Misallocation and the Early Career Consequences of Student Debt*

Titan Alon Natalie Cox Minki Kim

June 2026

Abstract

In the presence of credit frictions, student debt may prevent graduates from realizing the full returns to a college education by distorting their occupation choice and subsequent early career investments in human capital. This paper quantifies the aggregate size of these labor market distortions by computing the effect of large scale student debt forgiveness policies. The model's predictions are disciplined by new empirical evidence showing that more student debt leads to higher initial earnings, but lower returns-to-experience. The quantitative results suggest that rising student debt is having a substantial adverse effect on aggregate labor productivity and the occupational composition of employment.

Keywords: Student debt, occupation choice, wage profiles, credit constraints, misallocation of talent, college, higher education, labor productivity.

---

*We thank Nir Jaimovich, Alexander Ludwig, Simon Mongey, and seminar participants at Stanford, Princeton, Harvard, MIT Sloan, NYU, UCSD, USC, Dartmouth, Wharton, Tilburg, NBER SI, Keio, Chonnam National University, Frankfurt-Mannheim Macro Workshop, Vienna Macro Cafe, St.Louis Fed, and the Minneapolis Fed for their valuable comments and suggestions. We thank Marcelo Ferreira, Cristoforo Pizzimenti, and Clemens Lehner for outstanding research assistance. This research was conducted with restricted access to Bureau of Labor Statistics (BLS) data. The views expressed here do not necessarily reflect the views of the BLS. Support by the German Research Foundation (DFG) through CRC TR 224 (Project A03) is gratefully acknowledged. Alon: Department of Economics, University of California San Diego (email: talon[at]ucsd.edu), Cox: Bendheim Center for Finance, Princeton University (email: nmcox[at]princeton.edu), Kim: University of Mannheim (email: minki.kim[at]uni-mannheim.de).

1 Introduction

Total outstanding student loan debt reached $1.57 *trillion* in 2022, surpassing auto loans and credit card debt to become the second largest household financial liability after home mortgages.[1] On the one hand, the increase in student debt represents a policy success in which subsidized federal loans alleviated credit frictions to help broaden access to higher education. On the other hand, the rise in student debt may prevent many indebted students from fully realizing the benefits of a college education by exacerbating subsequent credit constraints after graduation. Surveys of *non-delinquent* borrowers from federal loan programs suggest such concerns are well founded, with 34% of borrowers reporting their student loans resulted in more hardship than anticipated; 54% reporting they would borrow less if they could repeat college; and nearly one-fifth reporting "significantly changing career plans because of student loan burdens."[2]

This paper provides a theoretical and empirical analysis of the aggregate welfare and productivity consequences of rising student debt. It focuses in particular on aggregate labor markets and the early career outcomes of college graduates. In the presence of credit frictions, economic theory predicts that student debt may distort household occupation choice and inhibit subsequent investments in human capital accumulation on the job. Exploiting exogenous variation in the composition of college funding, the paper provides empirical evidence consistent with these predictions. The results show that students graduating with more student debt have higher initial earnings, but lower returns to experience in the first ten years of their career. Employing a structural model calibrated to these empirical findings, the paper quantifies the scale of aggregate distortions by computing the impact of large scale student debt forgiveness policies.

While the paper's results apply more broadly to the interaction of credit market frictions and labor markets, there are several advantages to focusing in particular on the impact of student debt. First, student loan debt is among the largest and fastest growing forms of household credit. Second, it is primarily incurred early in life, before individual labor market experiences diverge, making its effects on lifecycle outcomes easier to measure and isolate. Figure 1 shows that student debt accounts for virtually all debt held by young college graduates and a substantial portion of the variation in their total net worth. Third, student debt is largely non-dischargeable, allowing the model to abstract from strategic delinquency and bankruptcy considerations. Empirically, households are

[1]See Federal Reserve Bank of New York (2022) for details on the composition of household debt.
[2]See Baum and O'Malley (2003) for details on the complete questionnaire and survey methodology.

Figure 1: Net Worth and Student Debt of Young Workers, Ages 22-25 with BA Degree

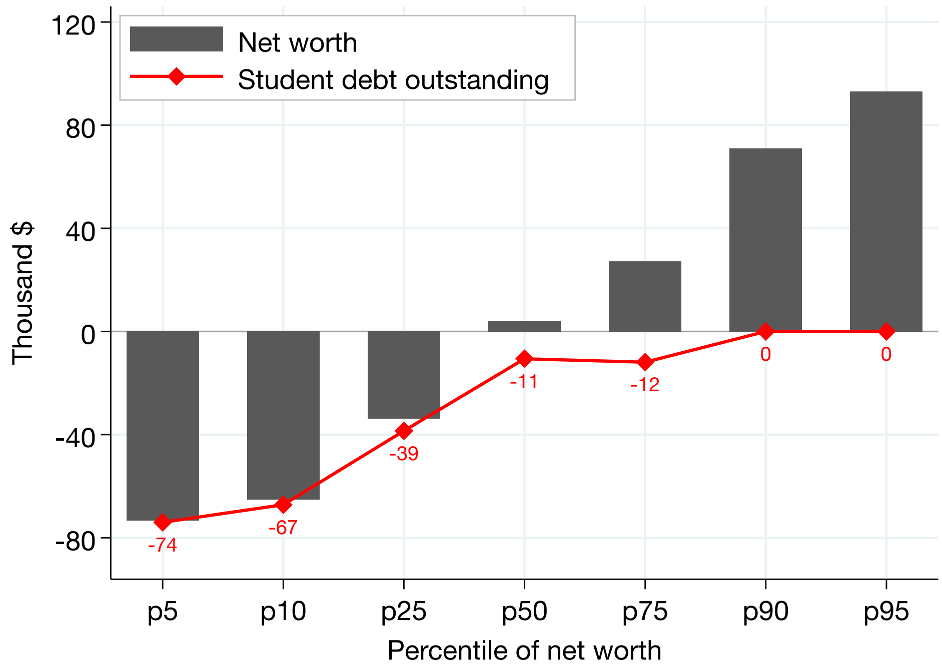

Notes: Source data from 2016 Survey of Consumer Finances. Student debt defined as the total market value of aggregate loan balance of education-purpose expenses. The sample is restricted to bachelor's degree holders between the ages of 22 and 25. Appendix Figure A1 reports the distribution of assets and student debt for the whole population between the ages of 22 and 25, regardless of educational attainment.

much less likely to become seriously delinquent (90-days or more) on their student debt than on other forms of credit, such as credit card debt or auto loans.[3] Finally, the federal government plays an out-sized role in student debt markets, accounting for 92.7% of the total outstanding student loan debt.[4] Changes in loan policies can therefore generate aggregate variation in household debt, providing a natural laboratory to study the causal effects on labor market outcomes. The dominant role of the federal government also means that there may be broad scope for policy improvements to deliver substantial aggregate welfare and productivity gains.

[3]In 2022Q3, only 1.04% of outstanding student debt became seriously delinquent, the lowest rate among all major categories of household credit except mortgages and home equity. See Federal Reserve Bank of New York (2022).

[4]Student loan debt held by the federal government is composed overwhelmingly of Direct Loans, which account for $1.4 trillion of the total. The remaining balance is made up mostly of Title IV loans issued through the FFEL Program and federal Perkins loans. See Consumer Financial Protection Bureau (2022) for further details.

Employing panel microdata from the National Longitudinal Survey of Youth (NLSY97), the paper provides empirical evidence on the effect of student debt on early career labor market outcomes. It instruments student debt levels using variation in the share of grant funding within a college and across cohorts. Constructing the instrument requires accessing restricted-use NLSY97 identifiers for the participants' educational institutions to merge in annual data on college-level loans and grants from the National Center for Education Statistics (NCES). The empirical results show that an additional $\$1000$ of student debt increases initial earnings by 1.50%, but *reduces* the returns to experience by 0.41 percentage points. The effect on initial earnings corresponds to an additional $\$342$ in annual earnings for every $\$1000$ in student debt, in line with similar estimates in the literature.[5] The effects of student debt on the returns to experience are statistically significant and sizable given that average annual earnings growth of individuals between the ages of 25 to 30 is estimated to be 7.75% (Guvenen, Karahan, Ozkan, and Song 2021).

To draw aggregate inference from the empirical evidence, the paper develops a dynamic model of lifecycle human capital accumulation and occupation choice in the presence of credit market frictions. When credit constraints bind, household discounting of future income streams is greater than prevailing market interest rates. The resulting intertemporal distortions lead households to decrease investments in human capital accumulation as an alternative form of consumption smoothing, reducing lifetime earnings and aggregate labor productivity. These adjustments can occur not only through reduced investment on the job, but also through changes in occupation choice. Indebted households disproportionately select into occupations with more front-loaded compensation schemes, even when their abilities make them a better match for other types of work. The result is a misallocation of talent that can compound the negative effects of student loan debt on aggregate labor productivity.

To study their macroeconomic implications, these mechanisms are embedded into a quantitative heterogeneous-agent, incomplete markets model that can be taken to the data. Individuals are born with heterogeneous family assets and occupation-specific abilities. They endogenously incur student debt when deciding whether or not to attend college, accounting for potential selection effects in the data. After graduation, households choose an occupation based on their innate abilities and financial assets, taking occupation wages and non-wage amenities as given. Earnings evolve endoge-

[5] For instance, see Rothstein and Rouse (2011), Chapman (2015), and Luo and Mongey (2024) for comparable estimates of the impact of student debt on initial earnings.

nously over their lifecycle as a consequence of costly investments in human capital and idiosyncratic labor market shocks. Households are also subject to progressive taxation, receive retirement benefits which depend on their earnings, and face realistic student debt repayment provisions and borrowing constraints. The calibration matches the aggregate earnings profile; the joint distribution of household assets, student debt, and college matriculation; and the employment shares, initial earnings, and returns to experience across 22 occupation groups. It also replicates, through indirect inference, the estimated marginal effects of student debt on an individual's lifecycle earnings.

Using the calibrated model, the paper quantifies the aggregate labor market distortions which result from student debt. To do so, it computes the impact of large scale student loan forgiveness policies on household welfare, earnings, college matriculation, labor productivity, and the occupational composition of employment. The first exercise analyzes the effect of a one-off federal student loan forgiveness program, similar to those being proposed by the current administration. The second evaluates the long-run consequences of making the proposed programs permanent. Conceptually, the second exercise is more an accounting exercise than a policy counterfactual, providing useful information about the aggregate size of labor market distortions from student debt.

The results show that both short-run and long-run policies increase labor productivity by stimulating human capital investment and improving the allocation of workers across occupations. The model predicts that recently proposed one-off federal student loan forgiveness policies could increase household lifetime earnings by up to 0.24% and raise aggregate labor productivity by up to 0.28% in the short-run. The policy also induces a moderate amount of occupational reallocation, with workers switching predominantly from high wage to high amenity occupations following debt relief, consistent with the predictions of Boar and Lashkari (2025) and Luo and Mongey (2024).

To quantify the potential size of aggregate labor market distortions and the misallocation of talent, the long-run exercise computes the impact of policies which permanently reduce the incidence of student debt among graduates.[6] Given the scope of these interventions, the computations allow occupational wages and college matriculation decisions to endogenously respond to the policy. The results suggest sizable macroeconomic effects. Average lifetime earnings increase by up to 6.24% under the long-run policy, driven by gains in labor productivity. The larger long-run productivity effect is due in

[6] Real world examples of such policies would include expanded education grants, reductions in federal student loan interest, and public tuition subsidies.

part to a 1.29 percentage point rise in the college matriculation rate that was absent in the short-run. It is also the result of changes in the scope and direction of occupational reallocation. Under the long-run policy, up to 5.63% of workers switch occupations. These reallocated workers flow predominantly to occupations with higher returns-to-experience, rather than amenities. These differences are due to the equilibrium response of occupational wages in the long-run, which induce second-order job reallocations and discourage flows toward high amenity occupations relative to the short-run policy. Consequently, long-run policies deliver larger gains in labor productivity through larger reductions in the aggregate misallocation of talent.

Related Literature. The findings contribute to the literature examining how credit market frictions affect labor market outcomes. Recent contributions have shown that access to consumer credit can affect household job search behavior with aggregate implications for the efficiency of worker sorting and business cycle volatility (Herkenhoff, Phillips, and Cohen-Cole 2016; Herkenhoff 2019). This paper focuses in particular on student debt and its effect on the early career outcomes of college graduates. As a result, it complements research on credit frictions in the financing of higher education (Andolfatto and Gervais 2006; Lochner and Monge-Naranjo 2012; Caucutt and Lochner 2020; Lochner, Stinebrickner, and Suleymanoglu 2021; Krueger, Ludwig, and Popova 2025) by investigating how student debt can affect post-graduation labor market outcomes. The mechanism may also help explain the substantial cross-country variation in the returns to college and the shape of lifecycle earnings profiles (Lagakos, Moll, Porzio, Qian, and Schoellman 2018; Martellini, Schoellman, and Sockin 2024).

Empirically, the paper presents new evidence on how student debt affects household earnings profiles and occupation choice. It contributes to a growing literature documenting the effect of student debt on lifecycle outcomes, such as entrepreneurship, homeownership, marriage, fertility, and graduate school (Rothstein and Rouse 2011; Zhang 2013; Goodman, Isen, and Yannelis 2018; Chakrabarti et al. 2022; Folch and Mazzone 2026). In particular, it provides evidence of how student debt affects lifecycle earnings through distortions to occupation choice (Herkenhoff, Phillips, and Cohen-Cole 2021; Boar and Lashkari 2025; Morazzoni 2023; Luo and Mongey 2024). The findings show that indebted graduates are more likely to select into occupations with front-loaded compensation schemes, mirroring the results in Hampole (2024) which show similar sorting patterns across college majors.[7] More broadly, the paper contributes to the literature on the misallocation of occupational talent by studying the role of debt, complementing earlier work focused on the role of taxation or race and gender based discrimination (Lockwood, Nathanson, and Weyl 2017; Hsieh et al. 2019).

Methodologically, the paper follows a common approach in the literature that analyzes the aggregate consequences of student debt using dynamic stochastic heterogeneous-agent lifecycle models with incomplete markets (Ionescu 2009; Abbott, Gallipoli, Meghir, and Violante 2019; Lee and Seshadri 2019, Fu, Lin, and Tanaka 2021). The model here builds in particular on the seminal approach in Huggett, Ventura, and Yaron (2011). It extends their framework along a number of dimensions by including occupational heterogeneity, multidimensional ability, student debt and college matriculation alongside a broader calibration strategy. To contextualize the effect of these extensions, the paper recomputes the main exercises in Huggett, Ventura, and Yaron (2011) on the sources of lifetime inequality to show how the model's predictions change. A comparison of the results shows a much more prominent role for initial wealth differences. While their original study finds that differences in human capital at age 16 account for the majority of lifetime inequality, the approach here finds that heterogeneity in initial assets is the dominant factor driving lifecycle inequality in earnings and welfare.

Finally, while the paper encompasses several mechanisms through which credit frictions can distort human capital investments, one important channel is absent: the role of parents. In particular, the model does not consider how parental savings for their children's education may respond to changes in education policies. This is relevant as several studies, such as Abbott et al. (2019), document a substantial crowding out effect. Even with a crowding out effect, much of the literature continues to find that education subsidies have a substantial positive effect on the human capital investments of high ability students from low income families (Abbott et al. 2019; Lochner, Stinebrickner, and Suleymanoglu 2021; Krueger, Ludwig, and Popova 2025). Hence, while our findings are consistent with the broader findings of this literature, the quantitative results here should be qualified and interpreted as net of any endogenous changes in parental savings that may feed back into student matriculation and investment decisions.

[7] For context outside the U.S., De Falco and Reichlin (2025) show similar sorting patterns among college students in Chile.

2 Model Environment

The economy is populated by a unit mass of forward-looking, heterogeneous households who make college matriculation, occupation choice, and on-the-job human capital investments subject to credit constraints and idiosyncratic labor market risks. Each period corresponds to five years. Agents begin life at age 18 when they are endowed with initial assets a and realize their occupation specific abilities $\Theta = \{\theta_0, \theta_1, \ldots, \theta_K\}$.

After observing their assets and abilities, individuals decide whether or not to attend college subject to a matriculation taste shock ζ capturing, among other things, the opportunity cost of a college education. Those who attend college will endogenously incur student debt d if they do not have sufficient funds to cover their education costs. College graduates enter the labor market at age 23 and choose a college occupation $k \in \{1, 2, \ldots, K\}$ to maximize their expected lifetime utility

$$k^* = \text{argmax} \; \{ \; V_1, V_2, ..., V_K \; \}$$

where V_k is the value function associated with occupation k. Those choosing not to attend college can only work in high school occupations $k = 0$. All households participate in the labor force for at least 40 years, retire by age 64, and pass away by age 78.

Households can be identified by the state vector (a, h, d, k, t, Θ) which summarizes their assets a, human capital h, student debt d, occupation k, age t, and innate abilities Θ. The decision problem of a working-age individual in occupation k can be expressed

$$V_k(a, h, d, t, \Theta) \; = \; \max_{c,\, s,\, a'} \; \left\{ \frac{c^{1-\rho} \, - \, 1}{1 - \rho} \; + \; \beta \mathbb{E} \left[V_k(a', h', d', t+1, \Theta) \right] \right\} \times \nu_k$$

subject to

$$c + a' = \mathbb{T} \left(w_k (1 - s) h \right) + (1 + r) a - \phi(a, h, d, t)$$

$$h' = exp(z') \mathrm{H}(s, h, \Theta)$$

$$d' = (1 + r_d) d - \phi(a, h, d, t)$$

$$a' \; \geq \; -\bar{a} \;\; , \;\; 0 \; \leq \; s \; \leq \; 1$$

where $\bar{a}$ is the exogenously given borrowing constraint and ν_k is a fixed occupation-specific non-wage amenity common to all workers.[8] The non-wage amenities account for factors–other than potential earnings–which influence household occupation choice. Recent work by Boar and Lashkari (2025) and Luo and Mongey (2024) has demonstrated that household substitution along this amenities margin is a relevant channel through which household liabilities affect their occupation choice.

As in Huggett, Ventura, and Yaron (2011) and Lee and Seshadri (2019), human capital accumulation is risky and modeled by a Ben-Porath (1967) learning technology subject to i.i.d. idiosyncratic log-normal shocks, z', so

$$\mathrm{H}(s, h, \Theta) = \theta^{1-\alpha}(sh)^{\alpha} + (1-\delta)h$$

where h is the individual's current human capital, θ is their occupation specific ability, and s is the share of time invested in skill acquisition on the job.[9] Parameter α summarizes the investment elasticity of the learning technology and δ captures the depreciation of skills. The functions $\mathbb{T}\left(w(1-s)h\right)$ and $\phi(a, h, d, t)$ represent the tax system and student debt repayment rule, respectively, which are described below.

Student Debt and College Matriculation. Households endogenously incur student debt when deciding whether or not to attend college. At age 18, every individual receives a college admission offer that allows them to enroll in college if they take on some amount of student debt $d \geq 0$. Admission offers can be freely accepted or rejected. Those who accept the offer take on student debt d and begin work in their preferred college occupation, k^*. Those who do not attend college enter the labor force directly to work in the high school sector, $k = 0$. High school graduates solve the same life-cycle problem as college graduates, except that they make no occupation choice and have no student debt. The high school graduate value function is therefore given by $V_0(a, h, 0, t, \Theta)$. Formally, the college matriculation problem can be expressed

$$\max \left\{ V_{k^*}(a, h, d, 1, \Theta)\ ,\ V_0(a, h, 0, 0, \Theta) + \zeta \right\}$$

[8]The exogenous occupation-specific amenities ν, common to all households within that occupation, do not play a meaningful role in the theory but are necessary quantitatively to match the patterns of occupation choice conditional on observed earnings in the exercises below. See Section 5 for details.

[9]The choice of s can also be thought of as the agent's chosen *career path* within an occupation, whether targeting advancement or maximizing current income.

where ζ is a college matriculation taste shock which accounts for unobserved idiosyncratic heterogeneity in college enrollment decisions not captured by the model.[10] It captures, among other things, the opportunity cost of attending college and the amenity value of higher education, both of which affect student matriculation decisions and may be shaped by family background.

To account for the variety of circumstances determining an individual's access to student financial-aid and family assistance, the model allows college admission offers to depend stochastically on household characteristics. Specifically, admission offers take the form $\{x, \tau\}$, which require students to take on debt $d = (1 - x)\tau$ to enroll. The Bernoulli random variable x controls the extensive margin of student debt, while $\tau \in \mathbb{R}^+$ determines the intensive margin of student loan sizes. In particular, the intensive margins of the college admission offer a household receives will depend on their initial assets, a_0.[11]

Since households only incur student loans when enrolling in college, the model's initial distribution of student debt depends on both the exogenous stochastic process generating offers as well as the endogenous choice to enroll in college. Formally, the initial distribution of student debt levels at age 23 is therefore defined implicitly by,

$$d(a, h, 0, 1, \Theta) \ = \begin{cases} (1 - x)\,\tau & \text{if} \quad V_{k^*}(a, h, d, 1, \Theta) \ \geq \ V_0(a, h, 0, 0, \Theta) + \zeta \\ 0 & \text{if} \quad V_{k^*}(a, h, d, 1, \Theta) \ < \ V_0(a, h, 0, 0, \Theta) + \zeta \end{cases}$$

Consequently, high school graduates have no student loans, nor does a mass of college graduates who received and accepted a debt-free college education.[12] The remaining population of college graduates has student loans which depend on realizations of τ.

The endogeneity of student debt means that college matriculation decisions will determine not only who attends college, but also which households end up taking on student loans, and how much. As a result, the initial distribution of student debt in the population will depend jointly on all the other household state variables which affect

[10]The matriculation taste shock helps the model match enrollment patterns by family background and improves computational tractability by smoothing the household value functions.

[11]Separately modeling the extensive and intensive margins of student debt realizations is necessary to match the large mass point of graduates without any debt observed in the empirical distribution of realized student debt. See Section 5 for details on the stochastic process generating admission offers.

[12]Note that some households may optimally choose not to attend college even though they received an offer allowing them to do so for free, e.g. $x = 1$. Furthermore, extending the model to allow for college drop-outs could generate a population of high school graduates with student debt. For a more complete theoretical and quantitative treatment of the implications of drop-out risk for the returns to college, see Hendricks and Leukhina (2017) and Hendricks and Leukhina (2018).

college matriculation. For instance, higher college ability and wealthier households are more likely to accept a given admission offer $\{x, \tau\}$ and enroll in college than are lower ability and less wealthy households. This interdependence helps match the patterns of selection into college by family background which are observed in the data and has theoretical and quantitative implications for the aggregate impact of policies which change the provision of college financial aid analyzed in the counterfactuals.

Student Debt Repayment. Households which borrow to attend college begin paying off their student debt after graduation. As in Luo and Mongey (2024), the benchmark repayment rule is modeled on the standard federal repayment plan whose provisions require fixed periodic payments that amortize the student loan over 10 years. This plan remains the most common provision among student loan borrowers, despite the rising popularity of income based repayment programs in recent years. In normal circumstances, an individual with outstanding student debt d would have to repay

$$\rho(d,t) = \left[\frac{r_d}{1-(1+r_d)^{-(\bar{T}-t+1)}}\right] d$$

in each period $t < \bar{T}$, where $\bar{T}$ is the repayment period and r_d is the student loan interest rate.[13]

Due to the stochastic risk in human capital accumulation, it is possible that some households will find themselves unable to make their student loan repayments $\rho(d,t)$. Households with few saved assets are particularly vulnerable to being unable to repay after sufficiently negative shocks to their human capital. Consistent with the "undue hardship" provisions of the standard federal repayment plan, households who find themselves in this circumstance can decrease or delay the size of their repayment obligations. Specifically, the student debt repayment rule is given by

$$\phi(a,h,d,t) = \min\left\{ \rho(d,t) \ , \ a + w(1-s)h - \bar{c} \right\}$$

In other words, households which cannot make their student loan payment receive a consumption floor $\bar{c}$ and dedicate their remaining assets and income to pay off their student debt. In this case, households will be delinquent and their repayment shortfall will be carried forward (with interest r_d) to their next period student debt balance, d'.

