Factors Affecting College Attainment and Student Ability in the U.S. since 1900*

Kevin Donovan
University of Notre Dame

Christopher Herrington
Virginia Commonwealth University

January 2017

Abstract

We develop a dynamic lifecycle model to study long-run changes in college completion and average ability of college students in cohorts born from 1900 to 1972. The model is disciplined in part by constructing a historical time series on real college costs from printed government documents covering this time period. The model captures nearly all of the increase in attainment for 1900 to 1950 cohorts. In counterfactual exercises we show that attainment would have been lower by almost half, on average, for 1925 to 1950 cohorts absent a large decrease in college costs relative to income. For post-1950 cohorts simultaneously rising college costs and education premia act in opposite directions to result in low college enrollment growth; however, endogenously declining average ability of college students lowers overall completion rates in the model. Furthermore, we find that economic factors have little impact on average student ability; rather, the precision of signals about true ability are the key driver of changes in average ability. We utilize historical data on the share of students who take SATs as a proxy for the increasing precision of ability signals, and show that this allows the model to match the aggregate ability sorting patterns.

JEL Classification Codes: I2, J24, N31

Keywords: College attainment; student ability; borrowing constraints; college costs; education premia

*Thanks to Eleanor Dillon, Todd Schoellman, Gustavo Ventura, and participants at Arizona State, the University of Richmond, Midwest Macro Meetings, and the AEA Annual Meeting for helpful comments. The usual disclaimers apply. Donovan (corresponding author): kdonova6@nd.edu; 434 Flanner Hall, Notre Dame, IN 46556. Herrington: herringtoncm@vcu.edu; 301 W. Main St, Richmond, VA 23284.

1 Introduction

During the twentieth century, higher education expanded dramatically in the United States. As shown in Figure 1, less than four percent of the 1900 birth cohort held a bachelor's degree by age 23, but that share increased to more than 30% by the 1972 cohort. Concurrent with the increase in college attendance, the gap in measured cognitive ability widened substantially between college students and those individuals who ended their formal education after high school graduation (i.e., "non-college" individuals). This pattern is seen in Figure 2a, which plots the average ability percentile of college and non-college individuals.[1] The average college student born around 1900 had measured cognitive ability only slightly above the average non-college individual. Yet over the next several decades, the average ability increased among college students and decreased among those with only a high school degree.

Figure 1: College Completion among U.S. birth cohorts 1900–1972

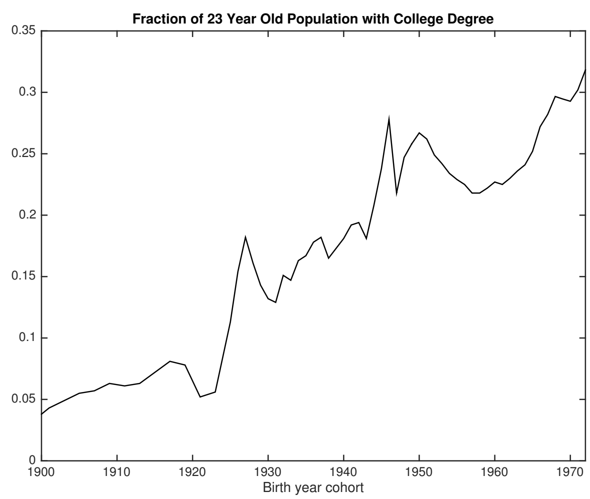

Figure 2b illustrates the pattern of increased ability sorting even more starkly by plotting the difference between average ability of college and non-college types shown in Figure 2a. The quadratic trend line indicates that the average ability difference more than doubled during the first half of the 20th century. The trend of increasing average ability among college students occurred even as the share of cohorts attending college grew steadily larger, implying imperfect selection into college on skills.

[1]Figure 2a is an updated version of the Figure 1b in Hendricks and Schoellman (2014), and we thank them for sharing their data. We will use the general term "ability" throughout the paper while recognizing that ability can measured in different ways in the data. For additional details see also Hendricks, Herrington and Schoellman (2016).

Figure 2: Ability Sorting among U.S. birth cohorts 1900–1972

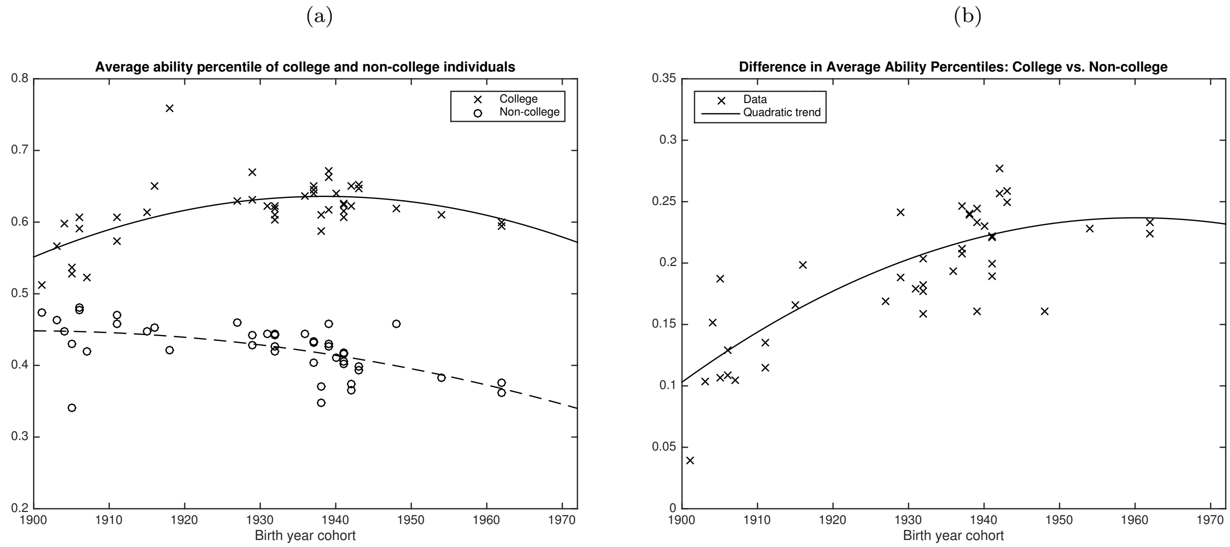

The goal of this paper is to understand the causes of these two empirical trends. This task is complicated by the vast number of changes in both the aggregate economy and education sector over this time period, as well as a lack of reliable data particularly during the first half of the century. We combat these issues by developing a lifecycle model that allows for several potential explanations to coexist within the model. First, each cohort is populated by an exogenous number of male and female high school graduates who, because they face different wage profiles, also face different education premia and opportunity costs. Endowed with financial assets and an ability to complete college - both of which are heterogenous - these high school graduates immediately decide whether to enter college or go to work. Moreover, ability is not perfectly known. Consistent with evidence in Cunha, Heckman and Navarro (2005), individuals observe only a noisy signal about true ability, which introduces a measure of uncertainty when combined with the limited borrowing opportunities for college. These features allow us to quantitatively compare a number of potential explanations within the framework of the model.

Naturally, the model predictions rely heavily on the costs and benefits associated with attending college, and we therefore take care to measure them properly. We construct the out-of-pocket costs for tuition and fees by combining data from a series of printed government documents dating back to 1916. This allows us to feed in time series data for college tuition costs, instead of relying on assumptions such as a constant tuition growth rate. The opportunity cost of attending college, along with the benefit of attaining a degree, is computed by estimating wage profiles from the 1940 to 2000 U.S. Censuses and 2006-2010 American Community Surveys. The wage profiles are allowed to depend on sex, age, and education to accurately capture the changing education earnings premia in the United States. These time series are exogenously fed into the model to capture the relevant tradeoffs faced by high school graduates when choosing to enter college.

After incorporating these time series data, we are left to calibrate the time-invariant parameters of the model. The lack of reliable data during the early 20th century poses a challenge, so we use more recent and reliable micro data from the 1979 National Longitudinal Survey of Youth and High School & Beyond Survey to calibrate the joint distribution of endowments and the riskiness of college. Individuals in these surveys were mostly born in the early 1960s, so they fall near the end of the 1900–1972 birth cohorts we consider. Taking these structural parameters as fixed, we then simulate the model beginning with the 1900 birth cohort (i.e. the 1918 high school graduation cohort) and feed in the constructed high school graduation rates, college costs, and life-cycle earnings profiles to assess the model's ability to replicate enrollment and average ability changes over time.

The model predicts almost the entire increase in college attainment from the 1900 to 1950 cohorts. Outside of over-predicting attainment during World War II, the time series tracks the data very well during this time period. An accounting decomposition of the college completion rate shows that the increase is driven almost exclusively by: (1) an increase in the fraction of individuals completing high school and (2) an increase in the fraction of individuals enrolling in college *conditional* on high school graduation. The former is exogenous to our model, and has been previously emphasized by Goldin and Katz (2010); however, the latter is endogenous. We therefore conduct a series of counterfactual experiments to understand the underlying forces driving the college enrollment increase, and we find that the direct costs of college are the main factor during this period. Relatively high costs depressed enrollment for pre-1920 cohorts, and falling costs boosted enrollment for later cohorts. In fact, our counterfactual results show that college attainment would have been lower by almost half, on average, for 1925 to 1950 cohorts absent the decrease in college costs relative to income.

When we consider cohorts born after 1950, the model initially predicts college completion rates higher than those observed in the data. This is a common issue (e.g. Card and Lemieux, 2001*a*), and is due in part to the large, perfectly forecasted increase in the college earnings premia experienced by these cohorts. However, rising college costs counteract this effect and result in declining college completion for cohorts born in the last decade of our sample, which is broadly consistent with evidence from Jones and Yang (2016). Our model also predicts endogenously decreasing average ability of college students for post-1950 birth cohorts, which leads to lower college graduation rates despite enrollment growth. This finding is consistent with evidence in Bound, Lovenheim and Turner (2010) and Castro and Coen-Pirani (2016), and supports the theory that variations in average ability across cohorts is important for understanding education decisions and outcomes in the second half of the century.

We lastly turn to model predictions on increased ability sorting over time. Through a series of counterfactual experiments, we consistently find that changes in economic factors (i.e., earnings premia, college costs, and opportunity costs) have little impact on ability sorting. Instead, the key feature in the model that accounts for this phenomenon is uncertainty about ability. We show that a decrease in the variance of ability signals can generate an increase in ability sorting similar to that in the data. We then link this decrease in variance to the share of high school graduates taking standardized SAT tests, which is increasing over this time period. Interestingly, the large increase in SAT prevalence matches up well with the required decrease in signal variance required to match sorting. This suggests an increase in standardized testing improved students' knowledge of their own ability relative to other students in their cohort, as discussed in Hoxby (2009), and is critical to understanding ability sorting in the early twentieth century.

1.1 Related Literature

This paper is primarily related to the literature on college attainment in the United States over time, which includes recent work by Garriga and Keightley (2007), Restuccia and Vandenbroucke (2013) Hendricks and Leukhina (2016), and Hendricks, Herrington and Schoellman (2016). These papers are complemented by others including Bound and Turner (2002) and Card and Lemieux (2001*b*) which examined the effects of policy changes like the G.I.

Bill and Vietnam war draft during particular points in history. More closely related to our work, Castro and Coen-Pirani (2016) ask whether educational attainment of the 1932–1972 cohorts can be explained by changes in the college earnings premia, tuition, education quality, and cohort average learning ability. Their complete markets model underpredicts college attainment for pre-1950 cohorts, while the combination of limited borrowing with our college tuition time series allows us to match it quite well. Hendricks and Schoellman (2014) also study early 1900 college attainment and ability sorting, but take college completion and student ability as given in order to understand the changes in the college earnings premium in a complete markets model. By contrast, we seek to understand the economic factors that affected college completion and average student ability.

