Misallocation and the Early Career Consequences of Student Debt*

Titan Alon Natalie Cox Minki Kim

July 2024

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, the Urban Institute, NBER SI, Keio, Chonnam National University, 2023 Frankfurt-Mannheim Macro Workshop, and the Minneapolis FRB 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

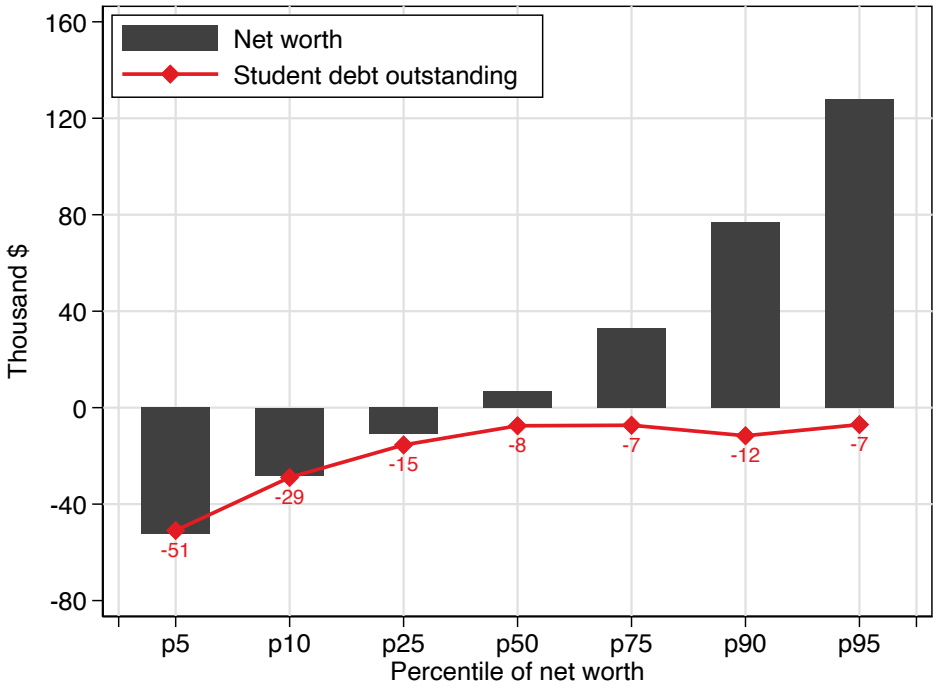
Total outstanding student loan debt reached $1.57 *trillion* in 2022, surpassing auto loans and credit card debt to become the second largest household financial liability after home mortgages.[1] On the one hand, the increase in student debt represents a policy success in which subsidized federal loans 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 workers and a substantial portion of the the variation in their net worth. Third, student debt is largely non-dischargeable, allowing the model to abstract from strategic delinquency and bankruptcy considerations. Empirically, households are much less

[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.



Notes: Source data from 2018 Survey of Consumer Finance. Student debt defined as the total market value of aggregate loan balance of education-purpose expenses. The sample is restricted to bachelor degree holders between the ages of 22 and 25.

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 is 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 of Education Statics (NCES). The empirical result show that an additional $\$1000$ of student debt increases initial earnings by 1.30%, but *reduces* the returns to experience by 0.50 percentage points. The effect on initial earnings corresponds to an additional $\$210$ 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 also sizable given that the average annual earnings growth of individuals between the ages of 25 to 30 are 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. A novel feature of the model is that 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 initial earnings and returns to experience.

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 once-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 exercises 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 once-off federal student loan forgiveness policies could increase household lifetime earnings by up to 0.28% and raise aggregate labor productivity by up to 0.31% 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 (2022) 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 sizeable macroeconomic effects. Average lifetime earnings increase by up to 7.93% under the long-run

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

policy, driven primarily by gains in labor productivity. The larger long-run productivity effect is due in part to a 1.52 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 re-allocation. Under the long-run policy, up to 7.18% of workers switch occupations. These re-allocated 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 re-allocations 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 effect labor market outcomes. Recent contributions have shown that access to consumer credit can effect 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 (Lochner and Monge-Naranjo 2012; Caucutt and Lochner 2020; Lochner, Stinebrickner, and Suleymanoglu 2021; Krueger, Ludwig, and Popova 2024) by investigating how student debt can effect 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 effects household earnings profiles and occupation choice. It contributes to a growing literature documenting the effect of student debt on lifecycle outcomes, such as homeownership, marriage, fertility, college major, and graduate school (Rothstein and Rouse 2011; Goodman, Isen, and Yannelis 2018; Chakrabarti et al. 2020; Hampole 2024). In particular, it provides evidence of how student debt effects lifecycle earnings through distortions to occupation choice (Boar and Lashkari 2022; Luo and Mongey 2024; Herkenhoff, Phillips, and Cohen-Cole 2021).

Methodologically, the paper follows a common approach in the literature that analyzes the aggregate consequences of student debt using dynamic stochastic heterogeneous-

agent models of lifecycle earnings with incomplete markets (Ionescu 2009; Abbott, Gallipoli, Meghir, and Violante 2019; Fu, Lin, and Tanaka 2021). A novel feature of the framework here is the inclusion of a non-trivial occupation choice, which endogenously interacts with student debt to determine lifecycle earnings and the occupational composition of employment. In this sense, the paper is closest to Boar and Lashkari (2022) and Luo and Mongey (2024) who show that debt affects occupation choice by altering household valuations of non-wage job characteristics. In particular, Luo and Mongey (2024) find that the positive response in initial wages to changes in student debt reflects an endogenous trade-off between job amenities and higher wages. Consistent with their findings, the computational results show that these effects dominate in the short-run, but may be moderated in the long-run by wage responses that redirect workers toward occupations with greater returns to experience.

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 63, 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

$$\begin{aligned} c + a' &= \mathbb{T}\left(w_k(1-s)h_k\right) + (1+r)a - \phi(a,h,d,t) \\ h_k' &= exp(z')\mathrm{H}(s,h,\Theta) \\ d' &= (1+r_d)d - \phi(a,h,d,t) \\ a' &\geq -\bar{a} \;\;,\;\; 0 \leq s \leq 1 \end{aligned}$$

where $\bar{a}$ is the exogenously given borrowing constraint and ν_k is a fixed occupation-specific non-wage amenity common to all workers.[7] The non-wage amenities account for factors–other than potential earnings–which influence household occupation choice. Recent work by Boar and Lashkari (2022) and Luo and Mongey (2024) has demonstrated that household substitution along this amenities margin is an important channel through which household liabilities effect their occupation choice.

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

$$\mathrm{H}(s,h,\Theta) = \theta_k^{1-\alpha}(sh_k)^{\alpha} + (1-\delta)h_k$$

where h_k is the individual's current human capital in occupation k, θ_k is their talent in occupation k, and s is the share of time invested in skill acquisition on the job.[8] The initial level of human capital is normalized to one for all households. Parameter α summarizes the investment elasticity of the learning technology and δ captures the depreciation of skills. The functions $\mathbb{T}\left(w_k(1-s)h_k\right)$ and $\phi(a,h,d,t)$ represent the tax system and student debt repayment rule, respectively, which are described below.

[7]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.

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

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, spend one period in college, and then 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 lifecycle problem as college graduates, except that they enter the labor force one period earlier, 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

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

where ζ is a college matriculation taste shock which accounts for unobserved idiosyncratic heterogeneity in college enrollment decisions not captured by the model.[9] It captures, among other things, the opportunity cost of attending college and the amenity value of higher education, both of which effect 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.[10]

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

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

[10] Separately modelling 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.

distribution of student debt levels at age 23 is therefore defined implicitly by,

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

Consequently, high school graduates have no student loans, nor does a mass of college graduates who received and accepted a debt-free college education.[11] 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 effect 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 important 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 modelled 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 circumstance, an individual with outstanding student debt d would have to repay

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

[11] 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 from the returns to college, see Hendricks and Leukhina (2017) and Hendricks and Leukhina (2018).

in each period $t < \bar{T}$, where $\bar{T}$ is the repayment period and r_d is the student loan interest rate.[12]

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

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

In other words, households which cannot make their student loan payments will consume fraction $(1 - \xi)$ of their net assets, and dedicate the remaining fraction ξ of assets and income to pay off 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'. All outstanding student debt must be repaid and no borrowers is permitted to defer repayments beyond maturity duration $\bar{T}$. The model therefore allows individuals to become delinquent, but not default on their student debt obligations – consistent with legal provisions that prevent discharging 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 $\psi(y)$ represent 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 A1. Accurately modelling effective marginal tax rates is important 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

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

human capital investment, such as the provisions of student loan programs.

Retirement and Retirement Benefits. Households retire deterministically by age 63 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 determine the composition of household investments across the demographic distribution. This distribution has important 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 (constant) 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

$$H_c = \left[\sum_k A_k H_{c,k}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}$$

where A_k are group specific productivity terms and η is the elasticity of substitution between college graduate specializations. The total production input for both high school labor and each group of specialized college labor depends on their total effective hours worked, which incorporates both their physical hours and human capital,

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

where $\mathcal{G}_k$ is the joint marginal cumulative distribution function of working age households (a, h, d, t, Θ) who optimally select into occupation $k \in \{0, 1, \ldots, K\}$. Modelling 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 when policies change the 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 dis-invest 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 $\Theta = \{\theta_1, \theta_2\}$.[13] 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 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,

[13] 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.