[13] The fact that obligations do not count toward the exogenous borrowing constraint $\bar{a}$ motivates why households do not prepay student loans.

All outstanding student debt must be repaid and no borrower is permitted to defer repayments beyond maturity duration $\bar{T}$. The model therefore allows individuals to become delinquent, while default remains a rare equilibrium outcome – consistent with legal provisions that strictly limit the discharge of student loan obligations.

The Tax System. $\mathbb{T}(y)$ is a function which represents the prevailing tax system and transforms gross household income y into after-tax income. In particular, $\mathbb{T}(y)$ takes the form of a step function

$$\mathbb{T}(y) = (1 - \psi(y)) \cdot y$$

where $\psi(y)$ represents the effective marginal tax rates for the tax bracket of individuals with income y. The brackets and marginal rates $\psi(y)$ are chosen to match the effective tax rates estimated by the Congressional Budget Office (CBO) displayed in Figure A2. Accurately modeling effective marginal tax rates is relevant since they influence the household's incentive to attend college and invest in human capital over their lifecycle (Saez, Slemrod, and Giertz 2012; Jones 2019). Accounting for these effects is also quantitatively important when assessing the impact of policies which change the net cost of human capital investment, such as the provisions of student loan programs.

Retirement and Retirement Benefits. Households retire deterministically by age 64 and continue to make consumption and savings decisions until they pass away. Retired households fund consumption out of their savings a and retirement benefits e, which they begin receiving after retiring. Following Huggett, Ventura, and Yaron (2011), retired households solve

$$V_R(a,t) = \max_{c,a'} u(c) \, + \, \beta V_R(a', t+1)$$

subject to

$$c + a' = e + (1+r)a$$
$$a' \geq -\bar{a}$$

where e is a social security benefit paid to all retired households, with the amount of payment depending on the last-period income to match the average annual social security transfer observed in the United States.

The retirement stage determines the model's terminal state which closes the household problem and influences the evolution of their lifecycle investment behavior. The fact that retired individuals do not work and instead rely on savings to finance consumption shapes household incentives to accumulate physical versus human capital as they age. The balance of these incentives determines the composition of household investments across the demographic distribution. This distribution has implications for the impact of policies which change the relative returns of human versus physical capital, both in aggregate and across cohorts.

Production Technologies and Firms. Production occurs in a competitive sector of firms which operate constant returns-to-scale technologies that employ both skilled and unskilled labor. As in Jones (2014), the production technology takes a nested CES structure

$$Y = \left[A_{hs} H_{hs}^{\frac{\sigma-1}{\sigma}} + A_c H_c^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$$

where Y is output, A are group specific productivity terms, and σ is the elasticity of substitution between "skilled" college labor and "unskilled" high school labor. The input of college labor depends on the imperfectly substitutable variety of occupational specializations k chosen by graduates so that

$$H_c = \left[\sum_{k=1}^{K} A_k H_k^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}$$

where A_k are group specific productivity terms and η is the elasticity of substitution between college graduate specializations. The effective labor input for each college occupation $k > 0$ depends on the total effective hours worked by individuals in that occupation, which incorporates both their physical hours and human capital,

$$H_k = \int (1 - s(a, h, d, t, \Theta)) h \, d\mathcal{G}_k(a, h, d, t, \Theta)$$

where $\mathcal{G}_k$ is the joint marginal cumulative distribution function of working age households (a, h, d, t, Θ) who optimally select into college occupation $k \in \{1, \dots, K\}$. The high school labor input H_{hs} (i.e., $k = 0$) is defined analogously to the other H_k, integrating over the distribution of workers who do not attend college. Modeling the production sector is necessary for the endogenous determination of the college wage premium and occupation specific wages that influence college matriculation decisions and the sorting of workers across occupations. Accounting for these effects is important when assessing the impact of credit frictions on human capital accumulation and the ultimate allocation of talent across education and occupation groups, particularly under policies which change relative wages across occupations.

3 Implications for Human Capital and the Allocation of Talent

This section presents a simplified version of the model to highlight its key mechanisms. When credit constraints bind, households invest less in human capital accumulation as an alternative mode of consumption smoothing. Section 3.1 describes the resulting intertemporal distortions to college matriculation and investments in human capital accumulation on the job. Section 3.2 shows how credit frictions may also distort occupation choice, pushing households towards occupations with more front-loaded compensation schemes over those which best match their innate abilities.

Consider a simplified lifecycle where all individuals live for only two periods – young (y) and old (o) – and there are only two occupations. Households are born young with initial assets a_y and innate abilities $\Theta = \{\theta_1, \theta_2\}$.[14] The decision problem of a young worker who chose occupation k can be expressed

$$V_k(a_y, \Theta) = \max_{c_y, c_o, a_o, s} u(c_y) + \beta u(c_o)$$

subject to

$$\begin{aligned} c_y &= w_k(1-s) + a_y - a_o \\ c_o &= w_k H(s,\Theta) + (1+r)a_o \\ a_o &\geq -\bar{a} \ , \ 0 \leq s \leq 1 \end{aligned}$$

where initial human capital is normalized to unity, $h_y = 1$. The household's human capital in old age depends on the investment s they made during their youth and their innate abilities, as determined by the simplified Ben-Porath technology $H(s,\Theta) = \theta_k^{1-\alpha} s^\alpha$ with

[14] Compared with the quantitative model, the simplified model abstracts from the role of the stochastic shocks, taxes, retirement, occupation specific amenities, the particular structure of student debt markets, and general equilibrium price effects.

full depreciation. Households choose their occupation at the beginning of their youth in order to maximize discounted lifetime utility,

$$k^*(a_y, \Theta) \;=\; \text{argmax}\; \{\; V_1(a_y, \Theta)\;,\; V_2(a_y, \Theta)\; \}$$

3.1 Intertemporal Distortions to Human Capital Accumulation

In the absence of credit constraints, households optimally invest in human capital accumulation until the marginal return on investment equals the return on physical capital,

$$\frac{\partial H(s^*, \Theta)}{\partial s} \;=\; 1 + r$$

where s^* is the optimal human capital investment in the absence of credit frictions. The expression is an example of the classic arbitrage condition derived from the first order conditions for household human (s) and physical (a_o) capital investment. The condition clarifies that in the absence of credit frictions, household investment in human capital depends only on their abilities Θ and not their initial assets a_y. To see this directly, substitute in the education technology to solve for the optimal human capital investment

$$s^*(\theta) \;=\; \left[\frac{\alpha}{1+r}\right]^{\frac{1}{1-\alpha}} \theta$$

which is determined only by the household's occupation-specific ability, θ.

When credit constraints bind, households discount future income streams at a shadow rate that is greater than the market interest rate $1+r$. Unable to borrow in financial markets, households instead invest less in human capital as an alternative form of consumption smoothing. Formally, when the borrowing constraint binds, households optimally invest in human capital until the marginal return equals the shadow rate of borrowing, $1 + r^c$, so that

$$\frac{\partial H(s^c, \Theta)}{\partial s} \;>\; 1 + r$$

where $s^c(a_y, \Theta)$ is the optimal investment for a constrained household.[15] Unlike the unconstrained case, the investment policy $s^c(a_y, \Theta)$ now depends on both individual abilities and initial assets. This is because initial assets partly determine the shadow

[15] The household shadow rate is given by the implied interest at which the household consumption-savings profile would lie on the Euler Equation. For constrained households, $1 + r^c \equiv \frac{u'(c_y)}{\beta u'(c_o)} > 1 + r$. As a result, the shadow rate varies across households and across occupations for a given household.

interest rate a household faces when borrowing constraints bind. Moreover, since the education technology is concave in its inputs, it is straightforward to show that this leads to a reduction in the overall investment in human capital, so that $s^c < s^*$.

This reduction in human capital investment acts as another mode of consumption smoothing by shifting the household lifecycle income profile to front-load earnings. Specifically, the decrease in human capital investment restores foregone earnings, resulting in higher current income

$$w_k(1 - s^c) \ > \ w_k(1 - s^*)$$

but lower returns to experience

$$\frac{H(s^c, \Theta)}{1 - s^c} \ < \ \frac{H(s^*, \Theta)}{1 - s^*}$$

In other words, reduced investment in human capital leads to a larger share of lifetime earnings to be realized during youth, even when borrowing constraints are binding. However, this method of consumption smoothing is much costlier than borrowing through financial markets (e.g. $r^c > r$) and results in foregone human capital accumulation that lowers labor productivity and overall lifetime earnings.

3.2 The Misallocation of Talent

Credit frictions may also inhibit human capital accumulation by distorting household occupation choice. When the ability to borrow against future income is limited, households may switch away from occupations with higher lifetime earnings in favor of those with more front-loaded compensation schemes. In other words, constrained households will favor high wage w occupations, which offer the greatest opportunity to maximize current income, over high ability θ ones providing the greatest opportunity for future income growth. The result is a *misallocation of talent* whereby credit constraints lead some workers to select into occupations that are not optimally matched to their abilities. The more constraints bind, the less the sorting of workers reflects comparative advantages in ability Θ, and the greater the loss in aggregate labor productivity and lifetime earnings.

For intuition, consider the household occupation choice problem when both occupations offer the same wage, $w_1 = w_2 = w$. In this case, households optimally sort into the occupation corresponding to their highest ability, so $k^*(\Theta) = \text{argmax } \{\theta_1, \theta_2\}$. This is true for both borrowing constrained and non-constrained households since, in the absence of occupational wage dispersion, there is no margin outside of individual ability θ that individuals can trade-off when moving across occupations. As a result, workers sort across occupations based strictly on their comparative advantage in ability Θ.

When occupation wages differ, unconstrained households will continue to sort based on their comparative advantage while accounting for the difference in market prices. In the absence of credit frictions, unconstrained households perfectly smooth consumption so that lifetime utility is monotonic in the present discounted value of lifetime earnings. As a result, the occupation which maximizes lifetime utility corresponds to the one which generates the greatest lifetime earnings. Given wages w_k and the investment policy $s^*(\theta_k)$, the optimal occupation choice for unconstrained households can be expressed

$$k^*(\Theta) \;=\; \text{argmax} \;\left\{\; w_1\left[1 - s^*(\theta_1) \;+\; \frac{h_o^*(\theta_1)}{1+r}\right] \;,\; w_2\left[1 - s^*(\theta_2) \;+\; \frac{h_o^*(\theta_2)}{1+r}\right] \;\right\}$$

where $h_o^*(\theta_k) \;=\; \left[\frac{\alpha}{1+r}\right]^{\frac{\alpha}{1-\alpha}} \theta_k$ is the human capital in old age for a household which invested optimally during their youth. The expression makes clear that the unconstrained household's occupation choice continues to depend only on their innate comparative advantage in abilities Θ, and not their initial assets a_y.

Constrained households additionally consider the timing with which income is realized over their lifecycle when choosing their occupation. As a result, occupations which yield the highest lifetime utility are not necessarily those which offer the highest lifetime earnings. For instance, constrained households may prefer an occupation with lower lifetime earnings provided income is concentrated earlier in life. The optimal occupation choice for the constrained household can be expressed

$$k^c(a_y,\Theta) \;=\; \text{argmax} \;\left\{\; u\left(c_y^c\left(a_y,\theta_1\right)\right) + \beta u\left(c_o^c\left(a_y,\theta_1\right)\right) \;,\; u\left(c_y^c\left(a_y,\theta_2\right)\right) + \beta u\left(c_o^c\left(a_y,\theta_2\right)\right) \;\right\}$$

where $c_j^c(a_y,\theta_k)$ is the optimal consumption policy of a financially constrained household of age j working in occupation k. The expression shows that, unlike unconstrained households, occupation choice $k^c(a_y,\Theta)$ depends on both household initial assets a_y as well as abilities Θ.

One implication of the different sorting rules is that credit constraints can give rise to a misallocation of talent that leads individuals to switch away from the occupation best suited to their abilities. To see this explicitly, substitute the optimal investment policy $s^*(\theta)$ into the unconstrained sorting rule $k^*(\Theta)$. Given a realization of occupation specific abilities $\Theta = \{\theta_1, \theta_2\}$, the condition reduces to choosing occupation 1 if

$$\theta_1 \; > \; \frac{w_2 - w_1}{\kappa w_1} + \frac{w_2}{w_1} \cdot \theta_2 \tag{1}$$

where $\kappa = \left(\frac{1}{1+r}\right)^{\frac{1}{1-\alpha}} \left[\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}\right] > 0$.[16] Similarly, evaluating the lifetime utility received in each occupation under constrained investment policy $s^c(a_y, \theta)$ shows that the sorting rule in $k^c(a_y, \Theta)$ reduces to choosing occupation 1 if

$$\theta_1 > \left(\frac{w_2 + a_y}{w_1 + a_y}\right)^{\frac{1+\alpha\beta}{(1-\alpha)\beta}} \cdot \frac{w_2}{w_1} \cdot \theta_2 \tag{2}$$

where notation is simplified by assuming logarithmic utility, $u(c) = \log(c)$.

Comparing the sorting rules for constrained and unconstrained households shows that there are some individuals $\Theta = \{\theta_1, \theta_2\}$ who would switch occupations when credit constraints bind. For instance, if occupation 2 offers a higher wage, $w_2 > w_1$, then there will be some workers who choose occupation 1 in the absence of credit constraints, but choose occupation 2 when financially constrained.[17] The constrained sorting rule also demonstrates how the level of initial assets shapes occupation choice within the constrained population: the lower their initial assets a_y, the more likely an individual is to select into the high wage occupation. For instance, when $w_2 > w_1$ a decrease in initial assets a_y will increase the right hand side of the constrained sorting rule, raising the ability threshold for households to select into the low wage occupation 1.

The sorting rules also demonstrate how the macroeconomic consequences for labor productivity and earnings profiles depend on the joint distribution of assets a_y and abilities Θ in the population. In part, this reflects the fact that whether or not a household is financially constrained may itself depend on their occupation choice. Recall that financial constraints bind when a household's desired level of borrowing, $a_o^*(a_y, \Theta)$, surpasses the exogenous borrowing constraints $\bar{a}$, which amounts to

$$\frac{\beta}{1+\beta}\left[a_y + w_k(1 - s^*)\right] - \frac{1}{(1+\beta)(1+r)}\, w_k h_o(s^*, \theta_k) \; < \; -\bar{a}$$

[16] The condition is economically intuitive: it is optimal to choose an occupation offering a lower wage only if one has sufficiently high ability to nevertheless generate greater lifetime earnings.

[17] To see this more directly, note that the occupation sorting rule for constrained households is a line passing through the origin with slope greater than w_2/w_1. This implies that the slope of the sorting rule in Θ space is steeper when individuals are constrained than when they are not constrained.

The expression shows that the extent to which credit constraints bind depends not only on a household's initial assets, a_y, but also on their occupational wage w_k which determines their current income, and on their occupational ability θ_k which determines their future income. This means a household can be financially constrained while working in one occupation, but not in another. The expression also implies that–for every asset level a_y–there exists an occupation-specific ability cutoff $\bar{\theta}_k$ such that all workers with ability $\theta_k > \bar{\theta}_k$ are credit constrained. Plugging in the optimal investment policy, the cutoff can be solved explicitly as

$$\bar{\theta}_k = \frac{\beta(1+r)^{\frac{1}{1-\alpha}}}{\alpha^{\frac{1}{1-\alpha}} + \alpha^{\frac{\alpha}{1-\alpha}}} \left(1 + \frac{a_y}{w_k}\right)$$

The cutoff reflects the fact that higher ability individuals expect faster income growth over their lifecycle, and hence have the greatest desire to borrow for consumption smoothing. It also shows that the lower a household's initial assets, a_y, the lower the ability cutoff $\bar{\theta}_k$ at which they become constrained. The effect is intuitive since a reduction in initial assets will increase a household's desire to borrow against future income, especially when that future income now constitutes a greater share of their lifetime wealth.

Figure 2 illustrates how the joint distribution of assets and abilities determines the aggregate misallocation of talent. Conditional on an initial level of assets a_y, it depicts the population of workers who switch occupations because of credit frictions. The cutoffs $\bar{\theta}_1$ and $\bar{\theta}_2$ are defined as above, indicating the regions where credit constraints bind in each occupation. The lower red border of the shaded region corresponds to the unconstrained occupation sorting rule in (1); the upper black border corresponds to the sorting rule in (2) when credit constraints bind. The shaded region represents the population of misallocated workers who switch from occupation 1 to occupation 2 in the presence of credit constraints. Consistent with the discussion above, misallocation is most prominent among the high ability population. In contrast, the low ability $\theta < \bar{\theta}_k$ population always choose their occupation according to the unconstrained rule, as they do not expect much earnings growth over their lifecycle.[18]

The misallocation region depicted in Figure 2 is conditional on a particular level of initial assets a_y. Moving across the household wealth distribution, both the cutoffs $\bar{\theta}_k$ and the frontier of the constrained sorting rule in (2) will shift, changing the size of the misal-

[18]Since $w_2 > w_1$, there is an additional population whose occupation choices are distorted because they become constrained in occupation 1 $\theta_1 > \bar{\theta}_1$ but not occupation 2, $\theta_2 < \bar{\theta}_2$, and so switch to the latter like other constrained households.

Figure 2: Misallocation of Talent

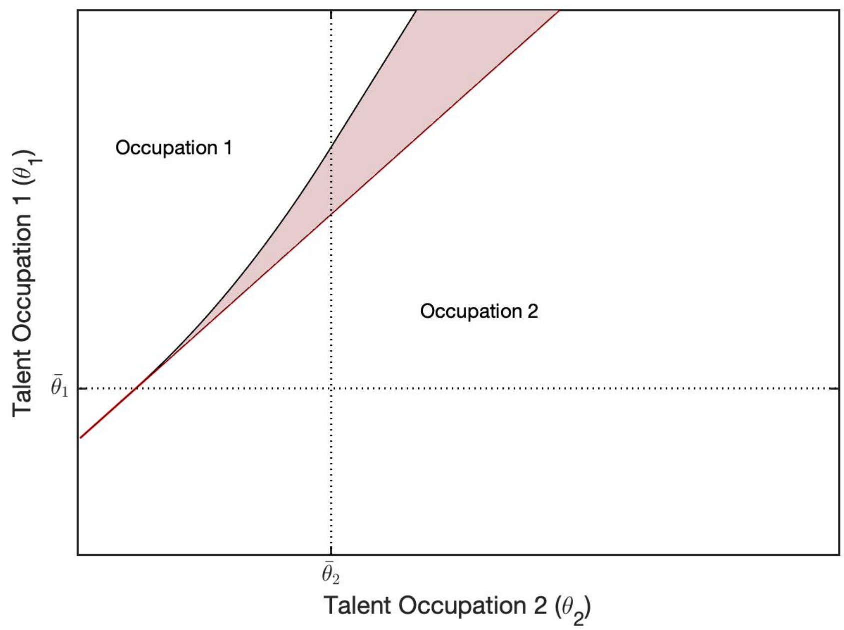

Notes: This figure illustrates how the misallocation of talent depends on an individual's abilities, for a given initial asset a_y. Occupation 2 is assumed to offer higher wages, $w_2 > w_1$.

located population. As discussed above, a reduction in wealth a_y will reduce the ability cutoffs $\bar{\theta}_k$ and shift out the constrained sorting frontier (the black line), both of which expand the region of misallocated workers. Computing the total population of constrained households therefore requires knowing the share of workers that fall into each of these misallocation regions, and then aggregating over the household wealth distribution.

The following sections turn to quantifying the size of the aggregate misallocation of talent in the full model. While the core mechanics are the same, the interactions of human capital accumulation and credit frictions are more detailed and nuanced in the quantitative model. One difference is that the distinction between constrained and unconstrained households is less stark in the full model. With stochastic human capital accumulation, all forward looking households will anticipate hitting the borrowing constraint with some probability. The behavior of households in the full model is therefore best thought of as a weighted average of the two stark types in the simplified model.

4 The Empirical Evidence

This section reviews direct evidence of the model's central mechanism. It provides reduced-form empirical evidence that the trajectory of lifecycle earnings is affected by the presence of student debt. The estimation employs panel data from the National Longitudinal Survey of Youth 1997 (NLSY 1997) and an instrumental variable design to estimate the impact of student debt on both initial earnings and returns to experience after graduation. The results provide identifying power for the model's parameters as they contain information on the marginal effect of varying a household's initial debt on their earnings profile. The calibration strategy in Section 5 combines these estimated marginal effects with information on the joint distribution of assets and abilities in the population to discipline the size and sensitivity of the treatable population which drives macroeconomic outcomes.

4.1 Data Sources

The primary dataset is the NLSY 1997, an individual-level panel dataset that contains information on higher education, student debt, and labor market outcomes. Our sample follows individuals from 1997 through 2015. Given the focus on earnings outcomes, the main analysis restricts attention to full-time, full-year workers whose highest level of education is a bachelor's degree.[19] Summary statistics are provided in appendix Table A1. The NLSY transcript surveys also provide detailed information on how individuals finance their time in college. Table A2 summarizes this data and compares the composition of education financing for students with and without student debt.

Using the NLSY, we instrument for student debt using variation in the share of grant funding across cohorts within a college, and measure how incremental debt impacts labor market decisions and lifetime earnings trajectories. To construct the instrument, we have accessed restricted-use data that identifies NLSY participants' educational institutions. Using the college identifier, we merge in information from the National Center for Education Statistics (NCES) on the amount of loans and grants awarded at a given college in a given year.[20]

[19]Full-time, full-year status includes those working at least 35 hours a week for at least 40 weeks a year. Robustness exercises in Section 4.3.1 report the sensitivity of the main results to these sample restrictions.

[20]For the details on how we constructed student debt and grant variables, see Appendix B.

4.2 Instrumental Variable Design

To estimate the effect of student debt on an individual's initial earnings and subsequent returns to experience, we employ the following empirical design

$$\begin{aligned} y_{it} \;=\; & \underbrace{\alpha_0 + \beta\,\mathrm{X}_{it}}_{\text{initial (log) earnings if no student debt}} + \underbrace{\alpha_1 \mathrm{Exp}_{it}}_{\text{returns to experience if no student debt}} + \\ & \underbrace{\alpha_2 \mathrm{SD}_i}_{\text{effect of student debt on initial (log) earnings}} + \underbrace{\alpha_3 \mathrm{SD}_i \times \mathrm{Exp}_{it}}_{\text{effect of student debt on returns to experience}} + \epsilon_{it} \end{aligned} \tag{3}$$

where y_{it} is annual log earnings of individual i in year t. The variable SD_i denotes total student debt at graduation and Exp_{it} denotes years of experience. The X_{it} represent additional controls, including year fixed effects; demographic fixed effects for race, age, and sex; initial occupation and occupation interacted with years of experience; and fixed effects for college attended and matriculation cohort.

The goal is to recover an unbiased and consistent estimate of α_2 and α_3. The effect of student debt on initial wages is measured by α_2. The effect of student debt on the returns to experience is measured by α_3. There are potential challenges to estimating equation 3 using OLS. For instance, there may be a correlation between the level of debt an individual takes on and the individual's unobservable ability or human capital. This bias can go either way. Individuals with high ability may expect to have higher future wage growth and so decide to borrow more today to smooth consumption over time, leading to an upward bias in α_2. On the other hand, debt may be positively selected. For instance, low ability individuals may come from low income households, who are unable to provide parental support for their child's education. This shows up as higher borrowing for the low ability individual, leading to a downward bias in α_2 and α_3.[21]

To address these identification challenges, we estimate the impact of student debt on earnings using a school-cohort level instrumental variable. The instrument construction follows that used in Luo and Mongey (2024) – it is defined as the share of grant funding, out of all grant and federal student loan funding, issued by a college c in a given year t.

[21] These challenges to identifying a causal impact of student debt on earnings are also highlighted by Field (2009), Rothstein and Rouse (2011), Chapman (2015), Luo and Mongey (2024). These papers use variation in forgiveness of debt and/or variation in grants within a school across cohorts to instrument for student debt. Like us, their identification comes from comparing cohorts within the same school, when these cohorts within the school differ in terms of grants received.

