

Debt, Human Capital, and the Allocation of Talent*

Titan Alon Natalie Cox Minki Kim Arlene Wong

Preliminary Draft
September 2023

Abstract

We empirically and theoretically analyze the impact of household assets and liabilities (namely student debt) on labor market outcomes. Exploiting exogenous variation in the composition of college funding, we find that more student debt leads to higher initial earnings but lower returns to experience. Initial occupation choice out of college plays an important role in driving the results. To explain the data, we develop and calibrate a quantitative model in which occupation choice and lifecycle human capital accumulation interact with credit constraints. Intertemporal distortions arising from credit limits cause households to dis-invest in human capital and switch to occupations with more front-loaded compensation schemes as alternative modes of consumption smoothing. Using the model, we analyze the aggregate productivity and welfare consequences of federal extended repayment and student debt forgiveness programs. The results show that while extended repayment policies always produce gains, the benefits of student debt forgiveness are non-monotonic due to the distortionary effects of redistributive taxation. Moreover, although the fraction of households induced to switch occupations is small, they account for almost half the aggregate gains in labor productivity as relaxed credit constraints induce workers to flow from high amenity to human capital intensive occupations.

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

*We thank seminar participants at Stanford, Princeton, MIT Sloan, NYU, UCSD, USC, Dartmouth, Wharton, the Urban Institute, NBER SI, and the Fed Minneapolis for their valuable comments and suggestions. We thank Marcelo Ferreira, Cristoforo Pizzimenti, and Clemens Lehner for outstanding research assistance.

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% reported they would borrow less if they could repeat college; and nearly one-fifth reported “significantly changing career plans because of student loan burdens”.²

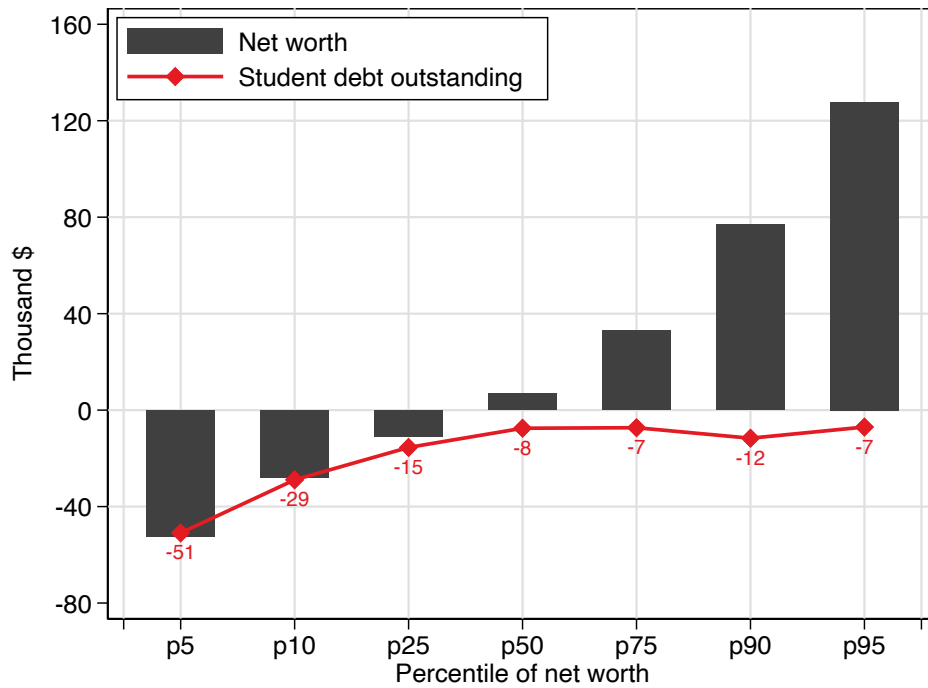
This paper provides a theoretical and empirical analysis of the welfare and productivity consequences of rising student debt. It focuses in particular on student debt’s effect on aggregate labor markets and the early career outcomes of college graduates. We show that students graduating with more student debt have higher initial earnings, but lower returns to experience in the first ten years of their career. Nearly half the effect is mediated by initial occupation choice after graduation, with more indebted students selecting into occupations with more front loaded compensation schemes. We develop a quantitative heterogeneous-agent, incomplete markets model to rationalize the findings and study their implications for federal student loan policies. The results show that budget-neutral reforms which alter the principal or duration on outstanding federal student loans can increase aggregate welfare and labor productivity. The gains arise from reducing credit constraints that inhibit human capital accumulation on the job and give rise to a misallocation of talent by reducing the efficiency of occupational sorting.

While our findings 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,

¹See [Federal Reserve Bank of New York \(2022\)](#) for details on the composition of household debt.

²See [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 total value of aggregate loan balance of education-purpose expenses. Sample is restricted to BA degree holders between age 22 and 25.

student debt is largely non-dischargeable, allowing us to largely abstract from strategic delinquency and bankruptcy considerations. Empirically, households are much less likely to become seriously delinquent (90-days or more) on their student debt than on other forms of credit, such as credit card debt or auto loans.³ Finally, the federal government plays an out-sized role in student debt markets, accounting for 92.7% of the total outstanding student loan debt.⁴ Changes in federal 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.