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 market interest rate $1+r$. Unable to borrow in financial markets, households instead dis-invest 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.[14] 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 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 dis-investment in human capital 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, dis-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

[14]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.

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 $\kappa = \left(\frac{1}{1+r}\right)^{\frac{1}{1-\alpha}} \left[\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}\right] > 0$.[15] 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

[15] 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.

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.[16] 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 selection 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 constrains bind when a household's desired level of borrowing, $a_o^*(a_y, \Theta)$, surpasses the exogenous borrowing constraints $\bar{a}$, which amounts to

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

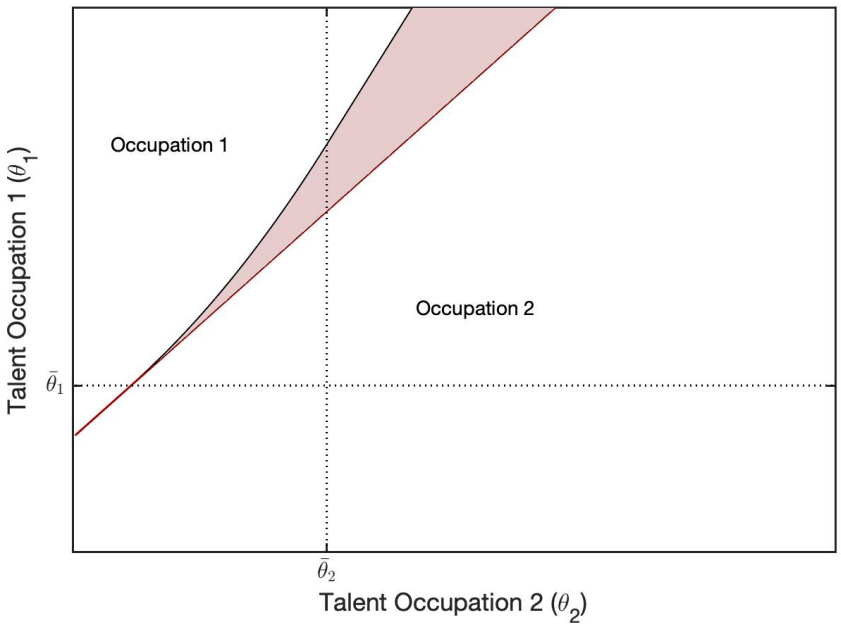
The expression shows that the extent to which credit constraints bind depends not only on a household's initial assets, a_y, but also on their occupational wage w_k which determines their current income, and on their occupational ability θ_k which determines their future income. This means a household 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

[16] 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 that the slope of the sorting rule in Θ space is steeper when individuals are constrained than when they are not constrained.

Figure 2: Misallocation of Talent



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

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.[17]

[17]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

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 misallocated 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 important 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 main mechanism. It provides reduced-form empirical evidence that lifecycle earnings are effected by the presence of student debt, and that part of the effect is mediated by initial occupation choice. The estimation employs panel data from the National Longitudinal Survey of Youth 1997 (NLSY) 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 important identifying power for the model's parameters since they contain information on the marginal effect of varying a household's initial assets. 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 will drive macroeconomic outcomes.

4.1 Data Sources

The analysis draws on several data sources. The primary dataset is the NLSY 1997, an individual-level panel dataset that contains information on higher education, student

other constrained households.

debt, and labor outcomes. It follows individuals from 1997 through 2015. Summary statistics are provided in Table A1 in the appendix. Given the focus on student debt and early career labor earnings outcomes, the analysis focuses on the full-time, full year employed population whose highest level of education is a bachelors degree.[18]

Using the NLSY, we instrument for student debt using variation in the share of grant funding within college and across cohorts, 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 institution. Using the college identifier, we then merge in information from the National Center for Education Statistics (NCES) on the amount of loans and amount of grants used at that given college in a given year.[19]

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

$$y_{it} = \underbrace{\alpha_0 + X_{it}\beta}_{\text{initial (log) earnings if no student debt}} + \underbrace{\alpha_1 \text{Exp}_{it}}_{\text{returns to experience if no student debt}} + \underbrace{\alpha_2 \text{SD}_{it}}_{\text{effect of student debt on initial (log) earnings}} + \underbrace{\alpha_3 \text{SD}_{it} \times \text{Exp}_{it}}_{\text{effect of student debt on returns to experience}} + \epsilon_{it} \quad (3)$$

where y_{it} is an outcome measure of individual i in year t for annual log earnings. The variable SD_{it} denotes the level of student debt. The variable Exp_{it} denotes the years of experience. The variable X_{it} includes additional controls school and individual controls, including dummies for race, gender, age at the year starting and finishing the BA degree, and whether the individual attended private or public college.

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

[18] Robustness exercises which additionally include those with associates degrees (AA) or those with post-graduate degrees are largely consistent with the findings of the BA-only population, which hold the majority of student debt. Nevertheless, the primary sample focuses on BA-only population since AA and post-graduate degree recipients make very different human capital investment decisions, having different forms of student debt, and have markedly different observed occupational choices after graduating. Full-time, full-year status includes those who work at least 35 hours a week and for at least 40 weeks a year.

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

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 quality or ability. 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.[20]

To address these identification challenges, we estimate the causal 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 in a given year. Specifically, the instrument is given by

$$Z_{c(i),j} = \frac{\text{total grants}_{c(i),j}}{\text{total grants}_{c(i),j} + \text{total loans}_{c(i),j}}$$

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 amount of student debt that students take out. Table A2 in the appendix shows a strong first stage effect of shifts in the college-year grant share on individuals' student debt. Importantly, the table also shows that changes in grant funding are

[20]These identification challenges to identifying a causal impact of student debt on earnings are also highlighted by Field (2009), Rothstein and Rouse (2011), Luo and Mongey (2024). These papers uses variation in forgiveness of debt (Field 2009) and variation in grants (Luo and Mongey 2024) within a school across cohorts to instrument for student debt. Their identification comes from comparing outcomes of cohorts within the same school, when cohorts within the school differ in terms of grants received.

compensated for almost *entirely* and *exclusively* by changes in student debt, not other sources of funding. Total funding for college remains mostly constant in response to one standard deviation increase in the college grant share. While the level of student debt decreases almost one-for-one with the increase in grant funding, family and work study aid remain constant. This precise, isolated substitution is important, because it allows us to study the impact of an increase in student debt on future earnings, absent of other confounding factors like more parental aid or increased work study while in college.

Another important concern is whether changes in the college grant share may also change other important education outcomes that could confound the results. Table A3 in the appendix investigates these concerns by checking if variation in the instrument is correlated with the probability of college completion, student ability, parental income, race, full time enrollment status, and age of enrollment. The results do not show any significant evidence of the instrument impacting enrollment or student selection on observable characteristics.

The second stage regression, which estimates the impact of instrumented student debt on lifecycle earnings, includes additional fixed effects for college type – for example, private, public, for-profit, etc. While it would be ideal to include fixed effects for each individual college, the small sample size does not allow this since there are very few instances in which more than a few students attended the same institution.

4.3 Estimated Impact of Student Debt on Lifecycle Earnings

This section employs the instrument to investigate whether those who graduate with more student subsequently choose jobs with significantly different earnings profiles. Table 1 reports the estimates coefficients α_2 and α_3 for log earnings which summarize the effect. Across nearly all specifications and robustness exercises, the estimated coefficient α_2 is positive, while the estimated α_3 are negative. The coefficients imply that an individual graduating with more student debt will have initially higher earnings, but subsequently lower returns to experience. The direction of the estimated effects are consistent with the theory on intertemporal distortions to human capital accumulation discussed in Section 3.

Table 1 reports the estimated coefficients constructively across specifications until the benchmark model in the final column. The first column reports the estimated coefficients without any controls (e.g. β=0). The second column adds the minimal college

Table 1: Estimated IV Coefficients of Student Debt on Earnings Lifecycle

Effect of Student Debt ($000s) on:	**Grants-based IV Specifications**			
	(1)	(2)	(3)	(4)
Log Initial Earnings	2.81%	2.87%	1.33%	1.30%
(p-value)	0.066	0.038	0.059	0.062
Mean Returns to Experience	-1.13%	-0.96%	-0.59%	-0.50%
(p-value)	0.024	0.009	0.001	0.003
Source of Variation	Across-cohort, within college-type (among all students)			
cohort and college type controls	no	yes	no	yes
job type FE and job x experience	no	no	yes	yes
number of observations	6,284	6,284	5,909	5,909

Notes: The table reports the instrumented estimates from regression 3. The IV utilizes changes in the college-year grant share, which in turn impacts the amount of student debt taken out by individual students. The dependent variable is log annual earnings. See text and data appendix B for additional details.

type and cohort controls, including fixed effects for the individual's enrollment cohort, age at first enrollment, and their BA award year. The purpose of these variables is to isolate the cross-cohort, within college variation in student debt from other cohort and school type variation. Comparing the specifications with and without these controls suggests that the results are not driven by changes in matriculation decisions or the sorting of students across different school types.