Specifically, the instrument is given by

$$Z_{c,t} = \frac{\text{total grants}_{c,t}}{\text{total grants}_{c,t} + \text{total loans}_{c,t}}$$

Each individual is then assigned an instrument $Z_i = \frac{1}{|T_i|}\sum_{t\in T_i} Z_{c,t}$ which averages over the years $t \in T_i$ that an individual was enrolled in college c.[22] The instrument utilizes the fact that students must fund their college tuition costs through a combination of parental funding, grants, work study aid, and student loans. While parental funding is specific and fixed at the student level, grant funding can vary significantly at the college-year level. As shown in Luo and Mongey (2024), variation in grant funding is substantial both across and within institutions and years.

Intuitively, the instrument captures the fact that when colleges have less to award to students in the form of grants, students must make up the remaining gap in funding using student loans. The exogeneity assumption relies on the fact that yearly variation in the total amount of grant funding available at a college is unrelated to the ability (or other unobserved characteristics) of any given student at that college. However, to meet the relevance assumption, this variation in grant funding must also create a meaningful change in the amount of student debt that students take out.

Table A3 shows how shifts in the grant share instrument impacts individual student debt and the composition of their college financing, conditional on college and cohort. The results show that a one standard deviation increase in the grant share is associated with a -$3,172 reduction in student debt, equivalent to roughly 15% of average debt balances. Importantly, changes in grant funding are compensated for almost entirely by changes in student debt, not other sources of funding. As the table shows, while the level of student debt decreases almost one-for-one with the increase in grant funding, other forms of financial support remain largely unchanged. Average tuition rates also remain largely the same. This precise substitution between loans and grants is important as it allows the estimation procedure to isolate the impact of an increase in student debt on future earnings, absent the potential confounding effects of other concurrent changes in financial aid or educational expenditures.

Another concern is that changes in the instrument may be correlated with other educa-

[22] If individuals attended more than one college c, the instrument averages over grant shares in both colleges based on the years of attendance at each. Formally, $Z_i = \frac{1}{|T_i|}\sum_{c\in\mathcal{C}_i}\sum_{t\in T_i(c)} Z_{c,t}$, where $\mathcal{C}_i$ is the set of colleges attended by individual i and $T_i(c)$ is the set of years individual i spent at college c.

tional inputs that can also influence student debt and earnings, such as student ability or parental characteristics. A key assumption of the empirical model is that, conditional on student and college characteristics, variation in the grant share is unrelated to these other inputs. This assumption rules out, for example, the sorting of low-ability (or high-ability) students into colleges with higher (or lower) grant shares. Table A4 investigates these concerns by checking if, conditional on college and cohort, variation in the college grant share instrument is correlated with student ability (ASVAB score), high school GPA, parental income and net worth. The results do not show any meaningful association between the instrument and these other factors.[23] The table also examines if variation in the instrument is correlated with selection into the main sample by influencing college completion rates, post-graduate educational attainment, or employment status. Again, the results show no significant evidence that the instrumental variable predicts selection into the sample. While not a definitive proof, the results suggest the model's key identifying assumptions are reasonably justified.

4.3 Estimated Impact of Student Debt on Lifecycle Earnings

This section employs the instrument to investigate how graduating with more student debt shapes an individual's lifetime earnings profile. Table 1 reports the estimated coefficients α_2 and α_3 which summarize the effect. In most specifications and robustness exercises, the estimated coefficient α_2 is positive and α_3 is negative. The coefficients imply that individuals graduating with more student debt have higher initial earnings, but subsequently lower returns to experience. In the fully specified model, a $\$1,000$ increase in student debt leads to a 1.50 percent increase in initial earnings and a -0.41 percentage point decline in the annual returns to experience. The directions of the estimated effects are consistent with the theory on student debt's intertemporal distortions to human capital accumulation discussed in Section 3. The magnitudes provide informative restrictions on the net effect of student debt in the quantitative model.

To better understand the role of the instrumental variable, Table 1 reports both the ordinary least squares (OLS) and instrumental variable (IV) estimated coefficients. To provide insight into the contribution of the controls, results are reported constructively by cumulatively adding covariates until the fully specified model in the final column (4).

[23] Only the ASVAB scores show a statistically significant association, and the magnitude of the effect is small, with a one standard deviation increase in the grant share being associated with a 4 percentage point increase in ASVAB percentile.

Table 1: Estimated Impact of Student Debt on Earnings

	Ordinary Least Squares				**Instrumental Variable**			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Log Initial Earnings (α_2)	-0.09	-0.14	-0.05	-0.14	2.48	2.06	1.50	1.50
	(0.05)	(0.03)	(0.06)	(0.07)	(2.40)	(1.73)	(1.47)	(0.50)
Returns to Experience (α_3)	-0.03	-0.02	-0.02	-0.03	-1.02	-0.94	-0.65	-0.41
	(0.01)	(0.01)	(0.01)	(0.02)	(1.01)	(0.84)	(0.23)	(0.21)
Number of Observations	6360	6319	5903	5633	5985	5946	5548	5548
R-squared \ F-statistic	0.26	0.30	0.46	0.63	23.30	9.73	4.52	10.90
Year Fixed Effects (FE)	x	x	x	x	x	x	x	x
Demographic FE		x	x	x		x	x	x
Occ and Occ x Experience FE			x	x			x	x
College and Cohort FE				x				x

Notes: *The dependent variable is log annual earnings. The coefficients report the effect of a $1000 increase in student debt. The sample includes all full-time, full-year workers with a bachelor's degree. The regression R-squared are reported for OLS results. The first stage F-statistics for student debt are reported with the IV results. For transparency, estimates are reported cumulatively as controls are added until the fully specified model in column (4). Table A5 reports results when covariates are included one at a time. See Appendix B for additional details.*

For completeness, appendix Table A5 reports results for specifications when each set of covariates is included one at a time, rather than cumulatively.

The first four columns of Table 1 report the OLS estimates of α_2 and α_3. In both the partially and fully specified regressions, the estimates show that more student debt is associated with small reductions in both initial earnings and the returns to experience. In the full specification (4), a $1000 increase in student debt is associated with a -0.14 percent reduction in initial earnings and a -0.03 percentage point decline in the returns to experience. The estimates are only marginally significant around the 10% threshold and their magnitudes appear economically inconsequential.

The remaining columns of Table 1 report the instrumental variable estimates of α_2 and α_3. Across the specifications, the IV coefficient estimates are an order of magnitude larger than the OLS results. The sign of the effect on initial earnings is also flipped from negative to positive. Similar changes in the sign and statistical significance of initial earnings effects after instrumentation also appear in several other empirical studies of student debt (for example, see Rothstein and Rouse 2011). To the extent that the IV is well-identified, the magnitude of the difference in the coefficient estimates suggests that the size of the bias in the OLS estimation may be quite large.

The first column reports coefficient estimates when the control variables X_{it} include only year fixed effects. The second column adds demographic controls by including fixed effects for individual age, race, and sex. Comparing columns (1) and (2) in the IV segment suggests the results are not driven by differences in the incidence of student debt across sex, race, or age sub-populations. Both yield estimates that show a positive effect on initial earnings ($\alpha_2 > 0$) and a negative effect on earnings growth ($\alpha_3 < 0$), though neither is statistically significant.

The third column adds fixed effects for initial occupation choice and the interaction of these fixed effects with years of experience. These occupational controls allow for the possibility that the estimated effects on earnings are the result of differences in the occupational sorting of students with and without student debt. Though not a formal decomposition, comparing columns (2) and (3) suggests that differences in occupation choice may account for roughly a quarter of the estimated effects.[24] The estimated effect on initial earnings falls from 2.06 to 1.50 while the impact on returns to experience declines from -0.94 to -0.65 and becomes statistically significant. The results suggest that the effect of student debt on lifecycle earnings profiles has both a within-occupation and a between-occupation component.

The final column (4) completes the empirical model by including college and college cohort fixed effects.[25] The purpose of these controls is to isolate variation in student debt that is independent from potentially confounding variation in college and student types. The aim is to rule out the possibility that the results are driven by the sorting of low-ability (or high ability) students into colleges with higher (or lower) grant shares.[26] The results show that including college and college cohort fixed effects strengthens the instrument and leads the estimated coefficients to become statistically significant. The outcome suggests that accounting for college heterogeneity and the sorting of students is crucial for addressing the underlying bias.

In the fully specified model (4), an additional $1,000 of student debt leads to a 1.50% increase in initial earnings and a -0.41 percentage point reduction in subsequent returns

[24]Table A5 shows that including the occupational controls has a similar proportional effect on the estimated α_2 and α_3 coefficients when they are considered in isolation, without other controls.

[25]Cohorts are defined by admission year, to account for the peers among which they were selected.

[26]For example, without these controls the estimation could attribute differences in the starting salaries of graduates from any two colleges, say Harvard and the University of California, entirely to differences in student debt across graduates from the two schools.

to experience. Both estimates are statistically significant at the five percent threshold. To interpret the magnitudes, the effect on initial earnings corresponds to an additional $342 in annual earnings upon graduation for every $1000 of additional debt. Despite working with a different sample and data source, the estimated effect of student debt on initial earnings is very similar to what others have documented in the literature. For example, Rothstein and Rouse (2011) and Luo and Mongey (2024) estimate an initial earnings effect around $200, while Chapman (2015) documents a $400 to $800 effect.

The estimated impact of student debt on annual earnings growth also appears consistent with the literature, though there are far fewer comparable studies. The most closely related estimates are provided by Folch and Mazzone (2026) who use a similar instrumental variable strategy to study long-run earnings outcomes in the Baccalaureate and Beyond Longitudinal Study. Their results imply that an additional $1000 in student debt leads to an average -0.25 percentage point decline in earnings growth, which lies at the lower end of the estimated confidence interval here.[27] The magnitude of the effect also appears significant given the average earnings growth of early career workers. For instance, Guvenen et al. (2021) employ IRS administrative data and estimate that the annual earnings growth of all 25 to 30 year-olds in the United States is 7.75%.[28] Within sample, the estimates here imply that the lower returns to experience erase the initial earnings advantage of those with student debt after roughly 5 years.

4.3.1 Additional Robustness

The benchmark sample places two strong restrictions on the population data by focusing on full-time, full-year workers whose highest level of educational attainment is a BA degree. These restrictions yield a well-behaved subpopulation representing the bulk of the college workforce that helps isolate the impact of student debt on human capital and occupation choice. However, they omit any effect that student debt may have on labor supply or graduate education. This raises concerns given that an existing literature shows that both margins can depend on debt in a variety of ways. In this case, the sample restrictions may lead to selection based on outcomes that are themselves influenced

[27] More specifically, Folch and Mazzone (2026) report that a $2,364 increase in student debt leads to a -5.3% reduction in earnings ten years after graduation. Rescaling the results implies an annual effect of $100 * (1 - (1.0224)^{1/9}) = -0.246$ percentage points per year. The authors also document that student debt leads to a 3.6% rise in initial earnings, roughly double the size of the effect estimated here.

[28] Though this number includes both high school and college graduates, which also likely exhibit significant between group variation in early career returns to experience. Within our sample, the average returns to experience is 6.59% for 25-30 year-olds and 11.87% over the entire 22-35 estimation window.

by student debt, potentially biasing the conclusions.

To investigate these concerns, Table A6 reports robustness estimates of the full empirical model when the sample restrictions on employment status are relaxed. The main estimation focuses on individuals who worked an average of at least 35 hours per week for at least 40 weeks per year. The table shows how the coefficients change when the restrictions on hours, weeks, or both are removed. It also reports results with tighter labor market restrictions on hours and weeks worked. Across the specifications, the results continue to show positive point estimates for initial earnings and negative effects on returns to experience. Estimates of the initial earnings coefficient range from 0.17 to 2.50, though their statistical significance is highly sensitive to lower bound restrictions on hours worked. In contrast, the estimated effect on the returns to experience appears to be relatively robust with statistically significant estimates ranging from -0.66 to -0.38.

Table A7 reports the coefficient estimates when sample restrictions on educational attainment are removed so as to include those with masters, professional, and PhD degrees in the sample. The table reports results for the benchmark model as well as for augmented specifications that include additional fixed effects for post-graduate education and its interaction with years of experience. The point estimates of the effect on returns to experience are nearly identical to the main results, but they are no longer statistically significant when including those with graduate and professional degrees. The initial earnings effect continues to be positive, but is an order of magnitude larger than in the benchmark sample.[29]

While broadly consistent with the main findings, these results should be interpreted with caution. First, the estimates are relatively imprecise in part because the post-graduate population in the data is relatively small given the richness of the empirical model (see Table A8). Second, the F-statistics indicate that the grant share is a weak instrument in the expanded sample with post-graduates. The instrument's weakness in the expanded sample is perhaps not so surprising given the substantial diversity in education financing arrangements across graduate programs and professional schools. Therefore while variation in a college's grant share may be a crucial determinant of student debt for undergraduates, it may be less relevant for those who enrolled in profes-

[29] One interpretation of this much larger initial earnings effect is that those who anticipate pursuing graduate degrees may accept jobs with initial lower earnings in order to prepare themselves for higher education. Since those with student debt are less likely to pursue graduate education (Zhang 2013; Chakrabarti et al. 2022; Folch and Mazzone 2026), the estimated effect of student debt on initial earnings now also captures early career earnings changes in anticipation of graduate enrollment.

sional or graduate programs. A more complete model would account for the sequential nature of degree decisions and the heterogeneity in financing and employment opportunities across branches of the higher education system.

A final concern is whether the set of colleges attended in the NLSY sample is representative of US higher education. Table A9 compares the characteristics of colleges attended in sample to the student-weighted US college averages. It shows that individuals in sample are more likely to attend smaller, private colleges which are more focused on undergraduate education. These colleges also provide slightly more student loans and grants overall compared to the U.S. average, but proportionally the college level grant shares are nearly identical.

A related question is whether, given the relatively small cross-section of graduates in the NLSY, the data provides sufficient power to estimate the college level fixed effects. For instance, 53.6% of survey respondents in the main sample – representing 19.6% of all observations – are the sole attendee at their college. To investigate the effect of this sparsity, Table A10 reports estimation results for the fully specified model and robustness variants after replacing the college fixed effects with a parsimonious set of college characteristics. These characteristics include an indicator for public or private status, college size measured by full-time equivalent (FTE) student enrollments, and the B.A. share of all degrees conferred to measure undergraduate institutional focus.[30] Re-estimating the fully specified model with college characteristics shows that a $1000 increase in student debt leads to a 2.27% rise in initial earnings and a -0.52 percentage point decline in the returns to experience. The estimates are statistically significant and comparable in magnitude to the fully-specified benchmark model, suggesting that the results are not unduly influenced by sparsity in the college data.

4.4 The Role of Occupation Choice

The instrumental variable analysis finds that student debt leads to more front-loaded earnings profiles, both within and between occupations. To better understand the latter, this section investigates how realized age-earnings profiles differ across occupation groups. Occupation-specific earnings trajectories are summarized using traditional

[30] These covariates capture the most differentiated margins between in-sample colleges and the average U.S. college (see Table A8) and have the added benefit of being plausibly independent of unobserved student characteristics within schools.

earnings regressions fit separately for each occupation where,

$$Earnings_{it} = \beta_{0,j} + \beta_{1,j}\text{Exp}_{it} + \beta_{2,j}\text{X}_{it} + \epsilon_{it}. \tag{4}$$

and all regression coefficients are allowed to vary across occupations j. The X_{it} include additional individual level control variables that may influence the realized cross-sectional income variation within and between occupations, including race, gender, region where the respondent resides, year and industry fixed effects.[31] The estimation encompasses the same sample population used in the regression analysis.

The results show a strong negative correlation between initial earnings and subsequent earnings growth across occupations. Figure 3 illustrates the point by plotting the initial earnings and returns to experience implied by the estimated coefficients $\beta_{0,j}$ and $\beta_{1,j}$. It shows that there exists a statistically significant negative relationship between an occupation's initial earnings after graduation and its average yearly growth rate.[32] The substantial cross-occupation variation in earnings profiles suggests that the marginal constrained household may have ample opportunity to intertemporally trade-off income by moving across occupations. The results are consistent with the conditions necessary for credit constraints to cause a misallocation of talent within the model.

It is worth noting again that these are equilibrium earnings trajectories, which encompass endogenous decisions such as labor supply, human capital investments, and occupational sorting. Nevertheless, while these age-earnings profiles are likely not exogenous (e.g. unaffected by the sorting of individuals with different debt or ability levels into certain categories), they do help explain the between occupation effect of student debt. The variation in occupational earnings profiles determines, in part, the scope of adjustment available to the marginal misallocated individual. Within the model, these realized cross-sectional differences in earnings must also be internally consistent with equilibrium occupational sorting and human capital investments. For these reasons, the calibration strategy includes these cross-occupation moments as targets for the quantitative model alongside the instrumental variable results estimated above.

[31] Appendix B provides additional details on the data and variable construction.

[32] The cross-sectional correlation does not meaningfully change if a quadratic term in experience is added to the wage regression. This is likely due in part to the fact that our sample contains only young workers under 40.

Figure 3: Initial Earnings and Earnings Growth by Occupation

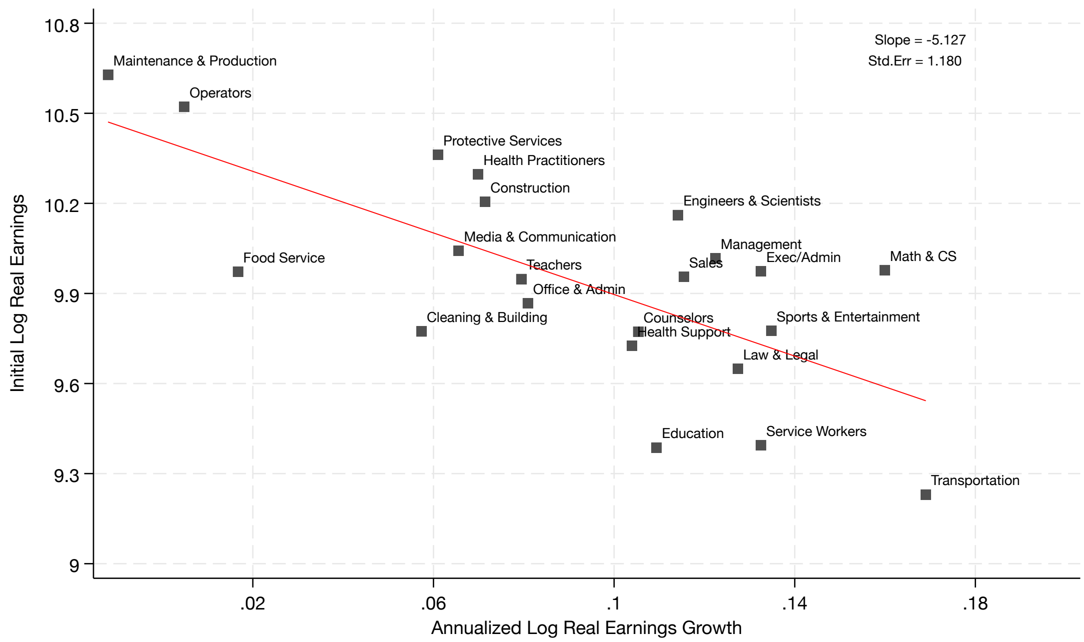

Notes: The figure plots the estimated earnings function coefficients from equation (4) for each three-digit occupation. Y-axis values correspond to the estimated $\beta_{0,j}$ and X-axis values correspond to $\beta_{1,j}$. Table A11 in the appendix reports the underlying regression coefficients.

5 Calibration Strategy and Model Fit

The goal of the quantitative model is to assess the aggregate consequences of intertemporal distortions to human capital and occupation choice resulting from student debt. As the discussion above explains, credibly doing so requires identifying the size and scope of the population holding student debt as well as the extent to which these debts affect their occupational choice and earnings decisions. The first requires matching the realized distribution of student debt across households, accounting for the fact that the population holding debt will not be random, but rather determined endogenously through the college matriculation choice. The second requires matching the estimated IV marginal effects of student debt on household earnings (i.e., α_2, α_3) and how workers differentially sort themselves across the heterogeneous occupations (i.e., Figure 3).

The following sections discuss how the model's parameters are tuned in accordance with this calibration strategy. While all the parameters will jointly determine the model's ability to match the data, each set of parameters is discussed in conjunction with their most closely associated data targets to help build intuition. Tables 2, A11, and 3 report the internal calibration targets and model fit. Table 2 summarizes college matriculation patterns and the distributional properties of student debt. Table A11 reports occupational employment shares and earnings heterogeneity across occupations. For ease of exposition, Figure 4 displays these targets graphically within the text. Table 3 contains the indirect inference targets and summarizes the model's ability to match the empirical evidence on the marginal effects of student debt estimated in Section 4. Table 4 summarizes the externally calibrated parameters.

College Matriculation and the Distribution of Student Debt. Modeling college matriculation is crucial since it will determine endogenously the population who choose to take on student debt and attend college. Identifying this population, and how it might change under various counterfactuals, is crucial for understanding the aggregate consequences of student debt. The model captures these mechanisms by jointly replicating in equilibrium how college matriculation rates and student debt vary across the household wealth distribution.

Recall that, to account for the variety of circumstances determining an individual's access to student financial-aid and family assistance, the model allows college admission offers to depend stochastically on household characteristics. After graduating high school, each household receives an *admission offer* of the form $\{x, \tau\}$, which require students to take on debt $d = (1 - x)\tau$ to receive their degree. Given that the empirical distributions of student debt and household assets appear log normal, the calibration parameterizes the distribution of latent admission offers τ by

$$\begin{pmatrix} a_0 \\ \tau \end{pmatrix} \sim LN\left[\begin{pmatrix} \mu_a \\ \mu_\tau \end{pmatrix}, \begin{pmatrix} \sigma_a^2 & \rho_{a\tau} \\ \rho_{a\tau} & \sigma_\tau^2 \end{pmatrix}\right]$$

where $\rho_{a\tau}$ is the correlation between a household's initial assets and the student debt it needs to take on in order to complete college. The extensive margin of student debt is captured by the Bernoulli random variable x which equals one with probability p_x.

The associated parameters $p_x, \mu_a, \mu_\tau, \sigma_a^2, \sigma_\tau^2, \rho_{a\tau}$ are jointly set in the internal calibration so that the model replicates (i) the marginal distribution of student debt (mean and variance) (ii) the marginal distribution of initial household assets (mean and variance) (iii) the share of college graduates without any student debt, (iv) the correlation between realized student debt levels and initial household assets. Panel (B) of Table 2 summarizes the internal calibration targets and model fit.

It is important to note again that these parameters generate the *latent* student debt distribution, while the realized student debt distribution will depend on households' endogenous selection into college after observing admission offers.[33] In addition to student debt, the decision to matriculate will depend on the returns to a college education. These in turn depend on market prices (e.g. the college wage premium), individual abilities Θ, human capital h_0, and whether the household has sufficient financial assets to avoid being constrained after graduation. The resulting selection into college will shape the model's equilibrium joint distribution of talents and assets, one of the key objects determining the aggregate consequences of student debt on earnings (see Section 3).

To discipline the selection into college, the calibration matches (i) college completion rates by household asset quintile and (ii) the average difference in lifecycle earnings for high school and college graduates (e.g. the college wage premium, returns to experience).[34] Panel (A) of Table 2 summarizes the internal calibration targets and model fit. The associated parameters include the aggregate productivity of high school and college labor, A_{hs} and A_c, the distribution of high school ability θ_0, and the distribution of the college taste shock ζ.