As it relates to post-1950 cohorts, the aforementioned Castro and Coen-Pirani (2016) show that cross-cohort variation in learning ability can alleviate the over-prediction of college attainment. Keller (2014), on the other hand, develops a model of college attendance and quality choice, and points to the slowdown of human capital rental rates as the cause of the post-1950 slowdown in attainment. Lochner and Monge-Naranjo (2011) emphasize student loan policies with limited commitment as the driving force behind post-World War II student ability sorting. Their focus on the latter half of the twentieth century forces us to exclude the student loan innovations they consider, since Figure 2b shows that a majority of the sorting occurred in cohorts before 1950.

A closely related strand of literature has also focused on the co-related roles of family background and borrowing constraints for college decisions. Keane and Wolpin (2001) finds that altering parental transfers can have significant effects on the overall schooling attainment of children when borrowing constraints are present; however, extending loans to ease borrowing constraints has little impact of college attendance decisions. Ionescu (2009) emphasizes initial human capital and learning ability of students as the primary determinants of college entry rather than family wealth or income during college. Belley and Lochner (2007) show that the role of family background has become increasingly important in recent decades, and that borrowing constraints are a crucial component for explaining this change. Finally, Athreya and Eberly (2016) show that college completion risk can dampen the expected return to college, and that this effect is magnified for low-wealth students who must borrow to finance the costs of college. Collectively these results inform our decision to include borrowing constraints, heterogenous parental transfers, and college completion risk.

The rest of this paper proceeds as follows. Section 2 begins by decomposing college attainment growth into changes in high school graduation rates and other factors. Section 3 then lays out a life-cycle model of college decisions which we use to discipline our analysis. Section 4 discusses the calibration, and Section 5 analyzes the quantitative performance of the baseline model. Section 6 conducts counterfactual exercises to isolate the key channels affecting college attainment, while Section 7 examines the factors that affect ability sorting over time. Finally, Section 8 concludes.

2 Accounting for College Attainment

We begin with an accounting exercise to identify the key channels that affect college completion over time. Our measure of college attainment for cohort t is the share of twenty-three year olds with a college degree.[2] This measure can be decomposed into three components:

$$\frac{N_t^{grad}}{N_t^{23}} = \underbrace{\left(\frac{N_t^{HS}}{N_t^{23}}\right)}_{\text{HS graduation rate}} \times \underbrace{\left(\frac{N_t^{enroll}}{N_t^{HS}}\right)}_{\text{college enrollment rate}} \times \underbrace{\left(\frac{N_t^{grad}}{N_t^{enroll}}\right)}_{\text{college graduation rate}} \tag{2.1}$$

N^{HS}, N^{enroll}, and N^{grad} are the number of individuals born in year t who complete high school, enroll in college, and graduate college. The first ratio on the right hand side – the HS graduation rate – is the fraction of all twenty-three year olds who have graduated high school. The latter two are the fraction of high school graduates who enroll in college (the college enrollment rate) and the fraction of college attendees who graduate college (the college graduation rate). We will refer to these ratios as such throughout the paper. Using identity (2.1), the growth rate of college attainment between cohorts born at t and t' is

[2] While educational attainment is often measured later in life to capture those who complete college at older ages, we prefer this series for two reasons. First, to our knowledge it is the only measure of college completion with consistent time series data for birth cohorts back to 1900. Second, our model is not constructed to evaluate college enrollment decisions of older students who: (i) are generally less financially-dependent upon parents when paying for education; (ii) face different opportunity costs of school after having been in the workforce for some time; and (iii) may anticipate different return on investment in education due to later-life completion.

Table 1: Growth rate decomposition

	Log differences			Annualized Percent Change		
Birth years	γ^{col}	γ^{hs}	γ^{endog}	γ^{col}	γ^{hs}	γ^{endog}
1900 - 1972	2.13	1.53	0.59	2.95%	2.13%	0.82%
1900 - 1920	0.54	1.13	-0.59	2.68%	5.64%	-2.95%
1920 - 1972	1.59	0.41	1.18	3.05%	0.78%	2.27%
1950 - 1972	0.18	-0.10	0.27	0.80%	-0.44%	1.24%

approximated by

$$\gamma^{col}_{t',t} := \log\left(\frac{N^{grad}_{t'}}{N^{23}_{t'}}\right) - \log\left(\frac{N^{grad}_{t}}{N^{23}_{t}}\right) = \gamma^{hs}_{t',t} + \gamma^{enroll}_{t',t} + \gamma^{grad}_{t',t}. \tag{2.2}$$

where γ^{hs}, γ^{enroll}, and γ^{grad} are the log differences of the three ratios on the right hand side of (2.1). Unfortunately we do not have sufficient data back to 1900 to separately compute γ^{enroll} and γ^{grad}. However, anticipating the model somewhat, both γ^{enroll} and γ^{grad} will be endogenous objects in our model, while γ^{hs} will be exogenous. We therefore rewrite equation (2.2) as

$$\gamma^{col}_{t',t} = \gamma^{hs}_{t',t} + \gamma^{endog}_{t,t'} \tag{2.3}$$

where $\gamma^{endog} := \gamma^{enroll} + \gamma^{grad}$ is the growth rate of factors endogenous to our model, and therefore potentially affected by changes in college costs and earnings premia. Since we have data on college attainment and high school graduation, the endogenous factors are defined as the residual that allows identity (2.3) to hold with equality. Table 1 decomposes γ^{col} for the entire time period (1900–1972 birth cohorts) and three sub-samples. For ease of interpretation, we also convert the log differences into annualized percent changes over each sample period.

Over the entire 1900 to 1972 time period, college attainment increased at an annual rate of nearly three percent. From a growth accounting perspective, over two-thirds of this would be attributable to the increase in high school completion, γ^{hs}, consistent with Goldin and Katz (2010). This result, however, masks large differences across sub-samples of the overall time period. For cohorts born 1900 to 1920 college attainment increased at 2.68% annually, but the sources of this growth were very different. These cohorts experienced high school graduation rate increases of 5.64% annually, but a substantial decrease in endogenous factors (-2.95% annually). From 1920 onwards, however, when the bulk of the college attainment increase occurs, three-fourths of the increase in attainment is accounted for by endogenous factors (i.e., 2.27% of the 3.05% annualized increase). Finally, we note that this decomposition also captures the well-documented slowdowns in both high school and college attainment for cohorts born after 1950. High school graduation rates fell by 0.44% annually, and growth in endogenous factors was only 1.24% per year. Combined, these resulted in college attainment growth of less than one percent annually since the 1950 cohort.

These decomposition results are consistent with other direct evidence over this time period as well. Snyder (1993) includes a measure of college completion conditional on high school graduation, defined as

$$\frac{\text{Number of college degrees award at year } t+4}{\text{Number of high school graduates at year } t}.$$

This captures roughly the same idea as the ratio N^{grad}/N^{HS} (e.g. the endogenous factors), though it requires assumptions on time-to-completion for the two to be exactly equivalent. Nevertheless, Figure 3 shows that this series follows the same path uncovered in Table 1, with a substantial decline for cohorts born in the early 1900s and then sustained growth afterward.

Taken together, these growth accounting exercises suggest that increases in high school graduation rates cannot be relied upon as the sole factor driving subsequent increases in college completion. Rather, it is vital to understand the factors affecting college enrollment decisions and completion rates for a full understanding of college attainment trends during the twentieth century. The remainder of this paper will develop and calibrate a model to quantitatively assess the main contributors to this growth.

3 Model

We now develop a lifecycle model to investigate the causes of increased college attainment and increased ability sorting over time. The relevant features include borrowing limits,

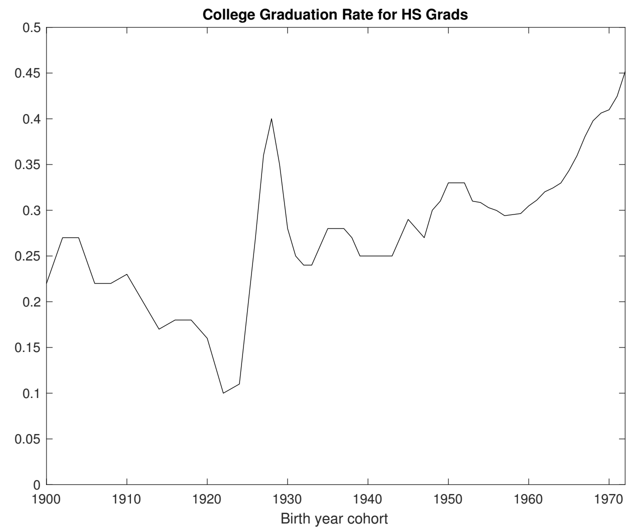

Figure 3: College Completion Rate Conditional on HS graduation

uncertain ability, and risky completion of college education. The notation introduced in this section is summarized in Table 6 of the Appendix.

Demographics and Preferences Time in the model is discrete, and a model period is one year. Each period, N_{mt} males and N_{ft} females are born, each of whom lives for a total of T periods. Let $a = 1, 2, \ldots, T$ denote age. Each individual maximizes expected lifetime consumption, given by

$$\mathbb{E}_0 \sum_{a=1}^{T} \beta^{a-1} \left(\frac{c_a^{1-\sigma} - 1}{1-\sigma} \right).$$

Endowments and Signals Individuals are ex-ante heterogeneous along three dimensions: their sex, m or f, initial asset endowment k_0, and ability to complete college, denoted α. The probability that an individual completes his or her current year of college is given by $\pi(\alpha)$, where $\frac{d\pi}{d\alpha} > 0$. Ability and initial asset holdings are drawn from a joint distribution with cumulative distribution function $F(\alpha, k_0)$.

While sex and asset endowments are perfectly observable, ability α is not. Instead, each individual receives a signal $\theta = \alpha + \varepsilon$ at the beginning of life. The error term is $\varepsilon \sim N(0, \sigma^2_{\varepsilon,t})$, so signal precision may vary across cohorts. Note that both the signal θ and asset endowment k_0 provide information about underlying ability, as assets and ability are potentially correlated. We therefore denote $\nu = (k_0, \theta)$ as the information an individual has about his or her true ability. After the initial college enrollment decision, ability α becomes publicly observable.

Education Decisions The population we consider is high school graduates, so that birth in this model translates to a high school graduation in the real world. At birth, every individual decides whether or not to enroll in college, given sex, asset endowment k_0, and ability signal θ. This is the only time this decision can be made. Once enrolled in college, individuals can exit college by either graduating or by failing out with annual probability $\pi(\alpha)$. After failure, individuals enter the labor force and may not re-enroll, consistent with the finality of dropout decisions discussed in Card and Lemieux (2001*a*). Graduating college requires C years of full-time education at a cost of λ_t per year. If an individual decides to not enter college, he or she immediately enters the labor market and begins to work.

Labor Market We adopt the common assumption that individuals of different ages, a, sex s, and education, e, are different inputs into a constant returns to scale production function that requires only labor. Therefore, wages depend on age a, year t, sex $s \in \{f, m\}$, education level $e \in \{0, 1, \ldots, C\}$, denoted by the function $w(a, t, e, s)$. While ability α has no direct effect on realized wages, it does affect expected wages because higher ability students are more likely to graduate college and earn higher wages.

Savings Each individual can borrow and save at an exogenous interest rate r. Individuals cannot die in debt, so $k_{T+1} = 0$. During life, borrowing is constrained to be a fraction $\gamma \in [0, 1]$ of expected discounted future earnings. Therefore, individuals must keep assets k_t each period above some threshold $\bar{k}$, where

$$\bar{k} = -\gamma \cdot \mathbb{E} \sum_{n=a}^{n=T} \frac{w(n, t, e, s)}{(1+r)^{n-a}}$$

Note that both the expectations operator and wage can depend on a number of factors, including ability α, age a, year t, education e, and sex s. Therefore, the borrowing constraint will be written as the function $\bar{k}(\alpha, a, t, e, s)$. In a slight abuse of notation, we will write $\bar{k}(a, t, e, s)$ when the borrowing constraint does not depend on ability α, as is the case once an individual finishes college.[3]

3.1 Timing and Recursive Problem

At the beginning of year t, N_{mt} men and N_{ft} women are born at age $a = 1$. Again, each individual is initially endowed with assets k_0, sex s, ability α, and a signal θ of true ability. Immediately, each individual decides whether or not to enroll in college. If she enrolls in college, true ability is realized, and the individual proceeds through college. In the case of leaving college due to failure or graduation, she proceeds to the labor market and works for the remainder of her life.[4] Individuals who do not enroll in college proceed directly to the labor market, where they receive the wage associated with age a, education $e = 0$, and sex s.