The third and fourth columns explore the role of occupation choice in explaining the earnings gap between those with different levels of student debt. They add fixed effects for the initial occupation and industry choice an individual chooses upon graduation. They also include the interaction of these fixed effects with the years of experience. The augmented specifications therefore allow differences in initial earnings and returns to experience to arise from differential job choice for those with student debt.[21] Comparing the specifications shows that these job type controls account for about half of the marginal effect of student debt. For instance, in the benchmark specification the effect on initial earnings is 55% lower (α_2 declines from 2.87% to 1.30%) and the effect on returns to experience declines by 48% (α_3 declines from 0.96% to 0.50%). The results

[21] In particular, it allows both initial earnings and returns to experience to vary across 30 occupations types and 15 industries. Data appendix B contains details on the occupation and industry definitions.

suggest that a substantial part of the difference in earnings profiles of those with and without student debt may be due to selection into different job types after graduation – suggestive of the misallocation effect discussed in section 3.2.

To interpret magnitudes, the coefficients of the fully specified model in column (4) imply that an additional $1,000 of student debt increases initial earnings by 1.30%. Within the data sample, this equates to to an additional $210 annual earnings upon graduation for every $1,000 of additional student debt. Despite analyzing a different sample from the existing literature, the resulting estimates on the effects for initial earnings are consistent with existing estimates. See, for instance, the initial earnings effects estimated in Luo and Mongey (2024), Rothstein and Rouse (2011), and Chapman (2015).

The results in column (4) also show that annual earnings grow by 0.50 percentage points *slower* per year of experience for every $1,000 of additional student debt. The effect is statistically significant and appears economically sizeable relative to the average earnings growth observed for early career workers. For instance, while the literature does not contain directly comparable instrumented estimates, Guvenen et al. (2021) employ IRS administrative data and estimate that the annual earnings growth of all 25 to 30 years old in the United States is 7.75% on average each year.[22]

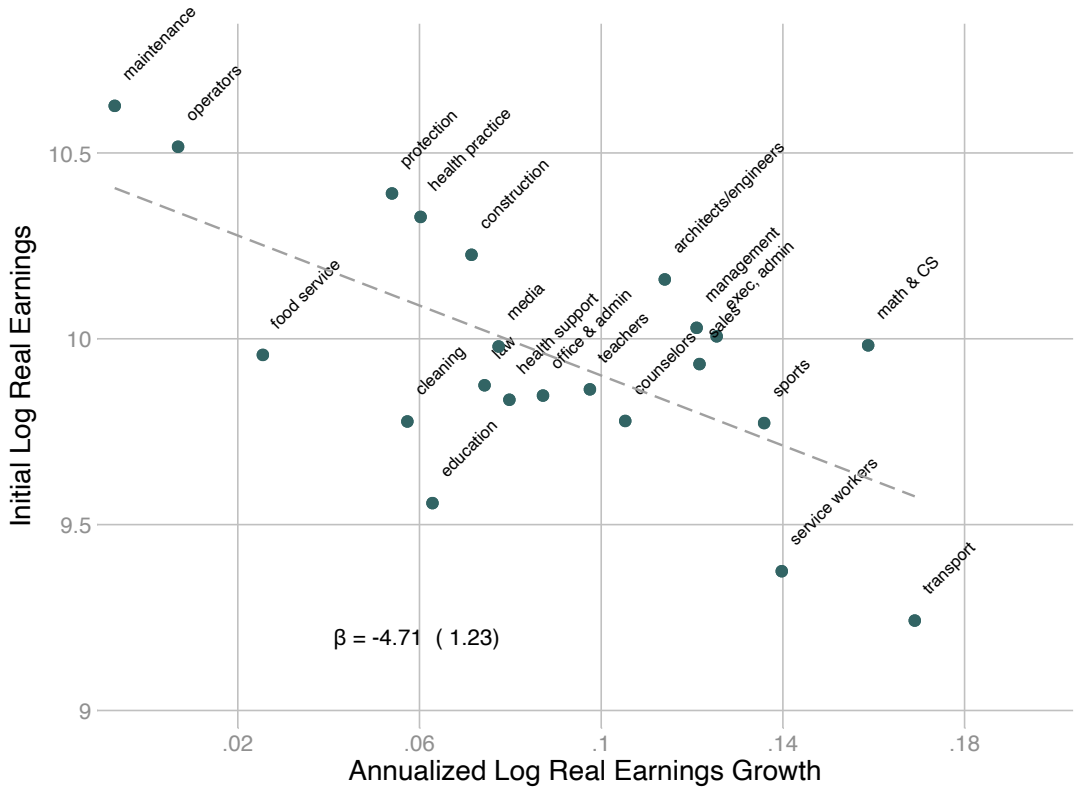
In summary, the findings suggest that individuals with high student debt trade off between current and future income early in life. Those with student debt have higher initial earnings upon graduation, but subsequent returns to experience are lower. Moreover, this trade-off appears to depend significantly on their *first* occupation and industry choice after graduation. Allowing initial earnings and returns to experience to vary with an individuals occupation and industry choice can explain roughly half of the marginal effects of student debt on early career earnings profiles.

4.4 The Role of Occupation Choice

The instrumental variable analysis finds that additional student debt leads to age-earnings profiles with initially higher earnings, but lower returns to experience. About half of this effect is explained by initial occupation choice, whereby individuals with higher student debt sort into professions with predictably front-loaded income trajectories. To better understand the results, this section investigates how realized age-earnings profiles differ across occupations. To estimate occupation-specific earnings trajectories, the

[22]Though this number includes both high school and college graduates, which also likely exhibit significant between group variation in early career returns to experience.

Figure 3: Initial earnings, earnings growth, and student debt by occupation



Notes: The figure plots the estimated earnings function coefficients for each 2-digit occupation.

NLSY data *for each* 2-digit sector is fit to the quadratic earnings model

$$Earnings_i = \alpha + \beta_1 Exp_i + \beta_2 Exp_i^2 + \beta_3 X_i + \epsilon_i$$

where all regression coefficients are allowed to vary across occupations. The sample includes individuals who are between 23–38 years old and have exactly a bachelors degree. The X_i include additional control variables which determine the cross-sectional income variation within and between occupations, including race, gender, region where the respondent resides, year and industry fixed effects.[23]

The result shows a strong negative correlation between initial earnings and subsequent earnings growth across occupations. The substantial cross-occupation variation in earnings profiles suggests that the marginal constrained household may have ample oppor-

[23] Appendix B.2 and B.3 contain the details about the occupation classification the data clearing and variable construction processes.

tunity to trade-off income today versus income tomorrow by moving across occupations. The results are consistent with the conditions necessary for credit constraints to cause a misallocation of talent within the model. They also provide important calibration targets which discipline the quantitative model's calibration in Section 5.

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 why occupational fixed effects have such a large impact on the regression results. They show that there is considerable heterogeneity across occupations in the level of earnings received upon graduation as well as the subsequent average earnings growth. Plotting the initial earnings and returns to experience implied by the estimated coefficients ($\beta_{1,j}, \beta_{2,j}$) makes the trade-off individuals face even more apparent. As shown in Figure 3, there exists a statistically significant negative relationship between an occupation's initial log earnings after graduation (y-axis), and their average yearly growth rate in the first 15 years (x-axis).[24] The implied variation in earnings profiles across occupations suggests a wide scope for households to adjust lifecycle earnings trajectories by moving across occupations, and highlights the slope-intercept earnings trade-off across occupations.

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 effect 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_1, α_2) 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

[24] This is the range that is consistent with the IV regression model since the NLSY 1997 sample used to estimated the model only provides data on the first 15 years of individual labor market outcomes.

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, 3, and 4 report the internal calibration targets and model fit. Table 2 summarizes college matriculation patterns and the distributional properties of student debt. Table 3 reports occupational employment shares and earnings heterogeneity across occupations. Table 4 contains the indirect inference targets and summarizes the model's ability to match the empirical evidence on the marginal effects of student debt in section 4. Table 5 summarizes the externally calibrated parameters.

College Matriculation and the Distribution of Student Debt. Modelling college matriculation is important 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 burdens 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 tuned with the internal calibration so that the 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.[25] In addition to student debt, the decision to matriculate will depends on the returns to a college education. These in turn depend on market prices (e.g. the college wage premium), individual abilities Θ, 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).[26] 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 44% college wage premium, as in Daruich and Kozlowski (2020).[27] Finally, given earnings, the college taste shock ζ is calibrated to match college completion rates. To capture the non-linearities in the data, the taste shock is allowed to depend on a household's initial assets via the polynomial $\zeta = b_0 + b_1 a_0 + b_2 (a_0)^2 + b_3 (a_0)^3 + \epsilon$, where ϵ is an idiosyncratic standard log-normal taste shock. The coefficients b_j are fit to match the college completion rates by household asset quintile.