To match these targets, the calibration abstracts from heterogeneity in high school ability and sets θ_0 to a common value for all households to match the average observed earnings growth for high school graduates. The aggregate productivity of high school labor, A_{hs}, is set to generate a 43% college wage premium, as in Daruich and Kozlowski (2020).[35] Finally, given earnings, the college taste shock ζ is calibrated to match college completion rates. To capture the non-linearities in the data, the college taste shock is allowed to depend on initial assets such that households in asset quintile i draw their taste shock from the distribution $N(0, b_i^2)$. Given the distribution of earnings and occupation choice, the standard deviations of these taste shocks b_i are chosen to exactly replicate college matriculation rates across quintiles of the household asset distribution.

[33] The exogenous parameterization of latent admission offers that shape equilibrium student debt is a convenient reduced form representation of the non-modeled process by which the government, family members, non-profit institutions, and education institutions determine tuition, grants, and financial aid.

[34] Note that without college drop-outs, there is no distinction between the college matriculation rate and college completion rate in the model.

[35] The college TFP term A_c is normalized to one, e.g. $A_c = 1$, without loss of generality as it is not separately identified from the level of occupation specific productivity terms A_k discussed below.

Table 2: Internal Calibration Targets on College Matriculation and Student Debt

	Data	Model
A. College Matriculation, Skill Premium, and Earnings Heterogeneity		
College wage premium	43%	42%
High school average returns to experience	4.16%	4.73%
Variance of log earnings at age 23	0.24	0.27
College completion rates by asset quintile	30.59%	30.69%
– First quintile (Q1)	18.40%	18.31%
– Second quintile (Q2)	18.52%	18.80%
– Third quintile (Q3)	26.55%	26.70%
– Fourth quintile (Q4)	35.29%	35.52%
– Fifth quintile (Q5)	54.51%	54.12%
B. Distributional Moments on Student Debt		
Mean level of initial assets	\$18,132	\$18,124
Standard deviation of initial assets	\$21,503	\$21,523
Mean level of student debt	\$21,843	\$21,788
Standard deviation of student debt	\$28,041	\$28,503
Correlation between initial assets and student debt	-0.15	-0.21
Fraction of BA graduates without SD	40.33%	41.65%

Notes: This table summarizes the internal calibration targets and model fit for the moments pertaining to college matriculation and the distribution of student debt. See Appendix B for additional details on data sources and variable construction.

Occupational Sorting and Earnings Heterogeneity. In addition to matching the average returns to college via the college wage premium, it is crucial that the calibrated model replicate the occupational sorting of workers and the resulting cross-occupation heterogeneity in lifecycle earnings. These moments will help identify the aggregate population of constrained households. They also have implications for the size, scope, and direction of occupational distortions resulting from credit frictions.

To capture these data, the calibration replicates in equilibrium the (i) initial earnings, (ii) average returns to experience, and (iii) employment share of each occupation. With 22 occupations, these data provide 66 additional calibration targets, to be matched primarily by parameters $\{A_k, \mu_k, \nu_k\}_{k=1,...,22}$. Specifically, the level of initial earnings in each occupation k is matched by occupation-specific productivities A_k. The returns to experience in each occupation are matched by the average ability level θ_k realized in the population (e.g. before selecting into occupations).[36] For tractability, the calibration parameterizes the population talent distribution for Θ with the log normal distribution

$$\begin{pmatrix} \theta_1 \\ \vdots \\ \theta_K \end{pmatrix} \sim LN \left[\begin{pmatrix} \mu_{\theta_1} \\ \vdots \\ \mu_{\theta_K} \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & \cdots & 0 \\ 0 & \ddots & 0 \\ 0 & \cdots & \sigma_\theta^2 \end{pmatrix} \right]$$

where μ_{θ_k} governs the average population ability in occupation k and σ_θ^2 the variance.[37]

Finally, given lifecycle earnings in each occupation, the corresponding amenity values ν_k, common to all workers within an occupation, are chosen to replicate the distribution of occupational employment shares. The presence of these occupational amenities is one of the reasons that the cross-occupation correlation between student debt and earnings profiles will not be perfectly monotone. This is because in addition to having different valuations of income today versus income tomorrow, constrained and unconstrained households also differ in the value they assign to non-wage amenities versus earnings, as discussed at length in Luo and Mongey (2024) and Boar and Lashkari (2025).[38]

Figure 4 summarizes the data targets and resulting model fit. The position of each occupation represents its initial earnings (y-axis), average returns to experience (x-axis), and employment share (bubble size) in the data and the calibrated model. Table A11 in the appendix reports the corresponding values for all 66 occupational data targets and model moments.[39] The result shows that the model does a good job overall at matching the sorting of workers and the heterogeneity in lifecycle earnings across occupations. It is worth noting again that, while there is expositional value in matching particular data targets and parameters, in practice all the internal calibration targets are jointly determined in equilibrium. The parameters $\{A_k, \mu_k, \nu_k\}_{k=1,\ldots,22}$ therefore do not map directly

[36]Recall that in the Ben-Porath (1967) model, the *level* of wages w_k do not determine optimal human capital investment for unconstrained households since they do not affect intertemporal trade-offs. However, for constrained households, optimal investments do depend on wage levels. See Section 3.

[37]The calibration of σ_θ^2 is discussed in the following section. Note that even though there is a common variance for all occupations, there will still be equilibrium heterogeneity in the variance of earnings across occupations. These differences will result primarily from the endogenous sorting of workers across occupations and the scope to which credit frictions affect investment within that occupation. They do not emerge from heterogeneity in the variance of abilities across occupations.

[38]Intuitively, households' marginal value of consumption, and hence earnings, increases as their debts rise while the utility value of occupational amenities, as often modeled, is invariant to household assets.

[39]Data targets correspond to those reviewed in Figure 3, augmented to also include employment shares.

Figure 4: Initial Earnings, Earnings Growth, and Employment Shares by Occupation

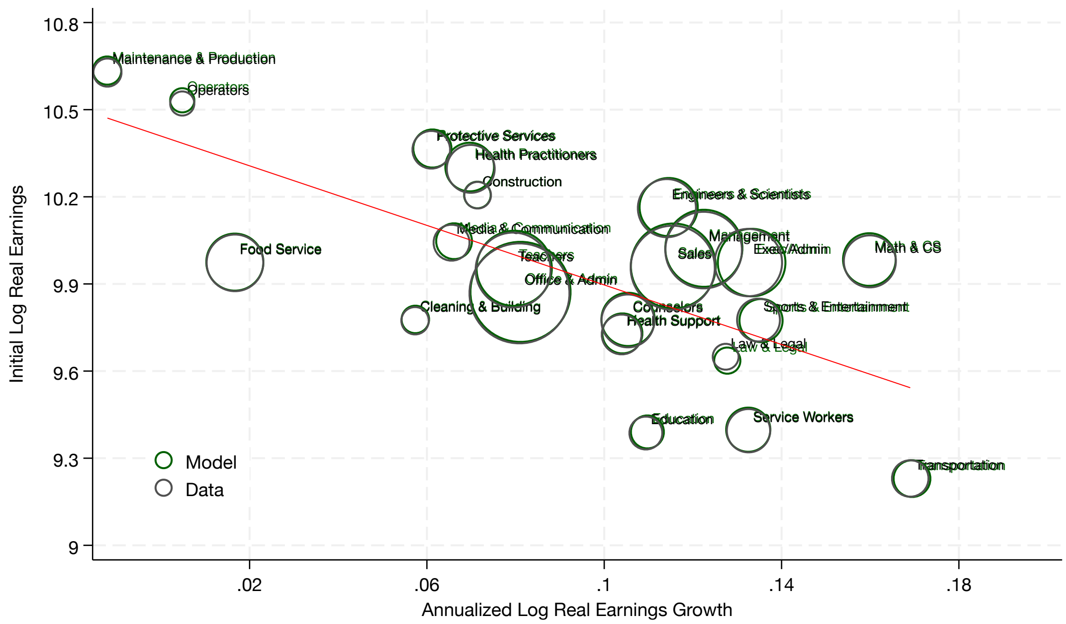

Notes: The figure plots the estimated earnings function coefficients for each occupation in the model and data. The dashed lines show the associated OLS coefficients summarizing cross-occupation correlations in earnings characteristics. Model-simulated average amount of student debt by occupation is represented using differently sized bubbles. Large bubbles have relatively high occupational employment shares. Appendix Table A11 reports the underlying data targets and corresponding model moments.

to the calibration targets. For instance, A_k influences the wage rate for occupation k both directly through labor demand and indirectly through its effect on the sorting of workers and human capital investment in the presence of credit constraints.

Evidence on the Marginal Effects of Student Debt. Until now, the calibration has sought parameters which, in equilibrium, replicate key distributional characteristics of the U.S. economy. These population moments are informative in that they place restrictions on the potential aggregate impact that student debt can have. In addition to these distributional moments, the calibration also targets the estimated marginal effects of student debt on earnings at the microeconomic level. Matching this evidence restricts how the individual households, which constitute the model's micro-foundation, respond to exogenous changes in student debt. In the context of this analysis, these moments correspond to the instrumental variable estimates in Section 4.

Unlike the distributional moments which correspond to equilibrium properties of the

Table 3: Internal Calibration Targets on the Marginal Effects of Student Debt

Target for Indirect Inference	**Coefficient**	**Data**	**Model**
IV Effect on Initial Earnings (IE)	$\hat{\alpha}_2$	1.50%	1.58%
IV Effect on Returns to Experience (RTE)	$\hat{\alpha}_3$	-0.41%	-0.31%

Notes: This table reports data targets for indirect inference. It summarizes internal calibration targets on the marginal effects of student debt and the model fit. The targets include the instrumental variable coefficients on initial earnings and returns to experience in column (4) of Table 1. The model entries correspond to the coefficients retrieved from running the benchmark empirical specification on the data produced by the model simulation. See Data Appendix B for additional details.

economy, the marginal effect targets correspond to household optimal responses to perturbations of the equilibrium. To match this evidence, we simulate the first ten years of each individual's lifecycle and then regress their log earnings on their student debt, their years of experience, and the interaction of the two. In other words, we replicate regression (3) using the simulated data and compare the coefficients to their empirical counterparts. Through indirect inference, the calibration seeks a parameterization which results in the simulated policy generating the same average marginal effects on initial earnings and the returns to experience as in Table 1. The results are presented in Table 3, which displays the original IV coefficients estimates in the fully-specified empirical model alongside the coefficients recovered from the model-based regression.

The empirical evidence primarily helps the calibration fit μ_{h0} and σ_θ^2 through indirect inference. This is because conditional on matching the other model observables – including occupational wages, matriculation patterns, and the joint distribution of assets and student debt – the average effect of exogenous changes in debt on lifecycle earnings will primarily depend on the household's initial human capital h_0 and the latent distribution of abilities θ. The dependence of the results on initial human capital h_0 is relatively straightforward to see. It primarily helps match the effect of student debt on initial earnings, summarized by α_2. Conditional on matching earnings, a larger h_0 implies that a larger share of household earnings is determined by human capital, as opposed to effective wages. As a result, changes in human capital investments induced by student debt will have a larger effect on initial earnings. Moreover, the prevailing level of human capital also directly influences the responsiveness of further human capital investments to debt given the dynamics of the Ben-Porath human capital technology.[40]

While simply adjusting the level of h_0 provides sufficient flexibility for the model to match the initial earnings effects of student debt, we follow Huggett, Ventura, and Yaron (2011) and allow for heterogeneity in initial human capital such that $h_0 \sim LN(\mu_{h_0}, \sigma^2_{h_0})$. This heterogeneity captures additional differences in household initial conditions and unobserved pre-college educational investments that are not otherwise accounted for by the model. To facilitate comparison (see Section 5.1), we fit the distribution of the initial heterogeneity in human capital, $\sigma^2_{h_0}$, to match the population variance of initial earnings at age 23, as in Huggett, Ventura, and Yaron (2011). The target and model fit are reported alongside the other earnings moments in Table 2. To fit the empirical evidence on the marginal effects of student debt on initial earnings, the calibration adjusts μ_{h_0}.

The population distribution of ability, σ^2_θ, primarily but not exclusively pins down the estimated effect of student debt on the returns to experience, α_3. Since individuals generally select into occupations in which they have higher abilities, it is the tail of the ability distribution implied by variance σ^2_θ which matters most for how human capital investments respond to tightening credit constraints. Comparing the optimal investment rules s^c and s^* in Section 3 illustrates how the average intertemporal distortions to human capital investments – and hence earnings growth – depends on the distribution of θ. Similarly, the distribution of θ also determines the mass of households which fall into the misallocation region displayed in Figure 2.

Externally Calibrated Parameters. Table 4 lists the model's externally calibrated parameters and their sources. These parameters predominantly govern (i) the strength of household consumption-savings incentives, (ii) structural characteristics of student debt repayment plans, and (iii) the risk and productivity of human capital investments over the working lifecycle. The household discount factor β and preference parameter ρ are set to common values from the literature in order for the model to generate reasonable consumption smoothing incentives. The risk free rate, r, is chosen to match the long-term U.S. interest rate and the interest rate on student debt, r_d, is set to the average interest rate on outstanding student loans according to the National Center for Education Statistics. The delinquency consumption floor, $\bar{c}$, is set to the Health and Human Services (HHS) poverty threshold, which results in an equilibrium default rate on stu-

[40] To see this more directly, divide through by current human capital in the Ben-Porath technology to express the growth rate of human capital as, $h_{t+1}/h_t = \theta^{1-\alpha} h_t^{\alpha-1} s_t^\alpha + (1-\delta)$.

Table 4: Externally Calibrated Parameters

Parameter	Interpretation	Value	Source
Household Preferences			
β	Discount rate	0.985	Standard
ρ	CRRA preference parameter	2	Standard
Student Debt Repayment Rule			
r	Risk-free interest rate	0.040	FRED
r_d	Student debt interest rate	0.042	NCES
$\bar{c}$	Delinquency consumption floor ($)	10,000	HHS poverty guideline
$\bar{T}$	Student debt repayment duration (years)	10	SFRP institutional feature
Human Capital Technology			
σ_z	Standard deviation of human capital shocks	0.111	Huggett et al. (2011)
α	Returns to scale in human capital tech.	0.7	Huggett et al. (2011)
δ	Human capital depreciation rate	0.029	Huggett et al. (2011)

Notes: All parameters are annualized. See Appendix B for additional details.

dent debt below 1.6%. The repayment duration $\bar{T}$ is set consistently with the standard federal repayment plan (SFRP) during the sample period.

The structural parameters governing the human capital technology are taken from the analysis in Huggett, Ventura, and Yaron (2011), who employ the same Ben-Porath formulation to study lifecycle earnings inequality in the United States. The riskiness of human capital accumulation, σ_z, is set to reflect the growing dispersion in earnings over the lifecycle. Parameter α controls the diminishing returns to human capital accumulation, and several studies find evidence for similar values in diverse settings (for example, see Ionescu 2009, Lee and Seshadri 2019). The human capital depreciation rate, δ, is set to match the decline in lifecycle earnings near the end of working life.

Finally, the counterfactual exercises allow individuals to re-optimize their human capital investments and occupational choices following student debt relief. To capture the idea that skills are imperfectly transferable across occupations, workers who switch occupations lose a fraction of their human capital acquired on the job. The calibration sets this penalty to 20%, the average estimated by Dvorkin and Monge-Naranjo (2019) across coarse job types in the United States. While a more complete treatment would parameterize the full bilateral matrix of human capital transferability across all 22 x 22 occupation pairs, such detailed estimates are currently unavailable in the literature.

5.1 Lifecycle Inequality and Connections to Huggett, Ventura, and Yaron (2011)

The model here builds on the seminal approach in Huggett, Ventura, and Yaron (2011). It extends their framework along a number of dimensions by including occupational heterogeneity, multidimensional ability, student debt and college matriculation. The approach here also pursues a broader calibration strategy that matches a larger set of moments and empirical evidence on the marginal effects of debt on earnings. One particularly consequential difference is in the calibration of initial household assets. Huggett, Ventura, and Yaron (2011) calibrate their model so that all young workers begin their life with zero assets.[41] In contrast, the calibration here jointly targets the initial distribution of household assets and student debt, as in Figure 1, alongside the other distributional targets. The distinction is important as, in both papers, credit frictions lead the initial distribution of assets to influence lifecycle earnings.

To contextualize the model extensions and differences in the calibration strategies, this section recomputes the main exercises in Huggett, Ventura, and Yaron (2011) and compares the results to their original findings. Their main exercise analyzes the sources of lifetime inequality–in welfare and lifetime wealth–by perturbing the underlying margins of individual heterogeneity and measuring how much lifetime outcomes change. Table 5 displays the results by reporting their original findings (HVY) alongside those derived from the model here (ACK). The table readily highlights their main finding that the largest factor contributing to lifetime inequality is differences in initial human capital. Heterogeneity in learning ability contributes much less to lifetime inequality and differences in initial assets make the smallest overall contribution.

In contrast, the model here finds that differences in initial assets are the most important factor explaining lifetime inequality. This is especially the case if one also includes the contribution of student debt. Initial differences in human capital are now the second most important factor, particularly for lifetime wealth. Finally, heterogeneity in learning ability is now the least important factor, although the overall magnitude of its contribution is relatively similar across the two models.

The more prominent role for initial assets and smaller contribution of initial human capital is likely the result of both differences in the calibration strategy and the structural

[41] As they write on page 2945, "The benchmark model abstracts from initial wealth differences. This is a potentially important omission, as wealth inequality is substantial among the young." In an auxiliary partial equilibrium exercise, they add wealth heterogeneity to their benchmark calibrated model ex-post and argue it does not impact their main conclusions. However, the exercise does not recalibrate the model parameters and the paper does not report how the model targets change following this modification.

Table 5: Sources of Lifetime Inequality

Variable	Change in variable	**Equivalent Variation (%)**		**Lifetime Wealth (%)**	
		ACK (2026)	HVY (2011)	ACK (2026)	HVY (2011)
Human capital	+1 st. deviation	6.9	39.3	15.5	47.5
	−1 st. deviation	-7.1	-28.3	-13.4	-31.7
Learning ability	+1 st. deviation	8.1	5.7	8.6	8.1
	−1 st. deviation	-7.7	-2.6	2.1	-3.9
Initial wealth	+1 st. deviation	14.0	7.1	71.1	5.0
	−1 st. deviation	-17.9	-1.6	-49.3	-1.3
Student debt	+1 st. deviation	-9.0	—	-10.9	—
	−1 st. deviation	8.7	—	17.6	—

Notes: The table states equivalent variations and the percentage change in the lifetime wealth associated with changes in each initial condition. The baseline initial condition is set equal to the median values of initial human capital, learning ability, wealth, and student debt. HVY columns report the results from Table 6 in Huggett, Ventura, and Yaron (2011). ACK columns correspond to the model developed here.

extensions to the model. For the calibration strategy, the most important difference is almost certainly the initialization of household assets. Huggett, Ventura, and Yaron (2011) calibrate their model assuming all households begin life with zero assets. Given the interaction between initial assets and human capital accumulation in both models, omitting heterogeneity in the former can necessitate a more prominent role for the latter to explain a given level of heterogeneity in lifecycle outcomes observed in the data. In contrast, the calibration strategy here matches the distribution of initial assets and student debt jointly with the model's other distributional targets. The impact of this asset heterogeneity on human capital accumulation is further disciplined by empirical evidence on the marginal effects of student debt on earnings. Accounting for the substantial wealth inequality among young workers (displayed in Figure 1) results in a parameterization that requires less initial heterogeneity in human capital to explain the same set of lifecycle earnings targets.

As with the calibration strategy, structural extensions to the model environment also increase the interdependency of assets and human capital. Chief among these is the inclusion of endogenous college matriculation and student debt. The strong positive correlation between initial assets and selection into college observed in the data further strengthens the connection between assets and human capital in the quantitative model.

Moreover, ex-post realizations of student debt also strongly depend on initial assets and are more subject to the distortionary effects of credit frictions given their short duration and stringent repayment rules. The inclusion of occupational heterogeneity also increases the responsiveness of human capital to debt by providing an additional margin along which households can adjust their earnings in response to credit frictions.

To summarize, while the model builds on the seminal framework of Huggett, Ventura, and Yaron (2011), its quantitative predictions on the role of household heterogeneity in determining lifecycle outcomes differ. While both models contain similar economic interdependencies between assets and human capital due to credit frictions, the quantitative strength of this channel differs due to differences in the calibration strategy and structural extensions to the model environment. The result is a more prominent role for the initial asset distribution of young workers in explaining inequality in their lifecycle welfare and earnings outcomes.

6 Macroeconomic Implications of Student Debt

This section reports the results of two computational exercises which quantify the aggregate impact of student debt on lifecycle earnings and the misallocation of talent. The first is a short-run exercise that computes the effect of a one-off student debt forgiveness program, similar to those being proposed by recent administrations.[42] The second is a long-run exercise examining the aggregate consequences of making the proposed education grants permanent. The computational results are reported in Tables 6 and 7.

Short-Run Student Debt Policies. The first exercise computes the short-run consequences of a one-off student debt forgiveness policy, holding occupational wages constant. It is a short-run policy in that it applies only to the current population of student debt holders at the time of the announcement. The model reports the impact of student debt forgiveness policies of up to 10K, 20K, and full-forgiveness. The first two reflect the size of programs proposed by recent administrations, while the latter provides a benchmark of the short-run aggregate distortions from student debt.

Table 6 reports the impact of each short-run policy on welfare and lifecycle earnings. Household consumption-equivalent welfare is decomposed into changes from lifecycle

[42] For details, see https://studentaid.gov/debt-relief-announcement

earnings and occupation specific amenities for those who switch jobs. Changes in household earnings are similarly decomposed into the contribution of occupational wages and effective hours, capturing the impact on labor productivity.[43] The final columns report the average level of student debt relief received under each policy and the average change in tax rates necessary to fund each program. Increases in taxation are implemented proportionally to existing tax rates so as to maintain the progressivity of the U.S. tax system. Online Appendix F reports results under a lump-sum taxation program. Outcomes are provided for the entire U.S. population and for the sub-populations of college graduates who either stay or are induced to switch jobs by the policy.

The results show that one-off student debt forgiveness programs would increase household welfare predominantly by increasing lifecycle earnings through a reduction in the distortions to human capital accumulation. Focusing on the full-forgiveness policy, average household welfare increases by 1.35%, with 1.08 percentage points of the rise coming from an increase in lifetime earnings, accounting for roughly 80% of the total effect. The increase follows from the improved efficiency of human capital investments over the working lifecycle. Under the full-forgiveness policy, total lifetime earnings increase by 0.24%, driven entirely by increases in human capital accumulation.

These population outcomes predominantly reflect a reduction of intertemporal distortions to human capital accumulation on the job, rather than a reduction in the misallocation of talent. This is because even under the full-forgiveness policy, only 0.59% of the population (1.94% of college graduates) is induced to switch jobs. Moreover, the results show that the direction in which the induced job switchers re-sort themselves across occupations runs *contrary* to the aggregate effects. While aggregate lifetime earnings increase, the population of induced job switchers experiences a large -21.32% reduction in lifetime earnings. The decline is driven by both their switching to jobs with lower wages and less scope for human capital accumulation.

It may be surprising, in light of the theory of misallocation developed in Section 3.2, that occupational reallocations are associated with *decreases* in lifetime income. This apparent contradiction is reconciled by the presence of heterogeneity in job-specific amenities, which also influences the direction of worker re-sorting. In response to debt relief–and the associated reduction in human capital distortions–some households find

[43] Labor productivity refers to labor's value-added adjusted to account for occupational wages (i.e. prices). In the model, labor productivity corresponds to the lifetime *effective* hours supplied by households which, given a fixed time endowment, varies due to investments in human capital accumulation over the lifecycle. See Appendix C for additional details on the earnings growth decomposition.