Recursive Problem for Worker For individuals currently not enrolled in college, their ability is irrelevant for their decision problem. Therefore, the value of entering year t at age a with assets k, years of college education e, and sex $s \in \{f, m\}$ is:

$$\begin{aligned} v^w_{a,t}(k,e,s) &= \max_{k'} \; u(c) + \beta v^w_{a+1,t+1}(k',e,s) \\ s.t. \qquad c + k' &= (1+r)k + w(a,t,e,s) \\ k' &\geq \bar{k}(a,t,e,s) \\ k_{T+1} &= 0 \end{aligned}$$

Recursive Problem for College Student If instead an individual is currently enrolled in college, she has already completed e years of his education and must pay λ_t in college costs for the current year. The probability that she passes and remains enrolled the next year, however, depends on her ability α.

The value of being enrolled in college at year t at age a, with assets k, ability α, e years

[3]Note that the assumption that ability uncertainty is immediately resolved after the college choice implies that the expectation over the borrowing constraint does not directly include uncertainty over about α, but only uncertainty about graduating college.

[4]We adopt "she" and "her" for simplicity, with the reminder that both sexes exist in the model.

of education completed, and sex $s \in \{f, m\}$ is:

$$\begin{aligned} v^c_{a,t}(k, \alpha, e, s) &= \max_{k'} \; u(c) + \beta\Big[\pi(\alpha) v^c_{a+1,t+1}(k', \alpha, e+1, s) + (1-\pi(\alpha)) v^w_{a+1,t+1}(k', e, s)\Big] \\ s.t. \quad c + k' - \lambda_t &= (1+r)k \\ k' &\geq \bar{k}(\alpha, a, t, e, s) \end{aligned}$$

The College Enrollment Decision Given the value of being enrolled in college and working, we can define the educational decision rule at the beginning of life. Recall that at this point, α is unknown, but each individual receives information $\nu = (k_0, \theta)$. Each individual then constructs beliefs over possible ability levels by using Bayes' Rule. Let $G(\alpha; k_0, \theta)$ be the cumulative distribution function of beliefs over ability levels. Given all this, an individual of sex s born in year t with assets k_0 and signal θ enters college if and only if the expected value of entering college is higher than the value of entering the workforce. The college decision is therefore represented by the value function

$$v_t(k_0, \theta, s) = \max\Big\{\int_\alpha v^c_{1,t}(k_0, \alpha, 1, s) G(d\alpha; k_0, \theta), v^w_{1,t}(k_0, 0, s)\Big\}$$

with associated decision rule $\phi_t(k_0, \theta, s) = 1$ if the individual decides to enroll, and $\phi_t(k_0, \theta, s) = 0$ otherwise.

We next turn to quantifying this model to assess the channels underlying changes in college attainment over time.

4 Calibration

Our goal is to use the model to assess how several economic features have affected college attainment and ability sorting over time. Toward that end we parameterize the model in several steps. First, the length of life, T, length of college, C, interest rate r, and discount factor β are all set to standard values. Next, we discuss model-independent time series data for high school graduates N_{mt} and N_{ft} and college costs λ_t that are fed directly into the model. Additionally, we use various micro data sources to estimate life-cycle wage profiles $w(a, t, e, s)$, the probability of passing college $\pi(\alpha)$, and the joint distribution of ability α

and parental transfers k_0. Finally, we impose a functional form on the time-varying ability signals $\sigma_{\varepsilon,t}$ and choose those parameters jointly with the borrowing constraint λ while solving the model.

4.1 Parameters Set to Standard Values

Parameters set exogenously prior to solving the model are: T, C, r, and β. We set the length of working life at $T = 48$, implying that individuals born into the model at age 18 would retire at age 65. The number of periods required to complete college is $C = 4$, so that all individuals in the model have post-secondary education $e \in \{0, 1, 2, 3, 4\}$. The real interest rate is set to $r = 0.04$ in all periods, and the discount rate is $\beta = 0.96$, a standard value in models with annual periods.

4.2 Historical Time Series Data

As previously mentioned, N_{mt} males and N_{ft} females are "born" into the model each year, meaning they graduate high school and enter the model eligible to make college enrollment decisions. We take high school completion, and thus the population of potential college enrollees, as exogenous. The series for N_{mt} and N_{ft} are from the Historical Statistics of the United States Millennial Edition Online, and we use linear interpolation to supply missing values.

Annual college costs per student, λ_t, are calculated as the average public tuition and fee expenses paid out-of-pocket by students each year.[5] Note that because we measure average *out-of-pocket* costs in the data, λ_t accounts for changes over time in the average amount of financial aid received by students in the form of public and private scholarships and grants. Full details of the data construction are relegated to Appendix A. Briefly, however, we compute λ_t each period as the total revenues from student tuition and fees received by all public institutions of higher education divided by the total number of full-time equivalent students enrolled in those institutions. The complete time series is constructed by splicing together data from historical print sources including the *Biennial Surveys of Education* (1900

[5] Additional student expenses, such as room and board, could also be included without altering results substantially. We choose to leave these out of the benchmark specification because such costs are usually more accurately classified as consumption rather than education expenses, and must be paid regardless of college enrollment status.

to 1958) and the *Digests of Education Statistics* (since 1962).

4.3 Life-Cycle Wage Profiles

Life-cycle wage profiles $w(a,t,e,s)$ are estimated using decennial U.S. Census data from 1940 through 2000, along with American Community Survey (ACS) data from 2006-2010. Each ACS data set is a 1 percent sample of the U.S. population, so that when combined they constitute a 5 percent of the U.S. population, similar to a decennial census. The data are collected from the Integrated Public-Use Microdata Series (IPUMS) (Ruggles et al., 2010), and include wage and salary income, educational attainment, age, and sex. From age and education data we compute potential labor market experience, x, as age minus years of education minus six. We assume that wages can be drawn from one of three education categories - high school, some college, or college. These correspond to $e = 0$, $e \in \{1, \ldots, C-1\}$ and $e = C$ in the model. For each education category, we estimate wage profiles for the non-institutionalized population between ages 17 and 65 who report being in the labor force using the following regression:

$$\log(w_{i,t}) = \sum_{j=1}^{4} \delta_j^s x_{i,t}^j + \delta_5 b_{i,t} + \delta_6 b_{i,t}^2 + \varepsilon_{i,t} \tag{4.1}$$

where i denotes individuals, b is birth-year cohort, s is sex, and x is potential labor market experience. In words, we regress log wages on a sex specific quartic function of experience and a quadratic function of birth year. This specification yields life-cycle wage profiles that vary by experience, education, and sex, as well as education premia that vary by birth cohort.

4.4 Probability of Passing College

We next need to set the annual probability of passing college by ability, $\pi(\alpha)$. Our proxy for ability in the NLSY79 data is an individual's AFQT percentile score. Note that $\pi(\alpha)$ is a reduced form way to capture college non-completion for any reason, including failure and voluntary drop-out, so we proceed by first calculating the *cumulative* probability of college completion among individuals in each AFQT percentile bin. From this we calculate the annual probability equal to the cumulative probability raised to the one-fourth power.

Thus, we assume there are four independent annual opportunities for failure, since the length of college is $C = 4$. Finally, we estimate the function $\pi(\alpha)$ by linear regression of the form $\pi(\alpha) = \beta_0 + \beta_1\alpha + \epsilon$, where α is the AFQT percentile. The estimated coefficients are $\beta_0 = 0.598$ and $\beta_1 = 0.004$, with $R^2 = 0.72$. Figure 4 shows the annual probabilities of passing college by AFQT percentiles in the NLSY79 data, and the linear regression line is the estimated function $\pi(\alpha)$.

Figure 4: Annual Probability of Passing College by AFQT Percentile

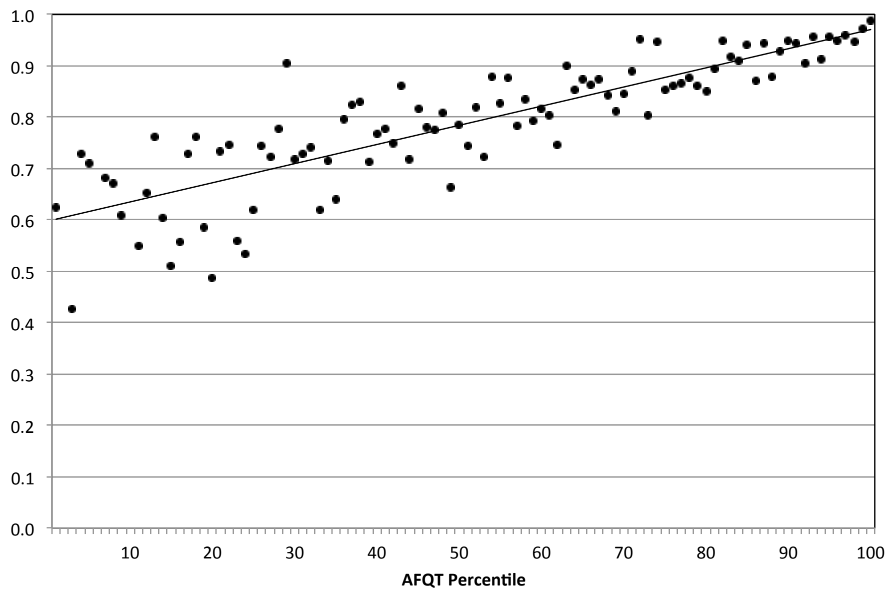

4.5 Joint Distribution of Assets and Ability

The last step in our model-free estimation is to construct a realistic joint distribution of initial assets and ability. While the estimation strategy is detailed completely in Appendix B, we briefly outline it here.

We require marginal distributions of ability and parental transfers, along with any potential correlation between the two. However, to our knowledge, no single data source contains information on both a measure of innate ability and parental transfers during the time period we consider. Our strategy is therefore to merge information from two separate datasets: the NLSY79 and the National Center for Education Statistics' High School & Beyond (HSB) dataset. NLSY79 includes information about students' grades and ability (i.e., AFQT score), but no information on parental transfers. HSB, on the other hand, includes information about grades and transfers. We utilize the fact that both datasets include high school grades. We first estimate the empirical link between grades and ability using information in the NLSY. We then impose that same relationship in HSB data, thus allowing us to construct a measure of predicted ability in the HSB. Finally, we derive the joint distribution of ability and parental transfers within the HSB data.

4.5.1 Ability Distribution

We first estimate the grade-ability relationship in the NLSY. We assume GPA is lognormally distributed, then use a linear regression to link log GPA to ability (proxied in the NLSY by AFQT score). By our distributional assumption on GPA, ability is therefore normally distributed as $\alpha^{NLSY} \sim N(\beta_0 + \beta_1 \mu_{g,NLSY},\ \beta_1^2 \sigma_{g,NLSY}^2 + \sigma_\varepsilon^2)$, where β_0, β_1, and σ_ε are the estimates and error standard deviation from the linear regression, and $\mu_{g,NLSY}$ and $\sigma_{g,NLSY}$ are the associated parameters of the GPA distribution.

Next, we project the estimated relationship between grades and ability onto grade information in HSB. Again assuming that HSB GPA is lognormal, this implies

$$\alpha_{HSB} \sim N(\beta_0 + \beta_1 \mu_{g,HSB},\ \beta_1^2 \sigma_{g,HSB}^2 + \sigma_\varepsilon^2). \tag{4.2}$$

Note that the parameters β_0, β_1, and σ_ε are estimated from the NLSY, but $\mu_{g,HSB}$ and $\sigma_{g,HSB}$ must be estimated from from the HSB grade data.

The estimation of $\mu_{g,HSB}$ is the last step of constructing the marginal ability distribution. HSB includes grade *bins*, not grades themselves, which requires us to estimate the distribution using masses in each discrete bin, not the underlying continuous distribution. Figure 5 shows that we match the distribution well with parameters $\mu_{g,HSB} = 4.404$ and $\sigma_{g,HSB} = 0.096$. Using these estimates in equation (4.2) gives us the implied marginal distribution of ability in the HSB dataset.

