in households that were previously at or below the enrollment margin. The two economies are therefore not comparable in levels. We therefore report the demographic experiment as a *response* within each economy — the change relative to that economy's own baseline path — and compare those responses rather than levels. We report the comparison twice: once with all parameters unchanged relative to the baseline calibration, and once with the comparison economy's utility cost of college recalibrated so that its benchmark enrollment matches that of the baseline economy.

Table 7: Response to the Demographic Cliff in 2086: Heterogeneous versus Representative College Market

Response to the cliff, 2086	Enrollment	Output (pc)	Lab. Tax Rate Δ	Interest Rate
Heterogeneous colleges (baseline)	8.69%p	2.30%	-1.15%p	-0.25%p
Representative college, no recalibration	3.81%p	-1.70%	0.08%p	-0.36%p
Representative college, recalibrated utility cost	4.01%p	-4.92%	1.30%p	-0.43%p

Notes: Each row reports the response of one economy to the same demographic cliff, measured in 2086 relative to that economy's own baseline path, as in Tables 4 to 6. Enrollment is the share of the youth cohort attending college (completers plus dropouts), and its entry is the change in that share. Output is expressed per capita by dividing by the adult population. Percentage changes are in log points; Δ entries are in percentage points. Both economies face the identical demographic scenario, so the differences across rows are the same whether the variables are measured in aggregate or in per-capita terms. Restricting the quality range raises attendance at unchanged preferences, so the two economies differ in levels and only their responses are comparable; the third row recalibrates the utility cost of college in the comparison economy so that its attendance matches that of the baseline economy.

Table 7 reports how much the response to the demographic cliff differs between the two economies. Restricting the quality range reduces the enrollment response from 8.69 to 3.81 percentage points and reverses the sign of the response of output per capita, from +2.30% to −1.70%; since the two economies face an identical demographic scenario, these gaps are the same whether the variables are measured in aggregate or per capita. The fiscal relief disappears as well. In the baseline economy the demographic cliff *lowers* the labor income tax rate by 1.15 percentage points, whereas in the comparison economy the same shock leaves the rate slightly higher, by 0.08 percentage points. College heterogeneity thus determines whether population aging is close to fiscally neutral or requires a substantial increase in the labor income tax rate.

The difference operates through the human capital that smaller cohorts acquire, not through capital deepening. The interest rate response differs by only 0.11 percentage points across the two economies, so the shrinking workforce raises the capital–labor ratio in much the same way; what differs is how much human capital the resulting increase in the return to schooling calls forth. With homogeneous colleges the enrollment response is substantially smaller. Measured against the baseline response of 8.69 percentage points, it is reduced by more than half whether or not the utility cost of college is recalibrated, and the margin along which students would otherwise respond by upgrading to better colleges is absent by construction. How much of population aging is absorbed by human capital is thus a property of the college market rather than of the demographic shock itself.

Why the response differs. The demographic cliff operates through the return to schooling: labor scarcity raises the college wage, and the decline in the pension replacement rate raises the value of own lifetime earnings. In the baseline economy a higher return calls forth two

responses. Students who would otherwise not have attended enroll — the extensive margin — and students who would have attended anyway move up the quality ladder and toward the STEM program — the intensive margin, or upgrading. The comparison economy weakens both, for different reasons. Appendix B.3 states the two margins formally, with propositions and proofs, in a simplified setting in which the household's type is a scalar and the tuition schedule is taken as given.

The intensive margin disappears outright. When a higher return raises the value of quality, every enrolled student in the baseline economy climbs the ladder a little, and these individually small movements aggregate over the entire enrolled population; with a single quality there is no ladder to climb, so the margin is absent by construction rather than merely weaker. The same holds for the major dimension: the reallocation toward the STEM program that accompanies the cliff in the baseline economy is unavailable once majors are equalized. The truncation of the quality distribution at the top, where the return to quality is steepest for high-ability students, is not a separate channel but the part of the intensive margin that is largest.

The extensive margin is dampened by the financing of the marginal student. In the baseline economy the marginal entrant sorts into the bottom of the quality ladder, the cheapest contract on the menu; in the comparison economy the only contract is the mid-quality bundle at mid-quality tuition. If financing costs were negligible at the margin, that substitution would be neutral: the bundle's higher tuition is matched by its higher quality, and entry would respond to the return in the same way under either menu. With borrowing constraints it is not, because marginal entrants are precisely the households for whom the cost of funds rises most steeply, so the financing burden of the compulsory bundle grows faster than its quality content. The homogeneous market thus confronts the marginal student with an entry ticket priced above what that student would choose to buy, and the mass of households near the enrollment margin is correspondingly thin. Selection cuts the other way: baseline entrants arrive at low quality and contribute few effective units each, whereas every entrant in the comparison economy carries the full bundle. Table 7 reports the combined effect of shutting down both margins rather than either on its own, but it shows that this advantage is not large enough to offset them: the same demographic shock produces a much smaller response of enrollment, output per capita and the labor income tax rate in the comparison economy.

One caveat. Collapsing the quality distribution also changes the tuition schedule itself, since pooling raises average ability at the single college and with it the resources required to deliver a given quality. This re-pricing is not a third margin — it operates through the two above, by moving the enrollment cutoff and the financing burden of the marginal student

— but our comparison does not separate it out. Doing so would require holding the tuition schedule fixed at its baseline values, which we leave to future work. Appendix B.3 also expresses the conditions behind both margins as bounds on observables, the tuition and earnings-value gradients across the quality menu, and verifies that in our calibration they hold with an order of magnitude to spare.

Because collapsing the quality distribution raises benchmark enrollment at unchanged preferences, part of the gap between the first two rows could reflect the higher starting level rather than the structure of the college market: the closer enrollment already is to one, the less room remains for it to rise. The third row addresses this by raising the utility cost of college in the comparison economy so that its benchmark enrollment matches that of the baseline economy. The gaps narrow but do not close: the enrollment response is still 4.7 percentage points smaller, output per capita still falls by a further 7.2 log points, and the fiscal gap remains 2.45 percentage points. The muted response is therefore not an artifact of the benchmark enrollment level.

We conclude that the structure of the higher education market is of first-order importance for evaluating the macroeconomic effects of demographic change. The same demographic shock, fed into two economies that differ only in the range of college contracts available to students, produces responses of aggregate labor input, output and the labor income tax rate that differ by an order of magnitude comparable to the effects themselves. Assessments of population aging that abstract from the heterogeneity of the higher education sector therefore risk substantially mismeasuring its fiscal consequences.

5.7 Transitional Dynamics

5.7.1 The Demographic Cliff

The results reported so far read the transition paths at a single date, 2086. This subsection follows the paths themselves through time, beginning with the demographic cliff, which is where the timing of the demographic change and of the college market's response to it becomes important. We solve the transition as an unanticipated (MIT) shock: households do not anticipate the demographic cliff until 2010; from 2010 onward they have perfect foresight of the new demographic path, and the wealth distribution accumulated up to that point is taken as given. The transition is initialized in 1950 with an artificial steady state, after which we feed in the population dynamics described in Section 4.1 and shown in Figure 4; the final steady state is reached around the year 2200. The transition induced by the subsidy cut alone is reported in Appendix D.

The single-date comparison and the full path answer different questions. The former isolates what the shock does to the economy by a fixed year, with the route by which it

got there left out of the picture. The path makes that route explicit, and the timing is of economic interest in its own right: higher education decisions and the macroeconomic effects of aging materialize at different dates, so the fiscal position of the economy in any given decade depends on how far each of the two has progressed. As argued in Section 5.5, this is what determines when education policy is most consequential, and a comparison of steady states cannot speak to it.

Figure 13 reports the aggregate response. By the end of the plotting window the cliff economy is close to half the size of the baseline: aggregate labor in efficiency units, output and capital are about 49 percent lower, essentially the decline in the adult population. Per capita, shown in the last panel of Figure 14, the two economies converge: by the end of the window output per capita is 1.2 percent higher under the cliff, and in the terminal steady state the gap vanishes (the remaining per-capita series are reported in Appendix D). In the long run the demographic cliff is therefore close to a pure scale effect, which is what one should expect given that the two scenarios converge to a common age structure. The transition, however, shows that the two economies are *not* the same per capita for a long time before that: the cliff economy runs *above* the baseline in output per capita from the late 2020s onward, with the gap reaching about 3.5 percent by mid-century and remaining of roughly that size into the early twenty-second century, before it closes.

That divergence comes from households' college decisions and from the shift in the distribution of students across college quality that they induce. College enrollment responds on impact at the announcement: in 2010 it stands 2.0 percentage points above the baseline, two decades before the demographic paths separate. The response is an anticipation effect: households that learn the cliff foresee the environment it implies — a rising relative price of college labor, and a pension system that absorbs aging through a lower replacement rate, shifting the financing of old-age consumption toward own lifetime earnings (Figure A4 in Appendix D).The gap then widens over the following decades, to 7.6 percentage points in 2050 and a peak of about 8.9 percentage points in the late 2070s (8.7 points in 2086), and closes only late in the horizon, with the cliff economy 0.4 percentage points above the baseline at the end of the window and level with it in the terminal steady state. Enrollment is thus higher under the cliff for roughly a century and a half, and essentially equal in the long run: the effect of the cliff on educational attainment is a transition phenomenon rather than a steady-state one. The associated paths of relative skill prices, the pension replacement rate and aggregate tuition are reported in Appendix D (Figure A4).

The fiscal path follows the same pattern. The labor income tax rate is identical in the two economies at the announcement and then falls faster under the cliff, running 0.9 percentage points below the baseline in 2050 and 1.2 percentage points below in 2086, with a maximum

saving of 1.3 percentage points in the early twenty-second century. The enrollment response, and the broader high-skill tax base it produces, more than offsets the smaller workforce for most of the transition.[10] The fiscal cushion is therefore concentrated in exactly the decades in which the demographic burden is building. The steady-state comparison of Section 5.5 already shows that the education response absorbs much of the fiscal cost of aging: a drop in the working-age population ratio of more than six percentage points leaves output per capita essentially unchanged and raises the labor income tax rate by only 0.46 percentage points. What the transition adds is the profile. Over roughly a century the cliff economy requires a labor income tax rate *below* the baseline rather than marginally above it, because the enrollment response arrives with the cohorts that learn about the cliff, whereas the increase in the dependency ratio arrives only as those cohorts age. The contrast is starkest around 2085, the year on which the stylized steady-state experiment is based: along the transition the cliff economy runs 1.2 percentage points below the baseline path in that year, against the increase of 0.46 percentage points that the steady-state comparison of the same demographic scenario requires. We conjecture that the differential timing is the reason. By 2085 the enrollment gap is near its maximum, while the rise in old-age dependency under the cliff is not yet complete, so the transition captures the economy at a point where the education response is fully in place and the demographic burden is not.

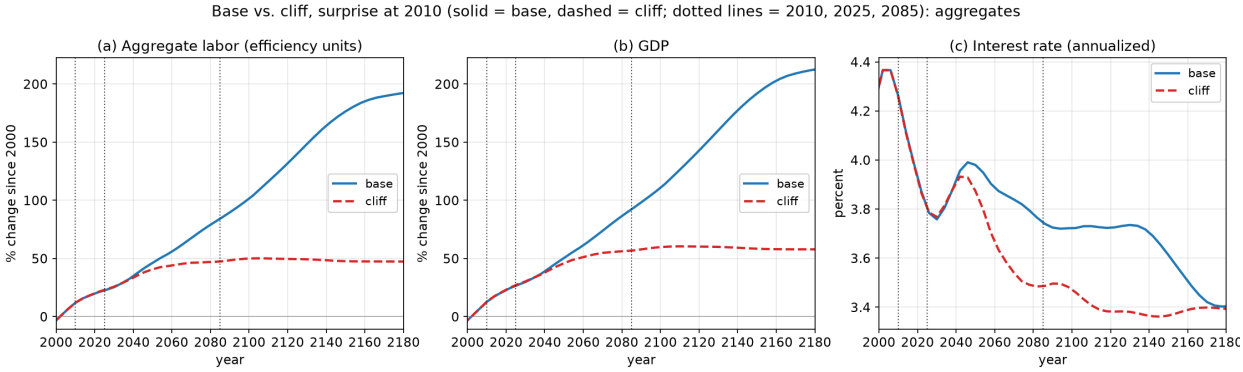


Figure 13: Aggregates along the surprise-cliff transition. Baseline (solid) versus demographic cliff (dashed); the cliff is learned in 2010. Panels: aggregate labor in efficiency units; GDP; and the interest rate. Labor and GDP are expressed as percentage changes since 2000 and share a common vertical axis; the interest rate is annualized and reported in percent. Dotted vertical lines mark 2010, 2025 and 2085; the plotting window is 2000–2180.

[10] Only at the very end of the horizon does the ordering reverse, with the cliff economy requiring a rate 0.4 percentage points higher, once the enrollment advantage has disappeared. The terminal comparison is a different object from the steady-state experiment of Section 5.5: at the end of the transition the two demographic scenarios share a common age structure by construction and differ only in the size of the population, whereas the steady-state experiment contrasts two stationary economies with different age structures.

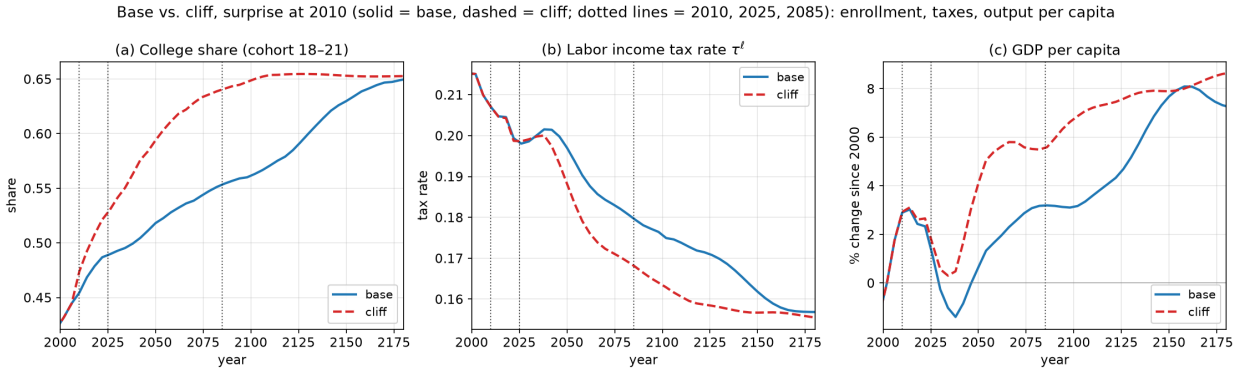


Figure 14: Enrollment, taxes and output per capita along the surprise-cliff transition. Panels: the college share of the 18–21 cohort (completers plus dropouts); the labor income tax rate $\tau^\ell = 1 - \lambda$; and output per capita, expressed as a percentage change since 2000. The first two panels are in levels. Conventions as in Figure 13. The remaining per-capita series are reported in Appendix D.