Table 6: Short-Run Student Debt Policies

Policy	Group	Welfare (CE%)			Lifetime Earnings (%)			Program Cost	
		Total	Amenities	Earnings	Total	Wages	Eff Hrs	Avg. Grant	ΔTax
	Population	1.35	0.27	1.08	0.24	-0.04	0.28	$24,694	
Full	Switchers (0.59%)	76.95	68.69	8.26	-21.32	-9.69	-11.62	$35,765	1.54 pp
	Stayers	5.84	0.00	5.84	1.84	0.00	1.84	$23,078	
	Population	0.74	0.13	0.61	0.15	-0.03	0.18	$14,595	
20K	Switchers (0.39%)	38.45	47.04	-8.59	-20.29	-9.56	-10.73	$18,049	0.93 pp
	Stayers	3.55	0.00	3.55	1.17	0.00	1.17	$14,272	
	Population	0.51	0.06	0.45	0.12	-0.02	0.14	$8,993	
10K	Switchers (0.23%)	25.85	37.54	-11.69	-18.86	-9.61	-9.25	$9,827	0.58 pp
	Stayers	2.63	0.00	2.63	0.86	0.00	0.86	$8,949	

Notes: Population outcomes include all college and non-college households. Switchers and stayers refer to conditional sub-groups of the population of college graduates. Note that population estimates include both the treated and untreated population in partial equilibrium, and so need not lie between the sub-group outcomes. Details on welfare and earnings decompositions are in Appendix C. Avg. Grant is the average amount forgiven, conditional on receiving forgiveness. Values correspond to percentage changes.

it optimal to give up earnings in order to move to occupations with higher non-wage amenities. This margin of occupational reallocation explains why induced job switchers derive enormous welfare benefits from the debt relief policies despite moving to occupations that reduce their lifetime earnings on average. These amenity-driven job reallocations align with the recent findings of Rothstein and Rouse (2011) and Luo and Mongey (2024), who show that higher student debt causes graduates to accept jobs with initially higher wages, but lower job satisfaction.[44]

While policy induced occupational reallocation does not have a large impact on aggregate labor productivity and earnings, it does play an important role in determining the heterogeneous policy outcomes across occupations. Table A12 summarizes the effect by reporting changes in the labor productivity (TFPR) of each occupation alongside changes in the composition of its workforce due to reallocation. It decomposes the average change in occupational talent within each occupation into contributions from new entrants and those who exit. The results indicate substantial heterogeneity in policy outcomes across occupations. Some occupations–such as Math & CS, Protective Services, and Engineers & Scientists– experience large increases in labor productivity from

[44]Relatedly, recent work by Boar and Lashkari (2025) documents a similar trade-off between wage and non-wage amenities by showing that children from high asset households are more likely to select into jobs with more desirable non-wage characteristics.

2.62% to 3.13% under the full forgiveness policy. Other occupations – such as Cleaning & Building, Transportation, and Education – experience large reductions in productivity ranging from -3.87% to -4.24%.

Since occupational wages are held constant in the short-run policy computations, the cross-occupation productivity changes are driven primarily by worker sorting. To highlight this point, Figure 5 plots the data from Table A12. It shows a strong positive correlation between the policy-induced change in an occupation's productivity (TFPR) and changes in the composition of its workforce–as measured by occupational talent. These changes in average occupational ability are the result of worker reallocation in response to the short-run debt forgiveness policy. As the figure shows, average talent increases in most occupations which saw an increase in productivity and decreases in most occupations where productivity falls. Consistent with the theory, the figure also shows that the impact of these reallocations on labor productivity was largest in occupations where workers initially had the highest levels of student debt.

Figure 5: Policy Induced Re-allocation and Productivity Growth

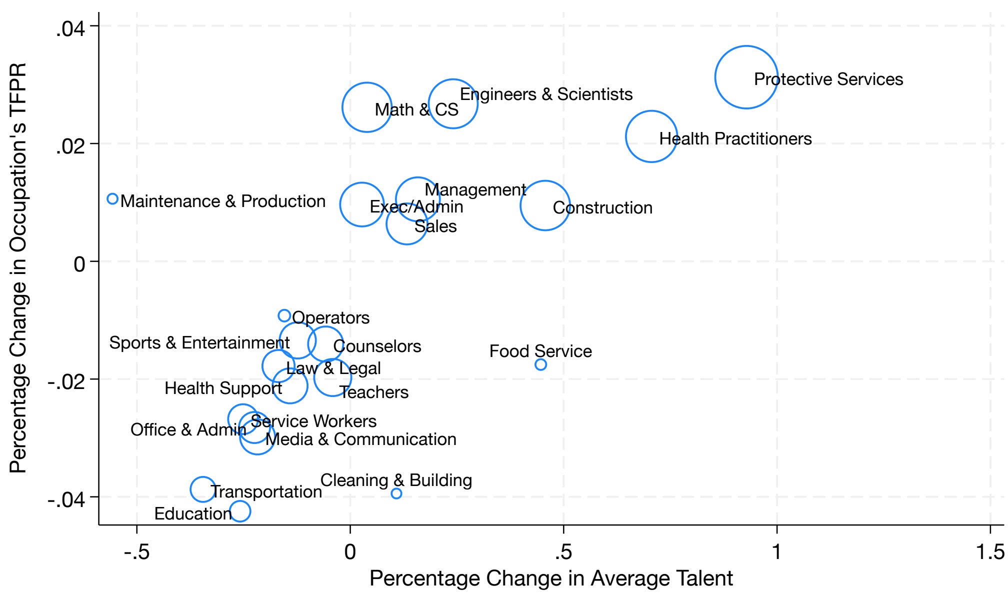

Notes: The figure displays the cross-sectional correlation between policy induced changes in an occupation's productivity (TFPR) and changes in the average talent of its workers. Increases in average talent within an occupation are indicative of a reduction in the misallocation of talent. The bubble sizes correspond to predicted student debt, as in Section 4. The changes in productivity and average talent correspond to the full-forgiveness policy reported in Table A12.

The decomposition in the latter columns of Table A12 also shows that the population of job switchers are, on average, lower ability than the job stayers in both their origin and destination occupations. This is intuitive, since the population of job switchers and job stayers is endogenous. The population which decides not to change jobs will therefore disproportionately include the unconstrained households who are already optimally matched to their best occupation. An implication of this endogeneity is that improvements in average occupational talent will be driven more by occupation exiters than entrants–as Table A12 shows. Consequently, reductions in the misallocation of talent will have a more muted effect on aggregate productivity since occupations with the largest increase in TFPR will also see their employment shares shrink (see Figure A3).

Finally, comparing the panels of Table 6 shows how the policy's impact varies with program size. The comparisons show that the macroeconomic effects of the policy are qualitatively similar across the programs. All three programs raise welfare predominantly by boosting earnings through increased human capital accumulation on the job and induce only a small reallocation of workers. While qualitatively similar, the quantitative impact of the policies appears highly non-linear in program size. One reason is that student debt burdens are continuously distributed in the population, so the average level of effective debt forgiveness (reported in the final column) does not vary one-for-one with the maximum education grant cap of each policy. For example, doubling the student debt forgiveness cap from 10k to 20k less than doubles the average level of debt relief received by households from \$8,993 to \$14,595. Similarly, the average increase in tax rates also less than doubles, rising from 0.58 to 0.93 percentage points.

Another reason for the non-linearity is the discrete nature of changes that come from occupation switching. Pushing a substantial population of workers over the job switching threshold (as illustrated in Figure 2) generally requires much larger transfers. In contrast, within occupation distortions to human capital investment respond continuously to a weakening of credit frictions. The distinction is again evident in the average size of debt relief received by those induced to switch jobs versus stayers. Under the full-forgiveness policy, the average person induced to switch jobs received \$35,765 in student debt relief, while job stayers received only \$23,078. These non-linearities therefore play an important role in determining the size and scope of occupational reallocation following policy interventions. The next section examines these effects in greater detail by analyzing the macroeconomic impact of larger, long-run debt forgiveness policies.

Long-Run Student Debt Policies. The second exercise computes the long-run consequences of making the policies permanent by offering student debt relief to both current and future generations. Given the size and scope of these policy changes, the long-run computations additionally allow occupational wages and college matriculation decisions to respond endogenously to the new policies.[45]

Table 7 displays the computational results. It reports steady state changes in the main outcome variables following the permanent implementation of each program. Relative to the short-run, it includes additional results for the policy-induced population of new college graduates (i.e., *matriculators*) which didn't exist in the previous exercises. At the aggregate level, the long-run policy outcomes are qualitatively similar to the short-run outcomes, but with larger magnitudes. Both sets of policies raise welfare predominantly by increasing lifetime earnings through more efficient human capital accumulation. In part, the larger aggregate effects are unsurprising, and follow mechanically from the greater scope of the long-run policies. Under the full-forgiveness policy, the average educational grant is $45,538, nearly double the size of the short-run, necessitating an additional 1.32 percentage point rise in marginal tax rates compared to the short-run.

While qualitatively similar, a closer examination reveals that there are notable differences between the short-run and long-run outcomes and the mechanisms which underlie them. For example, increases in lifetime earnings play a more important role in determining welfare gains in the long-run. As the decompositions in Table 7 show, earnings growth accounts for 92% of the rise in welfare under long-run policies versus 80% under the short-run policies. These differences are even starker among induced job switchers. Earnings growth accounts for 73% of the welfare gain for switchers under the long-run policy versus only 11% under the short-run policy. In other words, the long-run policies induce a reallocation of employment toward higher productivity occupations, while the short-run policies induce switching towards higher amenity jobs.

The greater gains in lifetime earnings under the long-run policies are the result of both (i) a much larger population of policy-induced job switchers and (ii) different patterns of job reallocation among them. Under the long-run full forgiveness policy, 5.63% of the population choose different occupations, compared with just 0.59% in the short-run. Moreover, while short-run job switchers experience an average *decline* in lifetime earnings, those induced to switch occupations under the long-run policy experience large *in-*

[45] Implicitly, the model holds constant any response in university pricing strategies following the change in student debt policies, which rules out the Bennett Hypothesis and related mechanisms.

Table 7: Long-Run Student Debt Policies

Policy	Group	Welfare (CE%)			Lifetime Earnings (%)			Program Cost	
		Total	Amenities	Earnings	Total	Wages	Eff Hrs	Avg. Grant	ΔTax
Full	**Population**	5.68	0.47	5.21	6.24	-0.18	6.42	$45,538	2.86 pp
	Switchers (5.63%)	53.34	14.20	39.14	27.77	-6.46	34.23	$47,724	
	Stayers	3.61	0.00	3.61	3.25	0.21	3.04	$45,362	
	Matriculators (1.29%)	23.11	12.96	10.15	152.22	-7.79	160.00	$69,580	
20K	**Population**	3.14	0.01	3.13	2.70	-0.04	2.74	$18,974	1.27 pp
	Switchers (3.58%)	36.02	2.22	33.80	22.52	-4.38	26.90	$18,854	
	Stayers	2.18	0.00	2.18	1.37	0.13	1.24	$18,980	
	Matriculators (0.54%)	15.11	13.53	1.58	128.92	-5.34	134.26	$19,597	
10K	**Population**	2.22	-0.04	2.26	1.47	-0.03	1.50	$9,945	0.67 pp
	Switchers (2.11%)	30.87	-1.44	32.31	21.78	-2.29	24.08	$9,930	
	Stayers	1.71	0.00	1.71	0.74	0.03	0.72	$9,945	
	Matriculators (0.30%)	18.73	17.21	1.53	115.28	-4.40	119.69	$9,979	

Notes: Population outcomes include all college and non-college households. Switchers and stayers refer to sub-groups of the population of college graduates. Matriculators include a new population of induced college graduates. Group percentages are with respect to the overall population. Details on welfare and earnings decompositions are contained in Appendix C. Avg. Grant is the average amount forgiven, conditional on receiving forgiveness. Values correspond to percentage changes.

creases in lifetime earnings driven by productivity growth. As Table 7 shows, long-run job switchers under the full forgiveness policy experience a 27.77% growth in lifetime earnings driven by a 34.23% increase in labor productivity. These large productivity gains, coupled with the larger population of policy induced job switchers, are the main driver of the policy differential.

The reason that similarly sized student debt relief programs deliver such different flows of worker reallocations in the short-run and long-run is due primarily to the endogenous response of occupational wages. As under the short-run policies, debt forgiveness leads some workers to switch to higher amenity occupations. In the long-run, the wages in high amenity occupations fall in response, discouraging further flows and moderating amenity driven reallocation. Figure 6 illustrates the effect by showing how, under the long-run policies, the endogenous wage declines are largest in the high amenity occupations. In contrast, worker flows into occupations that foster greater human capital accumulation lead to labor productivity increases that moderate the resulting decline in wages. These cross-occupation changes also trigger second-order labor reallocations, whereby some workers who may not have found it worthwhile to switch because of the debt forgiveness policy (such as those without student debt) may still be induced to

Figure 6: Policy Induced Wage Changes and Occupational Amenities

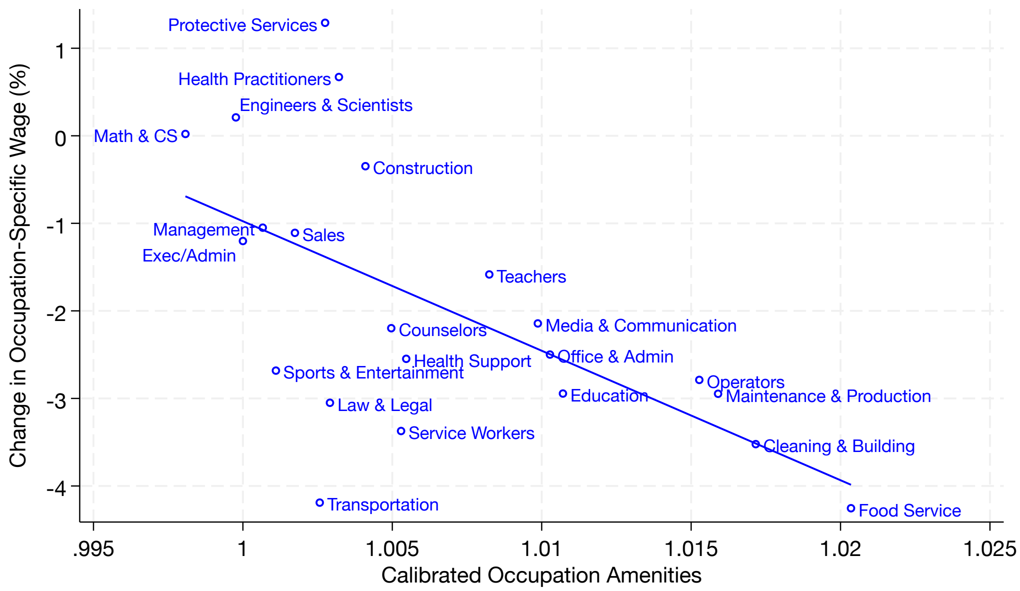

Notes: The y-axis reports changes in an occupation's wage under the long-run full student debt forgiveness policy. The x-axis reports the occupation's calibrated amenity value, re-indexed so executive administrators equal one. The corresponding occupation specific productivity and employment outcomes under the long-run full forgiveness policy are summarized in Table A13.

move by the resulting changes in occupational wages. These higher order wage effects partly explain the larger population of switchers under the long-run policy. The endogenous wage response also delivers modest wage gains to job stayers, who see their wages rise by 0.21% under the long-run full forgiveness policy – compared to no change in the short-run, where occupational wages are held fixed.

As a result of the equilibrium wage response, labor market reallocation under the long-run policy is more directed toward improving the allocation of talent rather than increasing occupational amenities. The change in flows is illustrated in Figure 7, which summarizes the reallocation of workers under the short-run and long-run policies along each of the three exogenous dimensions of occupational heterogeneity: amenities, wages (i.e. TFP), and the population endowments of occupational abilities. The top panel shows that the amenities-driven worker reallocation which characterized the short-run switchers is substantially moderated in the long-run. At the same time, the re-sorting of workers into higher talent occupations is much more pronounced. These changes in the flow of workers underlie the larger productivity gains and more modest welfare effects on job switchers under the long-run policies.

Table A13 summarizes the heterogeneous policy effects across occupations. As with the short-run policy, there is substantial variation in occupation level outcomes. While the magnitude of productivity and ability changes are much larger under the long-run policy, their distribution across occupations is similar. The cross-sectional correlation in occupational TFPR changes under the short-run and long-run policies is 0.69; the corresponding correlation of induced changes in worker ability across occupations is 0.81. The fact that the majority of occupations experience an increase in the average talent of their workforce shows that the long-run policy raises aggregate labor productivity by reducing the misallocation of talent. Decomposing flows into the effect of exiters and entrants, the results show that improvements in worker ability are again driven primarily by the exit of workers who are poorly matched – as in the short-run policies.

Finally, Table 7 also reports outcomes for the population of newly matriculated college graduates. Under the long-run policy, prospective students anticipate student debt forgiveness, which boosts enrollments by effectively reducing the cost of attending college. Surprisingly, the model predicts that the long-run policies will have only a small effect on college enrollments. Even under the full student debt forgiveness policy, only an additional 1.29% of the population is induced to enroll in college. The small response reflects the fact that college matriculation decisions in the calibrated model are driven more by the expected returns to a college education, than by the borrowing costs. These high returns, even at the margin, are evident in the large increases in lifetime earnings and productivity among the new college matriculators displayed in Table 7. Furthermore, while student debt relief always incentivizes additional matriculation, the resulting adjustments in occupational wages also disincentivize some workers who previously found it worthwhile to enroll.[46] As a result, while induced college matriculators experience some of the largest welfare and productivity gains, they have only a minor impact on aggregate outcomes given their small population size.

7 Conclusion

This paper provides a quantitative analysis of the macroeconomic consequences of rising student debt among college graduates. To do so, it develops a model of lifecycle human capital accumulation and occupation choice in the presence of credit frictions.

[46] However, the calibrated model suggests these effects are quantitatively negligible in practice, representing less than 0.05% of the college population even under the full forgiveness long-run policy.

Figure 7: Policy Induced Re-allocation of Workers by Occupational Characteristics

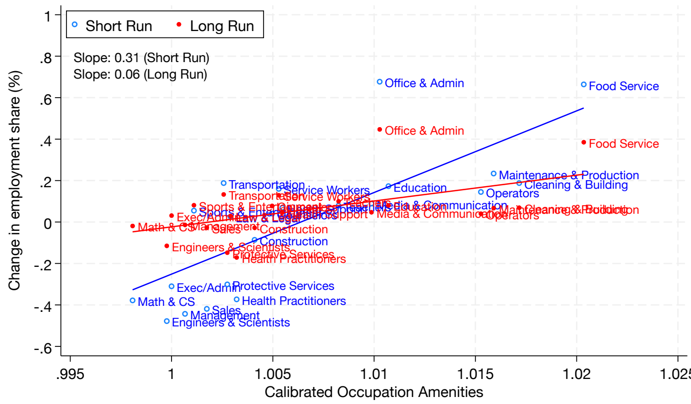

(a) Occupational Amenities

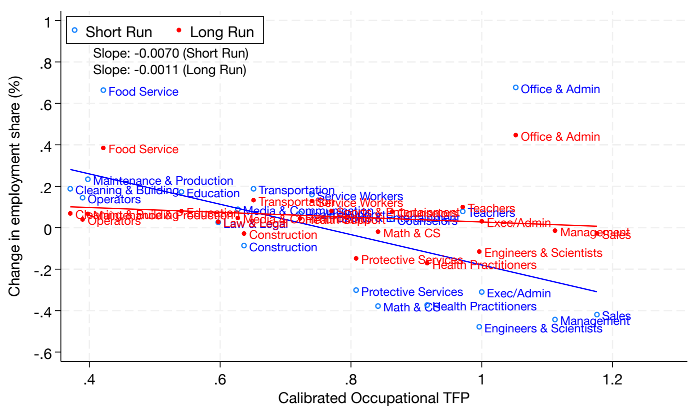

(b) Occupational TFP

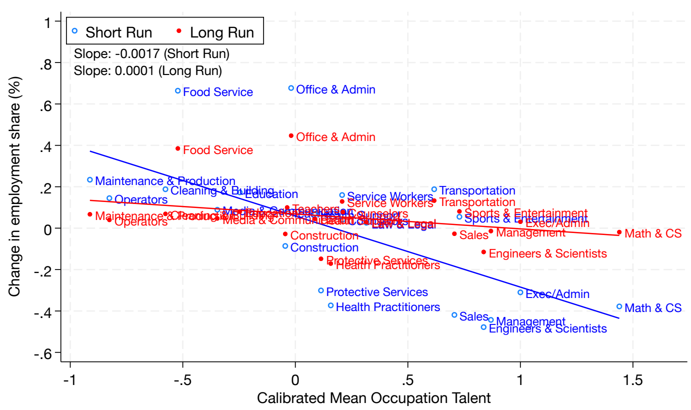

(c) Occupational Talent

Notes: Y-axis plots change in occupation employment shares (in percentage points) under short-run and long-run full student debt forgiveness policies. Each panel corresponds to one of the three dimensions on which the occupations are exogenously heterogeneous. The x-axis of each panel is normalized so that executive administrators equal one. Additional details are contained in the appendix.

When there is heterogeneity in occupational wages, the model shows how increasing student debt burdens can give rise to a misallocation of talent whereby workers sort into occupations that are not optimally matched to their skills, further inhibiting aggregate labor productivity. To quantify the effects, the calibrated model replicates both the aggregate distributional data–inclusive of student debt, earnings, assets, and occupation choice–as well as empirical evidence on the marginal effect of student debt on the early career labor market outcomes of college graduates.

The results of the computational analysis suggest that increases in the size and scope of student debt obligations in recent decades may be inhibiting the post-graduation, early career labor market outcomes of recent graduates. They show that both short-run and long-run student debt relief programs can increase labor productivity by stimulating human capital investment and improving the allocation of workers across occupations. Importantly, the model predicts meaningfully different short-run and long-run effects of student debt relief on the reallocation of labor. One-off student debt relief mostly induces workers to move toward higher amenity occupations, whereas the long-run policies induce larger flows into human capital intensive occupations – further augmenting labor productivity by reducing the misallocation of talent. Future work should assess the extent to which these gains could moderate, if not overcompensate, for the costs of replacing federal student loans with public education grants. The analysis should also be extended to more fully consider the sequential nature of educational attainment and how changes in undergraduate student debt policies can impact graduate education and its contribution to aggregate productivity and household welfare.

References

Abbott, Brant, Giovanni Gallipoli, Costas Meghir, and Giovanni L. Violante. 2019. "Education Policy and Intergenerational Transfers in Equilibrium." *Journal of Political Economy* 127 (6): 2569–2624.

Andolfatto, David, and Martin Gervais. 2006. "Human Capital Investment and Debt Constraints." *Review of Economic Dynamics* 9 (1): 52–67.

Baum, Sandy, and Marie O'Malley. 2003. "College on Credit: How Borrowers Perceive Their Education Debt." *NASFAA Journal of Student Financial Aid* 33 (12): 7–19.

Ben-Porath, Yoram. 1967. ""The Production of Human Capital and the Life Cycle of Earnings"." *Journal of Political Economy* 75 (4): 352–365.

Boar, Corina, and Danial Lashkari. 2025. "Occupational Choice and the Intergenerational Mobility of Welfare." Working paper 29381, National Bureau of Economic Research.

Caucutt, Elizabeth M, and Lance Lochner. 2020. "Early and Late Human Capital Investments, Borrowing Constraints, and the Family." *Journal of Political Economy* 128 (3): 1065–1147.

Chakrabarti, Rajashri, Vyacheslav Fos, Andres Liberman, and Constantine Yannelis. 2022. "Tuition, Debt, and Human Capital." *The Review of Financial Studies* 36 (4): 1667–1702 (09).

Chapman, Stephanie. 2015. "Student Loans and the Labor Market: Evidence from Merit Aid Programs."

Consumer Financial Protection Bureau. 2022, October. "Report of the CFPB Education Loan Ombudsman." Technical Report.

Daruich, Diego, and Julian Kozlowski. 2020. "Explaining intergenerational mobility: The role of fertility and family transfers." *Review of Economic Dynamics* 36:220–245.

De Falco, Adriano, and Yannick Reichlin. 2025. "Grants vs. loans: The Role of Financial Aid in College Major Choice." *The Economic Journal*, p. ueaf121.