4.5.2 Transfer Distribution

The empirical distribution of parental transfers is best matched with a gamma distribution. The shape and scale parameters are chosen to minimize the sum of squared errors between

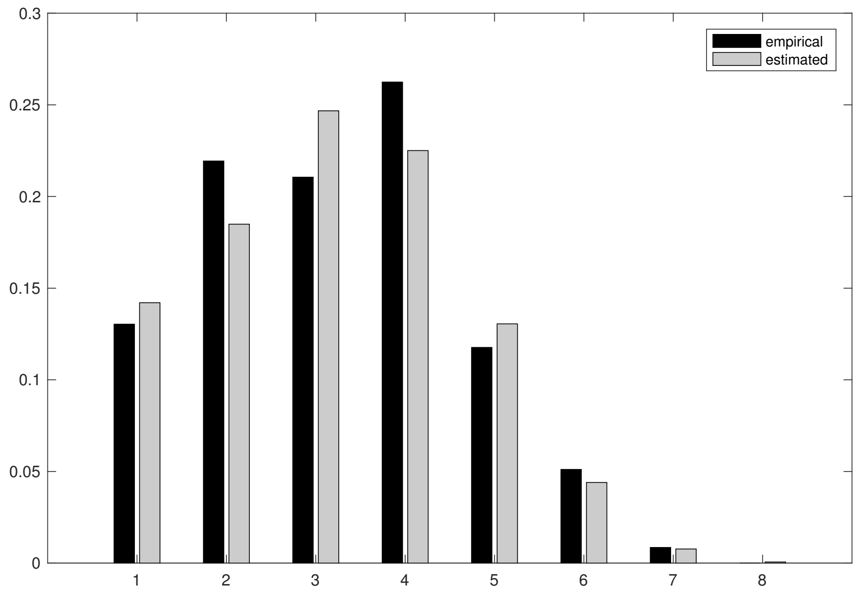

Figure 5: Data and Predicted Grade Bins in HSB

that empirical c.d.f. of average transfers (normalized by the mean) and that of the estimated gamma distribution. The best fit parameters are a shape parameter of 0.24 and a scale parameter of 4.44. Figure 6 plots the estimated and empirical distributions together, and shows that we match the data well. Lastly, we assume that the average transfer increases at the same rate as average disposable income in the United States.

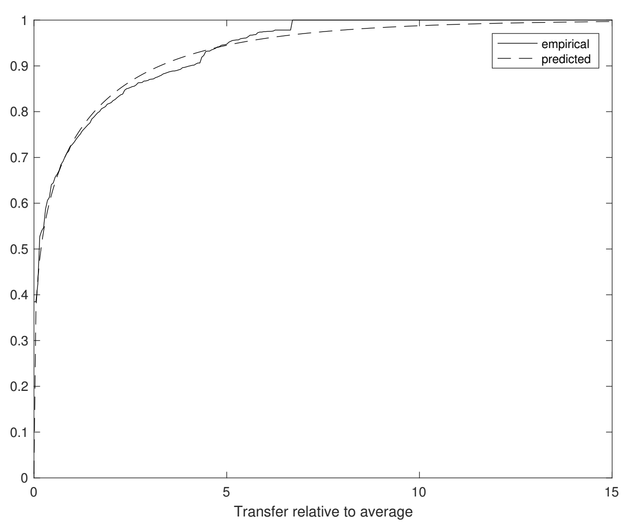

Figure 6: Cumulative Distribution Function for Transfers in HSB

4.5.3 Joint Distribution of Ability and Transfers

The last step is to compute the joint distribution of the two marginals we created above. We use a Frank copula to accomplish this, which takes the form

$$C(u,v) = \frac{-1}{\rho}\log\left(1 + \frac{(\exp(-\rho u)-1)(\exp(-\rho v)-1)}{\exp(-\rho)-1}\right).$$

The parameter ρ governs the dependence of draws. Our joint distribution of α and k_0 can therefore be written as

$$H(\alpha, k_0) = C\left[F(\alpha), G(k_0)\right]$$

where F and G are the cumulative distribution functions of the normal and gamma respectively. We are left to calibrate ρ, which roughly implies a positively correlation between the two series when $\rho > 0$. Note, however, that while our procedure gives the marginal distribution of ability in HSB, it does not provide individual-level estimates of ability. We therefore use the implied relationship to draw individual realizations of ability for students in the HSB data. We then compute the Kendall rank coefficient and repeat this simulation 1000 times. The average Kendall rank coefficient is 0.42, implying that high school graduates with higher ability on average have higher initial asset holdings. This implies a copula parameter of $\rho = 4.46$.

4.6 Variance of Ability Signals

We now proceed to impose some structure the process of ability signals across cohorts: $\{\sigma_{\varepsilon,t}\}_{t=1900}^{t=1972}$. Conceptually, the signals should be more precise (i.e., the variance should be smaller) when students have more and better information about their actual ability and likelihood of passing college. Therefore to discipline this series we seek data on how this information may have changed over time.

In the modern era an important signal of ability and college-readiness is a student's performance on standardized admissions tests like the Scholastic Aptitude Test (SAT). Introduced in 1926 (corresponding to the 1908 birth cohort), the SAT gained widespread use in the decades following. Beale (1970) reports that standardized testing grew throughout the first half of the 20th century, and by 1960 most colleges considered SAT scores "absolutely essential" to the admissions process. As standardized testing increased, students also gained better information about their own abilities, so we will utilize SAT testing data as a proxy to discipline the precision of ability signals over time.

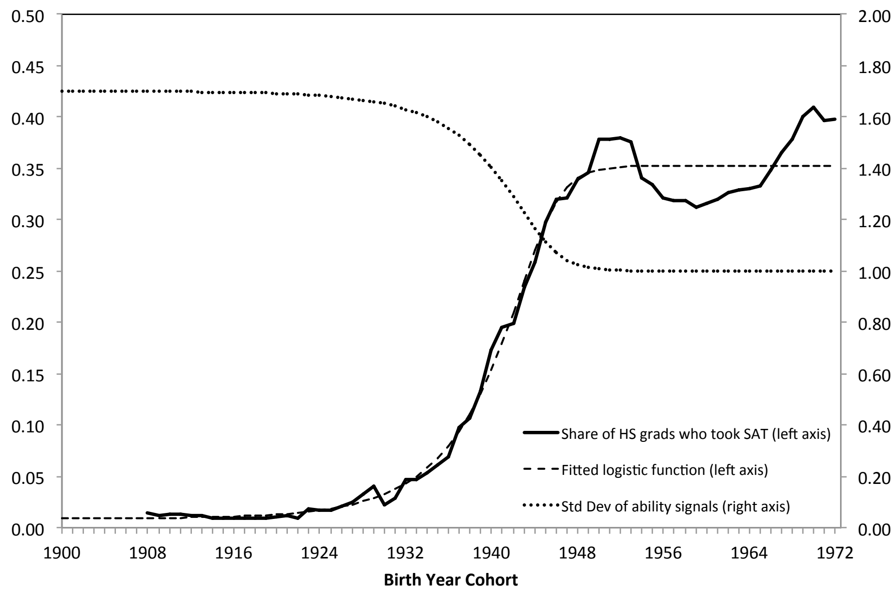

Figure 7: SAT Test Takers and Standard Deviation of Ability Signals

In Figure 7, the solid line shows the share of high school graduates who took the SAT in each birth cohort. The dashed line shows that this series is well-approximated by a generalized logistic function of the form $Y(t) = A_0 + \frac{A_\infty - A_0}{(1+e^{-B(t-M)})^{\frac{1}{V}}}$. This family of functions is often used to model the growth or diffusion of new technology, so it is a natural choice in this context. Parameters A_0 and A_∞ represent the asymptotic values at time zero and ∞, B controls the growth rate, M is the point in time when maximal growth occurs, and V governs the inflection point. Table 2 shows the estimated parameter values obtained via nonlinear least squares regression.

For the benchmark model, we impose this same functional form on $\sigma_{\varepsilon,t}$, and we maintain the parameters B, M, and V estimated in Table 2. Thus, we assume that the variance of ability signals changes concurrent with the rise in SAT testing, so that $\sigma_{\varepsilon,t}$ will have the same growth rate and inflection points as the curve fitted to SAT test data. However, we leave the asymptotic values A_0 and A_∞ as free parameters to match the magnitude of changes in ability sorting over time in the data.

Table 2: Parameters of logistic function estimated on SAT test takers

Parameter	Value
A_0	.009
	(.006)
A_∞	.352
	(.004)
B	.622
	(.222)
M	44.7
	(1.02)
V	3.43
	(1.70)
R^2	0.993
Obs.	65

4.7 Within-Model Calibrated Parameters

Finally, we are left to calibrate the asymptotic values A_0 and A_∞ for the function $\sigma_{\varepsilon,t}$, as well as the borrowing constraint λ. We choose these jointly so that the model matches as closely as possible the data shown in Figures 1 and 2b. Specifically, we target the time series of college completion for 23-year-olds and the quadratic trend line fitted to the difference in average ability of college and non-college individuals. We search over a parameter grid with spacing of 0.1 for A_0 and A_∞ and 0.001 for λ (one-tenth of one percent). The calibrated values that yield the best model fit are $A_0 = 1.7$, $A_\infty = 1.0$, and $\lambda = 0.0511$. The dotted line in Figure 7 illustrates that the path of $\sigma_{\varepsilon,t}$ is indeed an inverted version of the SAT test taking data with the same inflection points and growth rates. The shape reflects the intuition that the rise in standardized testing coincided with a decrease in the variance of ability signals.

5 Quantitative Results

With the calibrated model in hand, we now simulate the model for U.S. birth cohorts from 1900 through 1972 and verify that it replicates important features of the historical data.

5.1 Benchmark Model Fit

Figure 8a shows that, overall, the model replicates well the long-run trends in U.S. college attainment over much of the 20th century. Yet there are some notable deviations that warrant further discussion. First, the model under-predicts attainment by 2.8 percentage points, on average, for cohorts born between 1910 and 1920. In Section 6.1 we discuss evidence on relaxed borrowing constraints for these cohorts that can help reconcile this gap. In addition the model does not capture the initial decline and subsequent increase in college completion for cohorts born after 1950. Nevertheless, across the 1950–1972 cohorts, the average share who complete college by age 23 is 25.4% in the data, and 25.2% in our model. Despite missing some of the year-to-year variations, the model captures the long-run trends quite well. Furthermore, we note that some of the volatility in college attainment over this time period is related to non-economic factors like the Vietnam War, which our model is not designed to capture.