5.7.2 The Demographic Cliff and the Subsidy Cut

We next combine the two shocks in sequence: households learn the demographic cliff in 2010, as above, and are then surprised a second time by an elimination of public subsidies to higher education in 2026, with no anticipation of the cut before that date. The transition induced by the subsidy cut alone is reported in Appendix D. Figure 15 reports the paths of aggregate labor in efficiency units, output and the interest rate; Figure 16 those of college enrollment, the labor income tax rate and output per capita.

The path coincides with the cliff transition until 2026. When the cut arrives, enrollment collapses to 18.3 percentage points below the baseline path, and the fiscal position deteriorates along the same schedule as under the cut alone (Appendix D): the tax gap turns positive in the late 2040s, reaches 5.9 percentage points in 2086 and 9.4 percentage points by the end of the plotting window, and converges to the level of the pure subsidy-cut transition, since both experiments end in the same subsidy-free final steady state and the cliff's long-run contribution is a scale effect. Output per capita falls by 18.9 percent relative to the baseline path by 2086 and by 24.9 percent by the end of the window.

Two comparisons show what the demographic change contributes. Relative to the cut alone, the cliff keeps enrollment about 11 percentage points higher in 2086 and the labor income tax rate 1.5 percentage points lower. The education response to aging therefore operates also in the absence of subsidies, and it does so through three channels. Households learn about the cliff in 2010 and begin adjusting before the cut arrives, so part of the response is already in place by 2026. Aging raises the return to schooling directly, through the scarcity of labor and through a pension system that absorbs the demographic burden by lowering the

replacement rate. And smaller cohorts inherit more: the estates bequeathed at mid-century were accumulated by the larger pre-cliff generations, so bequests per entrant rise by roughly half.

The last of these channels is the one that depends on the subsidy regime. What changes with the regime is not so much the size of the inherited windfall as what it is used for. When subsidies are in place, college is cheap at the margin and a larger share of the additional resources is consumed rather than spent on attendance. When subsidies are withdrawn and tuition is paid in full, attendance is rationed by resources, and the same windfall instead clears the affordability threshold for students who would otherwise not enroll. The resource channel is therefore worth more where subsidies are absent, which is why the cliff's favorable effect is larger in the joint experiment than in isolation.[11]

None of this is enough to reverse the reform. The cliff lowers the joint fiscal cost by about a fifth, and output per capita remains about 19 percent below the baseline path. The demographic cliff is not enough to offset the long-run damage from the subsidy cut.

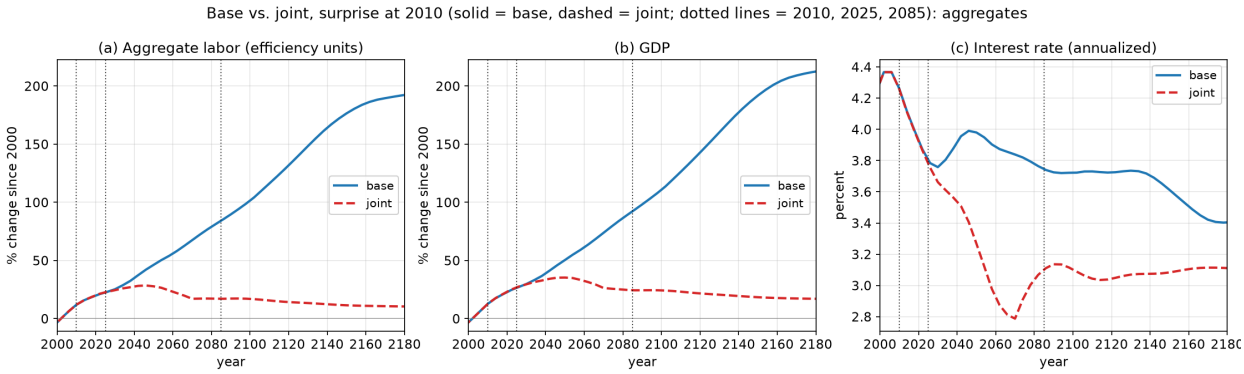


Figure 15: Aggregates along the joint transition: cliff learned in 2010, subsidy cut a second surprise in 2026. Baseline (solid) versus joint experiment (dashed). Panels: aggregate labor in efficiency units; GDP; and the interest rate (annualized, in percent). Labor and GDP are percentage changes since 2000 on a common vertical axis. Conventions as in Figure 13.

[11] At 2086 the cliff raises college enrollment by 10.6 percentage points on the subsidy-free path against 8.7 percentage points on the subsidised one.

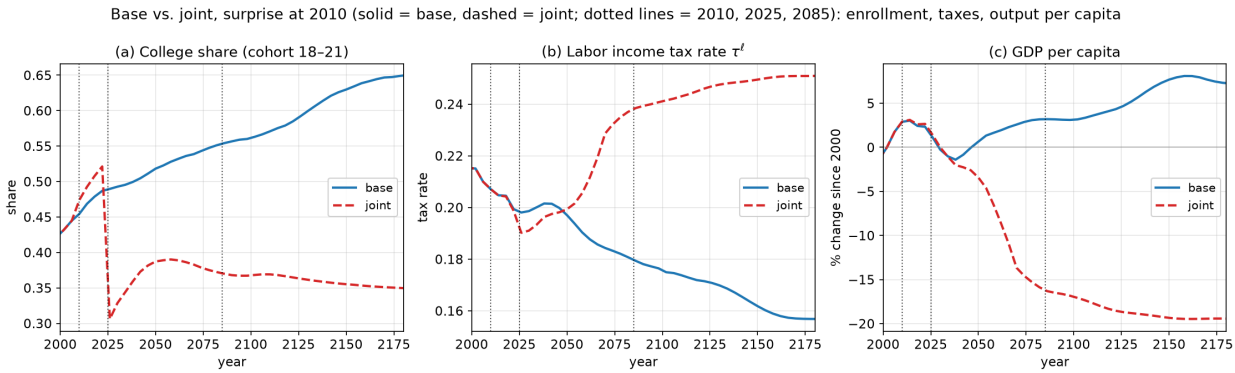


Figure 16: Enrollment, taxes and output per capita along the joint transition. Panels and conventions as in Figure 14. The remaining per-capita series and the college-market variables are reported in Appendix D.

6 Conclusion

In this paper we have constructed a general equilibrium life cycle model of the college market with heterogeneous colleges, student college quality and major choice, and subsequent labor market outcomes of workers in different occupations. We have applied the model to evaluate the aggregate and distributional consequences of the "demographic cliff" that reduce the number of high-school graduates in the next decades. We have then discussed the consequences of a funding cut towards public higher education akin to the recently imposed measures.

The model simulations show that demographic change fundamentally reshapes the college market and its interaction with public finance. While the "demographic cliff" reduces aggregate output through a smaller population, it simultaneously increases college enrollment rates and partially cushions fiscal pressures by expanding the share of high-skill workers. However, this buffer is fragile: when public subsidies to colleges are removed, enrollment falls substantially below baseline levels and the skill composition deteriorates. As a result, the combined fiscal impact of demographic change and subsidy cuts falls slightly short of the sum of their individual effects: the demographic change does more work for education where subsidies are absent, because the inheritance windfall it creates relaxes a constraint that public funding would otherwise have relaxed. The interaction is mild and leaves the policy reading intact. Withdrawing public funding is particularly costly against the backdrop of population aging, since it removes the education margin at the point where it would have done the most work.

How much of the demographic shock the education margin absorbs depends on the structure of the higher education market itself. Repeating the demographic experiment in a com-

parison economy in which colleges are homogeneous in quality and major, we find that the enrollment response, and with it the response of effective labor supply, is markedly smaller, and the labor income tax rate has to be raised substantially rather than remaining close to unchanged. A homogeneous system is closer to the more uniform public university systems of continental Europe, suggesting that the capacity of higher education to absorb population aging differs across countries with the institutional structure of their education sectors. Assessments of population aging that abstract from this heterogeneity risk substantially mismeasuring its fiscal consequences.

Our demographic cliff experiments so far were driven by a change in the domestic fertility and therefore domestic population growth rate of the economy. However, significant changes in immigration policy also will have a substantial impact on the age- and the skill distribution of the U.S. population. In a next step we will expand our demographic projections to include this force, with specific attention being placed on the reduction of high-skill migrations and against the backdrop of the quantitative importance of immigrants in STEM fields and innovation.

Second, when conducting our counterfactual experiments thus far, and even though different types of labor have a differential degree of substitutability with capital, we held technology constant (beyond a constant growth trend). Especially when it comes to the importance of STEM education, considering capital-skill-biased technological change (whether exogenously or, eventually, endogenously generated) is important for a full assessment of the changes in the higher education market, and the impact of public funding policies for that market. Whether college should be subsidized, and which programs should be publicly funded, will likely depend on the direction technological progress will take, and we plan to use our environment to shed light on this question. Finally, most of our analysis compares long-run steady states under different demographic and higher education policy scenarios. Sections 5.7 and 5.7.2 compute the transition paths induced by the demographic cliff and by its combination with the subsidy cut, each arriving as an unanticipated shock; the transition for the cut alone is reported in the appendix. Along this path the cliff economy requires a lower labor income tax rate than the baseline for roughly a century: the enrollment response sets in with the cohorts that learn about the cliff, whereas the old-age dependency ratio starts rising only as those cohorts age.[12]

[12]We report this single transition to emphasize the importance of the relative timing of the labor market effects of aging and of the college market response. The next iteration of the paper will contain the transition paths for all experiments, as well as the normative analysis, for which they are required: demographics and the cross-sectional distribution of higher education attainment are slow-moving objects, so the cohorts that bear the cost of a reform and those that benefit from it need not be the same.

References

Alon, T., D. Capelle, and K. Matsuda (2025). University research and the market for higher education. *American Economic Journal: Macroeconomics*.

Altonji, J., P. Arcidiacono, and A. Maurel (2016). Chapter 7 - the analysis of field choice in college and graduate school: Determinants and wage effects. Volume 5 of *Handbook of the Economics of Education*, pp. 305–396. Elsevier.

Altonji, J. G., E. Blom, and C. Meghir (2012, July). Heterogeneity in human capital investments: High school curriculum, college major, and careers. *Annual Review of Economics 4*(1), 185–223.

Altonji, J. G., L. B. Kahn, and J. D. Speer (2014, May). Trends in earnings differentials across college majors and the changing task composition of jobs. *American Economic Review 104*(5), 387–93.

Altonji, J. G. and S. D. Zimmerman (2018, October). The Costs of and Net Returns to College Major. In *Productivity in Higher Education*, NBER Chapters, pp. 133–176. National Bureau of Economic Research, Inc.

Arcidiacono, P. (2005). Affirmative action in higher education: How do admission and financial aid rules affect future earnings? *Econometrica 73*(5), 1477–1524.

Black, D. A. and J. A. Smith (2006). Estimating the returns to college quality with multiple proxies for quality. *Journal of Labor Economics 24*(3), 701–728.

Blair, P. Q. and K. Smetters (2021, September). Why don't elite colleges expand supply? Working Paper 29309, National Bureau of Economic Research.

Braxton, J. C. and B. Taska (2023, February). Technological change and the consequences of job loss. *American Economic Review 113*(2), 279–316.

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

Capelle, D. and K. Matsuda (2025). The great gatsby goes to college: Tuition, inequality and intergenerational mobility. Technical report.

Caunedo, J., D. Jaume, and E. Keller (2023, June). Occupational exposure to capital-embodied technical change. *American Economic Review 113*(6), 1642–85.

Cervellati, M., G. Meyerheim, and U. Sunde (2023). Human capital and the diffusion of technology. *Economics Letters 226*, 111108.

Chetty, R., J. N. Friedman, E. Saez, N. Turner, and D. Yagan (2020). Income segregation and intergenerational mobility across colleges in the united states. *Quarterly Journal of Economics 135*(3), 1567–1633.

College Board (2023). Trends in college pricing and student aid 2023. Technical report, College Board, New York.

Conesa, J. C., T. J. Kehoe, V. M. Nygaard, and G. Raveendranathan (2020). Implications of increasing college attainment for aging in general equilibrium. *European Economic Review 122*, 103363.

Deming, D. J. and K. Noray (2020, 06). Earnings dynamics, changing job skills, and stem careers. *The Quarterly Journal of Economics 135*(4), 1965–2005.

Dillon, E. W. and J. A. Smith (2020, None). The consequences of academic match between students and colleges. *Journal of Human Resources 55*(3), 767–808.

Epple, D., R. Romano, and H. Sieg (2006). Admission, tuition, and financial aid policies in the market for higher education. *Econometrica 74*(4), 885–928.

Freeman, J. A. and B. T. Hirsch (2008). College majors and the knowledge content of jobs. *Economics of Education Review 27*(5), 517–535.

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

Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal Tax Progressivity: An Analytical Framework. *Quarterly Journal of Economics 132*(4), 1693–1754.

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

Hendricks, L., T. Koreshkova, and O. Leukhina (2021, December). Causes and Consequences of Student-College Mismatch. Working Papers 2022-026, Federal Reserve Bank of St. Louis.

Hendricks, L., T. Koreshkova, and O. Leukhina (2024, Sep). College access and intergenerational mobility. Working Papers 2024-030, Federal Reserve Bank of St. Louis.