Dvorkin, Maximiliano A, and Alexander Monge-Naranjo. 2019. "Occupation Mobility, Human Capital and the Aggregate Consequences of Task-Biased Innovations." *FRB St. Louis Working Paper*, no. 2019-13.

Federal Reserve Bank of New York. 2022, November. "Quarterly Report on Household Debt and Credit." Technical Report.

Field, Erica. 2009. "Educational Debt Burden and Career Choice: Evidence from a Financial Aid Experiment at NYU Law School." *American Economic Journal: Applied Economics* 1 (1): 1–21.

Folch, Marc, and Luca Mazzone. 2026. "Go Big or Buy a Home: The Impact of Student Debt on Career and Housing Choices." *Review of Economic Dynamics* 59 (None): None (January).

Fu, Chao, Hsuan-Chih Lin, and Atsuko Tanaka. 2021. "College Loans and Post-Schooling Human Capital Investment." Working paper, University of Wisconsin-Madison.

Goodman, Sarena F., Adam Isen, and Constantine Yannelis. 2018, April. "A Day Late and a Dollar Short : Liquidity and Household Formation among Student Borrowers." Finance and economics discussion series 2018-025, Board of Governors of the Federal Reserve System (U.S.).

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2021. "What Do Data on Millions of US Workers Reveal about Lifecycle Earnings Dynamics?" *Econometrica* 89 (5): 2303–2339.

Hampole, Menaka. 2024. "Financial Frictions and Human Capital Investments."

Hendricks, Lutz, and Oksana Leukhina. 2017. "How Risky is College Investment?" *Review of Economic Dynamics* 26:140–163.

———. 2018. "The Return to College: Selection and Dropout Risk." *International Economic Review* 59 (3): 1077–1102.

Herkenhoff, Kyle, Gordon Phillips, and Ethan Cohen-Cole. 2016, May. "How Credit Constraints Impact Job Finding Rates, Sorting & Aggregate Output." Working paper 22274, National Bureau of Economic Research.

Herkenhoff, Kyle, Gordon M. Phillips, and Ethan Cohen-Cole. 2021. "The Impact of Consumer Credit Access on Self-Employment and Entrepreneurship." *Journal of Financial Economics* 141 (1): 345–371.

Herkenhoff, Kyle F. 2019. "The Impact of Consumer Credit Access on Unemployment." *The Review of Economic Studies* 86 (6): 2605–2642 (02).

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow. 2019. "The Allocation of Talent and U.S. Economic Growth." *Econometrica* 87 (5): 1439–1474.

Huggett, Mark, Gustavo Ventura, and Amir Yaron. 2011. "Sources of Lifetime Inequality." *American Economic Review* 101 (7): 2923–54 (December).

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, Benjamin F. 2014. "The Human Capital Stock: A Generalized Approach." *American Economic Review* 104 (11): 3752–77 (November).

Jones, Charles I. 2019, April. "Taxing Top Incomes in a World of Ideas." Nber working papers 25725, National Bureau of Economic Research, Inc.

Krueger, Dirk, Alexander Ludwig, and Irina Popova. 2025. "Shaping Inequality and Intergenerational Persistence of Poverty: Free College or Better Schools?" *Journal of Monetary Economics* 150:103694.

Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman. 2018. "Life Cycle Wage Growth across Countries." *Journal of Political Economy* 126 (2): 797–849.

Lee, Sang Yoon (Tim), and Ananth Seshadri. 2019. "On the Intergenerational Transmission of Economic Status." *Journal of Political Economy* 127 (2): 855–921.

Lochner, Lance, and Alexander Monge-Naranjo. 2012. "Credit Constraints in Education." *Annual Review of Economics* 4 (1): 225–256.

Lochner, Lance, Todd Stinebrickner, and Utku Suleymanoglu. 2021. "Parental Support, Savings, and Student Loan Repayment." *American Economic Journal: Economic Policy* 13 (1): 329–71 (February).

Lockwood, Benjamin B., Charles G. Nathanson, and E. Glen Weyl. 2017. "Taxation and the Allocation of Talent." *Journal of Political Economy* 125 (5): 1635–1682.

Luo, Mi, and Simon Mongey. 2024. "Assets and Job Choice: Student Debt, Wages and Amenities."

Martellini, Paolo, Todd Schoellman, and Jason Sockin. 2024. "The Global Distribution of College Graduate Quality." *Journal of Political Economy* 132 (2): 434–483.

Morazzoni, Marta. 2023. "Student Debt and Entrepreneurship in the US." *Available at SSRN:* http://dx.doi.org/10.2139/ssrn.4440926.

Rothstein, Jesse, and Cecilia Elena Rouse. 2011. "Constrained after College: Student Loans and Early-Career Occupational Choices." *Journal of Public Economics* 95 (1-2): 149–163.

Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz. 2012. "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." *Journal of Economic Literature* 50 (1): 3–50 (March).

Zhang, Lei. 2013. "Effects of College Educational Debt on Graduate School Attendance and Early Career and Lifestyle Choices." *Education Economics* 21 (2): 154–175.

A Tables and Figures Appendix

Figure A1: Net Worth and Student Debt of Young Workers, Ages 22-25.

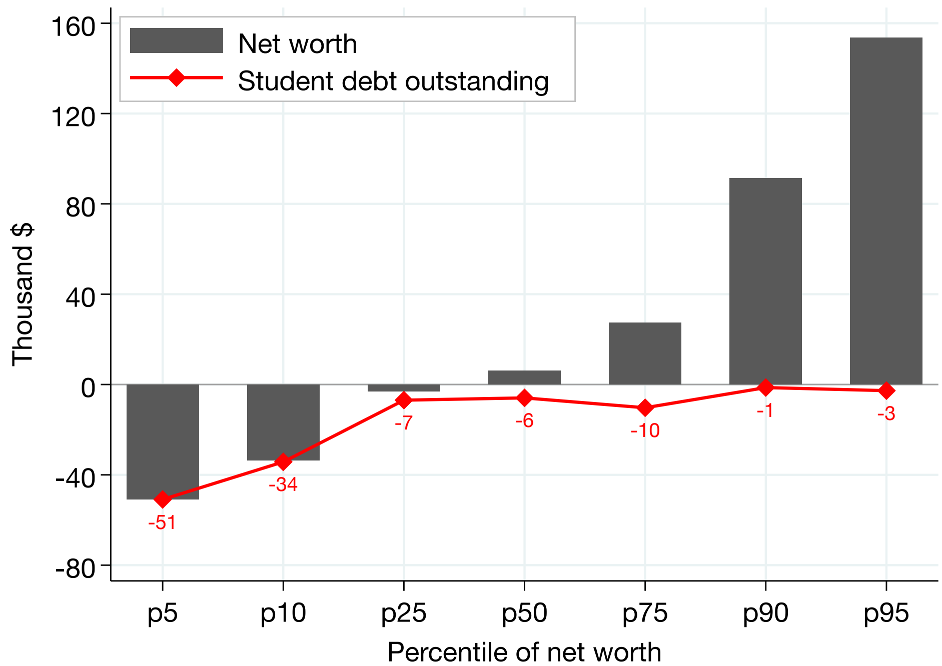

Notes: Source data from 2016 Survey of Consumer Finances. Student debt defined as the total market value of aggregate loan balance of education-purpose expenses. The sample is restricted to individuals between the ages of 22 and 25.

Figure A2: CBO Effective Tax Rates

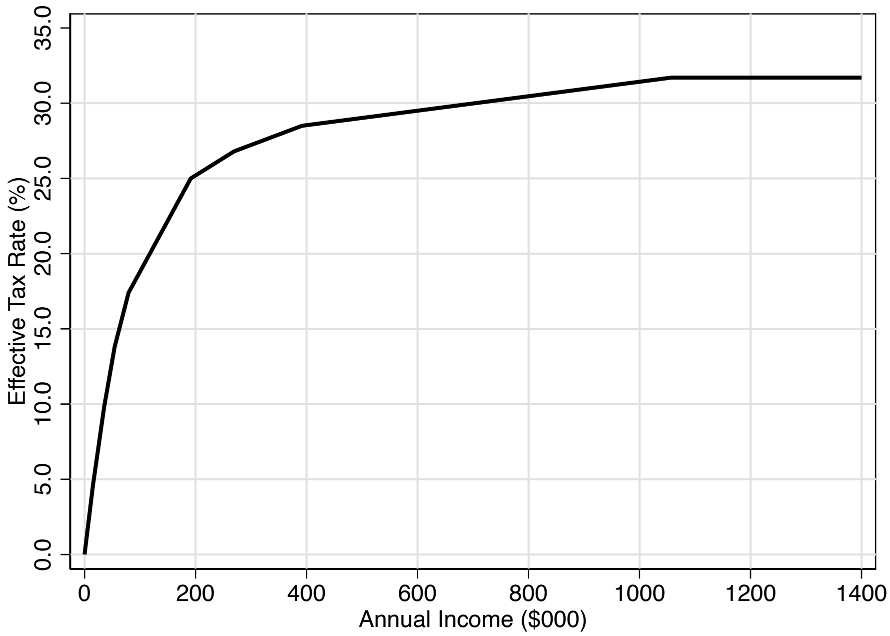

Notes: This figure plots effective tax rates by income categories. Effective tax rate includes individual income taxes, social security taxes, corporate income taxes, and excise taxes. Source: Congressional Budget Office, *Effective Federal Tax Rates, 1979–2004* (December 2006), Table 1.

Table A1: Individual Summary Statistics for NLSY 1997 Sample

	Population		Sub-Pop without Debt		Sub-Pop with Debt	
	Mean	**Std. Dev.**	**Mean**	**Std. Dev.**	**Mean**	**Std. Dev.**
Earnings (000s)	$42.8	$28.8	$46.7	$31.9	$40.6	$26.8
Student Debt (000s)	$14.3	$20.2	$0.0	$0.0	$22.1	$21.4
Percent with Student Debt	0.6	0.5	0.0	0.0	1.0	0.0
Average Age	27.0	3.3	26.9	3.3	27.1	3.3
Share Non-White	0.3	0.5	0.2	0.4	0.3	0.5
Share Female	0.5	0.5	0.5	0.5	0.6	0.5
Average Year BA	2005.3	2.3	2005.0	2.2	2005.5	2.3
Average Age BA	22.5	1.9	22.2	1.8	22.6	2.0
Initial Household Income	$72.9	$48.3	$90.9	$59.1	$63.5	$38.2
Initial Household Net Worth	$186.9	$181.9	$275.3	$205.0	$140.4	$149.2

Notes: *The table provides individual level summary statistics for the NLSY 1997 sample population and for the sub-populations with and without student debt.*

Table A2: Summary Statistics on Financing of Higher Education

	Population	**without Student Debt**	**with Student Debt**
Gross Tuition	$9,857	$9,762	$9,909
Net Tuition	$7,904	$7,802	$7,960
Student Debt	$14,296	$0	$22,362
Grants	$20,642	$13,274	$24,800
Family Loans	$1,898	$1,155	$2,317
Work Study	$865	$237	$1,220
Out-of-Pocket	$4,221	$3,865	$4,422
Other Support	$989	$1,147	$900

Notes: *The table provides summary statistics of individual financing of higher education and tuition paid for the sample population and sub-populations with and without student debt. Financing data come from the NLSY 1997 college transcript data and tuition is derived from NCES IPEDS after merging with the NLSY college-identifiers.*

Table A3: First Stage Financing Effects

	Student Debt	**Grants**	**Work Study**	**Family Loans**	**Out-of-Pocket**	**Other Support**	**Net Tuition**
Coefficient	-$3,172	$3,718	$205	$51	-$664	-$496	-$147
(se)	(1,895)	(2,352)	(216)	(885)	(677)	(669)	(136)
Observations	1358	1358	1358	1358	1358	1358	1362
R-squared	0.57	0.63	0.60	0.41	0.55	0.45	0.95

Notes: *Coefficients report the effect of a one standard deviation increase in the college grant share instrument on student debt as well as other sources of higher education financing reported in the NLSY 1997 and net tuition charged by the college. All regressions condition on college and student cohort fixed effects. Bracketed terms underneath the coefficients report robust standard errors.*

Table A4: Instrumental Variable Robustness

	Parent Income	**Parent Net Worth**	**ASVAB**	**HS GPA**	**Dropout**	**Post-Grad**	**FTFY Emp**
Coefficient	-0.89	16.02	4.13	0.00	0.01	0.00	0.00
(SE)	4.42	17.65	2.12	0.04	0.01	0.01	0.01
Obs	1342	950	1162	1068	5248	5248	5248
R2	0.52	0.67	0.69	0.70	0.39	0.42	0.39

Notes: *The coefficients report the effect of a one standard deviation increase in the college grant share on other educational determinants and characteristics which determine selection in the main sample. All regressions condition on college and student cohort fixed effects. Columns (1) - (4) report the correlation within the main estimation sample between the individual grant share instruments and the respondent's initial parental income (000s), parental net worth (000s), ASVAB score percentile, and high school GPA. Columns (5) - (7) examine the extent to which the grant share predicts selection into the estimation sample, based on educational attainment and employment status. Column (5) looks at dropout probability, defined by individuals who report having some college but no degree. Column (6) looks at probability of attaining a post-graduate degree and column (7) examines the probability of working full-time, full year. The sample for columns (5)-(7) includes all individuals 18 and older with at least some college (so they can be assigned a grant share) that could have been selected into the estimation sample.*

Table A5: Estimated Impact of Student Debt on Earnings

	Ordinary Least Squares					Instrumental Variable				
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
Log Initial Earnings (α_2)	-0.07	-0.09	-0.14	0.01	0.03	3.66	2.48	2.02	2.87	3.77
	(0.02)	(0.05)	(0.03)	(0.04)	(0.12)	(2.64)	(2.40)	(1.69)	(1.08)	(1.27)
Returns to Experience (α_3)	-0.03	-0.03	-0.02	-0.03	-0.04	-1.23	-1.02	-0.95	-0.89	-0.43
	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(1.13)	(1.01)	(0.85)	(0.29)	(0.24)
Number of Observations	6360	6360	6319	5935	6073	5985	5985	5946	5580	5985
R-Squared	0.22	0.26	0.30	0.40	0.44	8.67	23.30	6.49	5.86	2.25
Year Fixed Effects (FE)		x					x			
Demographic FE			x					x		
Occ and Occ x Experience FE				x					x	
College and College Cohort FE					x					x

Notes: *The dependent variable is log annual earnings. The coefficients report the effect of a $1000 increase in student debt. The sample includes all full-time, full-year workers with a bachelor's degree. The regression R-squared are reported for OLS results. The first stage F-statistics for student debt are reported with the IV results. Estimates correspond to partially specified models that only include a single set of covariates. Column (4) of Table 1 reports results for the fully specified model and the cumulative cases. See Appendix B for additional details.*

Table A6: Robustness with respect to Employment Status Restrictions

	(1)	(2)	(3)	(4)	(5)
Student Debt (α_2)	0.17	0.84	2.17	2.02	2.50
	(2.17)	(1.06)	(1.33)	(1.40)	(1.52)
Student Debt x Year Exp (α_3)	-0.56	-0.38	-0.66	-0.56	-0.53
	(0.25)	(0.16)	(0.22)	(0.19)	(0.18)
Number of Observations	6902	5548	5907	5672	4184
F-statistic	2.36	4.96	5.12	3.55	5.88

Notes: *This table reports the estimated coefficients of the fully specified IV model in column (4) of Table 1 when sample restrictions on employment status are removed. The benchmark sample corresponds to those who work at least 35 hours a week for at least 40 weeks per year. Column (1) removes all restrictions on weeks and hours worked. Column (2) removes restrictions on weeks worked, keeping restrictions on hours. Column (3) removes restrictions on hours worked, keeping restrictions on weeks. Column (4) includes part time workers with more than 20 hours a week. Column (5) tightens restrictions to those working at least 40 hours a week for at least 50 weeks a year.*

Table A7: Robustness with respect to Educational Attainment Restrictions

	(1)	(2)	(3)	(4)
Student Debt (α_2)	14.39	10.44	10.53	10.35
	(9.45)	(5.68)	(5.76)	(5.54)
Student Debt x Year Exp (α_3)	-0.50	-0.43	-0.42	-0.38
	(0.45)	(0.92)	(0.94)	(1.01)
Number of Observations	6977	8294	8294	8294
F-statistic	0.49	1.86	1.79	1.91

Notes: *This table reports the estimated coefficients of the fully specified IV model in column (4) of Table 1 when sample restrictions on educational attainment are removed. The benchmark sample includes only individuals whose highest degree is a BA. Column (1) removes sample restrictions on highest level of educational attainment. Column (2) includes post-graduate education labor market outcomes. Column (3) adds fixed effects for highest degree attained. Column (4) includes fixed effects for highest degree attained and its interaction with years of experience.*

Table A8: Educational Attainment and Student Debt

	Student Debt		Population Shares		
Highest Degree Attained	**Average Balance**	**% with SD**	**Total**	**with SD**	**without SD**
High school diploma (Regular 12 year program)	$13.9	18%	0.60	0.31	0.76
Associate/Junior college (AA)	$22.7	54%	0.09	0.14	0.06
Bachelor's degree (BA, BS)	$23.6	64%	0.21	0.38	0.12
Master's degree (MA, MS)	$20.1	62%	0.08	0.13	0.04
PhD	$20.3	59%	0.01	0.01	0.00
Professional degree (DDS, JD, MD)	$25.0	52%	0.01	0.02	0.01

Notes: *This table summarizes educational attainment (highest degree ever attained) and student debt in the NLSY 1997 sample. The first two columns report average student debt balances (for those with debt) and the share of individuals with positive student debt balances by educational attainment. The final three columns report the share of each educational attainment group within the sample population and within the sub-populations of those with and without student debt (i.e., columns sum to one). Note that some high school diploma holders have student debt if they enrolled in some college but dropped out before attaining a degree.*

Table A9: Summary Statistics on College Characteristics

	US Colleges FTE-weighted	**Sample Colleges Observation Weighted**	**Sample Colleges with Student Debt**	**Sample Colleges with Mult. Attendees**
Gross Tuition	8293	8964	9320	8583
Average Loan Amount	2630	3642	3702	3618
Average Grants	5973	8862	9010	8802
Grant Share	0.68	0.69	0.69	0.70
Public/Private Sector	0.44	0.64	0.61	0.74
FTE College Size (000s)	18.8	15.2	13.8	17.6
BA Degree Share of Total	0.52	0.64	0.64	0.66
FT Retention Rate	0.72	0.74	0.73	0.76
FT Faculty per 100 FTE	5.08	5.32	5.40	5.42

Notes: *The table compares college characteristics for the average US college, constructed by weighting outcomes by number of full-time equivalent students (FTE) to the in-sample college characteristics. The last two columns additionally report average college characteristics attended by in-sample respondents with student debt and for colleges that have more than one attendee within the estimation sample.*

Table A10: Comparing College Characteristics and College Fixed Effects Specifications

	College Characteristics				College Fixed Effects			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Student Debt (α_2)	2.27	2.90	3.64	7.90	1.50	0.17	10.44	12.34
	(1.24)	(1.10)	(1.57)	(5.91)	(0.50)	(2.17)	(5.68)	(6.24)
Student Debt x Year Exp (α_3)	-0.52	-0.85	-1.07	-2.14	-0.41	-0.56	-0.43	-1.59
	(0.22)	(0.24)	(0.38)	(1.72)	(0.21)	(0.25)	(0.92)	(1.11)
Number of Observations	5,531	6,881	8,272	10,573	5,548	6,902	8,294	10,599
F-stat	58.79	76.79	2.88	17.71	10.90	2.36	1.86	1.84

Notes: *The left panel reports regression results using college characteristics instead of college fixed effects. To ease comparison, the right panel reports the original results with college fixed effects. Column (1) corresponds to the main empirical specification in the final column of Table 1. Columns (2)-(4) correspond to robustness exercises. Column (2) removes restrictions on hours and weeks worked. Column (3) removes restrictions on educational attainment. Column (4) removes both restrictions on hours worked, weeks worked, and educational attainment.*

Figure A3: Policy Induced Changes in Productivity and Employment Shares

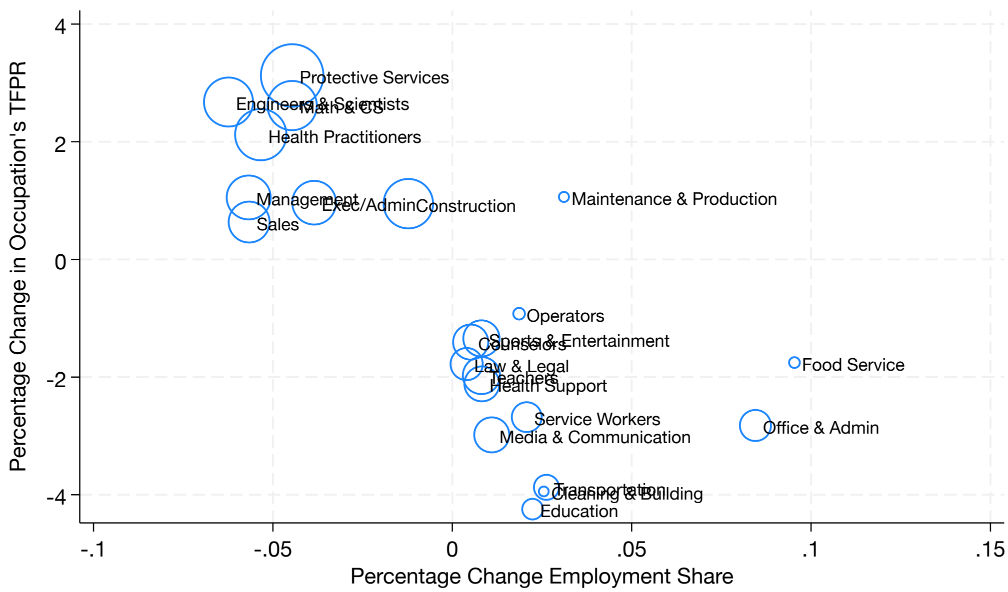

Notes: The figure displays the cross-sectional correlation between policy induced changes in an occupation's productivity (TFPR) and changes in an occupation's employment share. The bubble sizes correspond to predicted student debt, as in Section 4. The changes in productivity and employment shares correspond to the full-forgiveness policy reported in Table A12.

Table A11: Internal Calibration Targets on Occupational Heterogeneity

Occupation Group	Mean log Earnings ($)		Returns to Experience (%)		Employment Share (%)	
	Data	**Model**	**Data**	**Model**	**Data**	**Model**
Exec/Admin	9.97	9.97	13.25	13.33	7.39	7.35
Management	10.02	10.02	12.24	12.25	9.57	9.58
Math & CS	9.98	9.98	15.99	15.98	4.41	4.46
Engineers & Scientists	10.16	10.16	11.41	11.45	5.50	5.44
Counselors	9.77	9.78	10.53	10.54	4.55	4.57
Law & Legal	9.65	9.64	12.74	12.78	1.09	1.10
Teachers	9.95	9.95	7.94	7.97	9.16	9.10
Education	9.39	9.39	10.94	10.97	1.76	1.75
Sports & Entertainment	9.78	9.77	13.48	13.54	2.99	2.99
Media & Communication	10.04	10.05	6.56	6.61	2.10	2.06
Health Practitioners	10.30	10.30	6.99	6.97	3.73	3.81
Health Support	9.73	9.73	10.39	10.41	2.58	2.57
Protective Services	10.36	10.37	6.10	6.12	2.31	2.34
Food Service	9.97	9.98	1.67	1.67	5.22	5.24
Cleaning & Building	9.77	9.78	5.73	5.74	1.22	1.22
Service Workers	9.39	9.40	13.25	13.25	3.05	3.13
Sales	9.96	9.96	11.55	11.56	11.53	11.60
Office & Admin	9.87	9.87	8.09	8.11	16.28	16.14
Construction	10.21	10.21	7.15	7.14	1.15	1.14
Maintenance & Production	10.63	10.63	-1.21	-1.21	1.29	1.28
Operators	10.52	10.53	0.48	0.48	0.95	0.94
Transportation	9.23	9.23	16.90	16.94	2.17	2.18

Notes: This table summarizes the model fit of the 66 calibration targets governing occupational heterogeneity and the sorting of workers. Figure 4 plots the resulting targets and model moments to summarize the fit. Earnings and employment shares correspond to the population of BA graduates. See Data Appendix B for additional details on variable and occupation definitions.