Figure 8: Benchmark Model Results

(a)

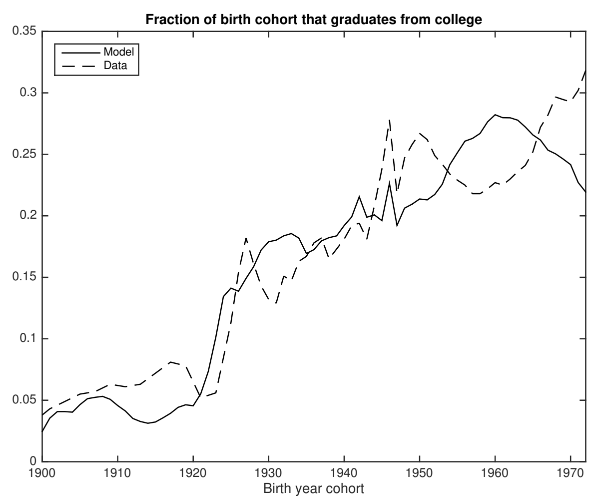

(b)

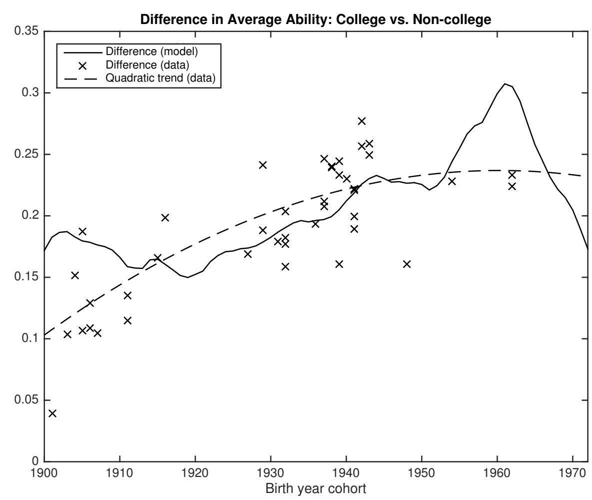

Figure 8b plots the average ability difference between college and non-college individuals in the data (shown earlier in Figure 2b) along with the corresponding model prediction. For cohorts born at the beginning of the 20th century, college and non-college students had similar ability on average, but the ability gap widened in the decades following. This general

Table 3: Growth rate decomposition in data versus model

Panel A: Data	Log differences			Annualized Percent Change		
Birth years	γ^{col}	γ^{hs}	γ^{endog}	γ^{col}	γ^{hs}	γ^{endog}
1900 - 1972	2.13	1.53	0.59	2.95%	2.13%	0.82%
1900 - 1920	0.54	1.13	-0.59	2.68%	5.64%	-2.95%
1920 - 1972	1.59	0.41	1.18	3.05%	0.78%	2.27%
1950 - 1972	0.18	-0.10	0.27	0.80%	-0.44%	1.24%
Panel B: Model	Log differences			Annualized Percent Change		
Birth years	γ^{col}	γ^{hs}	γ^{endog}	γ^{col}	γ^{hs}	γ^{endog}
1900 - 1972	2.08	1.53	0.55	2.89%	2.13%	0.76%
1900 - 1920	0.55	1.13	-0.57	2.77%	5.64%	-2.87%
1920 - 1972	1.53	0.41	1.12	2.94%	0.78%	2.15%
1950 - 1972	0.02	-0.10	0.12	0.09%	-0.44%	0.53%

pattern is also predicted by the model. However, the model over-predicts the gap among cohorts born between 1900 and 1910, while it does a better job capturing the increase in post-1920 cohorts. Overall, while the quadratic trend in the data exhibits an increase of about 13 percentage points (from 0.10 to 0.23) over cohorts born 1900 to 1950, the model predicts an increase of almost 6 percentage points (from 0.17 to 0.23) over the same time. Thus, the model captures approximately half of the ability sorting observed in the data. Note, however, that while we capture nearly the entire change between 1920 and 1950 cohorts, we over-predict the ability gap among pre-1910 cohorts. We return to this in Section 7, with a more detailed discussion of ability signals.

We next assess the model performance over the various sub-periods discussed in the growth accounting decomposition of Section 2. Table 3 provides this decomposition for both the model and data to facilitate comparison. Recall that high school graduation rates are exogenously fed into the model, hence the exact match of γ^{hs} between model and data. However, the endogenous factors (γ^{endog}) produced by the model also match the data quite well. We predict a decrease in these factors over the 1900-1920 birth cohorts, and then substantial growth for the 1920–1972 cohorts, consistent with observed trends in the data.

For further evidence that the "endogenous" portion of college completion matches well, we plot

$$\frac{\text{Number of college degrees awarded at year } t+4}{\text{Number of high school graduates at year } t}.$$

In the model, this is exactly equal to the fraction of high school graduates that eventually finish college (N_t^{grad}/N_t^{HS}), as all college completion is completed within four years. While this is not exactly true empirically, it provides a reasonable approximation for comparison. Figure 9 plots this time series in both the model and data, and shows that our model matches the empirical time series well. Indeed, consistent with the results in Table 3, there is a substantial decline for cohorts before 1920, an immediate increase between 1920 and 1930, and a relatively small increase for post-1930 cohorts.

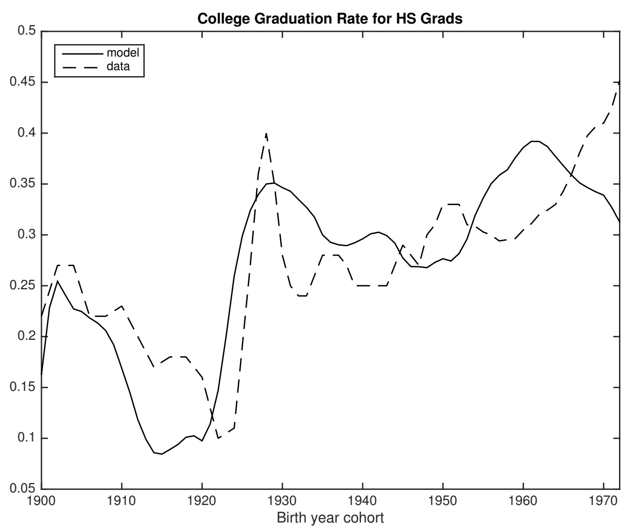

Figure 9: College Completion Rate Conditional on HS graduation

5.2 Decomposition of Endogenous Factors

Since the model matches the data well, we next turn to decomposing the endogenous factors: college enrollment of high school graduates and college graduation of high school enrollees. Unlike in the data, the model allows a full decomposition of equation (2.2), in which γ^{endog} is split into enrollment and graduation rates as $\gamma^{endog} := \gamma^{enroll} + \gamma^{grad}$. Table 4 shows results of this more detailed decomposition.

The striking result of this decomposition is the overwhelming importance of variation in college enrollment relative to variation in college graduation rates conditional on enrollment,

Table 4: Growth decomposition of endogenous factors in model

	Log Differences				Annualized Percent Change			
Birth years	γ^{col}	γ^{hs}	γ^{enroll}	γ^{grad}	γ^{col}	γ^{hs}	γ^{enroll}	γ^{grad}
1900 - 1972	2.08	1.53	0.64	-0.09	2.89%	2.13%	0.89%	-0.13%
1900 - 1920	0.55	1.13	-0.60	0.02	2.77%	5.63%	-2.98%	0.11%
1920 - 1972	1.53	0.41	1.23	-0.11	2.94%	0.78%	2.37%	-0.22%
1950 - 1972	0.02	-0.10	0.19	-0.07	0.09%	-0.44%	0.84%	-0.31%

as the latter accounts for only a small portion of changes in college attainment. Table 4 shows that during most of the period under consideration, the magnitude of γ^{enroll} dwarfs that of γ^{grad}. However, we note that for post-1950 cohorts declining average college student ability brought the graduation rate down by 0.31% annually. Combined with the 0.44% annual decline in high school graduation, these two changes almost completely offset the college enrollment growth in the model during 1950–1972.

In Figure 10a, we plot the results of several counterfactuals, in which we fix the enrollment and graduation rates at their 1900 level.[6] Indeed, the model with a fixed graduation rate is quite close to the baseline model, while fixing the enrollment rate generates substantial under-prediction of college attainment. Figure 10b shows why this is the case: variation in the graduation rate is substantially smaller than variation in the enrollment rate.

In the next section, we turn to understanding the underlying forces behind the dynamics of the enrollment and graduation rates.

6 Key Channels for College Enrollment and Graduation

In this section, we assess the role of college costs and the college earnings premium for understanding the results in the previous section, as both vary substantially over time. We isolate the effects of both by computing counterfactual economies in which each is held constant.

[6] These terms are defined in equation (2.1), but they are the fraction of high school graduates who enroll in college (N_t^{enroll}/N_t^{HS}) and the fraction of college enrollees who graduate (N_t^{grad}/N_t^{enroll}) respectively.

Figure 10: Decomposing Enrollment and Graduation

(a) Counterfactuals

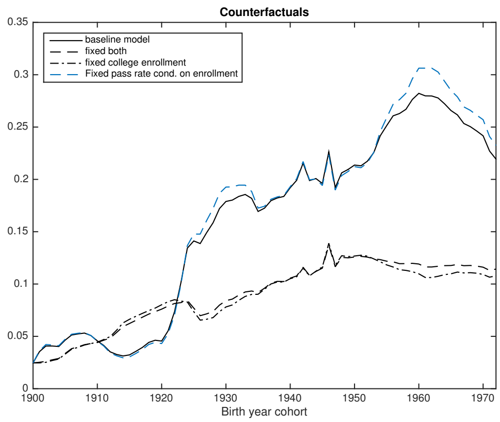

(b) Enrollment and Graduation Rates

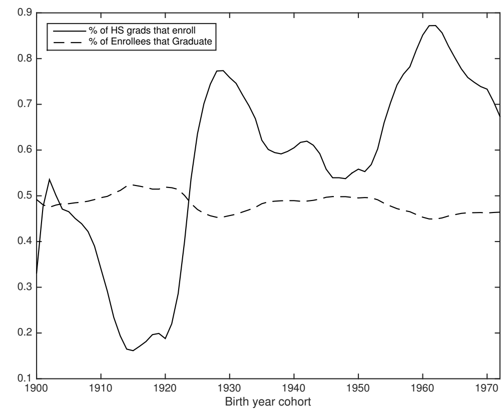

6.1 College Costs

We begin by considering the impact of college costs. Figure 11a plots the baseline level of tuition costs as a share average annual income in the year that cohort makes college decisions. This measure therefore gives a sense of the relative affordability of college across cohorts, and there is substantial variation. In particular, the relative cost spikes dramatically for cohorts born after 1912 until the early 1920s, as these cohorts made college decisions during the Great Depression when average incomes fell substantially. Shortly thereafter relative costs decrease sharply due to growing incomes and increased subsidies (recall these are out-of-pocket costs, net of scholarships, grants, and other aid). Finally, relative costs grow consistently for post–1930 cohorts. To assess the importance of these changes, we compute a counterfactual economy in which the direct cost of attending college is fixed as a constant fraction of average income, shown as the dashed line in Figure 11a. Figures 11b and 11c show the predicted time series for college attainment and ability sorting, as well as the baseline model predictions for comparison.

Consistent with the variation in costs, the counterfactual model predicts higher attainment for the 1910–1920 birth cohorts. Recall from Figure 8a that attainment in the baseline model was lower than the data for these cohorts, so the counterfactual economy actually fits better over this short time span. One potential reason for the baseline model to under-

Figure 11: Counterfactual economy with constant college costs

(a) College Costs

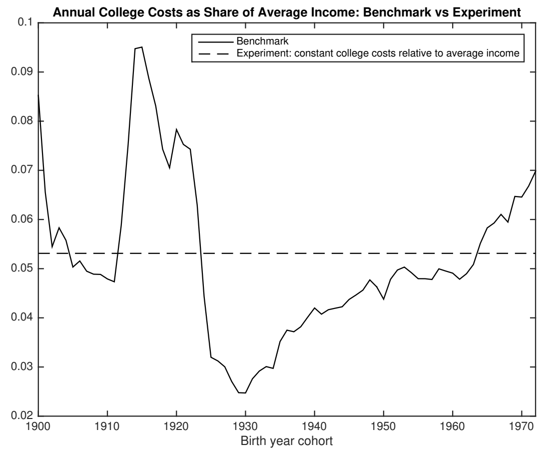

(b) Attainment

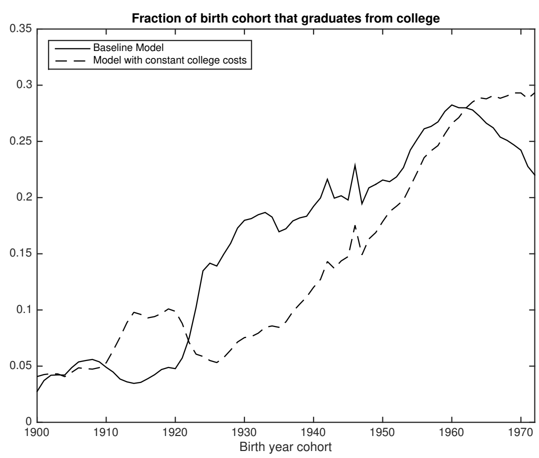

(c) Ability Sorting

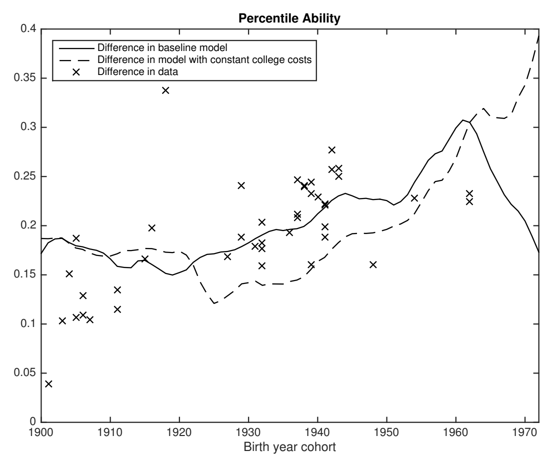

predict the data is that we hold the borrowing constraint fixed across cohorts, yet there is qualitative evidence that borrowing constraints may have actually eased somewhat during the Great Depression. The 1930-32 edition of the *Biennial Survey of Education* notes: "This question of student fees for the support of higher education is naturally made more acute by the depression. Students and parents have insufficient funds to meet high tuition rates. They offer notes in lieu of currency for their fees." In other words, some colleges and universities were extending credit directly to students and their parents. Additionally, the 1933 Federal Emergency Relief Administration extended temporary federal financial aid to college students.