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

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

[27] 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 and the Skill Premium		
College wage premium	43.0%	44.0%
High school average returns to experience	4.91%	4.88%
College completion rates by asset quintile	31.20%	30.56%
– First quintile (Q1)	18.20%	20.83%
– Second quintile (Q2)	22.75%	24.12%
– Third quintile (Q3)	26.65%	27.41%
– Fourth quintile (Q4)	37.70%	33.03%
– Fifth quintile (Q5)	50.70%	47.42%
B. Distributional Moments on Student Debt		
Mean level of initial assets	\$18,132	\$18,036
Standard deviation of initial assets	\$21,503	\$22,004
Mean level of student debt	\$21,843	\$22,022
Standard deviation of student debt	\$28,041	\$29,525
Correlation between initial assets and student debt	-0.15	-0.20
Fraction of BA graduates without SD	40.34%	42.53%

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 important 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 important 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 $\{A_k, \mu_k, \nu_k\}_{k=1,\ldots,22}$. Specifically, the level of initial earnings in each occupation k are 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).[28] 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.[29] Finally, given lifecycle earnings in each occupation, the corresponding amenity values v_k, common to all workers within an occupation, are chosen to replicate the distribution of occupational employment shares.

Table 3 summarizes the data targets and resulting model fit. For each occupation, it reports the initial earnings, average returns to experience, and employment share in the data and the calibrated model. Overall, the model appears to do a good job 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 $\{A_k, \mu_k, \nu_k\}_{k=1,\ldots,18}$ therefore do not map *directly* 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.

Causal Evidence on the Marginal Effects of Student Debt. Until now, the calibration has sought parameters which, in equilibrium, replicate important 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 is

[28] 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 effect intertemporal trade-offs. However, for constrained households, optimal investments do depend on wage levels. See section 3.

[29] 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 effect investment within that occupation. They do not emerge from heterogeneity in the variance of abilities across occupations.

Table 3: Internal Calibration Targets on Occupational Heterogeneity

Occupation Group	Mean log Earnings ($)		Returns to Experience (%)		Employment Share (%)	
	Data	**Model**	**Data**	**Model**	**Data**	**Model**
Executive and administrative	10.01	10.01	12.54	12.41	6.00	6.02
Management related	10.03	10.04	12.10	11.99	7.86	7.90
Mathematical and computer scientists	9.98	9.99	15.88	14.80	4.28	4.21
Engineers, architects, and other scientists	10.16	10.17	11.40	11.36	5.24	5.23
Counselors, social and religious	9.78	9.78	10.53	10.53	4.00	3.98
Lawyers, judges, and legal support	9.87	9.87	7.43	7.44	0.90	0.90
Teachers	9.86	9.87	9.75	9.71	9.93	9.97
Education, training, and library services	9.56	9.55	6.29	6.29	1.79	1.81
Entertainment and sports	9.77	9.78	13.59	13.36	3.52	3.56
Media and communications	9.98	9.98	7.74	7.71	2.76	2.77
Health diagnosing and treating practitioners	10.33	10.33	6.02	5.98	3.52	3.55
Healthcare technical and support workers	9.84	9.83	7.98	7.96	2.55	2.57
Protective Services	10.39	10.39	5.39	5.50	1.66	1.60
Food preparation and serving	9.96	9.95	2.55	2.56	6.48	6.50
Building and cleaning services	9.78	9.79	5.73	5.74	1.31	1.30
Service workers	9.37	9.38	13.97	13.94	3.38	3.37
Sales and related workers	9.93	9.94	12.16	12.07	12.76	12.79
Office and administrative support	9.85	9.84	8.72	8.73	16.14	16.11
Construction trades and extraction	10.23	10.21	7.15	7.43	1.31	1.24
Installation, maintenance, and repair	10.63	10.63	-0.71	-0.70	1.24	1.24
Production and operating workers	10.52	10.52	0.68	0.68	1.17	1.16
Transportation and material moving	9.24	9.24	16.90	16.82	2.21	2.21

Notes: This table summarizes the model fit of the 66 calibration targets governing occupational heterogeneity and the sorting of workers. Earnings and employment shares correspond to the population of BA graduates. See data appendix B for additional details on variable and occupation definitions.

important since it 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 economy, the marginal effect targets correspond to household optimal responses to perturbations of the equilibrium. Specifically, we use the simulated data and regress individual's log earnings on the years of experience, amount of initial student debt, and their interaction, using the first ten years of simulated life-cycles among college graduates. In other words, we replicate the same regression (3) with the simulated data and

Table 4: Internal Calibration Targets on the Marginal Effects of Student Debt

Target for Indirect Inference	Coefficient	Data	Model
IV Effect on Initial Earnings (IE)	$\hat{\alpha}_1$	1.30%	1.32%
IV Effect on Returns to Experience (RTE)	$\hat{\alpha}_2$	-0.50%	-0.32%

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 experimental data produced by the model simulation. See data appendix B for additional details.

compare the coefficients with the empirical counterparts. Through indirect inference it then seeks a parameterization which results in the stimulated policy generating the same average effect on initial earnings and the returns to experience in column (4) of Table 1. The results are presented in Table 4, which display the original IV estimated coefficients alongside the coefficients recovered from the model-based regression.

The empirical evidence primarily helps the calibration fit 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 debt relief on life-cycle earnings will primarily depend on the latent distribution of θ. More specifically, since households generally select into occupations in which they have higher abilities, it is the tail of the ability distribution implied by variance σ_θ^2 which matters most for the earnings response. Comparing the optimal investment rules s^c and s^* in Section 3 illustrates how the average intertemporal distortions to earnings will depend on the distribution of θ. Similarly, the distribution of θ will also determine the mass of households which fall into the misallocation region discussed in Figure 2.

Externally Calibrated Parameters. Table 5 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

Table 5: 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
$\bar{T}$	Student debt repayment duration (years)	10	SFRP institutional feature
ξ	Student debt delinquency garnishment	1.00	Luo and Mongey (2024)
Human Capital Technology			
σ_z	Standard deviation of human capital shocks	0.111	Hugget et al. (2011)
α	Returns to scale in human capital tech.	0.7	Hugget et al. (2011)
δ	Human capital depreciation rate	0.029	Hugget et al. (2011)

Notes: This table summarizes the externally calibrated parameters. All parameters are annualized. See appendix B for additional details.

consumption smoothing incentives. The risk free rate, r, is chosen to match 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 of Education Statistics. The repayment duration $\bar{T}$ and delinquency term ξ are set consistently with the standard federal repayment plan (SFRP), as discussed in Luo and Mongey (2024).

The parameters governing the human capital technology are taken from the study of 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 depreciate rate, δ, is set to match the decline in lifecycle earnings near the end of working life.

5.1 Validating the Mechanism

The internal calibration does a good job at replicating college matriculation patterns and important distributional properties of student debt. It is also consistent with the marginal effects of student debt on earnings by replicating the estimated IV coefficients.

The economic theory shows that these effects can arise through intertemporal distortions to human capital investment on the job and through a misallocation of talent. The presence of intertemporal distortions that reduce human capital investment at the intensive margin is readily evident in the model's specification of the human capital technology and continuous investment choice s. However, it is not immediately evident whether or not the resulting calibration also produces the conditions necessary for their to be a misallocation of talent.

To validate that this mechanism is operating, Figure 4 illustrates how the sorting of graduates interacts with earnings heterogeneity across occupations. In particular, it plots the calibrated model's output against the empirical graph in Figure 3. The position of each occupation is determined by its initial earnings (y-axis) and average returns to experience (x-axis). Consistent with the model fit in Table 3, the figure shows that the model accurately replicate the cross-occupation heterogeneity in lifecycle earnings. Moreover, the size of the bubble reflects the average level of student debt held by graduates who select into each occupation. Consistent with the presence of a misallocation of talent, graduates with more student debt cluster disproportionately into occupations with higher initial earnings and lower returns to experience. This occurs even though occupation-specific abilities are distributed independently of assets and student debt.