Table A12: Occupational Productivity and Re-allocation under Short-Run Policies

	TFPR	%p Emp share	%p Entrants	%p Leavers	% Average talent	Avg. talent of Entrants Relative to Baseline	Avg. talent of Leavers Relative to Baseline
Exec/Admin	0.96%	-0.04%	0.02%	0.06%	0.03%	-9.81%	-4.21%
Management	1.05%	-0.06%	0.02%	0.07%	0.16%	-11.29%	-8.74%
Math & CS	2.62%	-0.04%	0.02%	0.06%	0.04%	-9.45%	-3.76%
Engineers & Scientists	2.68%	-0.06%	0.00%	0.07%	0.24%	-7.85%	-6.32%
Counselors	-1.41%	0.01%	0.02%	0.02%	-0.06%	-12.37%	-11.29%
Law & Legal	-1.78%	0.00%	0.01%	0.00%	-0.17%	-9.33%	-4.97%
Teachers	-1.97%	0.01%	0.04%	0.03%	-0.04%	-12.76%	-12.34%
Education	-4.24%	0.02%	0.03%	0.00%	-0.26%	-7.41%	-14.65%
Sports & Entertainment	-1.34%	0.01%	0.02%	0.01%	-0.12%	-9.99%	-7.44%
Media & Communication	-2.98%	0.01%	0.02%	0.01%	-0.22%	-12.00%	-10.69%
Health Practitioners	2.12%	-0.05%	0.00%	0.05%	0.71%	-4.02%	-14.48%
Health Support	-2.12%	0.01%	0.02%	0.01%	-0.14%	-10.12%	-6.42%
Protective Services	3.13%	-0.04%	0.00%	0.05%	0.93%	-5.09%	-13.86%
Food Service	-1.75%	0.10%	0.10%	0.00%	0.45%	8.35%	34.58%
Cleaning & Building	-3.94%	0.03%	0.03%	0.00%	0.11%	0.04%	-13.48%
Service Workers	-2.68%	0.02%	0.03%	0.01%	-0.25%	-11.00%	-8.44%
Sales	0.64%	-0.06%	0.02%	0.08%	0.13%	-13.03%	-9.44%
Office & Admin	-2.82%	0.08%	0.11%	0.02%	-0.22%	-14.14%	-16.74%
Construction	0.95%	-0.01%	0.00%	0.01%	0.46%	-7.14%	-11.96%
Maintenance & Production	1.06%	0.03%	0.04%	0.01%	-0.56%	-9.24%	-13.96%
Operators	-0.92%	0.02%	0.03%	0.01%	-0.15%	-6.22%	-16.97%
Transportation	-3.87%	0.03%	0.03%	0.01%	-0.34%	-8.83%	-7.71%

Notes: Outcomes correspond to short run full forgiveness policy.

Table A13: Occupational Productivity and Re-allocation under Long-Run Policies

	TFPR	%p Emp share	%p Entrants	%p Leavers	% Average talent	Avg. talent of Entrants Relative to Baseline	Avg. talent of Leavers Relative to Baseline
Exec/Admin	7.33%	0.03%	0.36%	0.33%	-0.45%	-7.32%	-4.86%
Management	8.82%	-0.01%	0.39%	0.41%	0.10%	-8.13%	-8.59%
Math & CS	6.35%	-0.02%	0.34%	0.36%	-1.55%	-9.66%	-3.38%
Engineers & Scientists	11.98%	-0.11%	0.23%	0.35%	0.36%	-6.24%	-5.79%
Counselors	6.26%	0.08%	0.21%	0.13%	0.16%	-6.99%	-12.92%
Law & Legal	6.10%	0.03%	0.08%	0.05%	-0.20%	-5.03%	-6.50%
Teachers	6.26%	0.10%	0.35%	0.25%	0.25%	-7.74%	-13.78%
Education	-0.11%	0.08%	0.16%	0.08%	0.32%	-5.38%	-13.50%
Sports & Entertainment	4.61%	0.08%	0.19%	0.11%	-0.30%	-6.72%	-8.90%
Media & Communication	4.28%	0.05%	0.12%	0.07%	-0.09%	-8.64%	-13.37%
Health Practitioners	20.27%	-0.17%	0.13%	0.30%	4.27%	-3.22%	-15.39%
Health Support	7.28%	0.04%	0.14%	0.09%	0.33%	-4.99%	-10.52%
Protective Services	27.34%	-0.15%	0.08%	0.23%	5.93%	-0.62%	-14.81%
Food Service	-3.55%	0.38%	0.51%	0.12%	2.07%	9.44%	5.55%
Cleaning & Building	-1.18%	0.07%	0.16%	0.09%	2.04%	1.17%	-7.68%
Service Workers	2.53%	0.13%	0.22%	0.09%	-0.67%	-7.68%	-10.54%
Sales	9.30%	-0.03%	0.42%	0.45%	0.49%	-7.65%	-11.05%
Office & Admin	3.85%	0.45%	0.84%	0.40%	-0.21%	-9.53%	-17.49%
Construction	15.27%	-0.03%	0.05%	0.08%	2.36%	-3.81%	-12.18%
Maintenance & Production	-2.28%	0.07%	0.23%	0.16%	-1.45%	-9.62%	-9.49%
Operators	0.85%	0.04%	0.14%	0.10%	3.14%	-3.05%	-14.69%
Transportation	-0.50%	0.13%	0.22%	0.09%	-0.83%	-7.16%	-10.43%

Notes: Outcomes correspond to long run full forgiveness policy.

B Data Appendix

This section describes the data sets used in the analysis and provides additional details on variable construction and sample definitions. The main datasets used are the Survey of Consumer Finances (SCF), National Longitudinal Survey of Youth (NLSY 1997), and the National Center for Education Statistics (NCES) Integrated Postsecondary Education Data System (IPEDS).

B.1 Joint Distribution of Student Debt and Initial Assets in the Population

The initial distributions of student debt and initial assets, displayed in Figures A1 and 1, are derived from the 2007 sample extract of the Survey of Consumer Finances (SCF). Initial assets correspond to the total net worth reported by household heads.[47] The sample includes all household heads between the ages of 22 and 25, who are not married or living with a partner, and who do not have any children. The base year of 2007 is chosen for two reasons. First, it avoids the effect of the Great Recession on the household asset distribution. Second, it corresponds to the sample cohort in the National Longitudinal Survey of Youth (NLSY 1997) that is the focus of the analysis. The NLSY 1997 includes individuals who were between the ages of 12 and 16 in 1996, who would be roughly between the ages of 22 and 26 in 2007. Moreover, 2007 falls after the graduation year of the majority of respondents in the NLSY 1997, who report an average graduation year between 2005-2006. These moments also provide the population joint distribution of assets and student debt targeted in the model calibration.

B.2 Occupation and Industry Classification

Occupation and industry classifications correspond to the 2002 Census codes, as reported in the NLSY 1997. For both employee and self-employed jobs, respondents' verbatim descriptors of their occupations are coded using the three-digit Census code frame. Freelance jobs that do not qualify as self-employment are coded according to the type of work performed.[48] According to the 2002 Census occupation codes, there are 15 and 31 distinct three-digit non-military industries and occupations, respectively. For occupations, we drop agricultural occupations and combine several adjacent occupation codes when the sample size is too small in order to increase precision. The result converts the 31 occupation codes to a slightly more aggregated 22 occupation code system. Tables A14 and A15 summarize the industry and occupation classifications and the occupational grouping procedure.

[47] To see how the Survey of Consumer Finances defines net worth, assets, and student loans, see `https://www.federalreserve.gov/econres/files/Networth%20Flowchart.pdf`

[48] For details, see `https://www.nlsinfo.org/content/cohorts/nlsy97/other-documentation/codebook-supplement/attachment-1-census-industrial`

Table A14: Occupational Classification and Grouping

Census Occupation Code (N=31)	**Grouped Occupations (N=22)**
Executive, Administrative, and Managerial Occupations	Executive, Administrative, and Managerial Occupations
Management Related Occupations	Management Related Occupations
Mathematical and Computer Scientists	Mathematical and Computer Scientists
Engineers, Architects, and Surveyors	Engineers, Architects, and Other Scientists
Engineering and Related Technicians	
Physical Scientists	
Social Scientists and Related Workers	
Life, Physical, and Social Science Technicians	
Counselors, Social, and Religious Workers	Counselors, Social, and Religious Workers
Lawyers, Judges, and Legal Support Workers	Lawyers, Judges, and Legal Support Workers
Teachers	Teachers
Education, Training, and Library Workers	Education, Training, and Library Workers
Entertainers and Performers, Sports and Related Workers	Entertainers and Performers, Sports and Related Workers
Media and Communication Workers	Media and Communication Workers
Health Diagnosing and Treating Practitioners	Health Diagnosing and Treating Practitioners
Health Care Technical and Support Occupations	Health Care Technical and Support Occupations
Protective Service Occupations	Protective Service Occupations
Cleaning and Building Service Occupations	Building and Cleaning Services
Entertainment Attendants and Related Workers	
Funeral Related Occupations	Service Workers
Personal Care and Service Workers	
Sales and Related Workers	Sales and Related Workers
Office and Administrative Support Workers	Office and Administrative Support Workers
Farming, Fishing, and Forestry Occupations	–
Construction Trades and Extraction Workers	Construction Trades and Extraction
Installation, Maintenance, and Repair Workers	Installation, Maintenance, and Repair
Production and Operating Workers	
Setters, Operators, and Tenders	Operators and Tenders
Transportation and Material Moving Workers	Transportation and Material Moving Workers
Food Preparation and Serving Related Occupations	Food Preparation and Serving Related Occupations
Food Preparation Occupations	

B.3 Occupation IE, RTE, Employment share

Occupation-specific earnings, returns to experience, and employment shares are computed using the longitudinal data in the NLSY 1997. The benchmark sample includes individuals above age 18, whose highest educational attainment is a 4-year BA degree, and who work full-time, full year–defined as those who work at least 35 hours a week, for at least 40 weeks a year. The robustness exercises loosen the restrictions on educational attainment, hours, and weeks worked.

The occupation-specific initial earnings (IE) and average returns to experience (RTE) calibration targets are estimated by regressing log earnings on years of experience and a set of controls. The regressions are estimated separately for each occupation. The earnings variable is defined as

Table A15: Industry Classification

Census Industry Code
Agriculture, Forestry, Fishing and Hunting
Mining
Utilities
Construction
Manufacturing
Wholesale Trade
Retail Trade
Transportation and Warehousing
Information and Communications
Finance, Insurance, Real Estate, and Rental and Leasing
Professional, Scientific, Management, Administrative, and Waste Management Services
Educational, Health and Social Services
Arts, Entertainment, Recreation, Accommodations, and Food Services
Other Services (Except Public Administration)
Public Administration

the log of total annual income, converted into the real values using the consumer price index as a deflator. Years of experience is defined as the number of years since the individual obtained their BA degree, defined as zero in the initial year after graduation. The control variables include dummies for race, gender, geographic region, year and industry fixed effects. Initial earnings are defined as the predicted earnings with zero years of experience (i.e., the starting period) in each occupation. Annualized returns to experience are defined as the annual increase in earnings over each individual's working lifecycle. Table A11 reports the resulting estimated coefficients and employment shares for each occupation, which are also used to construct Figures 3 and 4.

The IE and RTE for high school graduates are computed in a similar fashion using the same regression specification pooling high school graduates into a single occupation. The sample is restricted to those who reported their highest degree as a high school diploma or equivalent (e.g., GED). We drop individuals who report hourly earnings less than the federal minimum wage in 2007, which was \$5.85. Specifically, as with the college sample, the high school sample includes the full-time, full year employed population who work at least 35 hours a week and for at least 40 weeks a year with annual earnings greater than \$8,190 (\$5.85 $\times$ 35 hours $\times$ 40 weeks).

B.4 Individual Student Debt and Grant Share Instrumental Variables

The instrumental variable regressions in Section 4 require computing two additional variables corresponding to individual level student debt at graduation and the institutional grant share corresponding to their college and attendance years. Individual student debt at the moment of graduation is computed from the NLSY 1997 transcript study by cumulating the student loans an individual received in each semester in the first roster college they report being enrolled in. The exact questionnaire in the NLSY97 reads as *"Other than assistance you received from relatives and friends, how much did you borrow in government-subsidized loans or other types of loans while you attended this school/institution this term?"* Hence, our student debt variable encompasses all formal loans taken both from the federal/government programs and private credit market. The transcript data is also used to construct the composition of higher education financing each individual receives throughout their time in college, which are summarized in Table A2, and also used in the first-stage robustness regressions reported in Table A3.

Annual information on college grant and loan funding, along with other college characteristics, is derived from the Integrated Postsecondary Education Data System (IPEDS) of the National Center for Education Statistics (NCES). These college-level data are merged to the NLSY 1997 using the confidential college identifiers in the restricted-use version of the dataset provided by the BLS. The grant share IVs assigned to each individual are defined as the average share of grant funding, out of all grant and loan funding, issued by their college during the years they were enrolled as students.

C Computational Appendix

The following sections provide additional computational details on the counterfactual exercises in Section 6 and the results reported in Tables 6 and 7.

C.1 Computation of Consumption-Equivalent Welfare Gains

Variables with superscript b denote the variables in the baseline calibrated economy at the steady state, and c in the counterfactual economy under the debt forgiveness policies. For example, c_t^b denotes the simulated time-t consumption in the baseline economy's steady state.

Lifetime utility can be decomposed into the utility from lifetime consumption and the occupation-specific amenity values. For instance, the welfare of an individual in the baseline and counterfactual economies can be expressed as:

$$\mathbb{W}_{\mathrm{b}} = \varepsilon_p^{\mathrm{b}} \sum_{t=1}^{T} \beta^{t-1} u(c_t^{\mathrm{b}})$$

$$\mathbb{W}_{\mathrm{c}} = \underbrace{\varepsilon_p^{\mathrm{c}}}_{\substack{\text{occ-specific} \\ \text{amenity}}} \underbrace{\sum_{t=1}^{T} \beta^{t-1} u(c_t^{\mathrm{c}})}_{\substack{\text{NPV utility from} \\ \text{lifetime consumption}}}$$

For those who do not switch occupations under a forgiveness policy (job stayers), $\varepsilon_p^b = \varepsilon_p^c$, so the entire welfare change for this population comes from changes in lifetime consumption. For those who switch occupations (switchers), there is an additional welfare change which comes from changes in occupation-specific amenities, as $\varepsilon_p^b \neq \varepsilon_p^c$.

Given these formulas, consumption-equivalent welfare gains λ for each individual can be defined as follows:

$$\varepsilon_p^{\mathrm{b}} \sum_{t=1}^{T} \beta^{t-1} u\left(c_t^{\mathrm{b}}(1+\lambda)\right) = \varepsilon_p^{\mathrm{c}} \sum_{t=1}^{T} \beta^{t-1} u(c_t^{\mathrm{c}})$$

The population average consumption-equivalent welfare gains Λ reported in the counterfactual tables are therefore computed by averaging over individuals so that:

$$\Lambda \equiv \frac{1}{N} \sum_i \lambda_i$$

C.2 Decomposition of Lifetime Earnings Gain

Earnings in a given period are defined as $w(1-s)h$, where w is the occupation wage and $(1-s)h$ is the effective hours (or efficiency units) the individual devotes to work. Lifetime earnings (LE) is defined as the discounted sum of per period earnings:

$$LE = \sum_{t=1}^{T} \left(\frac{1}{1+r} \right)^{t-1} w(1-s_t)h_t$$

Following the notation above, let LE^b be the lifetime earnings of an individual in the baseline calibrated economy in steady state and LE^c be the lifetime earnings under the counterfactual debt forgiveness policy. Then the change in lifetime earnings can be decomposed as follows:

$$\begin{aligned}
\Delta LE &\equiv \log LE^c - \log LE^b \\
&= \underbrace{\log w^c - \log w^b}_{\text{change in wage}} + \underbrace{\log \left[\sum_{t=1}^{T} \left(\frac{1}{1+r} \right)^{t-1} (1-s_t^c)h_t^c \right] - \log \left[\sum_{t=1}^{T} \left(\frac{1}{1+r} \right)^{t-1} (1-s_t^b)h_t^b \right]}_{\text{change in lifetime effective hours}}
\end{aligned}$$

In the short-run debt forgiveness policy, lifetime earnings gains come solely from the change in effective hours (efficiency units) for job stayers, as their wage rates w do not change. In the long-run debt forgiveness counterfactual, or for job switchers in the short-run debt forgiveness counterfactual, $w^c \neq w^b$ due to the general equilibrium effects on occupation wages or job switching. As a result, changes in lifetime earnings reflect the combined effect of changes in wages and effective labor supply.

ONLINE APPENDIX

D Extension: Occupation Switching in the Baseline Model

The baseline model restricts college graduates to choose an occupation only once, upon entering the labor market. While analytically and computationally convenient, this modeling assumption eliminates a potentially important channel of adjustment: dynamic occupation switching. If individuals are able to switch occupations during or after the student debt repayment period, they may be able to mitigate the misallocation of talent which arises from student debt. For example, highly indebted individuals may initially select into high-wage occupations to accelerate debt repayment and later move to occupations in which their abilities are a better match. The omission of this dynamic occupation switching channel in the baseline model suggests the results may be an upper bound on the true aggregate impact of student debt.

A complete accounting of the role of endogenous occupation switching would require substantial extensions to the model, both theoretically and computationally. The extent to which this type of mitigating endogenous occupation switching occurs would depend on many additional factors not specified in the current model, such as switching costs, labor market search frictions, and the bilateral portability of human capital across occupations. Even with these model extensions, solving such a model could be prohibitively costly in terms of computational time. The baseline model requires solving 22 occupational paths for each household type (a, h, d, Θ). In contrast, the model with endogenous switching would require solving $22^8 \approx 54.88$ billion possible occupational paths for each household type. Calibrating such a model would also require a richer set of data; the current data sources employed in the analysis are not nearly large enough to observe even a small fraction of the set of possible paths individuals could follow in the initial equilibrium or the counterfactual exercises.

While a complete treatment of endogenous switching therefore seems infeasible within the confines of this paper, we nevertheless consider a partial extension that allows limited switching during the student debt repayment period to consider the nature of occupation switching that might emerge and how it impacts the paper's main results. Specifically, we extend the model to allow individuals to switch occupations once, in the second period of life after labor market entry. Because each period corresponds to five years, this timing implies that switching occurs after individuals have partially repaid their student debt, but before repayment is complete. Given that human capital investment is most intensive in the first period, permitting switching in the second period offers a meaningful opportunity to realign workers with their comparative advantage and thereby attenuate the misallocation of talent generated by credit frictions.

Figure A4: Patterns of Planned Switchers on Talent and Wage Dimensions

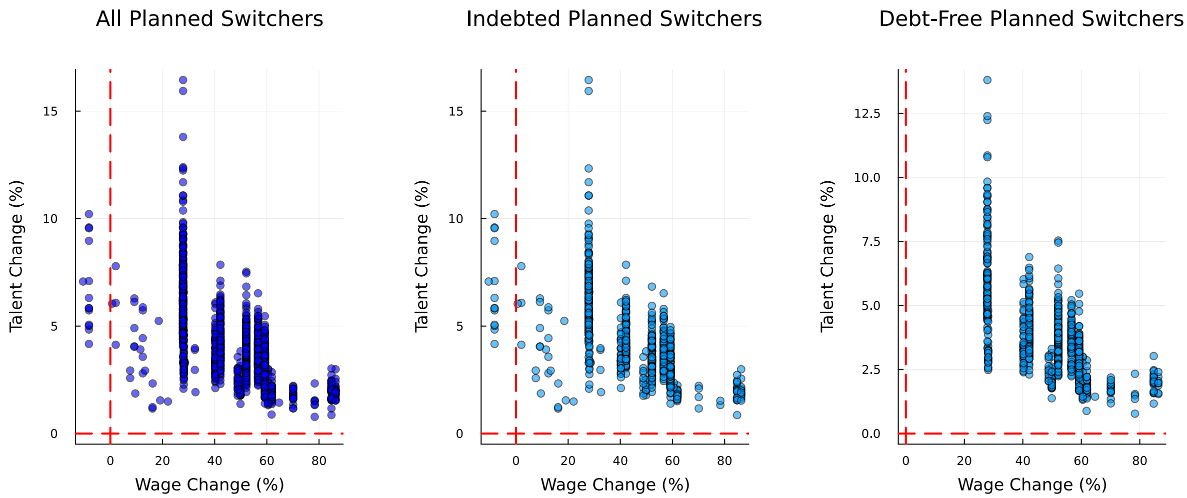

This simple extension allowing one opportunity to endogenously switch occupations midway through the student debt repayment period increases the number of occupational paths that must be computed for each individual (a, h, d, Θ) from 22 to $22^2 = 484$. In addition, the extended model continues to allow individuals to switch occupations at any stage of life t following student debt forgiveness in the counterfactuals. For tractability, rather than attempting to specify a complete bilateral matrix of human capital portability, we assume endogenous occupation switchers are subject to the same loss of accumulated human capital as in the benchmark counterfactuals. Furthermore, with endogenous occupation switching, individuals choose not only their initial occupation but also an occupational *path*: the occupation they enter and, if optimal, the occupation into which they will switch later in life. To distinguish these endogenous switchers from the job switchers induced by policy changes in the counterfactual exercises, we refer to them as "planned switchers", which are individuals who intentionally enter one occupation with the intention of moving to another later in their lifecycle.