Hendricks, L. and T. Schoellman (2023, January). Skilled labor productivity and cross-country income differences. *American Economic Journal: Macroeconomics 15*(1), 240–68.

Kelchen, R., D. Ritter, and D. Webber (2024, December). *Predicting College Closures and Financial Distress*. University of Chicago Press.

Kogan, L., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller (2023, November). Technology and labor displacement: Evidence from linking patents with worker-level data. Working Paper 31846, National Bureau of Economic Research.

Krueger, D. and A. Ludwig (2016). On the Optimal Provision of Social Insurance: Progressive Taxation versus Education Subsidies in General Equilibrium. *Journal of Monetary Economics 77*, 72–98.

Krueger, D., A. Ludwig, and I. Popova (2025). Shaping inequality and intergenerational persistence of poverty: Free college or better schools? *Journal of Monetary Economics 150*, 103694. Labor Market Dynamics: Drivers and Consequences.

Krusell, P., L. E. Ohanian, J.-V. Ríos-Rull, and G. L. Violante (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica 68*(5), 1029–1053.

Leukhina, O. (2023). The Changing Role of Family Income in College Selection and Beyond. *Federal Reserve Bank of St. Louis Review 105*(3), 198–222. Third Quarter.

Ludwig, A., T. Schelkle, and E. Vogel (2012, January). Demographic change, human capital and welfare. *Review of Economic Dynamics 15*(1), 94–107.

Marto, R. and Y. Wittmann (2024). Risky college admissions and the (mis)allocation of talent. Technical report.

Mendoza, E. G., A. Razin, and L. L. Tesar (1994). Effective Tax Rates in Macroeconomics. Cross-Country Estimates of Tax Rates on Factor Incomes and Consumption. *Journal of Monetary Economics 34*(3), 297–323.

National Center for Education Statistics (2023). Digest of education statistics 2023. U.S. Department of Education.

Sunde, U. and T. Vischer (2015). Human capital and growth: Specification matters. *Economica 82*(326), 368–390.

Trabandt, M. and H. Uhlig (2011). The Laffer Curve Revisited. *Journal of Monetary Economics 58*(4), 305–327.

Appendix

A Dynamic Programming Problems for $j \geq 1$

To be Completed

B Theoretical Results and Detailed Derivations

B.1 (In-)Efficiency of the Education Allocation

B.1.1 Simplified Model

Consider a unit mas of people with initial cross-sectional probability measure $\Phi(e, b)$. The can choose education program $p \in 0, 1, 2$ at resource cost $\kappa(p)$ and education p confers labor productivity $\gamma(p)$. Households value consumption $c(e, b)$ according to strictly concave utility function $u(c)$. Production takes place according to aggregate production function

$$Y = \left[\sum_p L(p)^\nu\right]^{1/\nu}$$

where $L(p) = \int \gamma(p)\mathbf{1}_{p(e,b)=p} d\Phi(e, b)$ and $p(e, b)$ is the education allocation function (chosen by the social planner or the household directly.) Social welfare is given by $\mathcal{W} = \int \psi(e, b)U(c(e, b))d\Phi(e, b)$ where $\psi(e, b)$ is the Pareto weight placed on household of type (e, b), Utilitarian social welfare is a special case with $\psi(e, b) \equiv 1$.

B.1.2 Social Planner Solution

The social planner solves

$$\max_{c(e,b)\geq 0, p(e,b)\in\{0,1,2\}} \mathcal{W} = \int \psi(e, b)U(c(e, b))d\Phi(e, b) \qquad \text{s.t.}$$

$$\int c(e, b)d\Phi(e, b) + \sum_p \int \kappa(p)\mathbf{1}_{p(e,b)=p} d\Phi(e, b) = \left[\sum_p \left(\int \gamma(p)\mathbf{1}_{p(e,b)=p} d\Phi(e, b)\right)^\nu\right]^{1/\nu} + \int b d\Phi(e, b)$$

Denote the Lagrange multiplier on the aggregate resource constraint by λ. The planner makes two choices, how to allocate a given amount of aggregate consumption among people, and how to allocate the population across education programs (and thus labor markets). We are mostly interested in the last choice, and in the social planner problem the education choice is independent of the consumption allocation choice. This will not be the case in the constrained planner problem and Ramsey-type tax problems.

Consumption allocation The optimality condition with respect to $c(e,b)$ is

$$\psi(e,b)U'(c(e,b)) = \lambda \tag{57}$$

for all (e,b). Thus individual consumption is increasing in that group's Pareto weight. And if social welfare is utilitarian, then consumption is equalized across all individuals, so that $c(e,b) = C$ for all (e,b).

Education allocation For the education decision, under the assumptions made it is indeterminate and irrelevant who goes to college since neither characteristics (e,b) matters. We will relax the assumption leading to this result below. The planner can then choose the share $\Phi(p)$ of the population going to each program $p \in \{0,1,2\}$ in order to maximize net output

$$\max_{\{\Phi(p)\}} \left[\sum_p \left(\gamma(p)\Phi(p)\right)^\nu\right]^{1/\nu} - \sum_p \kappa(p)\Phi(p) \qquad \text{s.t.} \qquad \sum_p \Phi(p) = 1$$

Note that for $\nu = 1$ we obtain the corner solution that $\Phi(p^*) = 1$ for $p^* = \arg\max_p\{\gamma(p) - \kappa(p)\}$. For $\nu < 1$ we have the optimality conditions

$$\left[\sum_p \left(\gamma(p)\Phi(p)\right)^\nu\right]^{1/\nu-1} \gamma(\hat{p})^\nu \left(\Phi(\hat{p})\right)^{\nu-1} - \kappa(\hat{p}) = \mu$$

for all $\hat{p}$. This implies

$$\left[\sum_p \left(\gamma(p)\Phi(p)\right)^\nu\right]^{1/\nu-1} \left[\gamma(\hat{p})^\nu \left(\Phi(\hat{p})\right)^{\nu-1} - \gamma(\tilde{p})^\nu \left(\Phi(\tilde{p})\right)^{\nu-1}\right] = \kappa(\hat{p}) - \kappa(\tilde{p})$$

for any two $\hat{p}, \tilde{p}$. These two conditions, together with $\sum_p \Phi(p) = 1$, determines the efficient education allocation.

B.1.3 Constrained-Efficient Allocation

B.1.4 Implementation

B.2 Theoretical Properties of College Market Equilibrium

The two theoretical properties of the college market equilibrium stated below follow directly from Cai and Heathcote (2022), extended to our setting in which college is an investment good – that is, college quality q and program choice p affect graduates' future labor market productivity and earnings – and in which colleges offer distinct programs $p \in \mathcal{P}$ alongside quality q. Since the college supply side operates through the same competitive market

structure and CRS quality production function (19), and since neither the investment nature of college nor the program dimension affect the colleges' within-period optimization over inputs $(i, \{\eta_e\})$, the arguments of Cai and Heathcote (2022) go through in our setting with (q, p) replacing q throughout.

Tuition Independence of Initial Resources. Any equilibrium with tuition schedule $\tilde{t}(e, r, b; q, p)$ can be supported by an equivalent schedule that does not depend on initial household resources b, specifically $t^*(e, r; q, p) = v(q, p, e) - s(e, r; q, p)$.

Proof. Fix a college type (q, p) and an ability level e. Let (b^*, r^*) be a revenue-maximizing pair,

$$(b^*, r^*) \in \arg\max_{b,r} \left\{ \tilde{t}(e, r, b; q, p) + s(e, r; q, p) \right\},$$

so that $v(q, p, e) = \tilde{t}(e, r^*, b^*; q, p) + s(e, r^*; q, p)$. Since college quality depends only on $\bar{e}$ and i, a college of type (q, p) can always replace a non-revenue-maximizing student of ability e with a revenue-maximizing one without affecting technological feasibility. Therefore, for any $b^{**} \neq b^*$ with $\tilde{t}(e, r^*, b^{**}; q, p) < \tilde{t}(e, r^*, b^*; q, p)$, colleges strictly prefer not to admit type (e, r^*, b^{**}), so college demand for such students at (q, p) is zero. By market clearing, supply of such students at (q, p) must also be zero. We can therefore raise the tuition of type (e, r^*, b^{**}) to $\tilde{t}(e, r^*, b^*; q, p)$ without affecting the equilibrium: supply remains zero since students now face higher tuition, and college demand remains zero since colleges still prefer type (b^*, r^*). Repeating this for all $b \neq b^*$ yields $t(e, r^*, b; q, p) = v(q, p, e) - s(e, r^*; q, p)$ for all b, which is independent of b.

An analogous argument applies to residence status: if $r^{**} \neq r^*$ generates strictly lower total revenue, demand for type (e, r^{**}) at (q, p) is zero, supply is zero by market clearing, and we can raise tuition until $t(e, r^{**}, b; q, p) = v(q, p, e) - s(e, r^{**}; q, p)$ for all b. The resulting schedule $t^*(e, r; q, p) = v(q, p, e) - s(e, r; q, p)$ supports the same equilibrium allocation and is independent of b. □

Revenue Schedule Linear in Ability. Any equilibrium can be supported by a revenue schedule that is linear in student ability e:

$$v(q, p, e) = b(q, p) - d(q, p) \left(e - e^l\right),$$

where $b(q, p) = v(q, p, e^l)$ is the sticker revenue from the lowest-ability student and $d(q, p) > 0$ is the equilibrium ability discount defined in (24).

Proof. By the previous result we work with the income-independent schedule, so it suffices to show $v(q,p,e)$ is linear in e. Fix (q,p) with $\chi(q,p)>0$ and let $A^+(q,p)=\{e:\eta_e(q,p)>0\}$ denote the set of ability types with positive enrollment. Suppose $A^+(q,p)$ contains more than one element, with maximum and minimum elements $e_{\max}$ and $e_{\min}$. Define

$$d(q,p)=-\frac{v(q,p,e_{\max})-v(q,p,e_{\min})}{e_{\max}-e_{\min}},\qquad b(q,p)=v(q,p,e^l)+d(q,p)(e^l-e_{\min}).$$

We claim $v(q,p,e)=b(q,p)-d(q,p)(e-e^l)$ for all e. Suppose not: then for some $e_j\in A^+(q,p)$ we have $v(q,p,e_j)>b(q,p)-d(q,p)(e_j-e^l)$, i.e., type e_j generates more revenue than the linear schedule would imply. A college could then increase profits by admitting more students of type e_j and substituting out a combination of $e_{\min}$ and $e_{\max}$ students chosen to keep $\bar{e}$ - and hence quality q - unchanged. This deviation is strictly profitable, contradicting zero profits in equilibrium. By a symmetric argument, if $v(q,p,e_j)<b(q,p)-d(q,p)(e_j-e^l)$ for some e_j, then $\eta_{e_j}(q,p)=0$, so $e_j\notin A^+(q,p)$; we can freely redefine $v(q,p,e_j)$ to lie on the linear schedule without affecting the equilibrium allocation. This establishes linearity for all active ability types and, by the redefinition argument, for all e.

Interior case. When $\eta_e>0$ for all e (so $A^+(q,p)=\mathcal{E}$), linearity follows directly from the first-order conditions. Letting $\lambda\geq 0$ denote the multiplier on (19) and μ the multiplier on $\sum_e\eta_e=1$, the first-order condition with respect to η_e gives

$$v(q,p,e)=\mu-\lambda\,\frac{\theta q}{\bar{e}}\,e,$$

using $\frac{\partial q}{\partial\bar{e}}=\frac{\theta q}{\bar{e}}$ from (19), which is immediately linear in e. In either case, setting $d(q,p)\equiv\lambda\frac{\theta q}{\bar{e}}$ and using the first-order condition with respect to i, which gives $\lambda=\frac{i}{(1-\theta)q}$, yields $d(q,p)=\frac{\theta}{1-\theta}\frac{i}{\bar{e}}$, consistent with (24). □

B.3 Why College Heterogeneity Shapes the Response to Aging: Two Margins

This appendix formalizes the claim in Section 5.6 that collapsing the quality and major dimensions of the college market dampens the economy's education response to demographic change through two distinct margins. Rather than restating the full dynamic problem, we condense the enrollment decision into a static one: the value of attending college at a given quality is summarized by a single surplus function, in which the scalar type x stands in for the household's multi-dimensional state (ability, parental resources, preference draws), $h(q)$ for the lifetime earnings value of college quality, $t(q)$ for the equilibrium tuition schedule, and R for the bundle of equilibrium price changes through which the demographic cliff operates (skill prices and the pension replacement rate). Each object below is thus a stand-in for an

equilibrium value of the quantitative model, and the propositions are comparative statics on this condensed problem: they isolate the mechanisms behind the quantitative comparison.

Setup

Households are indexed by a scalar type x with distribution function $F(x)$ and density $f(x)$, combining ability and parental resources.[13] A college of quality q charges tuition $t(q)$, with $t'(q) > 0$. The surplus from enrolling at quality q when the aggregate return to skill is R is

$$S(x, q; R) \;=\; R\, v(x)\, h(q) \;-\; \varphi\big(t(q) - b(x)\big), \tag{58}$$

where the first term is the labor-market payoff of attending and the second the cost of financing it. Specifically, $h(q)$ is the effective labor supplied by a graduate of a quality-q college, with $h'(q) > 0$: better colleges produce more effective labor. The function $v(x)$ scales this payoff by household type, with $v'(x) > 0$: better types benefit more from any given college. The price R is the aggregate return to skill at which effective college labor is valued. On the cost side, a student pays tuition $t(q)$ out of the resources $b(x)$ available at entry, with $b'(x) \geq 0$: better types come from weakly better-funded families. The utility cost of financing the difference is $\varphi\big(t(q) - b(x)\big)$, with φ increasing and weakly convex, and the value of not attending is normalized to zero. A linear φ describes financing at a constant marginal cost, under which a dollar of tuition costs a dollar whoever pays it; a strictly convex φ stands in for the borrowing limits and consumption smoothing of the quantitative model, under which the marginal cost of tuition rises as own resources are exhausted. The demographic cliff operates through the return to skill R: labor scarcity and the decline in the pension replacement rate raise the return to skill.