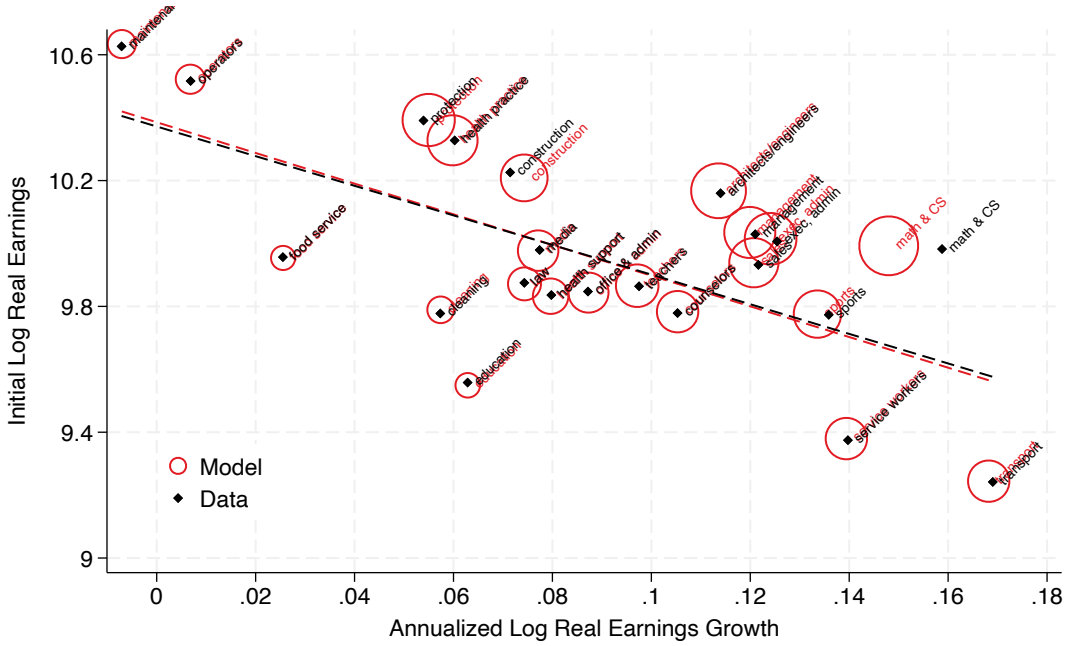
The cross-occupation correlation with student debt is not perfectly monotone in part due due to the presence of occupation-specific amenities, ν_k, which are also valued differently by constrained and unconstrained households. This differential preference for non-wage amenities, discussed at length in Luo and Mongey (2024) and Boar and Lashkari (2022), also impact the direction of occupational distortions which result from credit frictions in this model. The following section analyzes these effects more directly, alongside those discussed above, by using the calibrated model to assess the aggregate consequences of student debt for occupation choice and labor productivity.

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 misallocaiton 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 the current administration.[30]. The second

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

Figure 4: Initial earnings, earnings growth, and student debt by occupation



Notes: The figure plots the estimated earnings function coefficients for each occupation in the model and data. The dashed lines show the associate OLS coefficients summarizing cross-occupation correlations in earnings characteristics. Model-simulated average amount student debt by occupation is represented using differently sized bubbles. Large bubbles have relatively high predicted student debt levels compared to small bubbles.

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 once-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 being proposed by the current administration, while the latter provides a benchmark of the short-run aggregate distortions from student debt.

Table 6 reports the impact of each policy on household welfare and lifecycle earnings. Household consumption-equivalent welfare is decomposed into changes from lifecycle earnings and occupation specific amenities, for those who switched jobs. Changes in household earnings are similarly decomposed into the contribution of occupational

wages and effective hours, capturing the impact on labor productivity.[31] The final column reports the average level of student debt relief received under each policy. 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 once-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 1.65%, with 75% of the rise coming from increased earnings. The increase increase follows from the improved efficiency of human capital investments over the working lifecycle. For example, under the full-forgiveness policy, total lifetime earnings increases by 0.28%.

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 0.80% of the population (2.63% 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 -5.14% reduction in lifetime earnings. The decline is driven predominantly by their switching to jobs that on average have a -6.45% lower wage, which outweighs the modest 1.31% increase in lifecycle earnings.

The fact that induced job-switchers flow from high wage to low wage occupations that better match their skills and increase their lifetime productivity is consistent with the theory of misallocation developed in section 3.2. What is surprising is that these reallocations are associated with a *decrease* in lifetime earnings. 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 it optimal to give up earnings in order to move to occupations with higher non-wage amenities. This margin of occupational re-allocation explains why induced job switchers derive enormous welfare benefits from the debt relief policies despite moving on average to occupations that

[31] 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
	Population	1.65	0.41	1.24	0.28	-0.04	0.31	$45,418
Full	Switchers (0.80%)	127.05	86.05	41.00	-5.14	-6.45	1.31	$46,844
	Stayers	5.60	0.00	5.60	1.74	0.00	1.74	$36,372
	Population	0.57	0.18	0.39	0.07	-0.01	0.08	$18,947
20K	Switchers (0.56%)	48.89	48.73	0.15	-16.24	-2.44	-13.81	$18,959
	Stayers	2.25	0.00	2.25	0.76	0.00	0.76	$18,379
	Population	0.29	0.09	0.20	0.04	0.00	0.04	$9,943
10K	Switchers (0.34%)	34.78	41.16	-6.39	-16.52	-1.30	-15.22	$9,947
	Stayers	1.24	0.00	1.24	0.44	0.00	0.44	$9,902

Notes: Population outcomes include all college and non-college households. Switchers and stayers refer to sub-groups of the population of college graduates. Group percentages are with respect to the total population. 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.

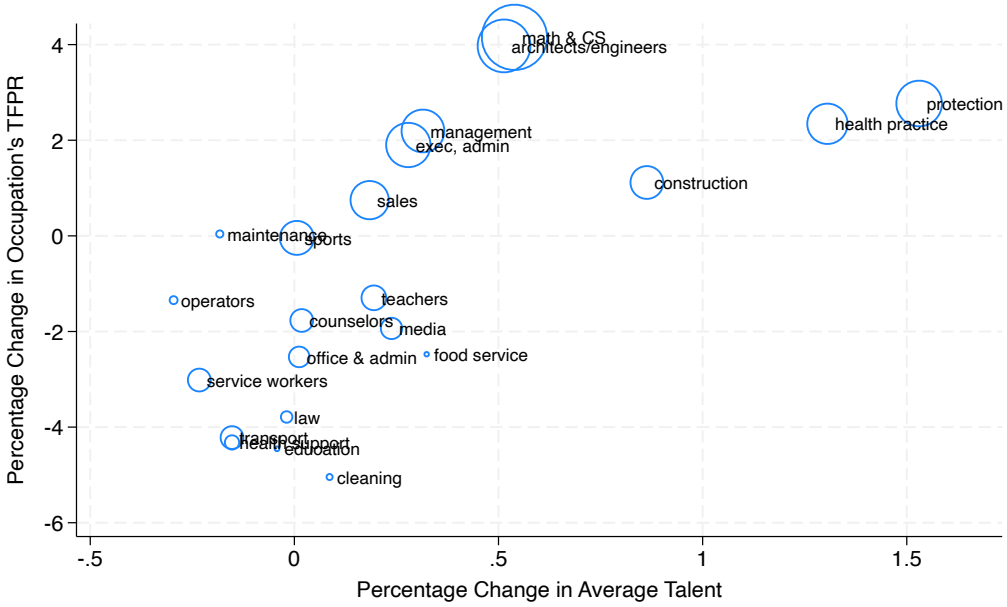
reduce their lifetime earnings. These amenity-driven job re-allocations align with the recent findings of Luo and Mongey (2024), who show that higher student debt causes graduates to accept jobs with initially higher wages, but lower job satisfaction.[32]

While policy induced occupational re-allocation 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 A2 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 Mathematics, Computer Science, Architects, and Engineering Professions– experience large increases in labor productivity of 3.97 to 4.15% under the full forgiveness policy. Other occupations – such as Building Services, Transportation, Education and Health Support – experience large reductions in productivity of nearly the same magnitude, ranging from -4.32% to -5.04%.

Since occupational wages are held constant in the short-run policy computations, the

[32]Relatedly, recent work by Boar and Lashkari (2022) 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.

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 is 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 A2.

cross-occupation productivity changes are driven primarily by worker sorting. To see this more directly, Figure 5 plots the data from Table A2. 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 re-allocation in response to the short-run debt forgiveness policy. The fact that average talent increases in virtually all occupations indicates a reduction in the aggregate misallocation of talent. Consistent with the theory, the figure also shows that the impact of these re-allocations was largest in occupations where workers initially had the highest levels of student debt.

The decomposition in the latter columns of Table A2 also shows that the population of job-switchers are, on average, lower ability than 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 op-

timally 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 A2 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 A4).

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 appear 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 $9,005 to $14,674.

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 responds 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. For example, under the full-forgiveness policy, the average person induced to switch jobs received $46,844 in student debt relief, while job stayers received only $36,372. These non-linearities also partly explain why only the largest short-run debt forgiveness policy is able to induce an increase in productivity among job switchers, whereas smaller programs induce primarily amenities driven switching. The next section examines this effect in greater detail by analyzing the macroeconomic impact of larger, long-run debt forgiveness policies.

Long-Run Student Debt Policies. The second exercises 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 de-

cisions to respond endogenously to the new policies.[33] In contrast to the short-run exercises, the long-run results should be thought of as accounting exercises rather than policy experiments. This is because there is much greater uncertainty regarding the role of omitted factors in quantitatively assessing the outcome of such large scale and permanent changes to education finance. Nevertheless, the computational exercises are useful in that they provide informative benchmarks as to the long-run aggregate distortions that can result from financing higher education through debt, rather than grants.

Table 7 displays the results of the long-run policies. It reports steady state changes in each of 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, except 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.

Upon closer examination, however, there are notable differences between the short-run and long-run outcomes and the mechanisms which underlie them. For instance, increases in lifetime earnings play a more important role in driving the welfare in the long-run. As the decompositions in Table 7 show, earnings growth accounts for almost 90% of the rise in welfare in the long-run, versus 75% in the short-run policies. Such differences are even starker among job switchers. Earnings growth accounts for 70% of the welfare gain for job switchers in the long run, versus 32% in the short run policies, implying that job switching is more labor productivity, not amenities in the long run policies.