³In 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 \(2022\)](#).

⁴Student 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), we provide empirical evidence on the effect of student debt on early career labor market outcomes. We instrument 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 3.14%, but *reduces* the returns to experience by 1.37%. The effect on initial earnings corresponds to an additional \$508 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 et al. 2021). Finally, nearly half the estimated effect on earnings is mediated by initial occupation choice after graduation, as more indebted students select into occupations with more front loaded compensation schemes.

To understand the data, we develop 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 dis-invest 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 effects of student loan debt on aggregate labor productivity.

To study the implications for federal policies, we embed these mechanisms 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 to attend college, accounting for any selection effects that may be in the data. After graduation, households

⁵For instance, see Rothstein and Rouse (2011), Chapman (2015), and Luo and Mongey (2019) for comparable estimates of the impact of student debt on initial earnings.

choose an occupation based on their innate abilities and financial assets, taking prevailing wages as given. Earnings evolve endogenously over their lifecycle as a consequence of costly investments in human capital and idiosyncratic labor market shocks. Households are also subject to progressive taxation, face idiosyncratic human capital shocks, receive retirement benefits which depend on their earnings, and face realistic student debt repayment provisions and borrowing constraints. The calibrated model matches the aggregate earnings profile and the joint distribution of family assets, student debt, and college matriculation. It also replicates the employment shares, initial earnings, and returns to experience in the 18 detailed occupation groups households choose between.

Using the quantitative model, we measure the aggregate welfare and productivity consequences of budget-neutral changes to the repayment duration or principal of federal student loans. The first exercise computes the effect of 2 and 5 year extensions to the standard federal repayment program. In each case, the interest rate on outstanding federal loans is recomputed to ensure that the net present value of individual student debt does not change. The model predicts these extensions would meaningfully boost welfare and labor productivity, with the gains increasing in the program's duration. For instance, a 5 year budget-neutral extension to the standard federal repayment program would increase aggregate labor productivity by 0.57% and consumption equivalent welfare by 0.45%, with benefits concentrated among low wealth households. The results also show that induced occupation switchers contribute negligibly to aggregate welfare gains, but account for one-third of the rise in aggregate labor productivity despite representing only 0.41% of the treated population. Decompositions of the aggregate productivity increase reveals that the out-sized contribution of induced-switchers is the result of a systematic pattern in the direction of occupation switching; the policy induces workers to flow predominantly from low-skill service jobs into more human capital intensive occupations, mainly Sales, Engineering, Math and Computer Science.

The second set of exercises examines the effect of reducing the principal of outstanding federal student loan debt via student debt forgiveness. It computes the aggregate and distributional effects of student debt forgiveness capped at the 10k, 50k, and 100k thresholds. Given the empirical distribution of student debt, the 100k-cap program effectively amounts to complete debt forgiveness. To ensure the program is budget-neutral, the debt forgiveness is funded by income taxes which are distributed across households with the same proportionality as the overall U.S. tax system.

The results show that the impact of debt forgiveness policies is non-monotonic in the

size of the program, with middle-sized programs performing the best. The 50k-capped program yields a modest 0.02% rise in consumption equivalent welfare and a 0.20% increase in aggregate labor productivity. In contrast, the 100k-capped program yielded only a 0.09% increase in labor productivity and a decline in consumption equivalent welfare, while the small 10k program lead to aggregate declines in both welfare and productivity. The non-monotonicity arises from the countervailing effects of debt forgiveness and the distortionary taxation which funds them. The small 10k program is insufficient to substantially alleviate credit constraints, despite still requiring a substantial rise in tax revenue. The 100k program virtually eliminates student debt, but requires large increases in distortionary taxation that reduce the benefits by discouraging human capital accumulation on their own. The model suggests that middle-sized programs which are large enough to alleviate credit constraints, but no so large as to substantially increase distortionary taxes, deliver the best returns. Overall, the results suggest that extended repayment duration policies are a better tool than student debt forgiveness. While both alleviate credit constraints, the former does not require distortionary taxes to support transfers across households.

Related Literature. Our 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. Our work complements research on credit frictions in the financing of higher education ([Lochner and Monge-Naranjo 2012](#); [Lochner, Stinebrickner, and Suleymanoglu 2021](#)) by investigating how student debt subsequently effects labor market outcomes after graduation.

Empirically, this 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 household lifecycle outcomes, such as homeownership, marriage, fertility, and attending graduate school ([Goodman, Isen, and Yannelis 2018](#); [Chakrabarti et al. 2020](#)). In particular, we provide additional evidence of how debt effects lifecycle earnings through distortions to occupation choice ([Rothstein and Rouse 2011](#); [Luo and Mongey 2019](#); [Herkenhoff, Phillips, and Cohen-Cole 2021](#)).