In the *heterogeneous college* economy households choose $q \in [q_\ell, q_h]$; enrollment is $N(R) = 1 - F(\bar{x})$, where the cutoff $\bar{x}$ satisfies $S(\bar{x}, q^*(\bar{x}); R) = 0$ and $q^*(x)$ is the chosen quality. In the *representative college* economy the menu is the singleton $\{\bar{q}\}$ at tuition $t(\bar{q})$. Under unchanged preferences the two economies do not enroll the same number of students: replacing the quality menu with a single mid-quality college changes who finds college worthwhile, so enrollment levels differ before any shock hits. Comparing responses across economies that start at different levels would confound the effect of the menu with a mechanical level effect: an economy whose enrollment is already higher has fewer households left near the margin, so

[13] The household type in the quantitative model is of course multi-dimensional. The scalar index should be read as movement along the direction in which both the value of college and the resources available to pay for it increase — the empirically relevant direction, since ability and initial resources are positively associated in our calibration. The collapse is for transparency only: with a multi-dimensional type the enrollment margin is a boundary in the state space rather than a cutoff, the decomposition (59) holds with the extensive term as a boundary integral and the intensive term unchanged, and both propositions apply pointwise along that boundary.

its enrollment responds less for that reason alone. We therefore adopt the same adjustment as the recalibrated row of Table 7. Let R_0 denote the benchmark value of the return to skill — its level before the demographic shock moves it. The representative economy's utility cost of college is recalibrated so that its enrollment cutoff at R_0 coincides with the heterogeneous economy's cutoff $\bar{x}$. This utility cost is an additive shifter of the surplus (58): recalibrating it aligns the cutoff but leaves every derivative of the surplus — and hence the financing wedge defined below — unchanged. The propositions then compare the two economies' responses to the same change in R from a common starting point: the same marginal household $\bar{x}$, and hence the same mass of households $f(\bar{x})$ at the enrollment threshold. Any difference in responses then reflects the structure of the menu rather than different starting levels of enrollment. Effective college labor is $L(R) = \int_{\bar{x}}^{\infty} h(q^*(x))\, dF(x)$, with $q^* \equiv \bar{q}$ in the representative economy. Throughout, the tuition schedule $t(\cdot)$ is held fixed across menus, the representative college charging the baseline schedule evaluated at $\bar{q}$; in the quantitative comparison the schedule itself re-prices in equilibrium, a channel that operates through both margins below and is discussed in Section 5.6.

The first step is that students sort: better types attend better colleges. The property of the surplus that delivers this is a complementarity between type and quality — the gain from attending a better college is larger for a better type. This is the familiar single-crossing condition, stated in its stronger differential form.

Lemma 1. *The surplus $S(x, q; R)$ has a strictly positive cross-partial derivative in type and quality,*

$$S_{xq} = R\, v'(x)\, h'(q) + \varphi''\big(t(q) - b(x)\big)\, b'(x)\, t'(q) \;>\; 0 :$$

the gain from attending a better college is strictly larger for a better type. Hence the chosen quality $q^*(x)$ is nondecreasing in type, and the marginal entrant chooses the bottom of the quality ladder attained in equilibrium.

Proof. The first term of S_{xq} is strictly positive since $v'(x) > 0$ and $h'(q) > 0$; the second is nonnegative since φ is weakly convex, $b'(x) \geq 0$ and $t'(q) > 0$.

To see that a higher type never chooses a lower quality, suppose the opposite: let $x_2 > x_1$ be two types whose optimal choices $q_2 = q^*(x_2)$ and $q_1 = q^*(x_1)$ satisfy $q_2 < q_1$. Because each type prefers its own choice, type x_2 weakly prefers q_2 to q_1 and type x_1 weakly prefers q_1 to q_2. Read together, these two statements say that the upgrade from the lower quality q_2 to the higher quality q_1 weakly benefits the *lower* type but weakly harms the *higher* type:

$$S(x_1, q_1) - S(x_1, q_2) \;\geq\; 0 \;\geq\; S(x_2, q_1) - S(x_2, q_2).$$

But the positive cross-partial says the opposite: the gain from the same quality upgrade is strictly *larger* for the higher type,

$$S(x_2, q_1) - S(x_2, q_2) \;>\; S(x_1, q_1) - S(x_1, q_2),$$

as follows from integrating $S_{xq} > 0$ over $[x_1, x_2] \times [q_2, q_1]$. The two displays contradict each other, so $q_2 \geq q_1$: the quality choice $q^*(x)$ is nondecreasing. Finally, since the marginal entrant $\bar{x}$ is the lowest type that enrolls, $q^*(\bar{x})$ is the lowest quality attended in equilibrium. □

Effective college labor responds to the return along two margins, which the following decomposition separates:

$$\frac{dL}{dR} \;=\; \underbrace{h\big(q^*(\bar{x})\big)\, f(\bar{x}) \left|\frac{d\bar{x}}{dR}\right|}_{\text{extensive}} \;+\; \underbrace{\int_{\bar{x}}^{\infty} h'\big(q^*(x)\big)\, \frac{\partial q^*(x)}{\partial R}\, dF(x)}_{\text{intensive}}. \tag{59}$$

The first term counts new students entering at the margin, each contributing the effective units of the quality they enter with; the second counts already-enrolled students upgrading along the quality menu. The decomposition itself is an accounting identity — Leibniz's rule applied to $L(R)$ — and in the representative college economy the intensive term is zero *by construction*: with a singleton menu there is no quality choice to adjust. What requires proof is that this term is not innocuous: in the heterogeneous college economy it is strictly positive, so collapsing the menu removes a strictly positive part of the response.

Proposition 3. *In the heterogeneous college economy, every student with an interior quality choice upgrades when the return rises: $\partial q^*(x)/\partial R = -S_{qR}/S_{qq} > 0$, and hence the intensive term in (59) is strictly positive. The major margin works in the same way: if the STEM premium rises with the return R — as capital–skill complementarity implies — students reallocate toward STEM, a response that equalizing majors removes as well.*

Proof. An interior quality choice satisfies the first-order condition $S_q\big(x, q^*(x); R\big) = 0$, which by (58) equates the marginal labor-market benefit of quality to its marginal financing cost:

$$R\, v(x)\, h'\big(q^*(x)\big) \;=\; \varphi'\big(t(q^*(x)) - b(x)\big)\, t'\big(q^*(x)\big).$$

Differentiating this condition totally with respect to R gives $S_{qq}\, \partial q^*/\partial R + S_{qR} = 0$, hence $\partial q^*/\partial R = -S_{qR}/S_{qq}$. The cross-derivative is $S_{qR} = v(x)\, h'\big(q^*(x)\big) > 0$: a higher return raises the marginal benefit of quality while leaving the financing side unchanged. The second-order condition of an interior optimum gives $S_{qq} < 0$. Hence $\partial q^*(x)/\partial R > 0$: every student with an

interior choice upgrades when the return rises. The intensive integrand $h'\big(q^*(x)\big)\,\partial q^*(x)/\partial R$ is therefore strictly positive for every such student, and since a positive mass of enrolled types chooses interior qualities in equilibrium — attended qualities are dispersed across the entire menu in the computed equilibria (Figure 6) — the intensive term in (59) is strictly positive. □

Proposition 4. *Let q_m denote the marginal entrant's bundle ($q_m = q^*(\bar{x})$ in the heterogeneous economy, $q_m = \bar{q}$ in the representative economy) and define the financing wedge per efficiency unit $\omega(q_m) \equiv \varphi'\big(t(q_m) - b(\bar{x})\big)/h(q_m)$. Then*

$$\left|\frac{d\bar{x}}{dR}\right| \;=\; \frac{v(\bar{x})}{R\,v'(\bar{x}) \;+\; \omega(q_m)\,b'(\bar{x})}. \tag{60}$$

(i) Neutrality benchmark. *The bundle q_m enters (60) only through the wedge term $\omega(q_m)\,b'(\bar{x})$. If this cost-side type gradient is negligible — either because the marginal financing cost $\varphi'\big(t(q_m) - b(\bar{x})\big)$ is negligible at the margin, or because own resources do not vary with type at the margin, $b'(\bar{x}) \simeq 0$ — then (60) reduces to $|d\bar{x}/dR| = v(\bar{x})/\big(R\,v'(\bar{x})\big)$, an expression that does not involve the bundle q_m at all: the enrollment response is then the same under any menu. This neutrality is a property of the multiplicative payoff form $R\,v(x)\,h(q)$: the benefit and its slope in type scale by the same factor $h(q_m)$, which therefore cancels from the ratio. The multiplicative form is not a simplification made for tractability here; the model's log-linear earnings function (5) delivers exactly this structure. A linear φ alone does not deliver neutrality: even at a constant marginal financing cost, family resources $b(x)$ still create the cost-side type gradient, so the wedge term remains. Since $\omega(q_m)$ then falls in $h(q_m)$, a compulsory mid-quality bundle would, if anything, amplify the enrollment response.* (ii) Dampening. *If the cost-side gradient is present, the representative economy has the smaller enrollment response if and only if*

$$\frac{\varphi'\big(t(\bar{q}) - b(\bar{x})\big)}{\varphi'\big(t(q^*(\bar{x})) - b(\bar{x})\big)} \;>\; \frac{h(\bar{q})}{h(q^*(\bar{x}))}, \tag{61}$$

that is, if the financing burden rises faster than the quality content between the entry rung and the compulsory bundle. Since $t(\bar{q}) > t(q^*(\bar{x}))$ by Lemma 1, the marginal financing cost $\varphi'(z)$ is evaluated at a larger financed amount under the compulsory bundle, and (61) holds whenever $\varphi'(z)$ rises sufficiently fast over the range $z \in \big[\,t(q^*(\bar{x})) - b(\bar{x}),\; t(\bar{q}) - b(\bar{x})\,\big]$. This is the empirically relevant case: in the quantitative model the marginal cost of tuition rises through the concavity of utility and, especially, through borrowing limits, which marginal entrants are at or near.

Proof. The cutoff is defined by zero surplus at the marginal entrant's bundle, $S(\bar{x}, q_m; R) = 0$. In the heterogeneous college economy the bundle $q_m = q^*(\bar{x})$ is itself optimally chosen, but by the envelope theorem its adjustment to R has no first-order effect on the maximized surplus, so we may differentiate holding q_m fixed. Differentiating the identity totally with respect to R gives $S_x\, d\bar{x}/dR + S_R = 0$, hence $d\bar{x}/dR = -S_R/S_x$. From the surplus (58), the two partial derivatives at the margin are

$$S_R = v(\bar{x})\, h(q_m), \qquad S_x = R\, v'(\bar{x})\, h(q_m) + \varphi'\big(t(q_m) - b(\bar{x})\big)\, b'(\bar{x}),$$

where the second term of S_x reflects that a higher type brings more own resources: the gap $t(q_m) - b(x)$ between tuition and own resources — the amount that must be financed — shrinks at rate $b'(\bar{x})$, which lowers the financing cost $\varphi\big(t(q_m) - b(x)\big)$ at the marginal rate $\varphi'\big(t(q_m) - b(\bar{x})\big)$. Both derivatives are positive, so $d\bar{x}/dR < 0$: a higher return moves the cutoff down and enrollment expands. Dividing numerator and denominator of $|d\bar{x}/dR| = S_R/S_x$ by $h(q_m)$ and using the definition of the financing wedge $\omega(q_m) = \varphi'\big(t(q_m) - b(\bar{x})\big)/h(q_m)$ yields (60).

For the comparison across menus, recall that by the recalibration convention the two economies share the same marginal household $\bar{x}$ and face the same return R, and that the recalibrated utility cost is additive and so appears in no derivative of the surplus. Every term of (60) is therefore identical across the two economies except the wedge: $\omega(q^*(\bar{x}))$ under the menu, $\omega(\bar{q})$ under the representative college. Since (60) is strictly decreasing in the wedge, the representative economy has the smaller enrollment response if and only if $\omega(\bar{q}) > \omega(q^*(\bar{x}))$; writing out the two wedges and rearranging gives exactly (61). □

Two further observations complete the argument. First, the two economies differ not only in how many students the return draws in, but in what each entrant contributes: under the menu, return-induced entrants arrive at the bottom of the quality ladder with few effective units, whereas every entrant in the representative economy carries the full bundle $h(\bar{q})$. The menu therefore wins on volume and loses on per-head units. Table efgerepq reports the combined effect of shutting down both margins rather than either on its own, but it shows that the per-head advantage is not large enough to offset them.

Second, Proposition 4 is a statement about the cost-side type gradient and financing frictions rather than preferences, and it is falsifiable within the model. Relaxing the borrowing limit in the representative-college comparison should attenuate the convexity-induced dampening of the *enrollment* response; the gap disappears where the cost-side gradient is

negligible. The gap in the *upgrading* response of Proposition 3, by contrast, survives by construction.[14]

Sufficient Statistics and Empirical Bounds

Each of the two propositions requires one substantive condition to hold, and both conditions can be stated as bounds on observables. Moreover, in our calibration both sides of each bound are disciplined by targeted or directly reported calibration moments. Whether the conditions hold is therefore determined by the data the model is calibrated to, not by a free modeling choice.

The intensive margin. Proposition 3 requires only that the quality choice responds to prices and returns for a positive-measure set of students: $\partial q^*/\partial R > 0$. Its exact counterpart is the within-student response of chosen quality with the enrolled population held fixed, an object the subsidy experiment does not deliver directly: the aggregate response of mean attended quality mixes within-student adjustment with the changing composition of the enrolled and the equilibrium re-pricing of the menu. The experiment is nonetheless informative. Eliminating the average 25% subsidy raises net tuition mechanically by 33% and lowers mean attended quality by about 12%, an elasticity of roughly -0.4. The composition effect runs the other way: by the sorting of Lemma 1, the students a price increase pushes out attend the lowest qualities, and their exit by itself *raises* mean attended quality. The observed decline therefore implies that within-student downgrading more than offsets the composition effect. Table 7 quantifies the contribution of this margin to the aggregate response: the representative-college economy differs from the baseline precisely in removing it.