The greater lifetime earnings gains in the long-run are the result of both (i) a much larger population of policy-induced job switchers and (ii) different patterns of job re-allocation among them. Under the long-run full forgiveness policy, 7.18% of college graduates choose different occupations, compared with just 0.80% in the short-run. Moreover, while short-run job switchers experienced on average *decline* in lifetime earnings and negligible productivity growth, those induced to switch occupations by the long-run policy experience large *increases* in lifetime earnings driven by productivity growth.

[33] 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

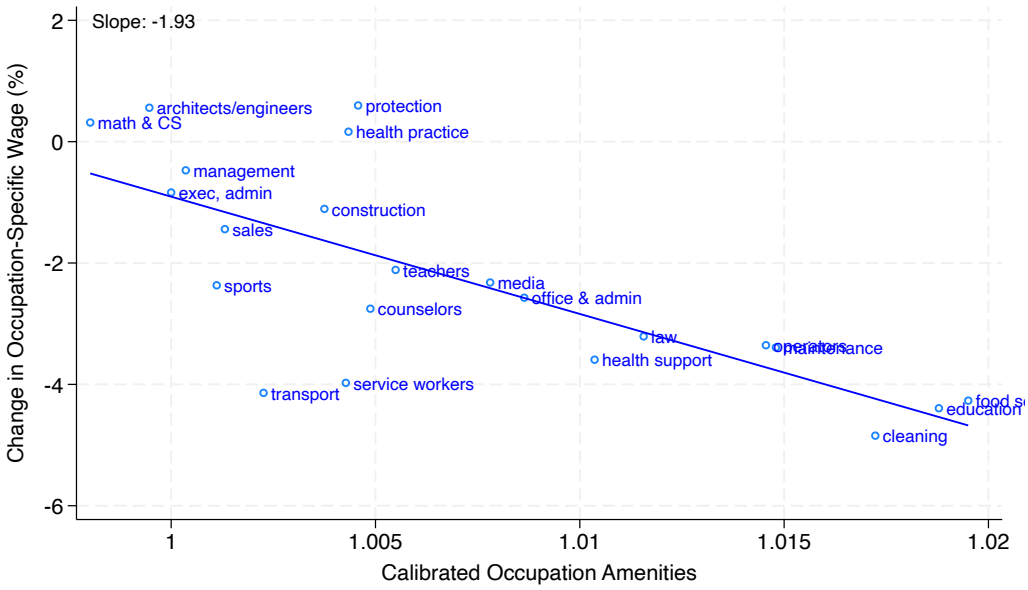
Policy	Group	Welfare (CE%)			Lifetime Earnings (%)			Program Cost
		Total	Amenities	Earnings	Total	Wages	Eff Hrs	Avg. Grant
Full	**Population**	8.27	0.90	7.37	7.21	-0.72	7.93	$45,417
	Switchers (7.18%)	78.28	23.73	54.55	36.41	-8.11	44.52	$53,206
	Stayers	4.36	0.00	4.36	3.09	0.00	3.09	$44,552
	Matriculators (1.52%)	35.68	17.73	17.95	149.67	-17.16	166.83	$69,938
20K	**Population**	3.33	0.02	3.31	3.17	-0.22	3.39	$18,947
	Switchers (4.90%)	38.42	2.03	39.85	26.72	-5.14	31.86	$19,139
	Stayers	1.74	0.00	1.74	1.29	0.08	1.21	$18,932
	Matriculators (0.65%)	28.73	19.68	9.04	123.98	-11.95	135.93	$19,754
10K	**Population**	1.79	-0.05	1.85	1.74	-0.30	2.04	$9,943
	Switchers (2.95%)	31.50	-1.55	33.05	23.71	-4.57	28.28	$9,956
	Stayers	0.98	0.00	0.98	0.73	-0.13	0.86	$9,942
	Matriculators (0.34%)	36.23	21.81	14.42	122.59	-13.75	136.34	$9,985

Notes: Population outcomes include all college and non-college households. Switchers and stayers refer to sub-groups of the population of college graduates. Matriculators includes 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.

As Table 7 shows, long-run job switchers under the full forgiveness policy experience a 36.41% growth in lifetime earnings driven by a 44.52% increase in labor productivity. These large productivity gains, coupled with the larger population of policy induced job switchers, is the main driver of the policy differential.

The reason that similarly sized student debt relief programs deliver such different flows of worker re-allocations in the short-run and long-run is due 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 re-allocation. Figure 6 illustrates the effect by showing how, under the long-run policies, the endogenous wage declines are largest in the high amenity occupations. These wage changes also trigger second-order labor re-allocations, 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 be induced to 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.

Figure 6: Policy Induced Wage Changes and Occupational Amenities

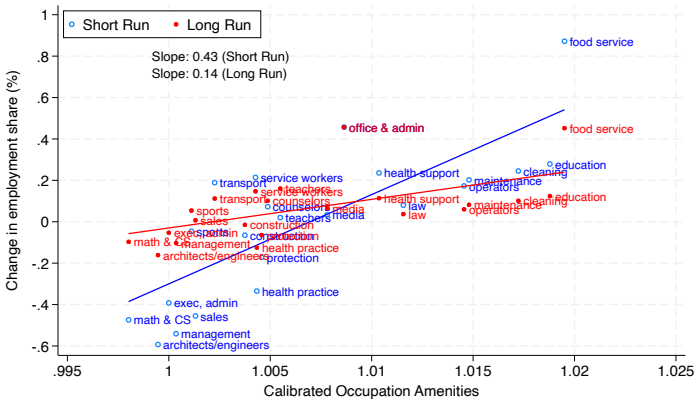


Notes: The y-axis reports changes in an occupation's wage under the long-run full student debt forgiveness policy. The x-axis reports the occupation's calibrated amenity value, re-indexed so executive administrators equal one. The corresponding occupation specific productivity and employment outcomes under the long-run full forgiveness policy are summarized in Table A3.

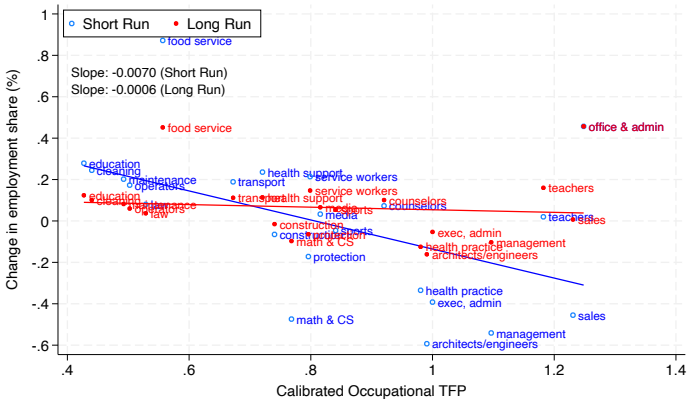
As a result of the equilibrium wage response, labor market re-allocation under the long-run policy are more directed toward improving the allocation of talent rather than increasing occupational amenities. The change in flows is illustrated in Figure 7, which summarizes the re-allocation of workers under 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 re-allocation which characterised the short-run switchers is substantially moderated in the long-run. At the same time, the re-sorting of workers into lower wage but higher talent occupations is much more pronounced. These changes in the flow of workers underlies the larger productivity gains and more modest welfare effects on job-switchers under the long-run policies.

Table A3 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

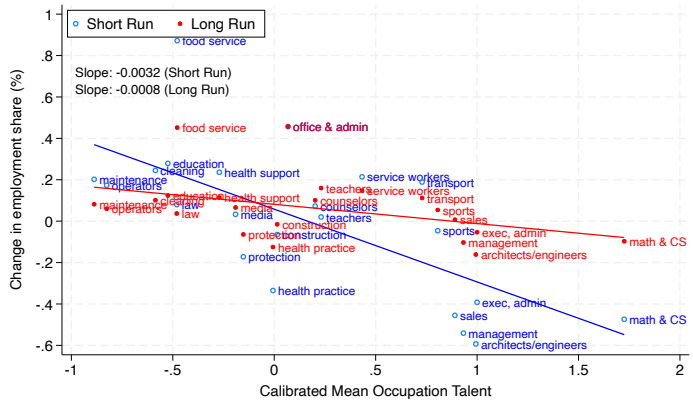
Figure 7: Policy Induced Re-allocation of Workers by Occupational Characteristics



(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 exogenous heterogeneous. The x-axis of each panel is normalized so that executive administrators equal one. Additional details are contained in the appendix.

occupational TFPR changes under the short-run and long-run policies is 0.82; the corresponding correlation of induced changes in worker ability across occupations is 0.92. The fact that the majority of occupations experience an increase in the average talent of their workforce shows that the long-run policy 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 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 1.52% of the population are 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, it is also possible that the resulting adjustments in occupational wages will disincentivize some workers who previously found it worthwhile to enroll. However, the calibrated model suggests these effects are quantitatively negligible in practice, representing less than 0.07% of the college population even under the full forgiveness long-run policy. As a result, while the population of induced college matriculators experience some of the largest welfare and productivity gains in the population, their small population size means that they have only a minor impact on aggregate outcomes. Instead, it is the occupational re-allocation of current college graduates, and the concomitant reduction in the misallocation of talent, that drives the long-run outcomes.