Figures A4 through A6 document the direction of planned endogenous switching across the wage, talent, and amenity characteristics of occupations. The figures demonstrate that, as with the induced switchers in the baseline model, the direction of occupational switching reflects a mix of amenity-driven and productivity-driven reallocation. Despite these aggregate similarities, additional distinctions exist wherein planned switchers appear to trade-off amenities rather than wages when initially choosing occupations less well matched to their abilities. To see this, note that while all planned switchers move toward occupations that are a better match for their talents (Figure A6), both indebted and non-indebted individuals also move toward occupations

Figure A5: Amenity difference between the 2nd and the 1st occupation

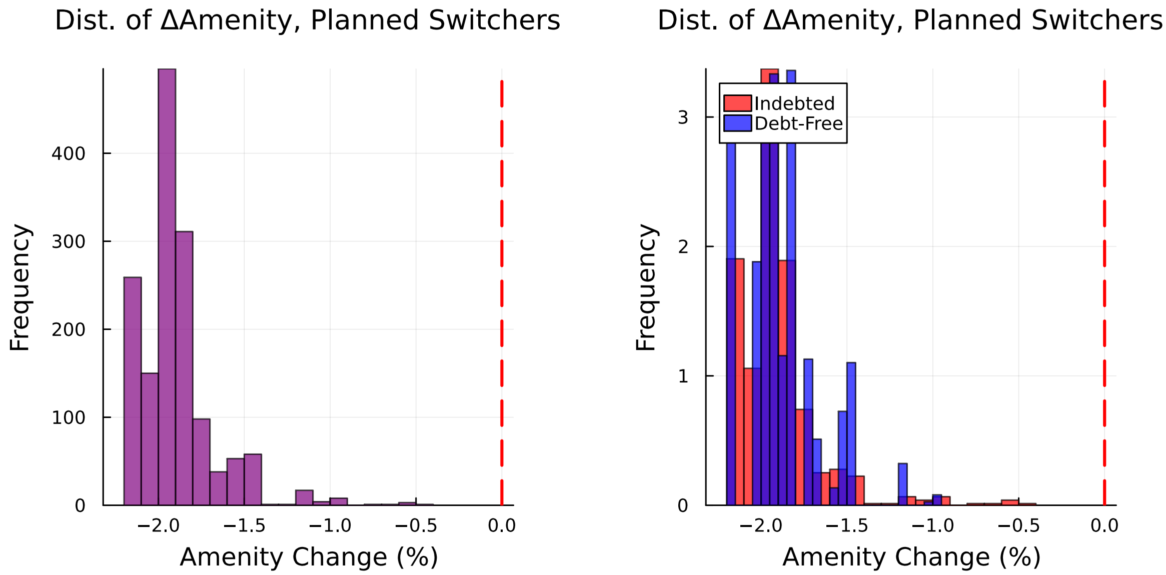

Figure A6: Talent difference between the 2nd and the 1st occupation

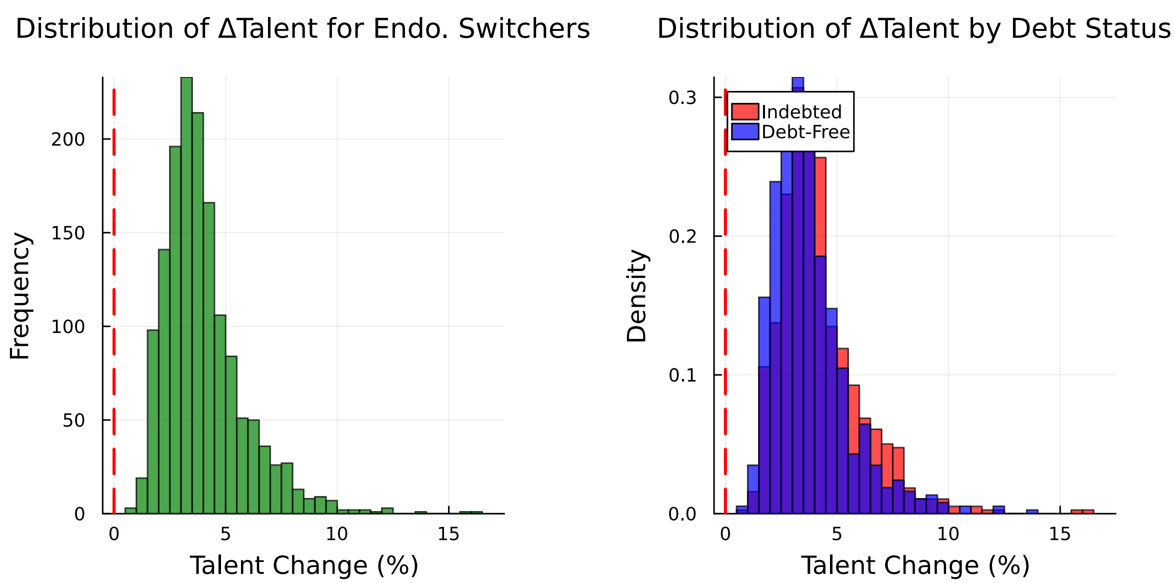

Table A16: Fraction of Planned Switchers among College Graduates

	Endo. Switching	No Amenity-Driven Switching Allowed	100% HC Loss Upon Switching
College Graduates	49.82%	10.06%	48.72%
Debt-Free College Graduates	60.69%	10.44%	54.00%
Indebted College Graduates	42.35%	9.80%	45.10%
Indebted (>100K) College Graduates	11.11%	9.26%	11.11%

Notes: This table reports the fraction of graduates who plan to switch occupations under three model specifications: (1) the baseline model with endogenous switching, (2) a model that shuts down the amenity channel for switching, and (3) a model which imposes a 100% human capital loss upon switching.

with *higher* occupational wages (Figure A4). This runs somewhat contrary to the pattern in the baseline model where indebted individuals start in high-wage occupations to pay off debt and then move to lower-wage occupations that better fit their abilities. Instead, the results suggest that planned switchers start in high amenity (rather than high wage) occupations and switch to occupations with lower amenities (Figure A5) but greater scope for human capital accumulation.

To highlight these distinctions more clearly, Table A16 summarizes the degree of planned occupation switching among indebted and non-indebted students and the extent to which it is driven by amenity versus human capital considerations. The results show that allowing endogenous switching generates a substantial amount of mobility, with nearly 50% of college graduates in the baseline economy choosing to become planned switchers. Moreover, the first column shows that the extent of planned switching is actually considerably more pronounced among individuals *without* student debt (60.69%) than among those with student debt (42.35%). The most indebted students, those with 100k or more in student debt, are the least likely to pursue planned career switches (11.11%). The fact that endogenous occupation switching is substantially more likely to be undertaken by those without student debt already suggests that its driving force is not necessarily the mitigation of the misallocation of talent.

Building on this observation, the remaining columns of Table A16 consider additional versions of the endogenous switching model which shut down the amenity value of switching occupations (column 2) and the portability of human capital (column 3). Comparing the degree of endogenous switching across the restricted models highlights how planned switches are driven primarily by amenities rather than human capital. When the amenity value of switching is turned off (column 2), the switching rates fall dramatically and there is little difference between indebted and debt-free individuals. In contrast, shutting off the portability of human capital (column 3) has little impact on either the scope or direction of occupational reallocation.

Taken together, the results suggest that the inclusion of endogenous switching primarily results in planned switchers who pursue high amenity occupations early in life, when they have little human capital, and then later switch to occupations with higher wages and more scope for human capital accumulation. Unsurprisingly, individuals without student debt are substantially more likely to pursue these new occupational paths than indebted students who mostly continue to pursue (and stay) in high wage occupations to meet their student debt repayment obligations. In other words, the inclusion of an endogenous switching channel is therefore more likely to benefit students *without* debt than those with debt. Consequently, relative to the baseline model, the extension introduces even larger negative welfare consequences of student debt as it also limits indebted graduates from taking advantage of these additional amenity driven occupational paths. As a result, the inclusion of endogenous occupation switching actually amplifies rather than dampens the impact of student debt. Table A17 illustrates this point by re-estimating the marginal effects of student debt using the simulated IV regressions in the model with endogenous switching and compares the result to the baseline model. The results show that the magnitude of both estimated coefficients is larger in the extended model, signaling an amplification of the role of student debt when endogenous occupation switching is permitted.

Table A17: Marginal Effects of Student Debt, with Endogenous Switching

Target for Indirect Inference	**Coefficient**	**Data**	**Model** (Baseline)	**Model** (Switching)
IV Effect on Initial Earnings (IE)	$\hat{\alpha}_2$	1.50%	1.58%	1.95%
IV Effect on Returns to Experience (RTE)	$\hat{\alpha}_3$	-0.41%	-0.31%	-0.36%

For a more complete picture of how these extensions influence the paper's main results, the following sections recompute the full-forgiveness counterfactual outcomes in the extended model to compare with the baseline. Doing so first requires some additional clarifications as the inclusion of endogenous switching complicates the decomposition strategies and the identification of induced job switchers. The following subsections clarify these issues before presenting the updated counterfactual results in the augmented model with endogenous occupation switching.

D.1 Calculation of Consumption-Equivalent Welfare Gain

When occupation switching is allowed, the welfare calculation must account for the possibility that an individual's occupation-specific amenity changes between the first and second period. A job switcher is now defined as any individual whose occupational *path* changes, i.e., who chooses a different occupation in at least one period. Accordingly, lifetime utility is decomposed as:

$$\mathbb{W}_{\text{b}} = \varepsilon_1^b u(c_1^b) + \varepsilon_2^{\text{b}} \sum_{t=2}^{T} \beta^{t-1} u(c_t^{\text{b}})$$

$$\mathbb{W}_{\text{c}} = \varepsilon_1^c u(c_1^c) + \underbrace{\varepsilon_2^{\text{c}}}_{\substack{\text{occ-specific}\\ \text{amenity}\\ \text{in second occ.}}} \underbrace{\sum_{t=2}^{T} \beta^{t-1} u(c_t^{\text{c}})}_{\substack{\text{NPV utility from}\\ \text{lifetime cons.}\\ \text{from } t=2 \text{ to } T}}$$

In other words, for those who do not switch occupations under a forgiveness policy (stayers), $\varepsilon_1^b = \varepsilon_2^b = \varepsilon_1^c = \varepsilon_2^c$, so the entire welfare change comes from changes in lifetime consumption. On the other hand, those who change occupations (induced switchers) incur additional welfare changes that come from the change in occupation-specific amenities.

Given this, consumption-equivalent welfare gain λ for each individual can be defined as follows,

$$\varepsilon_1^b u\left(c_1^b(1+\lambda)\right) + \varepsilon_2^b \sum_{t=2}^{T} \beta^{t-1} u\left(c_t^b(1+\lambda)\right) = \varepsilon_1^c u(c_1^c) + \varepsilon_2^{\text{c}} \sum_{t=2}^{T} \beta^{t-1} u(c_t^{\text{c}})$$

where average welfare gains in the population is the average of individual λs:

$$\Lambda = \frac{1}{N} \sum_i \lambda_i$$

D.2 Decomposition of Lifetime Earnings Gain

Note that earnings in a given period are defined as $w(1-s)h$, where w is the occupation wage and $(1-s)h$ is the effective hours (or efficiency units) the individual devotes to work. Lifetime earnings is defined as the discounted sum of period earnings:

$$LE = \sum_{t=1}^{T} \left(\frac{1}{1+r}\right)^{t-1} w_t(1-s_t)h_t$$

Following the notation above, let LE^b be the lifetime earnings of an individual in the baseline (no forgiveness) economy at the steady state, and LE^c be the lifetime earnings under the debt forgiveness policy. Then the change in lifetime earnings can be decomposed as follows:

$$\Delta LE \equiv \log LE^c - \log LE^b$$

$$
\begin{aligned}
= \quad & \underbrace{\log w_1^c - \log w_1^b}_{\text{change in wage in period 1}} \quad + \quad \underbrace{\log w_2^c - \log w_2^b}_{\text{change in wage in period 2}} \\
+ \quad & \underbrace{\log(1 - s_1^c)h_1^c - \log(1 - s_1^b)h_1^b}_{\text{change in effective hours in period 1}} \quad + \\
+ \quad & \underbrace{\log\left[\sum_{t=2}^{T}\left(\frac{1}{1+r}\right)^{t-1}(1 - s_t^c)h_t^c\right] - \log\left[\sum_{t=2}^{T}\left(\frac{1}{1+r}\right)^{t-1}(1 - s_t^b)h_t^b\right]}_{\text{change in effective hours from period 2 to } T}
\end{aligned}
$$

In the short-run debt forgiveness policy, lifetime earnings gains come solely from the change in effective hours (efficiency units) for job stayers, as their wages do not change. In the long-run debt forgiveness counterfactual, or for job switchers in the short-run debt forgiveness counterfactual, $w^c \neq w^b$ due to the general equilibrium effects on occupation wages or job switching. As a result, changes in lifetime earnings reflect the combined effect of both channels.

D.3 Student Debt Counterfactuals with Endogenous Occupation Switching

Finally, this section reports the results of the student debt forgiveness counterfactual exercises when (limited) endogenous occupation switching is allowed. The analysis focuses on the full-forgiveness policy and maintains the same proportional tax adjustment used to fund the program in the baseline model.

The results in Table A18 show that allowing for endogenous occupation switching does not alter the main qualitative patterns of the student debt forgiveness results. The prevalence of job switching increases relative to the baseline without switching, partly for mechanical reasons: a job switcher is now defined as any individual who chooses a different occupational path, including changes in either the initial occupation or the subsequent occupational transition, rather than only changes in the entering occupation. As in the baseline model without switching, the composition of occupational reallocation varies systematically across policy horizons. With short-run debt forgiveness, occupation switching is primarily driven by amenities, with individuals more likely to transition toward higher-amenity occupations. In contrast, under long-run forgiveness policies, earnings-driven switching becomes more dominant, although the welfare contributions from amenities remain positive. This reflects the fact that debt forgiveness relaxes constraints that previously prevented indebted individuals from selecting the new occupational paths which pursue high-amenity occupations at entry. With lower student debt, individuals who would otherwise have been locked into lower-amenity paths can now re-optimize toward occupations with higher amenity values, as those without student debt do in the extended model. Finally, consistent with the baseline results, induced matriculators remain a relatively small portion of

Table A18: Student Debt Forgiveness Policy with Endogenous Job Switching

Policy	Group	Welfare (CE%)			Lifetime Earnings (%)			Program Cost	
Panel A: Short-Run Policy									
		Total	Amenities	Earnings	Total	Wages	Eff Hrs	Avg. Grant	ΔTax
Full	**Population**	1.26	0.26	1.00	0.28	-0.72	1.01	\$24,028	1.55 pp
	Switchers (0.90%)	75.74	20.21	55.54	12.09	-110.84	122.92	\$46,533	
	Stayers	4.39	0.00	4.39	0.88	0.00	0.88	\$29,445	
Panel B: Long-Run Policy									
		Total	Amenities	Earnings	Total	Wages	Eff Hrs	Avg. Grant	
Full	**Population**	12.76	1.90	10.86	7.07	-9.16	16.23	\$38,731	2.79 pp
	Switchers (8.98%)	73.04	-1.70	74.75	38.38	-119.32	157.71	\$52,214	
	Stayers	6.84	0.00	6.84	1.86	1.02	0.84	\$29,205	
	Matriculators (1.49%)	11.21	-61.79	73.00	167.91	-77.45	245.36	\$69,372	

Notes: Population outcomes include all college and non-college households. Switchers and stayers refer to sub-groups of the population of college graduates. Matriculators include a new population of induced college graduates. Group percentages are with respect to the overall population. Details on welfare and earnings decompositions are contained in Appendix C. Avg. Grant is the average amount forgiven, conditional on receiving forgiveness. Values correspond to percentage changes.

the population, but continue to experience sizable earnings gains compared to others.

E Default/Non-Payment Options and Income-Based Repayment Plan

Another natural concern is whether the baseline model overstates the role of student debt because student debt repayment in practice is more flexible than the standard fixed amortization schedule. In reality, borrowers can opt into income-driven repayment plans, which tie required payments to current earnings. While widespread today, such programs are less relevant in our context as the overwhelming majority of individuals in our NLSY 1997 sample graduate before such plans became widely available. For instance, the majority of individuals in our sample graduate before 2006, while income-driven student debt repayment plans did not become widespread until 2010-2015 with the passage of the PAYE and REPAYE programs under the Obama administration. Nevertheless, understanding the impact of these income-driven repayment plans can provide useful insight into how increased student debt repayment flexibility might affect the model's core quantitative implications.

Towards this end, we consider an additional model extension which maximally loosens the structural constraints on student debt repayment. Specifically, we consider an extension that allows individuals to freely choose how much of their student debt to repay in each period of their lifecycle, subject only to the constraint that all obligations must be paid off before retirement. This alternative repayment scheme provides maximal flexibility, above and beyond what would be allowed under any income-driven repayment plan. While loosening the repayment schedule on student debt, the model extension continues to enforce the borrowing constraint that prohibits individuals from borrowing to finance consumption.

To understand how expanded repayment flexibility influences the model's predictions, Table A19 recomputes the marginal effects of student debt using the within-model IV regression and compares the results to the estimated coefficients in the baseline model. These results show that the added flexibility substantially reduces the IV coefficients, though they remain statistically significant. The extension therefore suggests that more flexible repayment attenuates, but does not fully eliminate, the distortions highlighted in the baseline model. The effect of any income-driven repayment provision then would likely lie in between the baseline model and the maximal repayment flexibility results reported here.

Table A19: Marginal Effects of Student Debt with Flexible Repayment

Target for Indirect Inference	**Coefficient**	**Data**	**Model** (Baseline)	**Model** (Flexible)
IV Effect on Initial Earnings (IE)	$\hat{\alpha}_2$	1.50%	1.58%	0.47%
IV Effect on Returns to Experience (RTE)	$\hat{\alpha}_3$	-0.41%	-0.31%	-0.06%

Notes: This table reports the simulated IV coefficients under the baseline model and the flexible repayment extension. The Data column reports the empirical IV estimates from the fully specified model in Table 1. The Baseline column reports the coefficients from the calibrated model. The Flexible column reports the coefficients when individuals may spread repayment over the entire lifecycle.

F Alternative Public Financing of Debt Forgiveness Policies

The benchmark counterfactual results in Tables 6 and 7 finance the student debt policies through proportional increases in existing marginal tax rates, thereby preserving the progressivity of the U.S. tax system. This section considers an alternative in which the counterfactual policies are instead funded by lump-sum taxation. Under lump-sum taxation, every working-age individual pays the same per-capita amount regardless of income, while retirees are exempt. Although lump-sum taxes avoid distortions to marginal incentives, they are regressive in that the tax burden as a share of income falls disproportionately on lower-income households – many of whom do not even go to college. Comparing the two financing schemes therefore helps isolate the role of the tax structure in determining the welfare effects of the counterfactual student debt policies.

Table A20: Short-Run Student Debt Policies with Lump-Sum Tax

Policy	**Group**	**Welfare (CE%)**			**Lifetime Earnings (%)**			**Program Cost**	
		Total	Amenities	Earnings	Total	Wages	Eff Hrs	Avg. Grant	ΔTax (LS)
Full	**Population**	1.36	0.02	1.34	0.58	-0.01	0.59	$24,694	$1,852
	Switchers (0.55%)	59.59	8.75	50.83	56.35	-1.83	58.17	$39,689	
	Stayers	6.34	0.00	6.34	2.02	0.00	2.02	$23,309	
20K	**Population**	0.75	-0.01	0.76	0.38	0.01	0.38	$14,595	$1,363
	Switchers (0.40%)	27.87	-5.34	33.21	55.10	1.87	53.23	$18,881	
	Stayers	3.75	0.00	3.75	1.24	0.00	1.24	$14,322	
10K	**Population**	0.52	-0.05	0.57	0.31	0.01	0.30	$8,993	$1,092
	Switchers (0.26%)	12.31	-30.02	42.33	77.86	7.20	70.66	$9,932	
	Stayers	2.75	0.00	2.75	0.90	0.00	0.90	$8,962	

Notes: Population outcomes include all college and non-college households. ΔTax (LS) shows the per-capita lump-sum tax levied on all working-age individuals per period. Retirees are exempt. Values correspond to percentage changes.

Table A20 reports the short-run results under lump-sum taxation. The qualitative patterns are similar to the proportional tax benchmark in Table 6: welfare and lifetime earnings increase across all three program sizes, driven predominantly by gains in human capital accumulation. The magnitudes are also comparable. Under the full-forgiveness policy, average welfare increases by 1.36% with lump-sum taxation versus 1.35% under proportional taxation, while lifetime earnings rise by 0.58% versus 0.24%. The per-capita lump-sum tax ranges from $1,092 under the 10K program to $1,852 under full forgiveness.

Table A21 reports the long-run results, where the choice of tax instrument has more consequential implications. While lifetime earnings effects remain comparable to the proportional tax benchmark in Table 7– increasing by 6.24% under full forgiveness in both cases – population welfare now *declines* under all three long-run policies. Under the full-forgiveness policy, average welfare falls by -1.95%, in contrast to the 5.68% welfare gain under proportional taxation. The reversal arises because lump-sum taxation imposes a flat per-capita burden that disproportionately impacts lower-income households, particularly those with only a high school education who do not benefit from the debt forgiveness policies but still bear a substantial share of its cost.

Figure A7 illustrates these distributional implications by plotting the present discounted value of lump-sum taxes paid against the average amount of debt forgiveness received, by group. Households with only a high school degree, who comprise the majority of the population, pay substantial taxes under the new policies but receive no student debt reductions. Under the long-

Table A21: Long-Run Student Debt Policies with Lump-Sum Tax

Policy	Group	Welfare (CE%)			Lifetime Earnings (%)			Program Cost	
		Total	Amenities	Earnings	Total	Wages	Eff Hrs	Avg. Grant	ΔTax (LS)
Full	**Population**	-1.95	-1.47	-0.48	6.24	-0.24	6.49	$45,538	$2,696
	Switchers (5.44%)	32.57	-34.68	67.25	61.60	-0.06	61.66	$51,261	
	Stayers	-3.73	0.00	-3.73	2.11	-0.16	2.27	$45,181	
	Matriculators (1.03%)	38.30	14.40	23.90	148.62	-8.42	157.05	$75,452	
20K	**Population**	-4.05	-1.10	-2.95	1.95	0.17	1.78	$18,974	$1,535
	Switchers (3.66%)	18.08	-34.48	52.56	46.61	3.79	42.82	$19,103	
	Stayers	-4.84	0.00	-4.84	0.07	0.05	0.02	$18,969	
	Matriculators (0.29%)	39.92	18.06	21.85	107.13	-2.53	109.67	$19,958	
10K	**Population**	-4.76	-0.92	-3.85	0.53	0.94	-0.42	$9,945	$1,114
	Switchers (2.34%)	10.86	-44.65	55.51	45.60	6.80	38.81	$9,943	
	Stayers	-5.14	0.00	-5.14	-0.59	0.81	-1.39	$9,945	
	Matriculators (0.07%)	106.33	30.28	76.06	85.74	-1.15	86.89	$10,000	

Notes: Population outcomes include all college and non-college households. Switchers and stayers refer to sub-groups of the population of college graduates. Matriculators include a new population of induced college graduates. ΔTax (LS) shows the per-capita lump-sum tax levied on all working-age individuals per period. Retirees are exempt. Values correspond to percentage changes.

run full-forgiveness policy, they experience a large welfare loss of -7.5%, far exceeding that of any other group. This contrasts with the proportional tax results in Table 7, where the progressive tax structure ensures that the cost of the program is borne more heavily by higher-income households. This structure mitigates the distributional consequences of the counterfactual policies as higher-income households are also more likely to attend college and therefore disproportionately benefit from the reductions in student debt.

Figure A7: Tax Burden and the Amount of Forgiveness by Group

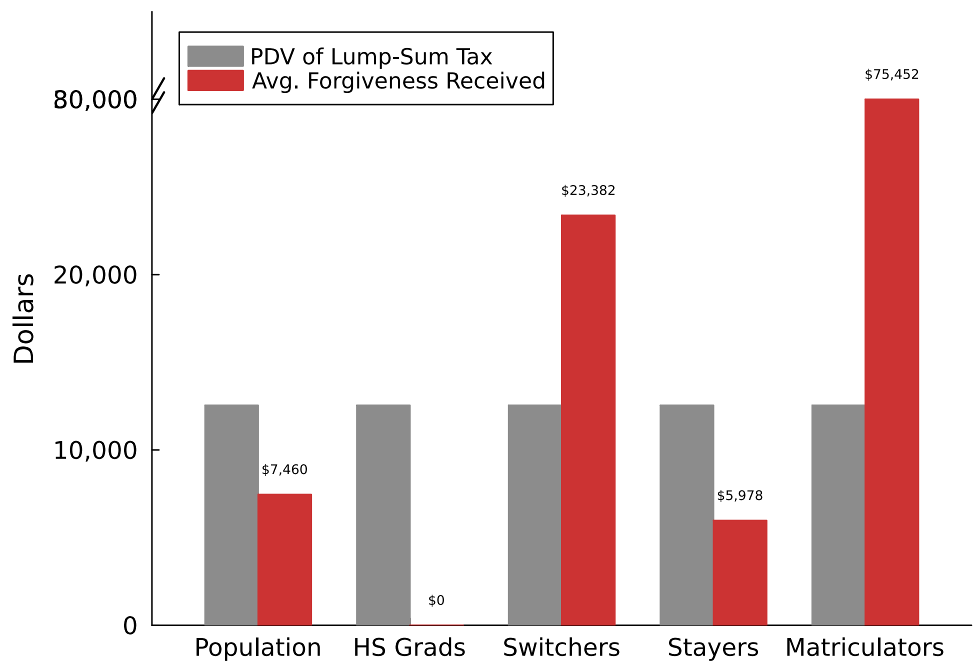

Notes: Gray bars show the present discounted value (PDV) of the lump-sum tax paid over the working life, computed as $\sum_{t=0}^{T} \tau/(1+r)^t$, where τ is the per-period lump-sum tax and T is the number of working periods. Red bars show the average student debt forgiveness received by individuals in each group. The average forgiveness amount was calculated including individuals with no student debt at the baseline. All statistics are computed for the age-23 entry cohort in the long-run full-forgiveness counterfactual.