The extensive margin. The dampening condition (61) compares the growth of the financing burden with the growth of the earnings value between the marginal entrant's rung q_e and the compulsory bundle $\bar{q}$. Two ratios summarize this comparison. The first is the out-of-pocket cost ratio $\big(t(\bar{q}) - b\big)/\big(t(q_e) - b\big)$: how much more of their own financing the mid-quality bundle requires from the marginal entrant than the entry rung does, where b denotes the entrant's mean initial transfers. The second is the value ratio $h(\bar{q})/h(q_e)$: how much more effective labor the bundle delivers than the rung. The dampening condition asks whether the first ratio, raised to the curvature of the financing cost, exceeds the second. To make the curvature a single number, take an isoelastic marginal financing cost, $\varphi'(z) \propto z^{\varepsilon_\varphi}$

[14] The same financing margin is what gives the demographic cliff's bequest channel its force in Section 5.7.2. Smaller cohorts inherit more per head, and the windfall relaxes precisely the constraint $\varphi'\big(t(q) - b(x)\big)$ that Proposition 4 identifies as the source of the dampened enrollment response. This is also why the cliff's enrollment response is larger when subsidies are absent: with tuition paid in full, more marginal students sit against the constraint that the inherited windfall loosens, so the same windfall buys more enrollment.

with $\varepsilon_\varphi \geq 0$: the exponent ε_φ is the elasticity of the marginal financing cost with respect to the amount financed — how fast the burden of one more tuition dollar escalates as the financed amount grows. At $\varepsilon_\varphi = 0$ the marginal financing cost is constant — φ is linear — and the condition below fails: dampening requires a marginal financing cost that rises with the amount financed. If the financing cost is the utility cost of forgone consumption under CRRA preferences with risk aversion σ, then ε_φ equals σ times the ratio of the financed amount to the consumption that remains, so log utility with tuition absorbing half of remaining resources gives $\varepsilon_\varphi \approx 0.5$, and ε_φ grows without bound as a borrowing limit is approached. Condition (61) then reads

$$\left(\frac{t(\bar{q}) - b}{t(q_e) - b}\right)^{\varepsilon_\varphi} > \frac{h(\bar{q})}{h(q_e)} \qquad \Longleftrightarrow \qquad \varepsilon_\varphi > \frac{\log\big(h(\bar{q})/h(q_e)\big)}{\log\big((t(\bar{q}) - b)/(t(q_e) - b)\big)}. \tag{62}$$

Both ratios are observable, and in the model both are disciplined by targeted or directly reported calibration moments. On the cost side, evaluating the equilibrium tuition schedule in 2025 at the lowest attended qualities ($q_e \in [0.66, 0.92]$) and at mean attended quality ($\bar{q} \approx 2.0$) gives a gross tuition ratio of 2.2–2.6; net of mean initial transfers ($\bar{b} = \$61{,}200$, the PSID moment) the out-of-pocket ratio is 3.3–4.7; applying the average subsidy first — students pay 75% of gross tuition — raises this ratio further, to 4.2–7.6. On the value side, the earnings function (5) with the calibrated saturation ($\beta_1 = 0.55$, $\beta_3 = 1.5$, disciplined by the average college premium and the bottom-versus-top-quartile premium gap) implies $h(\bar{q})/h(q_e) = 1.12$–1.20. The bound (62) therefore requires only $\varepsilon_\varphi \gtrsim 0.15$ — far below the curvature implied by log utility, let alone by binding borrowing limits. The dampening result is thus not knife-edge: tuition rises across the low-to-mid quality range by a factor of three to five in out-of-pocket terms while the earnings value rises by less than a fifth, and both gradients are features of the U.S. college market that discipline the calibration, not free choices.[15] Both sides of the bound also line up with their direct empirical counterparts, which we record here for later reference. On the cost side, published tuition and fees in 2023–24 average \$3,990 at public two-year colleges, \$11,260 at public four-year institutions (in-state), and \$41,540 at private nonprofit institutions (College Board, 2023): the gradient from the entry rung to the average attended four-year institution is a factor of 2.8 to 4.8 in sticker prices, if anything steeper than the model's 2.2–2.6, and subtracting a common family transfer from both prices raises the out-of-pocket ratio further, exactly as in the model. On the value side, quasi-experimental estimates of the causal effect of college quality

[15] This comparison speaks to the entry margin — why the compulsory bundle burdens the marginal entrant — and not against the upgrading margin of Proposition 3: an enrolled student equates *marginal* benefit and *marginal* cost at the chosen quality, so a higher return induces upgrading regardless of how steep the schedule is, as the quality-price elasticity above confirms quantitatively.

on earnings are on the order of 4 to 9 percent per standard deviation of a quality index (Black and Smith, 2006; Dillon and Smith, 2020); over the roughly one to one-and-a-half standard deviations separating the entry rung from mean attended quality, this implies a value ratio of about 1.06 to 1.15, close to the model's 1.12–1.20. Descriptive earnings gaps across college tiers are considerably larger (Chetty et al., 2020), as they include the sorting of students across colleges; since $h(q)$ prices quality conditional on type, with sorting modeled separately, the causal estimates are the appropriate comparison. In sum, the calibrated model reproduces the two features of the U.S. college market that drive this mechanism — the steep tuition gradient and the flat earnings-value gradient across the lower half of the quality distribution — and condition (62) holds comfortably whether evaluated at the calibrated or at the measured values. The dampened enrollment response under a representative college of Proposition 4 is thus grounded in the measured structure of the market.

The remaining requirement — that enrollment at the margin responds to net price at rates consistent with the quasi-experimental financial-aid evidence — is the standard aid semi-elasticity, which in the model is disciplined jointly by the enrollment-rate and tuition-spending targets and is visible in the on-impact response of the subsidy-cut transition.

C Quantitative Model Calibration Appendix

C.1 Demographic Model

Derivation of Equation (55). Use (1) in (54) to get

$$N_{t+1,0} = \bar{\zeta} \sum_{j=0}^{J} \zeta_j \prod_{i=0}^{j-1} \psi_i N_{t-j,0}. \tag{63}$$

Next, assume a steady state with population growth rate n to get

$$\begin{aligned} N_{t+1,0} &= \bar{\zeta} \sum_{j=0}^{J} \zeta_j \prod_{i=0}^{j-1} \psi_i \frac{N_{t+1,0}}{(1+n)^{1+j}} \\ \Leftrightarrow \qquad \bar{\zeta} &= \frac{1}{\sum_{j=0}^{J} \zeta_j \prod_{i=0}^{j-1} \psi_i \frac{1}{(1+n)^{1+j}}}. \end{aligned}$$

C.2 Asymmetric Effects of Technological Change on College Programs: Mapping College Majors to Tasks

Each college program $p \in \mathcal{P}$ can be characterized by a bundle of tasks — for example, abstract, routine, manual, and interpersonal tasks. We construct a task profile for each program by mapping its curricular content to occupational task requirements, drawing on task

measures from sources such as DOT/O*NET. Comparing the task demands of occupations with the curricular content of college majors to establish such a mapping has precedent in the literature; see, for example, Freeman and Hirsch (2008) and Altonji et al. (2014). Relative to an approach based on the occupational destinations of graduates, a curriculum-based mapping more directly captures the skills a program imparts, abstracting from selection into particular occupations.

This mapping allows us to introduce asymmetric effects of technological change across college programs parsimoniously, without altering the model structure. Skill-biased or STEM-biased technological change — including the emergence of AI — manifests as a shift in the productivity of certain task bundles. Because programs differ in their task profiles, such shifts generate differential changes in labor market returns across majors, affecting both college program choices and the equilibrium distribution of college quality. The model's supply and demand structure remains unchanged; only the calibration of program-specific productivity parameters varies across technological scenarios.

D Additional Results

This appendix reports the transition induced by the subsidy cut alone, and collects additional series along the transitions of Section 5.7.

Transitional Dynamics: Subsidy Cut

This section reports the transition induced by the subsidy cut alone, under the conventions of Section 5.7: public subsidies to higher education are eliminated in 2026 (the model period closest to 2025) as an unanticipated (MIT) shock on the baseline demographic path, and households have perfect foresight thereafter. Figure A1 reports the paths of aggregate labor in efficiency units, output and the interest rate; Figure A2 those of college enrollment, the labor income tax rate and output per capita.

The enrollment response is immediate and large: the college share of the entering cohort falls by 21.9 percentage points on impact, as the direct price channel operates within the entry period. The gap widens to 29.0 percentage points by 2086 and 31.9 percentage points by the end of the plotting window. The fiscal path, in contrast, is the mirror image of the timing pattern of Section 5.7: on impact the budget *improves* — the labor income tax rate runs 0.9 percentage points below the baseline path, since the subsidy outlay disappears at once while the missing graduates reach the labor market only cohort by cohort. The tax gap turns positive around 2040, reaches 1.2 percentage points in 2050, 7.4 percentage points in 2086, and 10.8 percentage points by the end of the plotting window. Output per capita falls by 7.6 percent relative to the baseline path by 2050, 21.4 percent by 2086, and 27.5 percent

by the end of the window. The subsidy cut is thus self-financing over a political horizon and costly over a generational one: the fiscal benefits are front-loaded, the damage arrives with the smaller and less skilled cohorts.

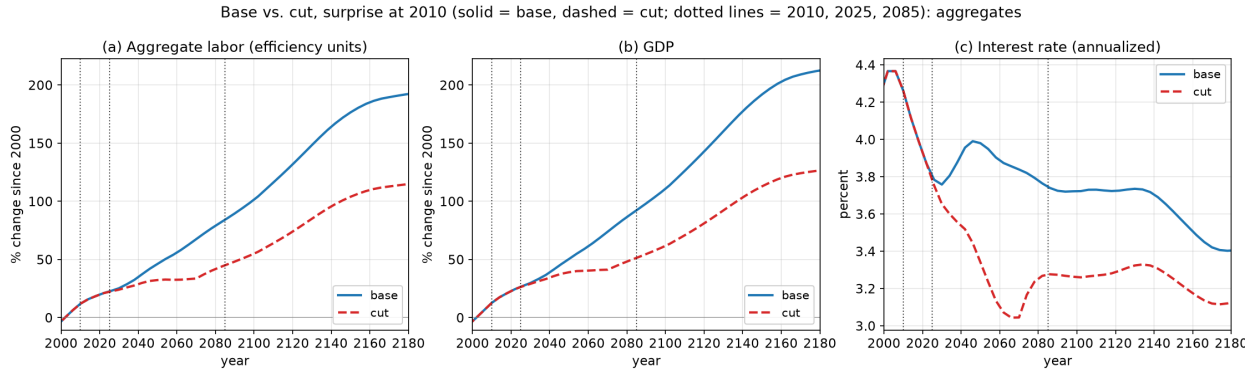


Figure A1: Aggregates along the surprise subsidy-cut transition. Baseline (solid) versus subsidy cut (dashed); the cut is learned and implemented in 2026. Panels: aggregate labor in efficiency units; GDP; and the interest rate. Labor and GDP are expressed as percentage changes since 2000 and share a common vertical axis; the interest rate is annualized and reported in percent. Dotted vertical lines mark 2010, 2025 and 2085; the plotting window is 2000–2180.

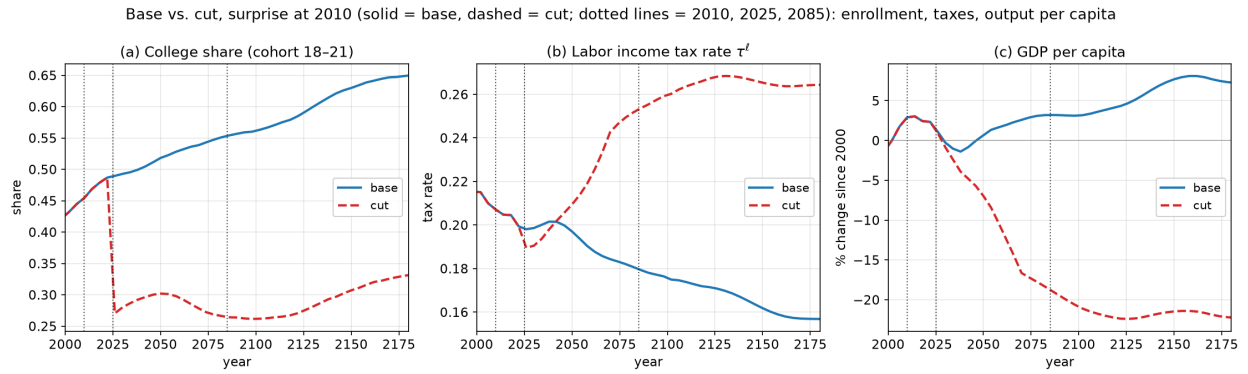


Figure A2: Enrollment, taxes and output per capita along the surprise subsidy-cut transition. Panels: the college share of the 18–21 cohort (completers plus dropouts); the labor income tax rate $\tau^{\ell} = 1 - \lambda$; and output per capita, expressed as a percentage change since 2000. Conventions as in Figure A1. The remaining per-capita series and the college-market variables are reported below.

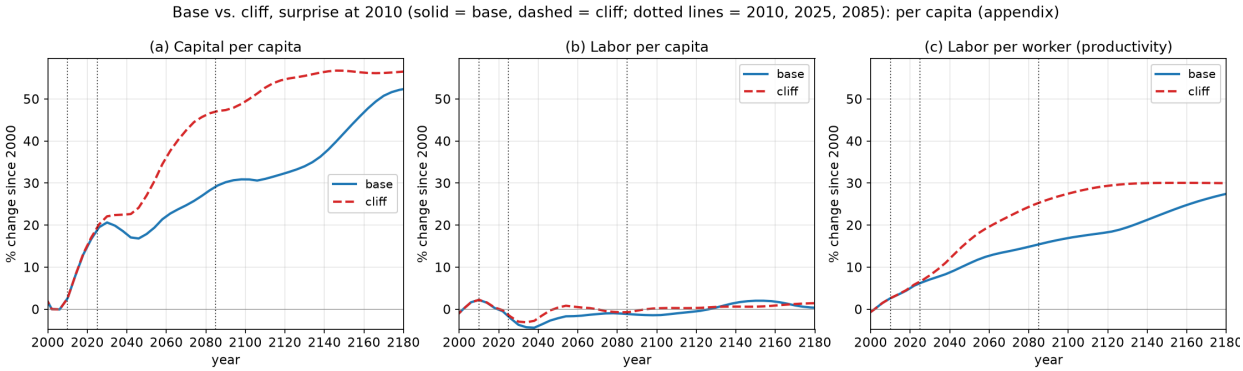


Figure A3: Additional per-capita series along the surprise-cliff transition. Baseline (solid) versus demographic cliff (dashed). Panels: capital per capita; labor per capita (aggregate efficiency units divided by the adult population); and labor per worker (efficiency units divided by the working-age population, i.e. average efficiency per worker). All three are expressed as percentage changes since 2000 and share a common vertical axis. Conventions as in Figure 13.