Looking beyond these early cohorts, we see that the counterfactual exercise under-predicts attainment for cohorts born after the early-1920s until 1960. In other words, college enrollment rates in the baseline model were boosted substantially by the large decline in relative college costs experienced by pre-1930 cohorts and the relatively low costs that sustained through subsequent decades. In fact, absent these changes our counterfactual model predicts that college enrollment and completion would have been lower by almost half, on average, for the 1925-1950 cohorts, compared to the baseline economy. As seen in Figure 11b, the largest gaps are for the cohorts born during the early 1930s; college attainment is depressed by more than 10 percentage points in the counterfactual economy for these cohorts.

For cohorts born after 1960, the counterfactual model predicts college attainment that is actually closer to the data than the benchmark model. Higher college attainment in this exercise is due to lower direct college costs, but the same effect would be generated by relaxing borrowing constraints. And in reality these cohorts did experience greater ability to borrow as Federal Student loan limits expanded in the late 1970s and 1980s. By contrast, attainment in the baseline model falls over the last decade as costs are rising and the borrowing limit is fixed.

Finally, we note in Figure 11c that this counterfactual model predicts a fairly similar trajectory of ability sorting compared to the baseline, at least until the 1960-1970 cohorts, when they diverge substantially. There is a small level difference from the mid-1920s through the 1950s, but the change in ability sorting is similar over this time period in both models. What is remarkable, however, is that the large fluctuations in relative college costs for pre-1930 cohorts had only minor impact on ability sorting.

Table 5, Panel B provides a decomposition of this counterfactual economy into the same components and sub-periods previously studied. The baseline results are also reported again for ease of comparison in Panel A. Over the entire sample period, growth in college attainment is quite close to the baseline results. The overall patterns in the sub-periods are consistent as well, but with some differences in magnitude. The largest difference between the two is seen in the most recent period. Comparing the annualized percent changes over the 1950–1972 cohorts, the baseline model predicts a substantial slowdown in college attainment, but this slowdown does not occur in the counterfactual with constant college costs. Clearly, rising

Table 5: Model decomposition in baseline versus counterfactual exercises

Panel A: Baseline Model

	Log Differences				Annualized Percent Change			
Birth years	γ^{col}	γ^{hs}	γ^{enroll}	γ^{grad}	γ^{col}	γ^{hs}	γ^{enroll}	γ^{grad}
1900 - 1972	2.08	1.53	0.64	-0.09	2.89%	2.13%	0.89%	-0.13%
1900 - 1920	0.55	1.13	-0.60	0.02	2.77%	5.63%	-2.98%	0.11%
1920 - 1972	1.53	0.41	1.23	-0.11	2.94%	0.78%	2.37%	-0.22%
1950 - 1972	0.02	-0.10	0.19	-0.07	0.09%	-0.44%	0.84%	-0.31%

Panel B: Counterfactual with constant college costs

	Log Differences				Annualized Percent Change			
Birth years	γ^{col}	γ^{hs}	γ^{enroll}	γ^{grad}	γ^{col}	γ^{hs}	γ^{enroll}	γ^{grad}
1900 - 1972	1.97	1.53	0.56	-0.12	2.74%	2.13%	0.78%	-0.17%
1900 - 1920	0.89	1.13	-0.27	0.03	4.43%	5.63%	-1.35%	0.15%
1920 - 1972	1.09	0.41	0.83	-0.15	2.09%	0.78%	1.60%	-0.30%
1950 - 1972	0.50	-0.10	0.75	-0.16	2.25%	-0.44%	3.43%	-0.73%

Panel C: Counterfactual with constant earnings premia

	Log Differences				Annualized Percent Change			
Birth years	γ^{col}	γ^{hs}	γ^{enroll}	γ^{grad}	γ^{col}	γ^{hs}	γ^{enroll}	γ^{grad}
1900 - 1972	0.09	1.53	-1.46	0.02	0.13%	2.13%	-2.03%	0.02%
1900 - 1920	1.38	1.13	0.29	-0.03	6.92%	5.63%	1.46%	-0.17%
1920 - 1972	-1.29	0.41	-1.70	0.05	-2.49%	0.78%	-3.37%	0.10%
1950 - 1972	-1.27	-0.10	-1.12	-0.05	-5.78%	-0.44%	-5.11%	-0.24%

college costs relative to income play a substantial role in dampening enrollment for these cohorts. Finally, comparing the γ^{enroll} and γ^{grad} columns, we again see that all of these changes in the counterfactual economy are driven primarily by changes in enrollment rather than graduation rates, as was the case in the baseline model.

6.2 Earnings Premia

As documented by Goldin and Katz (2010), the college earnings premium followed a U-shaped pattern over much of the 20th century. Consistent with their evidence, our estimated wage profiles show that education premia bottomed out for birth cohorts in the 1920s, and begin rising rapidly for cohorts born post-1930. To assess the importance of changes in the education premia over time, we simulate a counterfactual economy in which the education premia (for both college and some-college education) are constant across time.[7] Figures 12a and 12b show the estimated college earnings premium relative to high school graduates for men and women, along with the counterfactual fixed levels. The premium for some college looks very similar, so we omit it to save space.

Figure 12: College Earnings Premia for Men and Women

(a) Men

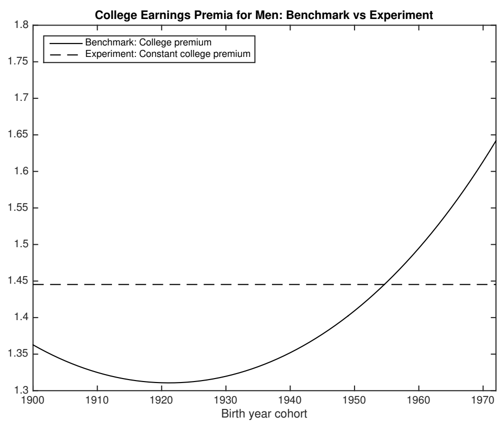

(b) Women

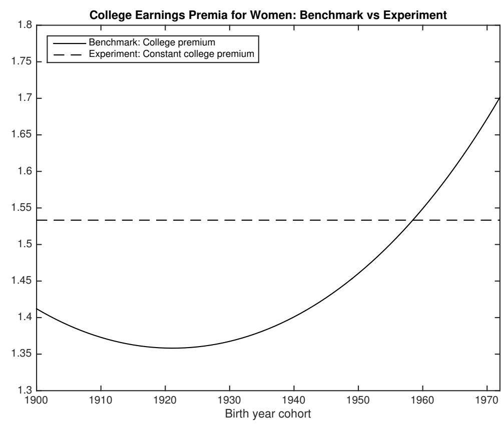

(c) Attainment

(d) Ability Sorting

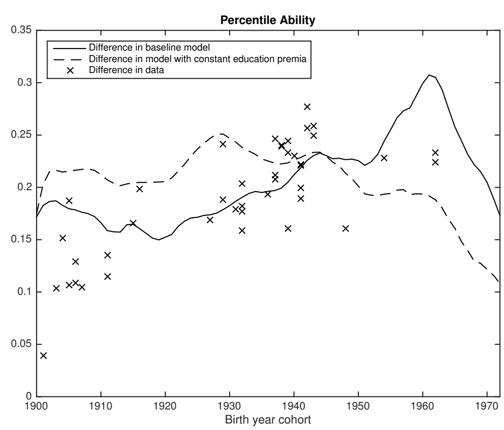

Fixing the education earnings premia over time implies that the counterfactual model will overshoot college attainment for early cohorts, and then undershoot for later cohorts

[7]For this experiment we also re-calibrate the borrowing constraint to match college attainment for the 1900 cohort, which results in $\lambda = 0.046$. Absent this adjustment the counterfactual college premium for pre–1940 cohorts is large enough, and the resulting borrowing limit is loose enough, that nearly all high school graduates are induced to attempt college.

in the absence of the large increase in the education premia. Furthermore, because college costs are simultaneously increasing over this time, the counterfactual model implies a large drop in college attainment starting around the 1940 cohort. Figure 12c shows this this result clearly, and Panel C in Table 5 provides the decomposition of this counterfactual model. Compared to the baseline in Panel A, we see drastic differences over the entire 1900–1972 sample period; however, most of this is due to the post–1940 effects just discussed. Also, as in the baseline model and the counterfactual with constant college costs, we again find that changes in graduation rates are relatively less important than changes in college enrollment.

Examining the 1950–1972 sub-period, we find that college enrollment and completion both fall by more than 5% annually in this counterfactual exercise. Recall from earlier that the average college completion rate among 1950–1972 cohorts is is 25.4% in the data, and 25.2% in the baseline model, yet in this counterfactual it falls to 10.3%. Clearly the rising college premium was a dominant force driving higher college attainment for these cohorts, even as they faced simultaneously rising college costs.

Finally, turning to ability sorting in Figure 12d we see that holding the earnings premia constant results in almost no increase in ability sorting for pre–1940 cohorts, and a substantial decrease in the average ability difference post–1940. The ability difference is higher early on in this counterfactual exercise because college enrollment is much higher compared to the baseline. Individuals who choose not to attempt college despite the higher premium do so because their ability (and therefore probability of passing) is very low.

Taken together, these counterfactual results demonstrate an important role for the both the education premium and college costs in accounting for college attainment over time. Large declines in relative college costs were an important factor for increasing enrollment, particularly among pre–1930 cohorts. Increases in education premia had greater impact on post–1940 cohorts, but these effects were dampened by rapidly rising college costs later on. Next we assess the critical model aspects that drive increased ability sorting over time.

7 Sorting and Ability Signals

While the model predicts an increase in ability sorting over time, it only captures about half of the difference found in the data. In this section, we pose two questions. First, what drives the predicted increase in ability sorting in the model? As shown in the counterfactual experiments, both the earnings premium and college costs play a minor role in ability sorting; however, we show here that the critical model component generating the sorting is a decrease in the variance of ability signals. Therefore, we turn to a second question: what time path of signal variance would be required to match the entire increase in the data?

7.1 Constant Variance of Ability Signals

The previous section showed that the economic factors we considered have relatively little impact on ability sorting over time. Rather, students' information about their own ability (i.e., the noisiness of their ability signal) is a key determinant of ability selection in the college decision. To demonstrate this, we conduct a counterfactual exercise in which the standard deviation of the ability signal is held constant at the initial level of $A_0 = 1.7$. Figure 13a shows that the signal variance in the the baseline and counterfactual economies begins diverging around the 1920 birth cohort. Figures 13b and 13c compare college attainment and ability sorting in this counterfactual model with the baseline model. First, note that the signal variance has almost no affect on overall attainment, as seen in Figure 13b. However, Figure 13c shows that as the signal variance between the two models begins to diverge, so does the predicted ability difference. If ability signals had not become more precise over time, in fact, this model predicts that the difference in average ability would have fallen between the 1930 and 1950 cohorts, rather than risen. This drop in signal variance is therefore critical for the baseline model to generate the increasing ability gap observed in the data.