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. 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 for the 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 causal evidence on how exogenous variation in student debt effects the early career labor market decisions 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 re-allocation 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 high ability occupations – further augmenting labor productivity by reducing the misallocaiton 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.

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.

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. 2022. "Occupational Choice and the Intergenerational Mobility of Welfare." Nber working papers 29381, National Bureau of Economic Research, Inc.

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. 2020. "Tuition, Debt, and Human Capital." Staff reports 912, Federal Reserve Bank of New York.

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.

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.

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).

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. 2024. "Shaping inequality and intergenerational persistence of poverty: Free college or better schools." Technical Report, National Bureau of Economic Research.

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." *Annu. Rev. Econ.* 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).

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

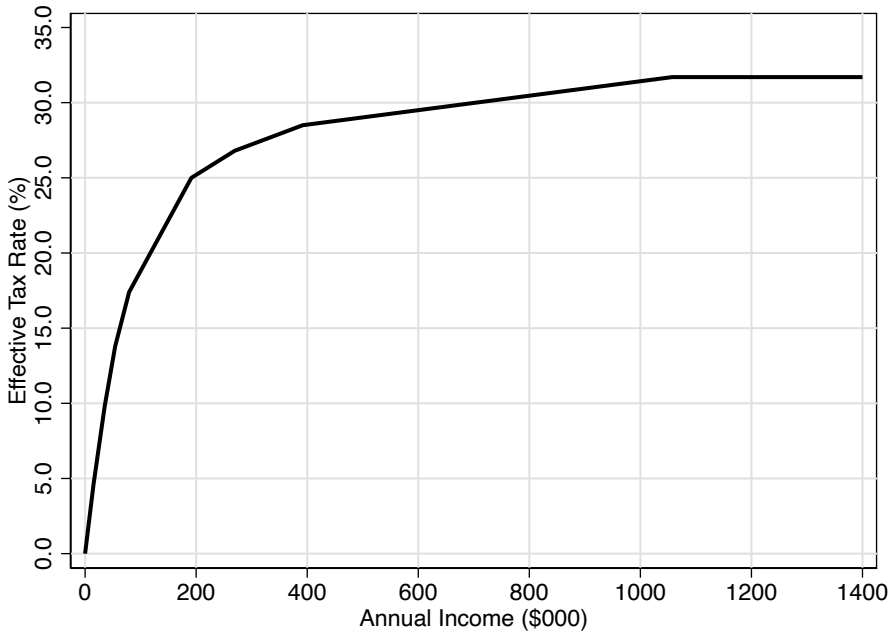
Martellini, Paolo, Todd Schoellman, and Jason Sockin. 2024. "The Global Distribution of College Graduate Quality." *Journal of Political Economy* 132 (2): 434–483.

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 (February).

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).

A Tables and Figures Appendix

Figure A1: 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 exercise taxes. Source: Congressional Budget Office, *Effective Federal Tax Rates, 1979–2004* (December 2006), Table 1.

Table A1: Summary Statistics for NLSY Sample

All Individuals Used in IV Sample							
Variable	Mean	Std. Dev.	Median	Min.	P25	P75	Max.
% Male	0.40	0.49	0.00	0.00	0.00	1.00	1.00
% White	0.70	0.46	1.00	0.00	0.00	1.00	1.00
Age at BA	23.16	2.73	22.00	19.00	21.00	24.00	34.00
Year of BA	2006	3	2006	2001	2004	2007	2015
HH Networth in 1997	138,384	134,914	95,375	250	33,000	197,751	599,001
Avg. HH Income	69,890	48,552	59,676	30	36,253	90,254	285,805
Ability Quartile	3.27	0.86	4.00	1.00	3.00	4.00	4.00
$ Student Loans	17,990	25,203	11,500	-	-	25,750	351,000
Conditional on Positive Student Debt							
Variable	Mean	Std. Dev.	Median	Min.	P25	P75	Max.
% Male	0.38	0.49	0.00	0.00	0.00	1.00	1.00
% White	0.67	0.47	1.00	0.00	0.00	1.00	1.00
Age at BA	23.26	2.71	22.00	19.00	21.00	24.00	34.00
Year of BA	2006	3	2006	2001	2004	2008	2015
HH Networth in 1997	116,109	115,873	79,620	250	27,500	162,500	588,000
Avg. HH Income	62,417	40,421	55,200	30	34,000	80,350	285,805
Ability Quartile	3.24	0.87	3.00	1.00	3.00	4.00	4.00
$ Student Loans	27,259	26,643	20,975	300	12,000	35,000	351,000

Notes: This table provides summary statistics for the NLSY97 population sample that we use in our instrumental variables regression. The top panel includes all individuals in the sample, while the bottom panel includes only those with positive student debt.

Figure A2: IV First Stage Estimates

Effect of 1sd (10ppt) increase in college grant share on:	Total Funding	Grants	Debt	Family Aid	Work Study Aid	Tuition costs
Coefficient	-$160	$7,670	-$5,076	-$23	$197	-$863
(pvalue)	0.95	0.00	0.00	0.43	0.09	0.56

Figure A3: IV Robustness Estimates

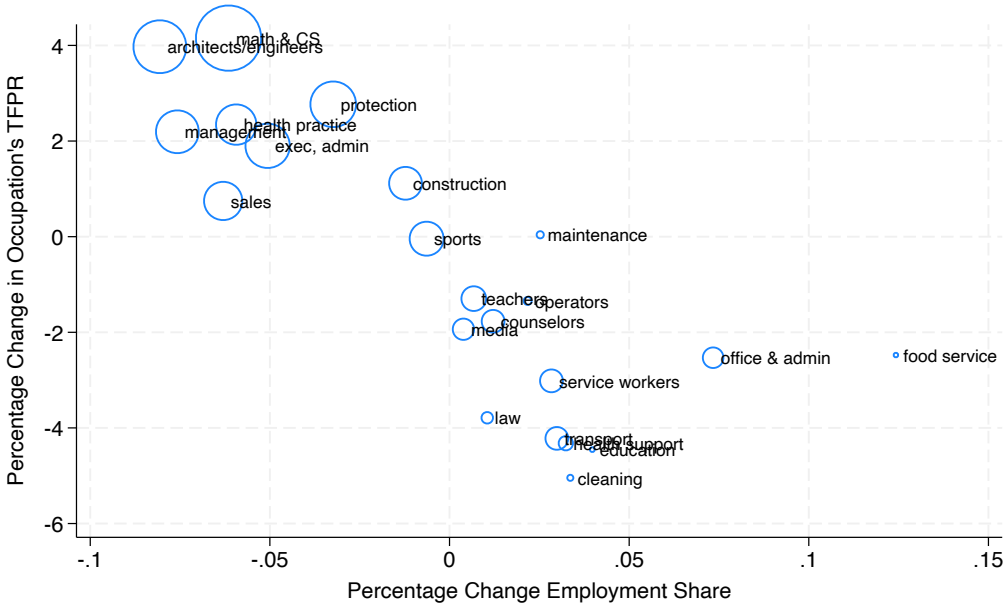
Effect of 1sd (10ppt) increase in college grant share on:	Years at college	Completion rate	1(Full-time)	Age starting college	Ability (percentile)	Parental income	1(White)
Coefficient	0.11	0.14%	0.01	-0.02	1.72%	$1,145	0.38
(pvalue)	0.35	0.28	0.21	0.71	0.23	0.60	0.11

Table A2: 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
Executive and administrative	1.90%	-0.05%	0.02%	0.07%	0.28%	-5.35%	-8.49%
Management related	2.19%	-0.08%	0.01%	0.09%	0.31%	-6.29%	-9.18%
Mathematical and computer scientists	4.15%	-0.06%	0.03%	0.09%	0.54%	-5.96%	-9.06%
Engineers, architects, and other scientists	3.97%	-0.08%	0.01%	0.09%	0.51%	1.29%	-8.34%
Counselors, social and religous	-1.77%	0.01%	0.03%	0.02%	0.02%	-5.83%	-11.22%
Layers, judges, and legal support	-3.79%	0.01%	0.01%	0.00%	-0.02%	-2.77%	-11.82%
Teachers	-1.30%	0.01%	0.06%	0.05%	0.19%	-3.49%	-15.52%
Education, training, and library services	-4.45%	0.04%	0.04%	0.00%	-0.04%	-1.97%	-22.21%
Entertainment and sports	-0.04%	-0.01%	0.03%	0.03%	0.01%	-8.42%	-6.93%
Media and communications	-1.94%	0.00%	0.02%	0.02%	0.24%	-1.04%	-13.18%
Health diagnosing and treating practitioners	2.34%	-0.06%	0.02%	0.08%	1.31%	7.85%	-15.96%
Healthcare technical and support workers	-4.32%	0.03%	0.04%	0.01%	-0.15%	-5.12%	-12.68%
Protective services	2.77%	-0.03%	0.01%	0.04%	1.53%	10.02%	-15.32%
Food preparation and serving	-2.48%	0.12%	0.12%	0.00%	0.32%	5.58%	42.23%
Building and cleaning services	-5.04%	0.03%	0.04%	0.00%	0.09%	-1.02%	-19.40%
Service workers	-3.02%	0.03%	0.04%	0.01%	-0.23%	-8.68%	-8.66%
Sales and related workers	0.75%	-0.06%	0.04%	0.10%	0.18%	-11.17%	-11.22%
Office and administrative support	-2.53%	0.07%	0.12%	0.04%	0.01%	-5.51%	-16.43%
Construction trades and extraction	1.12%	-0.01%	0.01%	0.02%	0.86%	1.38%	-16.94%
Installation, maintenance, and repair	0.04%	0.03%	0.03%	0.01%	-0.18%	-7.11%	-18.89%
Production and operating workers	-1.34%	0.02%	0.03%	0.01%	-0.30%	-8.25%	-18.16%
Transportation and material moving	-4.22%	0.03%	0.04%	0.01%	-0.15%	-5.08%	-8.90%

Notes: Outcomes correspond to short run full forgiveness policy.