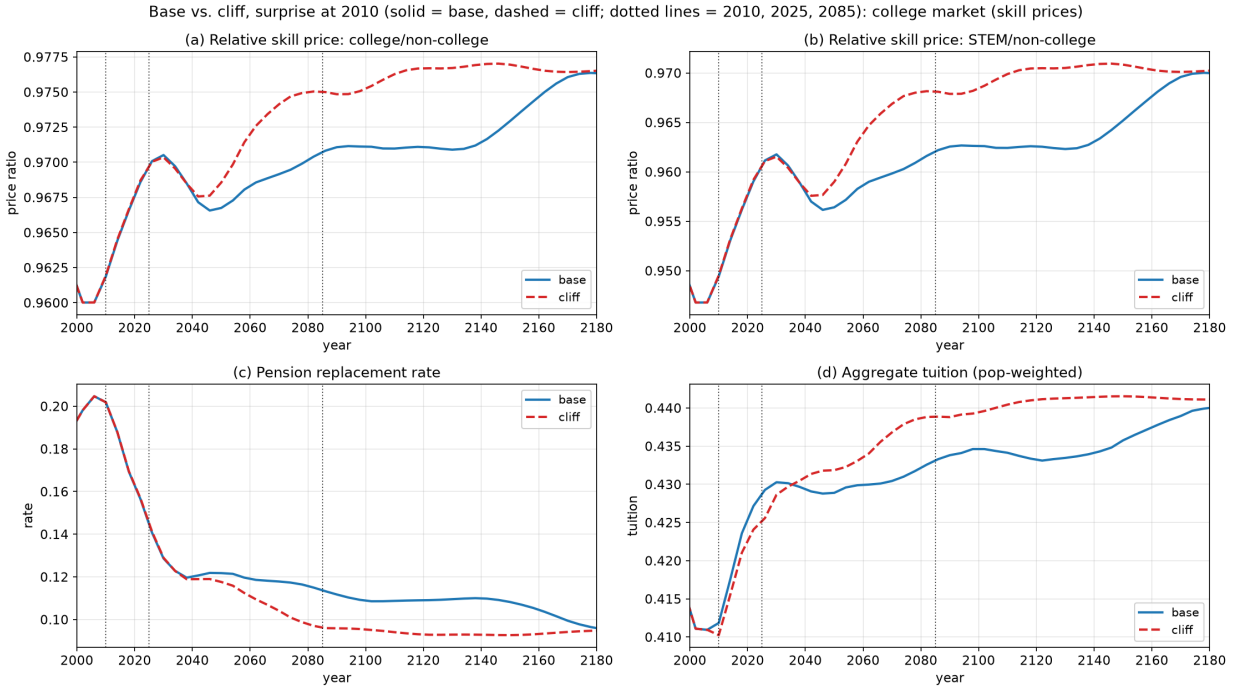


Figure A4: The college market along the surprise-cliff transition. Baseline (solid) versus demographic cliff (dashed). Panels: the relative skill price of college non-STEM labor (per efficiency unit, relative to non-college labor); the relative skill price of STEM labor; the pension replacement rate; and population-weighted aggregate tuition. Skill prices and the replacement rate are converged equilibrium objects of the transition solution. Conventions as in Figure 13. The two paths coincide before 2010 by construction, since the cliff is unanticipated until then.

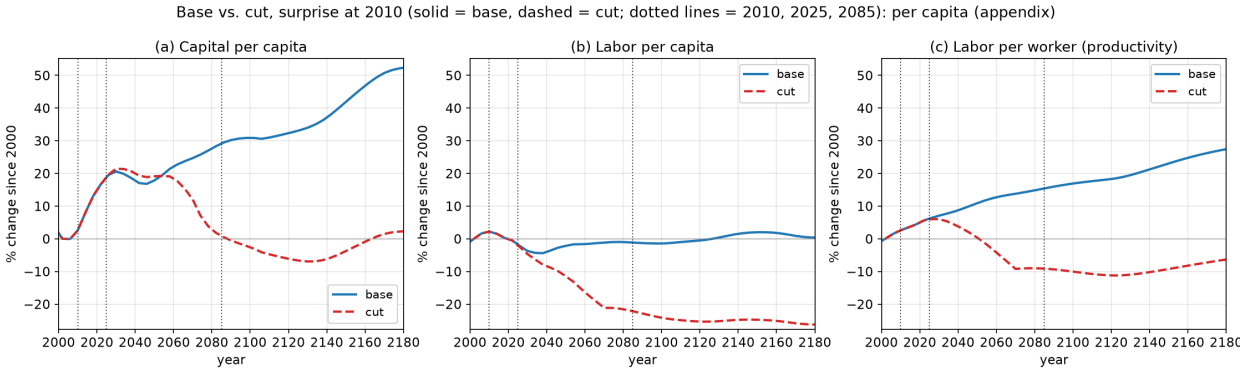


Figure A5: Additional per-capita series along the subsidy-cut transition. Panels and conventions as in Figure A3.



Figure A6: The college market along the subsidy-cut transition: relative skill prices, the pension replacement rate, and population-weighted aggregate tuition. Conventions as in Figure A4.

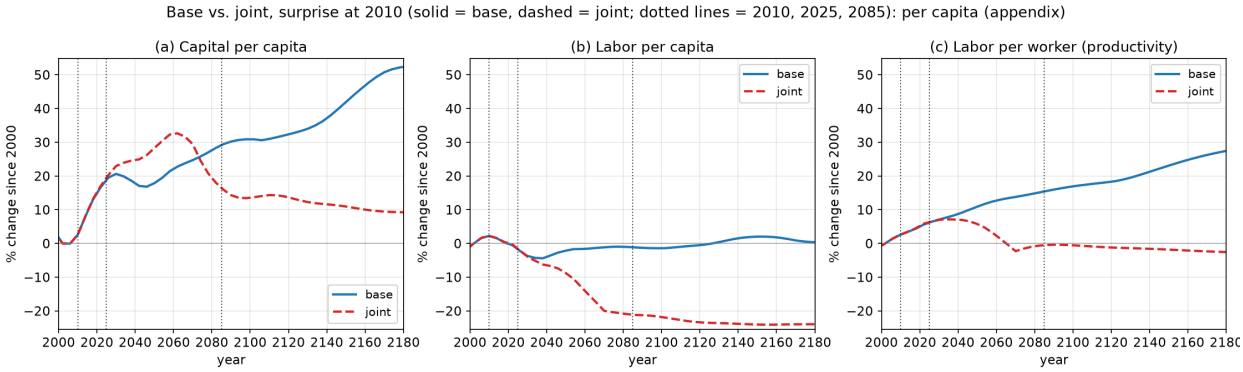


Figure A7: Additional per-capita series along the joint transition. Panels and conventions as in Figure A3.

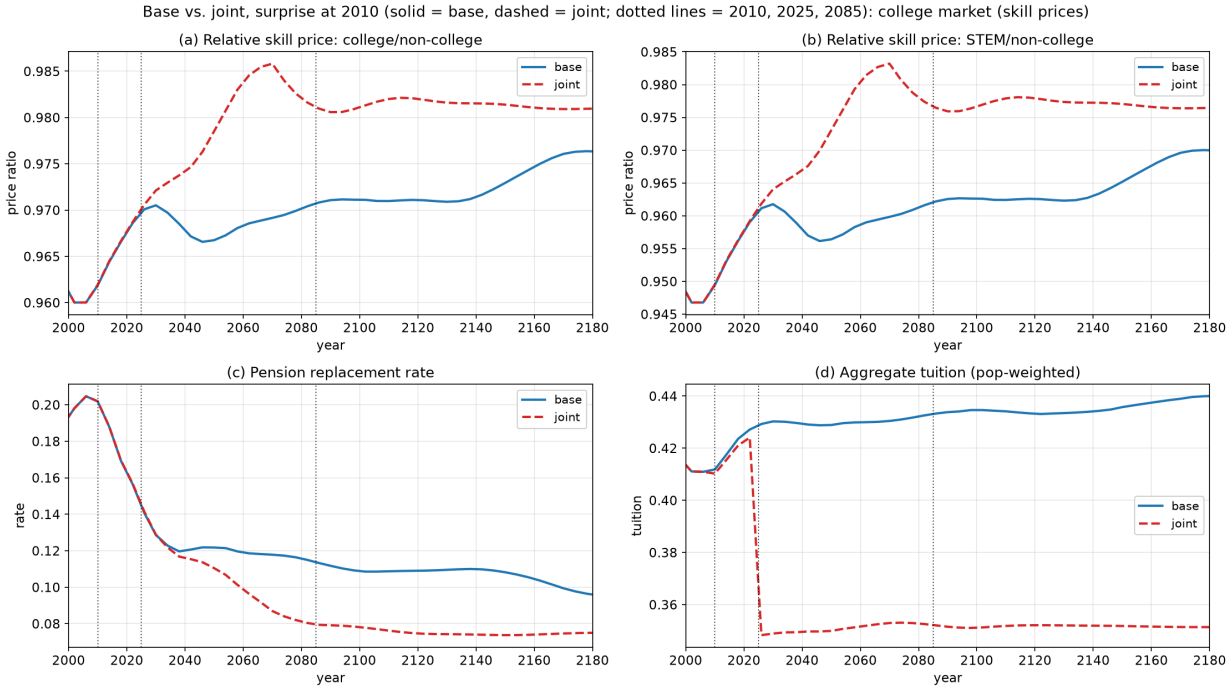


Figure A8: The college market along the joint transition: relative skill prices, the pension replacement rate, and population-weighted aggregate tuition. Conventions as in Figure A4.

E Computational Appendix

E.1 Solution Algorithm: College Market Equilibrium

In the current discrete implementation, college quality q and program p are both discrete, and households choose the pair (q, p) jointly. The key equilibrium pricing object is the

per-unit ability discount

$$d(q,p) > 0, \tag{64}$$

which determines the slope of the revenue schedule in ability. By the college-side optimality conditions, revenues are affine in ability:

$$v(q,p,e) = b(q,p) - d(q,p)e, \qquad e \in \mathcal{E}, \tag{65}$$

where $b(q,p)$ is an intercept pinned down by zero profit.

The pass-through from college subsidies to students' net tuition payments is 100% by construction. Accordingly, the household problem is solved using the gross tuition schedule $v(q,p,e)$, while the college-side update is most conveniently expressed in terms of the discount function $d(q,p)$.

Overview. The algorithm has three blocks:

1. solve the household problem after education decisions are made;
2. given a candidate discount schedule $d(q,p)$, recover the college-side objects and construct the gross tuition schedule;
3. solve the household college/program choice problem and update the discount schedule using the implied allocation.

Step 1: Solve the household problem after college entry. Solve the household dynamic programming problem by backward induction from the terminal age J to age $j = j_a$, i.e. up to the initial period at the beginning of which all education-related choices are made.[16] We use an adaptation of the endogenous grid method.

In period j_a, the household problem is solved on *after-tuition-financial-aid* savings and cash-at-hand/assets grids. After solving for optimal consumption as a function of end-of-period savings:

1. construct beginning-of-period resources net of tuition and financial aid:

$$b^{\text{totnet}} = b - \mathbb{1}_{q>0}\Big[t(b,e,r;q,p)\big(1-\zeta(b)\big)\Big]; \tag{66}$$

[16] From age $j_a + 1$ onward, the household problem is independent of college-market variables such as tuition and financial aid. The remaining household decisions depend only on labor-market prices and policy parameters.

2. for each (q, p), ability e, and residence status r, store the age-j_a value function

$$V_{j_a}(b^{\text{totnet}}, e, r; q, p),$$

defined on the endogenous grid for net initial resources b^{totnet}.

Step 2: Iterate on the discount schedule. The iteration computes the fixed point of the mapping

$$d(q,p) \;\to\; (\bar{e}(q,p), i(q,p), v(q,p,e)) \;\to\; \pi(q,p|b,e,r) \;\to\; \bar{e}^{new}(q,p) \;\to\; d^{new}(q,p).$$

Take as given the age-j_a value functions $V_{j_a}(b^{\text{totnet}}, e, r; q, p)$ computed in Step 1.

1. For each college cell (q, p), initialize a candidate discount

$$d(q, p) > 0.$$

2. Given $d(q, p)$, recover instructional spending and target average student ability from the college-side optimality conditions:

$$i(q,p) = \left[\frac{1-\theta}{\theta} d(q,p)\right]^{\theta} \frac{q}{A_q} \left(\bar{e}^{\text{avg}}\right)^{\theta} \left(i^{\text{avg}}\right)^{1-\theta}, \tag{67}$$

$$\bar{e}(q,p) = \left[\frac{1-\theta}{\theta} d(q,p)\right]^{\theta-1} \frac{q}{A_q} \left(\bar{e}^{\text{avg}}\right)^{\theta} \left(i^{\text{avg}}\right)^{1-\theta}. \tag{68}$$

Let $\{\eta_e(q,p)\}_{e\in\mathcal{E}}$ denote ability shares consistent with the college-side target ability in cell (q, p). These shares satisfy

$$\sum_{e\in\mathcal{E}} \eta_e(q,p) = 1, \qquad \bar{e}(q,p) = \sum_{e\in\mathcal{E}} \eta_e(q,p) e. \tag{69}$$

These target shares are implied by the current candidate discount schedule $d(q, p)$ and are used to construct the tuition schedule entering the household problem.