7.2 Targeting Ability Sorting Directly

Because signal precision plays an important role in understanding ability sorting, we lastly ask what pattern of signals would be required to match ability sorting exactly. To do so, we remove the functional form assumption that we placed on $\sigma_{\varepsilon,t}$ in the baseline model and

Figure 13: Counterfactual model with constant ability signals

(a) Signal Variance Time Series

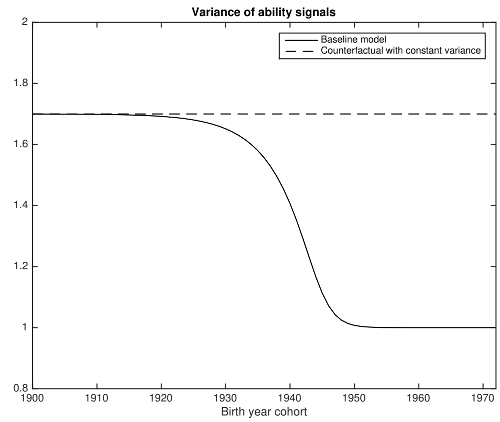

(b) Attainment

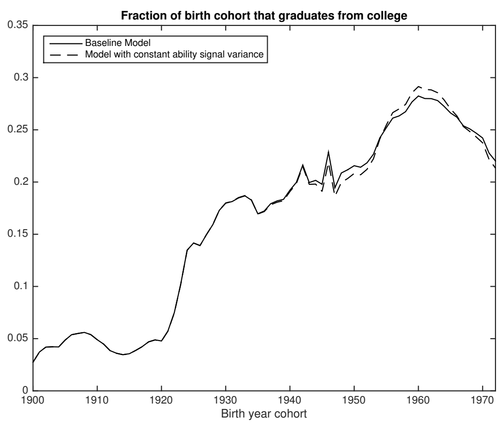

(c) Ability Sorting

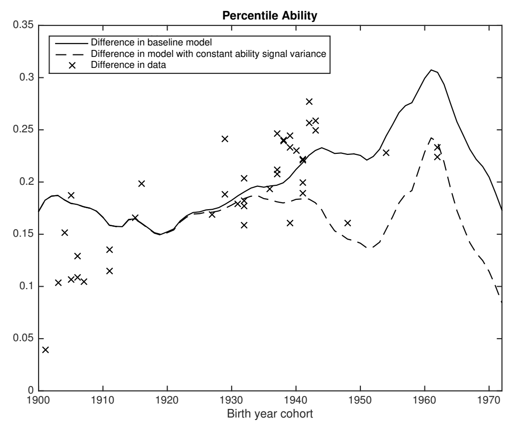

instead allow the ability signals to adjust freely across cohorts. We then search for values of $\sigma_{\varepsilon,t}$ at each cohort so that we match exactly the quadratic trend line shown in Figure 2b. The estimated time series for $\sigma_{\varepsilon,t}$ is shown in Figure 14a along with the baseline case (where we assumed a logistic functional form) for comparison. Additionally, Figures 14b and 14c show the predicted college attainment and ability sorting in this counterfactual exercise.

Several results are notable. First, Figure 14a shows that the counterfactual model requires a substantially larger signal variance for pre–1910 cohorts in order to match the small difference in average ability early on.[8] Intuitively, as variance increases individuals become

[8] To best match sorting, the model requires a 167 percent decrease in $\sigma_{\varepsilon,t}$ from 3 to 1 over time. This is a substantial drop, though we admittedly have little evidence to quantify its relative magnitude. Better understanding the underlying sources of this movement or its interaction with other theories of ability sorting, especially among early 1900s cohorts, is an interesting topic we leave unanswered here.

Figure 14: Counterfactual economy with ability signals as free parameter

(a) Variance of Ability Signal

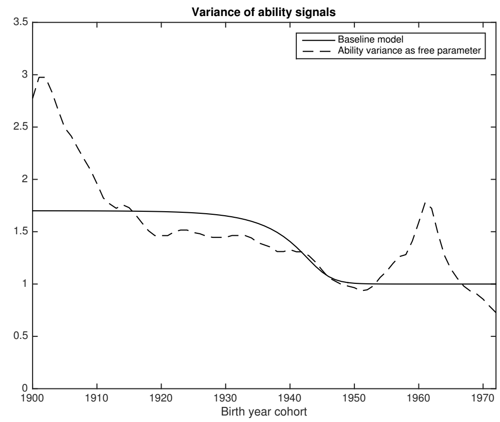

(b) Attainment

(c) Ability Sorting

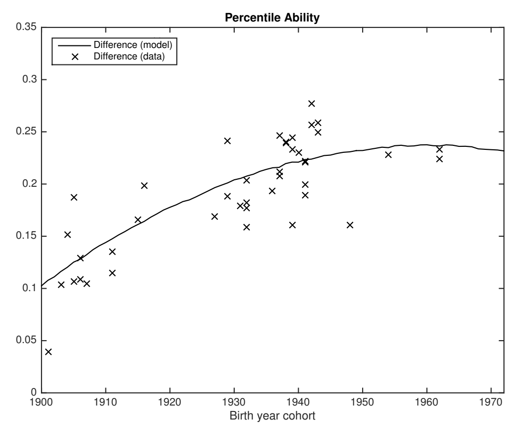

less certain of their true ability to pass college, which drives down enrollment among high ability students and drives up enrollment among lower ability students. This mechanism is certainly plausible during the earliest parts of the 20th century before widespread standardized testing gave students much more precise information about their true ability. Second, we also see in Figure 14a that the logistic functional form utilized in the baseline is remarkably close to the freely chosen series for $\sigma_{\varepsilon,t}$ during the 1910–1950 cohorts, which is when the most significant ability sorting occurred. Third, we observe in Figure 14b that the variance of ability signals continues to have little impact on overall attainment, consistent with the previous section. Signal precision seems to be a unique feature in this model (unlike college costs or education premia) that can affect selection in and out of college along the margin of student ability without having a significant impact on overall attainment.

8 Conclusion

We develop a life-cycle decision model to investigate long-run trends in college completion and ability sorting for the 1900–1972 birth cohorts. Key features of the model included unobservable ability, heterogenous parental transfers, and borrowing constraints. To discipline our model, we utilize historical data series including statistics on high school graduation rates, college costs, education earnings premia, and standardized college admissions testing. The model matches the overall trends in college attainment well. Furthermore, when we decompose the endogenous components of college completion, we find that changes in completion are driven almost exclusively by changes in college enrollment, not in the graduation rate of college students.

However, the economic factors that drive attainment – including real college costs and education wage premia – have little quantitative role in increasing ability sorting over time. Instead, signals about underlying ability are the key factor for sorting in this model. In particular, a decrease in the variance of ability signals can properly match ability sorting over time, a trend which we attribute to increases over time in standardized testing. One notable shortcoming of the model is that it predicts a larger average ability gap between college and non-college individuals born prior to 1910. Reconciling model and data for these cohorts would require either higher variance ability signals for early 1900 cohorts or some other change to the economy that is not considered here. One promising approach would also consider ability selection at high school graduation, which we do not consider here. The rapid growth of high school completion, particularly during the 1900–1920 cohorts, almost certainly resulted in changes in the average ability of high school graduates relative to the total population. Such changes would subsequently affect college student ability, but they are difficult to document empirically. We leave these topics for future research, with the goal to better understand who chooses to attend college, and why.

References

Athreya, Kartik, and Janice Eberly. 2016. "Risk, the College Premium, and Aggregate Human Capital Investment."

Beale, Andrew V. 1970. "The Evolution of College Admission Requirements." *Journal of the National Association of College Admissions Counselors*, 15(3): 14–15.

Belley, P., and L. Lochner. 2007. "The Changing Role of Family Income and Ability in Determining Educational Achievement." *Journal of Human Capital*, 1(1): 37–89.

Bound, John, and Sarah Turner. 2002. "Going to war and going to college: Did World War II and the GI Bill increase educational attainment for returning veterans?" *Journal of Labor Economics*, 20(4): 784–815.

Bound, John, Michael F. Lovenheim, and Sarah Turner. 2010. "Why have college completion rates declined? An analysis of changing student preparation and collegiate resources." *American Economic Journal: Applied Economics*, 2(3): 129.

Card, David, and Thomas Lemieux. 2001*a*. "Dropout and enrollment trends in the postwar period: What went wrong in the 1970s?" In *Risky behavior among youths: An economic analysis*. 439–482. University of Chicago Press.

Card, David, and Thomas Lemieux. 2001*b*. "Going to college to avoid the draft: The unintended legacy of the Vietnam War." *American Economic Review*, 91(2): 97–102.

Castro, Rui, and Daniele Coen-Pirani. 2016. "Explaining the Evolution of Educational Attainment in the United States." *American Economic Journal: Macroeconomics*, 8(3): 77–112.

Cunha, Flavio, James Heckman, and Salvador Navarro. 2005. "Separating uncertainty from heterogeneity in life cycle earnings." *Oxford Economic Papers*, 57(2): 191–261.

Garriga, Carlos, and Mark P. Keightley. 2007. "A General Equilibrium Theory of College with Education Subsidies, In-School Labor Supply, and Borrowing Constraints." Federal Reserve Bank of St. Louis Working Paper.

Goldin, Claudia, and Lawrence F. Katz. 2010. *The Race between Education and Technology.* Harvard University Press.

Hendricks, Lutz, and Oksana Leukhina. 2016. "The Return to College: Selection Bias and Dropout." *Working Paper.*

Hendricks, Lutz, and Todd Schoellman. 2014. "Student abilities during the expansion of US education." *Journal of Monetary Economics*, 63: 19–36.

Hendricks, Lutz, Chris Herrington, and Todd Schoellman. 2016. "Family Background, Academic Ability, and College Decisions in the 20th Century US." *Working Paper.*

Hoxby, Caroline M. 2009. "The Changing Selectivity of American Colleges." *Journal of Economic Perspectives*, 23(4): 95–118.

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

Jones, John Bailey, and Fang Yang. 2016. "Skill-Biased Technical Change and the Cost of Higher Education." *Journal of Labor Economics*, 34(3).

Keane, Michael, and Kenneth Wolpin. 2001. "The effect of parental transfers and borrowing constraints on educational attainment." *International Economic Review*, 42(4): 1051–1103.

Keller, Elisa. 2014. "The slowdown in American educational attainment." *Journal of Economic Dynamics and Control*, 46: 252–270.

Lochner, Lance J., and Alexander Monge-Naranjo. 2011. "The Nature of Credit Constraints and Human Capital." *American Economic Review*, 101(6): 2487–2529.

Restuccia, Diego, and Guillaume Vandenbroucke. 2013. "The evolution of education: A macroeconomic analysis." *International Economic Review*, 54(3): 915–936.

Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Geoken, Matthew B. Schroeder, and Matthew Sobek. 2010. "Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]." Minneapolis: University of Minnesota.

Snyder, Thomas D., ed. 1993. *120 Years of American Education: A Statistical Portrait.* United States Department of Education.

Appendices

A College Cost Data

This appendix details the construction of annual college costs, which we take as exogenous to the model. The primary sources of the underlying data are the Biennial Surveys of Education (BSE), which were published every other year from 1918 through 1958, and the Digests of Education Statistics (DES), which have been published annually since 1962. In some cases data were revised in later publications, so we take the latest published estimates available.

We calculate average out-of-pocket college costs per student each year, λ_t, equal to the total current fund revenue from student tuition and fees divided by the number of full-time equivalent students enrolled.[9] Importantly, we consider only public colleges and universities in this calculation. Additional costs paid for private tuition or fees could be interpreted as payment for higher quality education, consumption, or to satisfy personal preferences for education at a particular institution. In any case, such marginal costs are not relevant for the binary decision of whether to attend college, so we exclude them from consideration.