Figure A4: 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 A2.

Table A3: 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
Executive and administrative	11.51%	-0.05%	0.38%	0.43%	0.70%	-5.82%	-7.98%
Management related	11.85%	-0.10%	0.37%	0.48%	0.83%	-6.74%	-9.33%
Mathematical and computer scientists	13.39%	-0.10%	0.44%	0.54%	1.29%	-7.83%	-9.27%
Engineers, architects, and other scientists	15.42%	-0.16%	0.32%	0.48%	2.15%	-3.85%	-8.99%
Counselors, social and religious	6.06%	0.10%	0.25%	0.15%	0.59%	-4.74%	-13.33%
Layers, judges, and legal support	2.60%	0.04%	0.09%	0.05%	1.54%	-0.28%	-9.91%
Teachers	6.89%	0.16%	0.56%	0.40%	1.23%	-4.17%	-15.74%
Education, training, and library services	-1.39%	0.12%	0.21%	0.09%	0.92%	-0.21%	-7.53%
Entertainment and sports	7.95%	0.05%	0.26%	0.20%	-0.01%	-6.58%	-8.30%
Media and communications	5.07%	0.07%	0.20%	0.13%	1.60%	-1.76%	-13.81%
Health diagnosing and treating practitioners	18.69%	-0.13%	0.25%	0.38%	6.53%	1.21%	-15.81%
Healthcare technical and support workers	2.89%	0.11%	0.23%	0.11%	0.71%	-3.81%	-13.46%
Protective services	19.19%	-0.06%	0.12%	0.19%	7.28%	2.46%	-15.02%
Food preparation and serving	-2.05%	0.45%	0.65%	0.20%	1.63%	7.16%	3.53%
Building and cleaning services	-2.43%	0.10%	0.18%	0.08%	2.42%	2.16%	-9.66%
Service workers	4.35%	0.15%	0.32%	0.17%	-0.19%	-6.57%	-10.89%
Sales and related workers	10.05%	0.01%	0.63%	0.62%	0.58%	-8.30%	-12.00%
Office and administrative support	4.44%	0.46%	0.94%	0.48%	0.71%	-4.85%	-17.48%
Construction trades and extraction	13.70%	-0.01%	0.08%	0.10%	4.38%	0.22%	-16.65%
Installation, maintenance, and repair	-3.76%	0.08%	0.22%	0.14%	-0.54%	-10.30%	-14.61%
Production and operating workers	-0.22%	0.06%	0.18%	0.12%	0.40%	-9.48%	-15.66%
Transportation and material moving	2.86%	0.11%	0.25%	0.14%	0.56%	-4.15%	-10.66%

Notes: Outcomes correspond to long run full forgiveness policy.

B Data Appendix

This section describes data sets used in empirical and calibration steps and provides additional details on variable construction and sample used. The main datasets used are the Survey of Consumer Finance (SCF) and the National Longitudinal Survey of Youth (NLSY97).

B.1 Student Debts and Initial Assets

We use the assets and education loans data for household heads between 22 and 25 years old from the Survey of Consumer Finance (SCF) to calculate statistics associated with assets and students debts in the model.[34]. We use SCF over NLSY because assets in the NLSY are household assets measured at the beginning of the survey (year 1997), and hence largely reflects parental assets by the time individuals were co-residing with their parents. Furthermore, for married individuals, assets include partner's assets. Since both our empirical and theoretical analysis are on the individual level, SCF provides better counterpart to the initial assets upon graduating from college. Nevertheless, Table A1 shows that the statistics on student debts calculated from SCF and NLSY are quite similar. Those without any educational loans are identified as without student debts when calculating the fraction of indebted population.

We use the summary extract version of 2007 survey and further restrict the sample to those with a BA degree, between 22 and 25 years old, not married or living with a partner, and without children. The selection criteria are motivated by several considerations. Year 2007 is chosen because to avoid any effects from the Great Recession, and also it coincides the survey year when NLSY97 collected information about the amount of student loans. The restrictions on age, marriage, and child status are chosen to provide better counterparts for the initial assets (a) and student debts (d) in the model, which are defined at the individual-level.

B.2 Occupation and Industry Classification

NLSY97 codes the respondents' industry and occupation using the 2002 Census industry and occupation codes. For both employee and self-employed jobs, respondents' verbatim descriptors of their occupations are coded using a three-digit Census code frame. Freelance jobs that do not qualify as self-employment are coded according to the type of work performed.[35] 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

[34]To see how the Survey of Consumer Finance defines net worth, assets, and student loans, see `https://www.federalreserve.gov/econres/files/Networth%20Flowchart.pdf`. Net worth is used to for Figure 1.

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

adjacent occupations together with small sample sizes in order to increase precision. Table A4 and A5 summarize the industry and occupation classifications and the occupational grouping procedure.

Table A4: 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's Workers	Installation, Maintenance, and Repair
Production and Operating Workers	
Setter, Operators, and Tenders	Production and Operating Workers
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

We use the NLSY97 to calculate occupation-specific initial earnings, returns to experience, and the employment shares. We restrict the sample to individuals above age 18, whose highest degree is a 4-year BA degree, and working on the full-time, full year, which is least 35 hours a week and for at least 40 weeks a year. We identify an individual's first occupation as the first full-time,

Table A5: 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

full-year occupation that he/she chooses after obtaining the highest degree. The employment share of each occupation is calculated as the share of each first occupation.

To estimate the occupation-specific initial earnings (IE) and the average returns to experience (RTE), we regress the log earnings on the years of experience for each occupation along with control variables. Earnings is defined as log of total annual income, converted into the real values using the annual CPI. Years of experience is defined as the number of years since obtaining BA degree, hence zero in the starting year. The control variables include dummies for race, gender, region where the respondent resides (Northeast, North Central, South, or West), year and industry fixed effects. Then initial earnings are the predicted earnings with zero years of experience (i.e., the starting period). Annualized returns to experience is predicted annual increase in earnings.

The IE/RTE for high school graduates are calculated in a similar way with the same regression specification, but now with the additional occupation fixed effect. The sample is restricted to those who who reported their highest degree as high school diploma. We drop any individuals who reported earnings less than the federal minimum wage in 2007, which was \$5.85. Since the analysis focuses on the full-time, full year employed population who work at least 35 hours a week and for at least 40 weeks a year, observations with annual earnings less than \$8,190 (\$5.85 $\times$ 35 hours $\times$ 40 weeks) are excluded.

B.4 IV regression

The two important variables in the IV regression are students debts and grant share. Student debt refers to the total loans an individual received for their first college, surveyed in the year 2007. 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.

Annual information on the grant funding, along with other college characteristics, is taken from the Integrated Postsecondary Education Data System (IPEDS) of the National Center of Education Statistics (NCES). These college-level information are merged to the NLSY97 using the confidential college identifiers. Specifically, grant share is defined as the share of grant funding, out of all grant and federal student loan funding, issued by the first college a person attended in a given year.

C Computational Appendix

C.1 Calculation of Consumption-Equivalent Welfare Gains

Variables with superscript b denotes the variables in the baseline (no forgiveness) economy at the steady state, and c in the counterfactual economy under the debt forgiveness policy. For example, c_t^b denotes the simulated time-t consumption at the baseline economy.

Lifetime utility can be decomposed by the utility from lifetime consumption and the occupation-specific amenity values. For instance, the welfare of an individual at the baseline and counterfactual (under forgiveness) economies can be expressed as the following:

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

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

For those who do not switch occupations under a forgiveness policy (stayers), $\varepsilon_p^b = \varepsilon_p^c$, so the entire welfare changes comes from changes in lifetime consumption. For those who switch occupations (switchers), on the other hand, additional welfare change comes from the change in occupation-specific amenities, because $\varepsilon_p^b \neq \varepsilon_p^c$ for them.

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

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

Average welfare gain is the average of individual λs:

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

C.2 Decomposition of Lifetime Earnings Gain

Note that earnings in a given period is 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 is devotes to work. Lifetime earnings is defined as the discounted sum of period earnings:

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

Following the notation above, let LE^b be the lifetime earning 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:

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

In the short-run debt forgiveness, lifetime earnings gains solely comes from the change in effective hours (efficiency units) for job stayers their wage 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.