3. Using the affine pricing rule

$$v(q,p,e) = b(q,p) - d(q,p)e, \tag{70}$$

and the zero-profit condition

$$\sum_{e\in\mathcal{E}} \eta_e(q,p)v(q,p,e) - i(q,p) - \kappa(p) = 0, \tag{71}$$

recover the intercept

$$b(q,p) = i(q,p) + \kappa(p) + d(q,p)\bar{e}(q,p). \tag{72}$$

Therefore, the full ability-specific revenue schedule is

$$v(q,p,e) = i(q,p) + \kappa(p) + d(q,p)\big(\bar{e}(q,p) - e\big), \qquad e \in \mathcal{E}. \tag{73}$$

Special case: two ability types. When $\mathcal{E} = \{e^l, e^h\}$ with $e^h > e^l$, the college-side allocation is summarized by the share of high-ability students $\eta(e^h; q, p)$. In this case,

$$\bar{e}(q,p) = \eta(e^h;q,p)e^h + \big(1 - \eta(e^h;q,p)\big)e^l, \tag{74}$$

so

$$\eta(e^h;q,p) = \frac{\bar{e}(q,p) - e^l}{e^h - e^l}. \tag{75}$$

The low-type revenue level is

$$b(q,p) = v(q,p,e^l) = \kappa(p) + i(q,p) + \eta(e^h;q,p)d(q,p)(e^h - e^l), \tag{76}$$

and the two ability-specific revenue levels are

$$v(q,p,e^l) = b(q,p), \tag{77}$$
$$v(q,p,e^h) = b(q,p) - d(q,p)(e^h - e^l). \tag{78}$$

Equivalently,

$$v(q,p,e^l) = \kappa(p) + i(q,p) + \eta(e^h;q,p)d(q,p)(e^h - e^l), \tag{79}$$
$$v(q,p,e^h) = \kappa(p) + i(q,p) - \big(1 - \eta(e^h;q,p)\big)d(q,p)(e^h - e^l). \tag{80}$$

4. Let $\Phi(b \mid e)$ denote the exogenous distribution of initial resources before tuition and financial aid. Construct an exogenous grid

$$\mathcal{G}^b = \{b^1, b^2, \ldots, b^n\}$$

over pre-aid resources b.

5. Given the implied tuition schedule $v(q, p, e)$, solve the household's joint college/program choice problem for each state (b, e, r):

   (a) obtain residence-specific net tuition

$$t(e, r; q, p) = v(q, p, e) - s(e, r; q, p), \tag{81}$$

   where $s(e, r; q, p)$ is the subsidy schedule;

   (b) for each college cell (q, p), compute net initial resources after tuition and aid:

$$b^{\text{totnet}} = b - \mathbb{1}_{q>0}\Big[t(e, r; q, p)\big(1 - \zeta(b)\big)\Big]; \tag{82}$$

   (c) interpolate the stored age-j_a value function

$$V_{j_a}(b^{\text{totnet}}, e, r; q, p)$$

   at the corresponding value of b^{totnet};

   (d) solve the household's joint choice over (q, p). In the implementation, the choice is smoothed using Type-I extreme value shocks, so the ex-ante value is

$$U(b, e, r) = \sigma_\eta \log\left(\sum_{q=0}^{q_{\max}} \sum_{p=1}^{P} \exp\left(\frac{V_{j_a}(b^{\text{totnet}}, e, r; q, p)}{\sigma_\eta}\right)\right); \tag{83}$$

   this step yields joint choice probabilities

$$\pi(q, p \mid b, e, r).$$

6. Using the exogenous distribution of household states and the implied joint choice probabilities, compute enrollment in each college cell (q, p), denoted by $\chi(q, p)$. For each

active college cell, compute the associated ability shares

$$\eta_e^{new}(q,p) = \frac{\sum_r \int \pi(q,p \mid b,e,r)\, d\Phi_0(b,e,r)}{\chi(q,p)}, \qquad e \in \mathcal{E}, \tag{84}$$

where

$$\chi(q,p) = \sum_{e\in\mathcal{E}} \sum_r \int \pi(q,p \mid b,e,r)\, d\Phi_0(b,e,r). \tag{85}$$

Hence, average student ability in cell (q,p) is

$$\bar{e}^{new}(q,p) = \sum_{e\in\mathcal{E}} \eta_e^{new}(q,p) e. \tag{86}$$

7. In equilibrium, the household allocation must be consistent with the college-side objects $\bar{e}(q,p)$, $i(q,p)$, and $v(q,p,e)$ implied by $d(q,p)$. The discount schedule therefore adjusts until the demand-implied composition matches the college-side target.
Update the discount schedule using the college-side optimality condition

$$d^{\text{new}}(q,p) = \frac{\theta}{1-\theta} \frac{i^{new}(q,p)}{\bar{e}^{new}(q,p)}, \tag{87}$$

where updated instructional spending is implied by the quality production function (using the demand-side average ability):

$$i^{new}(q,p) = \left(\frac{q}{A_q}\right)^{\frac{1}{1-\theta}} i^{\text{avg}} \left(\bar{e}^{\text{avg}}\right)^{\frac{\theta}{1-\theta}} \left(\bar{e}^{new}(q,p)\right)^{-\frac{\theta}{1-\theta}}. \tag{88}$$

Equivalently, substituting out $i^{new}(q,p)$,

$$d^{\text{new}}(q,p) = \frac{\theta}{1-\theta} \frac{i^{\text{avg}}}{\bar{e}^{\text{new}}(q,p)} \left(\frac{q}{A_q}\left(\frac{\bar{e}^{\text{avg}}}{\bar{e}^{\text{new}}(q,p)}\right)^{\theta}\right)^{\frac{1}{1-\theta}} \tag{89}$$

$$= \frac{\theta}{1-\theta} \left(\frac{q}{A_q}\right)^{\frac{1}{1-\theta}} \left(\bar{e}^{\text{avg}}\right)^{\frac{\theta}{1-\theta}} \imath^{\text{avg}} \left(\bar{e}^{\text{new}}(q,p)\right)^{-\frac{1}{1-\theta}}. \tag{90}$$

Special case: two ability types. When $\mathcal{E} = \{e^l, e^h\}$, the ability shares are summarized by the high-ability share

$$\eta^{new}(e^h;q,p) = \frac{\sum_r \int \pi(q,p \mid b,e^h,r)\, d\Phi_0(b,e^h,r)}{\chi(q,p)}, \tag{91}$$

and

$$\bar{e}^{new}(q,p) = \eta^{new}(e^h;q,p)e^h + \big(1-\eta^{new}(e^h;q,p)\big)e^l. \tag{92}$$

Therefore,

$$d^{\text{new}}(q,p) = \frac{\theta}{1-\theta}\left(\frac{q}{A_q}\right)^{\frac{1}{1-\theta}} (\bar{e}^{\text{avg}})^{\frac{\theta}{1-\theta}}\, i^{\text{avg}} \left[\eta^{\text{new}}(e^h;q,p)e^h + \big(1-\eta^{\text{new}}(e^h;q,p)\big)e^l\right]^{-\frac{1}{1-\theta}}. \tag{93}$$

8. Compute the sup norm distance in the discount schedule:

$$\Delta = \max_{q,p} \left| d^{\text{new}}(q,p) - d(q,p) \right|. \tag{94}$$

If Δ exceeds a chosen tolerance threshold, update the discount schedule and return to Step 2(b). Otherwise, the algorithm has converged.

Remark. The key feature of the implementation is that, given a candidate discount schedule $d(q,p)$, the college-side optimality conditions determine instructional spending $i(q,p)$, the college-side target for average student ability $\bar{e}(q,p)$, and the affine ability-specific tuition schedule $v(q,p,e)$. The household block then determines the allocation of students across college cells (q,p), and the discount function adjusts until these objects are jointly consistent with market clearing in every college sub-market.

E.2 Solution Algorithm: College Market Equilibrium with Continuous College Quality

In the continuous implementation, college quality q is continuous while program p remains discrete, so colleges are indexed by pairs (q,p). As in the discrete implementation, the key equilibrium pricing object is the per-unit ability discount

$$d(q,p) > 0, \tag{95}$$

which determines the slope of the revenue schedule in ability. By the college-side optimality conditions, revenues are affine in ability:

$$v(q,p,e) = b(q,p) - d(q,p)e, \qquad e \in \mathcal{E}, \tag{96}$$

where $b(q,p)$ is an intercept pinned down by zero profit.

To obtain a finite-dimensional fixed-point problem, we approximate the discount schedule $d(q,p)$ using cubic B-splines in q. For each program p,

$$d(q,p) = \sum_{m=1}^{M} \theta_{m,p} B_m(q), \tag{97}$$

where $\{B_m(q)\}_{m=1}^{M}$ are spline basis functions defined on the admissible quality interval $[\underline{q}, \overline{q}]$.

The pass-through from college subsidies to students' net tuition payments is 100% by construction. Accordingly, the household problem is solved using the gross tuition schedule $v(q,p,e)$, while the college-side update is most conveniently expressed in terms of the discount function $d(q,p)$.

Quality grids. The computation uses two separate grids for quality:

1. a coarse grid $\mathcal{Q}^c$, used to solve and store household value functions;
2. a dense grid $\mathcal{Q}^d$, used to solve the household college choice problem and to update the discount schedule.

Objects stored on the two grids. The dense quality grid $\mathcal{Q}^d$ is used only for college-market objects needed to update the discount schedule, namely the college-cell enrollment measure $\chi(q,p)$, the demand-implied ability shares $\eta_e(q,p)$, and the corresponding average ability $\bar{e}(q,p)$. All other household-side objects — including value functions, policy functions, and the household cross-sectional distribution over individual states — are stored only on the coarse quality grid $\mathcal{Q}^c$. Hence, the dense grid is used only within the college-choice block, whereas the rest of the household problem is represented on the coarse grid.

Overview. The algorithm has three blocks:

1. solve the household problem after education decisions are made;
2. given a candidate discount schedule $d(q,p)$, recover the college-side objects and construct the gross tuition schedule;
3. solve the household college/program choice problem on the dense grid and update the discount schedule using the implied allocation.

Step 1: Solve the household problem after college entry. Solve the household dynamic programming problem by backward induction from the terminal age J to age $j = j_a$, i.e. up

to the initial period at the beginning of which all education-related choices are made.[17] We use an adaptation of the endogenous grid method.

In period j_a, the household problem is solved on *after-tuition-financial-aid* savings and cash-at-hand/assets grids. After solving for optimal consumption as a function of end-of-period savings:

1. construct beginning-of-period resources net of tuition and financial aid:

$$b^{\text{totnet}} = b - \mathbb{1}_{q>0}\Big[t(b,e,r;q,p)\big(1-\zeta(b)\big)\Big]; \tag{98}$$

2. for each $q \in \mathcal{Q}^c$, program p, ability e, and residence status r, store the age-j_a value function

$$V_{j_a}(b^{\text{totnet}}, e, r; q, p),$$

defined on the endogenous grid for net initial resources b^{totnet}.

Step 2: Iterate on the discount schedule. Take as given the age-j_a value functions $V_{j_a}(b^{\text{totnet}}, e, r; q, p)$ computed in Step 1.

1. For each program p, initialize spline coefficients θ_p, which imply a candidate discount schedule

$$d(q,p) = \sum_{m=1}^{M} \theta_{m,p} B_m(q), \qquad q \in \mathcal{Q}^d. \tag{99}$$

2. Given $d(q,p)$, recover instructional spending and target average student ability from the college-side optimality conditions:

$$i(q,p) = \left[\frac{1-\theta}{\theta} d(q,p)\right]^{\theta} \frac{q}{A_q} \left(\bar{e}^{\text{avg}}\right)^{\theta} \left(i^{\text{avg}}\right)^{1-\theta}, \tag{100}$$

$$\bar{e}(q,p) = \left[\frac{1-\theta}{\theta} d(q,p)\right]^{\theta-1} \frac{q}{A_q} \left(\bar{e}^{\text{avg}}\right)^{\theta} \left(i^{\text{avg}}\right)^{1-\theta}. \tag{101}$$

[17] From age $j_a + 1$ onward, the household problem is independent of college-market variables such as tuition and financial aid. The remaining household decisions depend only on labor-market prices and policy parameters.

Let $\{\eta_e(q,p)\}_{e\in\mathcal{E}}$ denote ability shares consistent with the college-side target ability in cell (q,p). These shares satisfy

$$\sum_{e\in\mathcal{E}} \eta_e(q,p) = 1, \qquad \bar{e}(q,p) = \sum_{e\in\mathcal{E}} \eta_e(q,p)e. \tag{102}$$

These target shares are implied by the current candidate discount schedule $d(q,p)$ and are used to construct the tuition schedule entering the household problem.

3. Using the affine pricing rule

$$v(q,p,e) = b(q,p) - d(q,p)e, \tag{103}$$

and the zero-profit condition

$$\sum_{e\in\mathcal{E}} \eta_e(q,p)v(q,p,e) - i(q,p) - \kappa(p) = 0, \tag{104}$$

recover the intercept

$$b(q,p) = i(q,p) + \kappa(p) + d(q,p)\bar{e}(q,p). \tag{105}$$

Therefore, the full ability-specific revenue schedule is

$$v(q,p,e) = i(q,p) + \kappa(p) + d(q,p)\big(\bar{e}(q,p) - e\big), \qquad e \in \mathcal{E}. \tag{106}$$

Special case: two ability types. When $\mathcal{E} = \{e^l, e^h\}$ with $e^h > e^l$, the college-side allocation is summarized by the share of high-ability students $\eta(e^h; q, p)$. In this case,

$$\bar{e}(q,p) = \eta(e^h;q,p)e^h + \big(1 - \eta(e^h;q,p)\big)e^l, \tag{107}$$

so

$$\eta(e^h;q,p) = \frac{\bar{e}(q,p) - e^l}{e^h - e^l}. \tag{108}$$

The two ability-specific revenue levels are

$$v(q,p,e^l) = \kappa(p) + i(q,p) + \eta(e^h;q,p)d(q,p)(e^h - e^l), \tag{109}$$

$$v(q,p,e^h) = \kappa(p) + i(q,p) - \big(1 - \eta(e^h;q,p)\big)d(q,p)(e^h - e^l). \tag{110}$$

4. Let $\Phi(b \mid e)$ denote the exogenous distribution of initial resources before tuition and financial aid. Construct an exogenous grid

$$\mathcal{G}^b = \{b^1, b^2, \ldots, b^n\}$$

over pre-aid resources b.