We consider birth cohorts 1900 – 1972, so these students graduated high school and made college decisions 1918 – 1990. Because college takes four years, the last cohort faced college costs through the 1993-94 academic year. Therefore, we construct a time series for annual college costs that runs from 1918 – 1993 in the model. The numerator for λ_t is constructed as follows:

- 1961 – 1993: current revenues from student tuition and fees for all public 2-year and 4-year institutions of higher education, taken from various issues of the DES. Data for 1962, 1964, 1966, and 1978 are missing, so we linearly interpolate these values.
- 1957 – 1961: missing values are interpolated between the 1961 data described above, and the 1957 data described below.
- 1918 – 1957: current revenues from student tuition and fees for public colleges, univer-

[9] Alternatively, one could use the total current expenditures rather than revenues, but this makes little difference quantitatively because revenues and expenditures track each other quite closely. In addition, the revenue data is preferable because it allows us to determine how much of costs are paid out-of-pocket by students for tuition and fees, and how much comes from other sources such as state, local, and federal governments, private gifts, endowment earnings, auxiliary enterprises (athletics, dormitories, meal plans, etc.), and other sources.

sities, and professional schools, taken from various issues of the BSE. Data are generally available every other year, and missing years are linearly interpolated.

The denominator for λ_t is constructed as follows:

- 1967 – 1993: number of full-time equivalent students enrolled, taken from various issues of the DES.
- 1953, 1955, 1957, 1963, and 1969: number of full-time equivalent students calculated from the following disaggregated data series on student enrollment: full-time, part-time, first professional, graduate, unclassified, and extension students. Consistency is verified by comparing the calculated versus published data in 1969, which are within one-half of one percent of each other. Missing years are linearly interpolated.
- 1918 – 1953: No published data are available for full-time equivalence, nor have we been able to obtain disaggregated data to estimate full-time equivalence as above. Rather, we use the published data for "resident college enrollment" in public colleges, universities, and professional schools, taken from various issues of the BSE. Data are generally available every other year, and missing years are linearly interpolated.

B Estimating the Joint Distribution of Assets and Ability

This appendix provides details on the procedure used to estimate the joint distribution of initial assets (parental transfers) and ability using the NLSY79 and HSB data sets.

B.1 Ability Distribution

Grade-Ability Relationship in NLSY We begin by considering the relationship between grades and ability in the NLSY. We first construct the variable GPA which is student high school grades on a $0-100$ scale. We assume the underlying distribution of GPA is lognormally distributed with mean and variance μ_g and σ_g. We next link this to ability, which in the NLSY is AFQT scores. We impose a linear relationship between log(GPA) and AFQT scores, and estimate the parameters of this relationship with the regression

$$AFQT_i = \alpha_0 + \alpha_1 log(GPA)_i + \varepsilon_i \tag{B.1}$$

which returns coefficient estimates $\alpha_0 = -8.00$ and $\alpha_1 = 1.91$, and error term standard deviation $\sigma_\varepsilon = 0.88$. By our assumption of lognormal GPA, it follows that $AFQT$ is normally distributed as $AFQT \sim N(\alpha_0 + \alpha_1 \mu_{g,NLSY},\ \alpha_1^2 \sigma_{g,NLSY}^2 + \sigma_\varepsilon^2)$.

Imposing the Relationship in HSB With the grades-ability link now estimated in the NLSY, we next impose this relationship on the High School and Beyond (HSB) data. This step is required because HSB includes student grades, but no information on underlying ability. In principal, we could compute the underlying AFQT score in HSB by directly imposing the estimates derived from equation (B.1). After assuming HSB GPA is lognormal, and computing $\mu_{g,HSB}$ and $\sigma_{g,HSB}$, this would imply

$$AFQT_{HSB} \sim N(\alpha_0 + \alpha_1 \mu_{g,HSB},\ \alpha_1^2 \sigma_{g,HSB}^2 + \sigma_\varepsilon^2)$$

where α_0, α_1 and σ_ε^2 are estimated from equation (B.1). Unfortunately, HSB does not ask for grades directly, but instead what bins they fall into. The bins are (1) 90-100, (2) 85-89, (3) 80 - 84, (4) 75-79, (5) 70-74, (6) 65-69, (7) 60-64, (8) lower than 60. We therefore must estimate the mean and variance of grades from these bins.

To this end, let $\bar{g}_j$ and $\underline{g}_j$ be the maximum and minimum grades in any grade bin j. Letting F denote the cumulative distribution function of the underlying (lognormal) grade distribution, each bin j includes mass

$$\widetilde{M}_j = F(\bar{g}_j) - F(\underline{g}_j).$$

Let M_j be the empirical mass in each grade bin calculated in the HSB data. We therefore choose $\mu_{g,HSB}$ and $\sigma_{g,HSB}$ to minimize the sum of the squared errors $(\mu_{g,HSB}, \sigma_{g,HSB}) \in \arg\min \sum_j (\widetilde{M}_j - M_j)^2$. The estimated parameters are $\mu_{g,HSB} = 4.404$ and $\sigma_{g,HSB} = 0.096$. Figure 15 plots the estimated bins along with their empirical counterparts, and shows that they match well. The sum of squared errors between the two discrete distributions is 0.004.

We now have the underlying grade distribution in HSB, which is distributed $logN(\mu_{g,HSB}, \sigma_{g,HSB})$. Imposing the estimated link between ability and grades in the NLSY onto the HSB data, we get that $AFQT_{HSB} \sim N(\alpha_0 + \alpha_1 \mu_{g,HSB},\ \alpha_1^2 \sigma_{g,HSB}^2 + \sigma_\varepsilon^2)$.

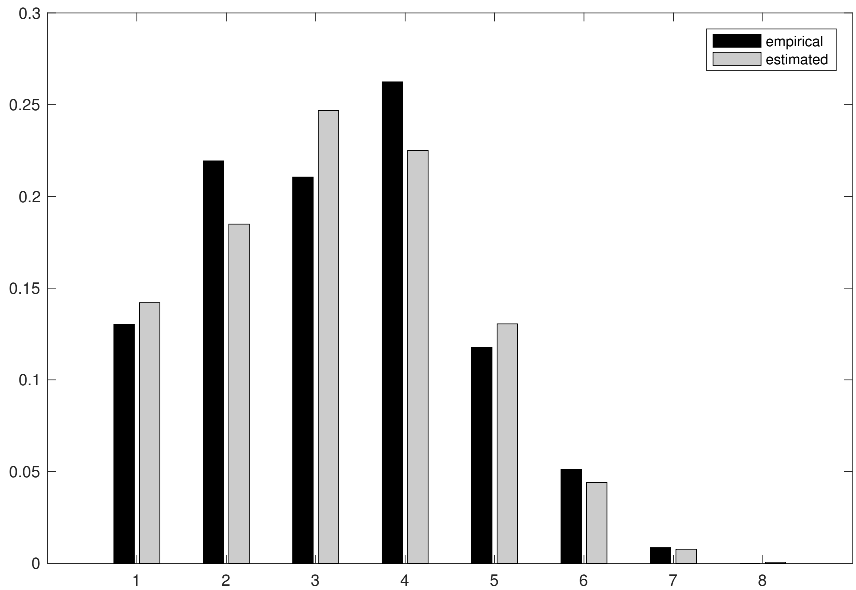

Figure 15: Data and Predicted Grade Bins in HSB

B.2 Transfer Distribution

Since we have the marginal distribution of $AFQT_{HSB}$, we now need to construct the unconditional distribution of transfers k_0. The empirical distribution of transfers has a large mass near zero. Therefore, we assume the marginal distribution of transfers follows a gamma distribution. We compute the distribution of unconditional transfers from parents to high school graduate children, then normalize by the average transfer. Denote $\widetilde{k}_0$ as the transfer normalized by the mean. The shape and scale parameters of the gamma distribution are chosen to minimize the sum of squared errors between that empirical c.d.f. and that of the estimated gamma distribution. The best fit parameters of the gamma distribution are a shape parameter of 0.24 and a scale parameter of 4.44. Figure 16 plots both together, and shows that the estimated distribution matches the data well.

B.3 Joint Distribution of Ability and Transfers

The last step is to compute the joint distribution of the two marginals we created above. We use a Frank copula to combine these two marginal distribution into the joint distribution required by the model. The Frank copula takes the form

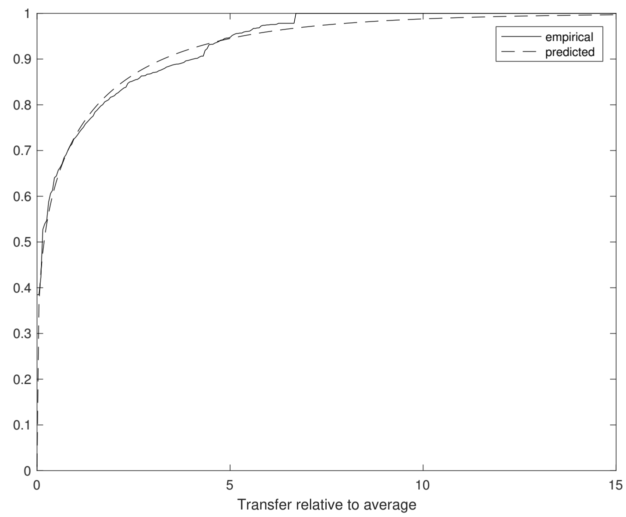

Figure 16: Cumulative Distribution Function for Transfers in HSB

$$C(u,v) = \frac{-1}{\rho} \log \left(1 + \frac{(\exp(-\rho u) - 1)(\exp(-\rho v) - 1)}{\exp(-\rho) - 1} \right)$$

where ρ governs the dependence of draws. Our joint distribution of α and k can therefore be written as

$$H(\alpha, k) = C\left[F(\alpha), G(\widetilde{k}_0)\right]$$

where F and G are the cumulative distribution functions of the normal and gamma respectively. We are therefore left to calibrate ρ, which roughly implies a positively correlation between the two series when $\rho > 0$. Note, however, that while our above procedure gives the marginal distribution of AFQT in HSB, it does not provide individual-level estimates of AFQT. Moreover, we only have grade bins, not actual grade realizations. We therefore proceed as follows. We first assign each individual a random grade realization that is consistent with (1) the bin of their grades and (2) the estimated lognormal distribution of underlying grades. Conditional on that realization, we then use the implied relationship between grades and AFQT to draw an individual realization of AFQT. We then compute the Kendall rank coefficient. We repeat this simulation 1000 times, and the average Kendall rank coefficient is 0.42, implying that high school graduates with higher ability on average have higher initial asset holdings. This implies a copula parameter of $\rho = 4.46$.

B.4 Growth in Transfers over Time

Since the HSB data is for one cohort of high school graduates, we lastly need to make an assumption on the growth of transfers over time. We fix the parameters governing the joint distribution constructed above, then allow the actual dollar amount of transfers to evolve according to

$$
\begin{aligned}
k_{0it} &= \left(\frac{\overline{k}_{0,1962}}{y_{1962}}\right) y_t \widetilde{k}_{0i} \\
&= \left(\frac{2225}{34473}\right) y_t \widetilde{k}_{0i} \\
&= 0.065 y_t \widetilde{k}_{0i}
\end{aligned}
$$

where y_t is average disposable income in year t and $\overline{k}_{0,1962}$ is the average transfer in the HSB data. That is, we assume that the average transfer is always equal to 6.5 percent of average disposable income (as in HSB), and therefore scales proportionally with average income over time.

Table 6: Summary of Model Notation

Symbol	Definition
N_{mt}	Number of males born into model (i.e., graduating high school) each year
N_{ft}	Number of females born into model (i.e., graduating high school) each year
a	Age of individual, where $a = 1, 2, ..., T$
s	Sex of individual, where $s \in \{f, m\}$
k_0	Initial asset endowment
α	Ability endowment
$\pi(\alpha)$	Annual probability of passing college, given ability α
θ	Signal of true ability, where $\theta = \alpha + \varepsilon$
ε	Error term on signal of true ability
σ_ε	Standard deviation of ability signals
ν	Vector of variables that are informative about true ability, where $\nu = (k_0, \theta)$
C	Number of years required to graduate from college
e	Years of education completed by individual, where $e \in \{0, 1, ..., C\}$
λ_t	Annual cost of college in year t
γ	Individuals may not borrow more than a fraction γ of expected discounted future earnings
$\bar{k}$	Minimum asset level for individual, given age, sex, education, and γ