5. Given the implied tuition schedule $v(q, p, e)$, solve the household's college/program choice problem for each state (b, e, r) on the dense quality grid:

   (a) obtain residence-specific net tuition

$$t(e, r; q, p) = v(q, p, e) - s(e, r; q, p), \tag{111}$$

   where $s(e, r; q, p)$ is the subsidy schedule;

   (b) for each dense-grid quality point $q \in \mathcal{Q}^d$ and each program p, compute net initial resources after tuition and aid:

$$b^{\text{totnet}} = b - \mathbb{1}_{q>0}\Big[t(e, r; q, p)\big(1 - \zeta(b)\big)\Big]; \tag{112}$$

   (c) interpolate the stored age-j_a value function

$$V_{j_a}(b^{\text{totnet}}, e, r; q, p)$$

   at the corresponding value of b^{totnet}, interpolating across both net resources and quality q from the coarse grid $\mathcal{Q}^c$ to the dense grid $\mathcal{Q}^d$;

   (d) conditional on each quality level q, solve the program choice problem. In the implementation, program choice is smoothed using Type-I extreme value shocks:

$$W(b, e, r; q) = \sigma_\eta \log\left(\sum_{p=1}^{P} \exp\left(\frac{V_{j_a}(b^{\text{totnet}}, e, r; q, p)}{\sigma_\eta}\right)\right). \tag{113}$$

   This step yields conditional program choice probabilities

$$\pi(p \mid b, e, r; q).$$

(e) given $\{W(b,e,r;q)\}_{q\in\mathcal{Q}^d}$, solve the college-quality choice problem by direct comparison on the dense grid:

$$q^*(b,e,r) = \arg\max_{q\in\mathcal{Q}^d} W(b,e,r;q). \tag{114}$$

Since quality is treated as continuous and approximated on a dense grid, no additional extreme-value shocks are introduced at the quality-choice stage.

6. Using the exogenous distribution of household states and the implied policy functions, compute enrollment in each college cell (q,p), denoted by $\chi(q,p)$. For each active college cell, compute the associated demand-implied ability shares

$$\eta_e^{\text{new}}(q,p) = \frac{\sum_r \int \mathbb{1}\{q^*(b,e,r)=q\}\pi(p \mid b,e,r;q)\, d\Phi_0(b,e,r)}{\chi(q,p)}, \qquad e\in\mathcal{E}, \tag{115}$$

where

$$\chi(q,p) = \sum_{e\in\mathcal{E}}\sum_r \int \mathbb{1}\{q^*(b,e,r)=q\}\pi(p \mid b,e,r;q)\, d\Phi_0(b,e,r). \tag{116}$$

Hence, demand-implied average student ability in cell (q,p) is

$$\bar{e}^{\text{new}}(q,p) = \sum_{e\in\mathcal{E}} \eta_e^{\text{new}}(q,p)e. \tag{117}$$

Only the college-market aggregates $\chi(q,p)$, $\eta_e^{\text{new}}(q,p)$, and $\bar{e}^{\text{new}}(q,p)$ are retained on the dense grid after solving the household problem.

7. In equilibrium, the demand-implied composition must coincide with the college-side target composition in every college cell. The discount schedule therefore adjusts until the demand-implied average ability matches the college-side target. Using the demand-implied average ability, update instructional spending via the quality production function:

$$i^{new}(q,p) = \left(\frac{q}{A_q}\right)^{\frac{1}{1-\theta}} i^{\text{avg}} \left(\bar{e}^{\text{avg}}\right)^{\frac{\theta}{1-\theta}} \left(\bar{e}^{new}(q,p)\right)^{-\frac{\theta}{1-\theta}}. \tag{118}$$

Then update the discount schedule using the college-side optimality condition:

$$d^{\text{new}}(q,p) = \frac{\theta}{1-\theta}\frac{i^{\text{avg}}}{\bar{e}^{\text{new}}(q,p)}\left(\frac{q}{A_q}\left(\frac{\bar{e}^{\text{avg}}}{\bar{e}^{\text{new}}(q,p)}\right)^{\theta}\right)^{\frac{1}{1-\theta}} \tag{119}$$

$$= \frac{\theta}{1-\theta}\left(\frac{q}{A_q}\right)^{\frac{1}{1-\theta}}(\bar{e}^{\text{avg}})^{\frac{\theta}{1-\theta}}\,\imath^{\text{avg}}\,(\bar{e}^{\text{new}}(q,p))^{-\frac{1}{1-\theta}}\,. \tag{120}$$

Two-type update. When $\mathcal{E} = \{e^l, e^h\}$, the ability shares are summarized by the high-ability share

$$\eta^{\text{new}}(e^h; q, p) = \frac{\sum_r \int \mathbb{1}\{q^*(b, e^h, r) = q\}\pi(p \mid b, e^h, r; q)\, d\Phi_0(b, e^h, r)}{\chi(q,p)}, \tag{121}$$

and

$$\bar{e}^{\text{new}}(q,p) = \eta^{\text{new}}(e^h; q, p)e^h + \big(1 - \eta^{\text{new}}(e^h; q, p)\big)e^l. \tag{122}$$

Therefore,

$$d^{\text{new}}(q,p) = \frac{\theta}{1-\theta}\left(\frac{q}{A_q}\right)^{\frac{1}{1-\theta}}(\bar{e}^{\text{avg}})^{\frac{\theta}{1-\theta}}\,i^{\text{avg}}\left[\eta^{\text{new}}(e^h; q, p)e^h + \big(1 - \eta^{\text{new}}(e^h; q, p)\big)e^l\right]^{-\frac{1}{1-\theta}}. \tag{123}$$

8. For each program p, update the spline coefficients by projecting the dense-grid discount update onto the B-spline basis. Let $\mathcal{Q}^{d,\text{act}}(p) \subseteq \mathcal{Q}^d$ denote the set of dense-grid quality points with non-negligible enrollment mass in program p. Then

$$\theta_p^{\text{new}} = \arg\min_{\theta_p} \sum_{q \in \mathcal{Q}^{d,\text{act}}(p)} w(q,p)\left[d^{\text{new}}(q,p) - \sum_{m=1}^{M}\theta_{m,p}B_m(q)\right]^2, \tag{124}$$

where the weights $w(q,p)$ are proportional to enrollment mass $\chi(q,p)$.

9. Compute the sup norm distance between the updated and current spline coefficients:

$$\Delta = \max_p \max_{m=1,\ldots,M} \left|\theta_{m,p}^{\text{new}} - \theta_{m,p}\right|. \tag{125}$$

10. If Δ exceeds a chosen tolerance threshold, optionally damp the coefficient update

$$\theta_p \leftarrow \lambda\theta_p^{\text{new}} + (1-\lambda)\theta_p, \qquad \lambda \in (0,1], \tag{126}$$

and return to Step 2(b). Otherwise, the algorithm has converged.

Remark. The continuous-quality implementation follows the same logic as the discrete one. Given a candidate discount schedule $d(q,p)$, the college-side optimality conditions determine instructional spending $i(q,p)$, the college-side target for average student ability $\bar{e}(q,p)$, and the affine ability-specific tuition schedule $v(q,p,e)$. The household block then determines the allocation of students across college cells (q,p) on a dense quality grid, and the discount function adjusts until these objects are jointly consistent with market clearing in every college sub-market.

E.3 Solution Algorithm: Capital and Labor Market Equilibrium

The optimality condition of the firm with respect to capital $K(p)$ used in the p-specific output production is given by

$$\begin{aligned} r+\delta &= \Psi Y^{1-\nu}\lambda(p)\alpha(p)\left(\alpha(p)\left(\Psi K(p)\right)^{\rho(p)}+(1-\alpha(p))\left(\Gamma(p)L_t(p)\right)^{\rho(p)}\right)^{\frac{\nu-\rho(p)}{\rho(p)}}\left(\Psi K(p)\right)^{\rho(p)-1} \\ &= \Psi Y^{1-\nu}\lambda(p)\alpha(p)\left(\alpha(p)+(1-\alpha(p))\left(\frac{1}{k(p)}\right)^{\rho(p)}\right)^{\frac{\nu-\rho(p)}{\rho(p)}}\left(\Psi K(p)\right)^{\nu-1} \\ &= \Psi\left(\frac{Y}{\Psi K(p)}\right)^{1-\nu}\lambda(p)\alpha(p)\left(\alpha(p)+(1-\alpha(p))\left(\frac{1}{k(p)}\right)^{\rho(p)}\right)^{\frac{\nu-\rho(p)}{\rho(p)}}, \end{aligned} \tag{127}$$

where $k(p)=\frac{\Psi K(p)}{\Gamma(p)L(p)}$ is the effective capital-labor ratio Similarly, the first order condition with respect to labor $L(p)$ is

$$\begin{aligned} w(p) &= \Gamma(p)Y^{1-\nu}\lambda(p)(1-\alpha(p))\left(\alpha(p)\left(\Psi K(p)\right)^{\rho(p)}+(1-\alpha(p))\left(\Gamma(p)L(p)\right)^{\rho(p)}\right)^{\frac{\nu-\rho(p)}{\rho(p)}}\left(\Gamma(p)L(p)\right)^{\rho(p)-1} \\ &= \Gamma(p)Y^{1-\nu}\lambda(p)(1-\alpha(p))\left(\alpha(p)+(1-\alpha(p))\left(\frac{1}{k(p)}\right)^{\rho(p)}\right)^{\frac{\nu-\rho(p)}{\rho(p)}}\left(\Psi K(p)\right)^{\nu-1+1-\rho(p)}\left(\Gamma(p)L(p)\right)^{\rho(p)-1} \\ &= \Gamma(p)\left(\frac{Y}{\Psi K(p)}\right)^{1-\nu}\lambda(p)(1-\alpha(p))\left(\alpha(p)+(1-\alpha(p))\left(\frac{1}{k(p)}\right)^{\rho(p)}\right)^{\frac{\nu-\rho(p)}{\rho(p)}}k(p)^{1-\rho(p)} \end{aligned} \tag{128}$$

Using (127) in the above equation we obtain:

$$w(p)=\frac{\Gamma(p)}{\Psi}\frac{1-\alpha(p)}{\alpha(p)}\cdot(r+\delta)\cdot k(p)^{1-\rho(p)} \tag{129}$$

which gives the capital intensity $k(p)$ as a function of p-specific wages and interest rates:

$$k(p) = \left(\frac{w(p)}{\frac{\Gamma(p)}{\Psi} \frac{1-\alpha(p)}{\alpha(p)} \cdot (r+\delta)} \right)^{\frac{1}{1-\rho(p)}} \tag{130}$$

This suggests the following algorithm:

1. Guess aggregate interest rate r, wage normalization factor $\Gamma(1)$ and p-specific wages $w(p)$ (for $p > 0$).

2. Solve the household problem (normalization the non-college wage to 1)

   - Compute aggregate assets K^s
   - Compute p-specific labor supply $\{L^s(p)\}_{p=0}^{P}$

3. Use eq. (129) to compute $k^d(p)$:

$$k^d(p) = \left(\frac{w(p)}{\frac{\Gamma(p)}{\Psi} \frac{1-\alpha(p)}{\alpha(p)} \cdot (r+\delta)} \right)^{\frac{1}{1-\rho(p)}} \tag{131}$$

   where $\Gamma(p) = \Gamma(1)\tilde{\Gamma}(p)$ with $\tilde{\Gamma}(p)$ being exogenous scaling parameters.

4. Compute aggregate capital demand: $K^d = \sum_p K^d(p) = \frac{\sum_p k^d(p) L^s(p) \Gamma(p)}{\Psi}$

5. Compute wage normalization factor $\Gamma(1)$:

$$\Gamma(1) = \frac{1}{\left(\frac{Y}{\Psi K^d(0)}\right)^{1-\nu} \lambda(0)(1-\alpha(0)) \left(\alpha(0) + (1-\alpha(0)) \left(\frac{1}{k^d(0)}\right)^{\rho(0)}\right)^{\frac{\nu-\rho(0)}{\rho(0)}} k^d(0)^{\rho(0)-1}} \tag{132}$$

6. For $p > 0$, compute p-specific wages consistent with p-specific household labor supply $L^s(p)$:

$$\tilde{w}(p) = \Gamma(p) \left(\frac{Y}{\Psi K^d(p)}\right)^{1-\nu} \lambda(p)(1-\alpha(p)) \left(\alpha(p) + (1-\alpha(p)) \left(\frac{1}{k^d(p)}\right)^{\rho(p)}\right)^{\frac{\nu-\rho(p)}{\rho(p)}} k^d(p)^{\rho(p)-1} \tag{133}$$

   Observe that both aggregate output Y and p-specific capital demand $K(p)$ are a function of p-specific labor supply by households $L^s(p)$.

7. Iterate on r and $w(p)$ for $p > 0$ and $\Gamma(1)$ until convergence, ensuring that the aggregate capital market clearing condition and p-specific labor market clearing conditions hold.

Special Case of Perfect Substitutes. With $\nu = 1$, equation (127) simplifies to

$$r + \delta = \Psi\lambda(p)\alpha(p)\left(\alpha(p) + (1-\alpha(p))\left(\frac{1}{k(p)}\right)^{\rho(p)}\right)^{\frac{1-\rho(p)}{\rho(p)}} \tag{134}$$

from which we can pin down in closed form $k(p)$ in each p as

$$k(p) = \left[\frac{1-\alpha(p)}{\left(\frac{r+\delta}{\Psi\lambda(p)\alpha(p)}\right)^{\frac{\rho(p)}{1-\rho(p)}} - \alpha(p)}\right]^{\frac{1}{\rho(p)}}. \tag{135}$$

Likewise, (128) simplifies to

$$w(p) = \Gamma(p)\lambda(p)(1-\alpha(p))\left(\alpha(p)k(p)^{\rho(p)} + (1-\alpha(p))\right)^{\frac{1-\rho(p)}{\rho(p)}}. \tag{136}$$

and we may now use (135) in (136) to see how wages are pinned down from $r + \delta$.