Our analysis employs a dynamic stochastic heterogeneous agent model of lifecycle earn-

ings with incomplete markets. Following the literature, we analyze the effect of policy reforms by calibrating the model with microdata on household assets, student debt, education, and labor market outcomes (Ionescu 2009; Huggett, Ventura, and Yaron 2011; Abbott et al. 2019; Fu, Lin, and Tanaka 2021). A novel feature of our framework is the inclusion of 18 distinct occupations which, based in part on the endogenous sorting of workers, exhibit different earnings and returns to experience. In this sense, our paper is most similar to Luo and Mongey (2019) who develop a quantitative model of how student debt effects household earnings and occupation choice. We build on their contribution by showing how student debt reduces the returns to experience by incentivizing households to sort into occupations with front loaded compensation schemes. As a result, while Luo and Mongey (2019) find that reducing student debt leads workers to sort into occupations with greater amenity value, we find that workers flow into occupations with greater scope for earnings growth. Taken together, our results suggest that greater scope for human capital accumulation on the job may be one particular non-wage amenity driving the sorting patterns they uncover.

The remainder of the paper is organized as follows. Section 2 presents a simple model to show how credit constraints lead to intertemporal distortions that can reduce human capital investment and give rise to a misallocation of talent. Section 3 presents empirical evidence on the effect of student debt lifecycle earnings profiles and occupation choice. Section 4 describes the quantitative model, calibration strategy, and reports the model fit. Section 5 presents the computational results. Section 6 concludes.

2 An Illustrative Model

This section introduces a simple model of human capital accumulation and occupation choice in the presence of credit constraints. When constraints bind, households disinvest in human capital accumulation as an alternative mode of consumption smoothing. Section 2.1 describes the intertemporal distortions to human capital accumulation on the job. Section 2.2 shows how credit constraints also effect aggregate human capital by distorting occupation choice, pushing households towards occupations with front loaded compensation schemes over those which are best matched to their abilities.

Households live for two periods, young (y) and old (o), and choose consumption, savings, an occupation, and how much to invest in human capital accumulation. Each household is endowed with one unit of time in each period and begins life with initial assets a_y . The problem of a household in occupation k is given by

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

subject to

$$c_y = w_k(1 - s) + a_y - a_o$$

$$c_o = w_k h(\theta_{ik}, s) + (1 + r)a_o$$

$$a_o \geq -\bar{a} \quad , \quad s \in [0, 1]$$

where $h(s, \theta)$ is the human capital in adulthood of a household with talent θ_i who invested s time in human capital accumulation. Human capital during youth is normalized to one. The human capital technology $h(s, \theta)$ is an increasing and concave function of time invested,

$$\frac{\partial h}{\partial s} > 0, \quad \frac{\partial^2 h}{\partial s^2} < 0$$

To interpret θ_{ik} as an index of individual talent which makes it easier for individual i to accumulate human capital in occupation k , we impose that

$$\frac{\partial h}{\partial \theta} > 0, \quad \frac{\partial}{\partial \theta} \frac{\partial h}{\partial s} > 0, \quad \frac{\partial}{\partial \theta} \frac{\partial^2 h}{\partial s^2} < 0$$

Most widely used models of human capital accumulation satisfy these criteria, such as [Ben-Porath \(1967\)](#) and [Mincer \(1974\)](#). Households realize their occupation specific abilities $\Theta = \{\theta_k\}$ as soon as they enter the labor market and, taking wages w_k as given, choose the occupation which maximizes their discounted lifetime utility. Formally, the optimal occupation choice is

$$k^* = \operatorname{argmax} \{ V_1, V_2, \dots, V_K \}$$

The credit constraints appear in the parameter \bar{a} which limits the amount households can borrow against future income. The key economic friction is that households cannot collateralize their human capital in order to loosen present day borrowing constraints.

2.1 Intertemporal Distortions to Human Capital

In the absence of frictions, households invest in human capital accumulation until the marginal return of further investment equals the return on physical capital. The optimal

investment in human capital s^* is given by,

$$\frac{\partial h(\theta_{ik}, s^*)}{\partial s} = 1 + r \quad (1)$$

When borrowing constraints bind, households discount future income streams relative to the present at a rate that is greater than the market interest rate $1 + r$. As a result, households dis-invest in human capital as an alternative form of consumption smoothing. When credit constraints bind, households invest in human capital accumulation until the marginal return equals the shadow interest rate, $1 + r^c$, so that

$$\frac{\partial h(\theta_{ik}, s^c)}{\partial s} > 1 + r$$

Consequently, credit constraints lead households to invest less in human capital accumulation, $s^c < s^*$, leading to higher initial earnings

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

but lower returns to experience

$$\frac{h(\cdot, s^c)}{(1 - s^c)} < \frac{h(\cdot, s^*)}{(1 - s^*)}$$

Moreover, total lifecycle human capital accumulation is reduced for workers facing credit constraints, leading to a reduction in aggregate labor productivity alongside changes in the structure of lifecycle earnings.

2.2 The Misallocation of Talent

In addition to reducing investment on the job, households can also respond to credit constraints by switching occupations. Limited ability to borrow against future income leads some households to switch away from occupations that offer opportunities for human capital accumulation on the job to those with more front-loaded compensation schemes. As a result, occupation choice will depend on initial household assets, giving rise to a misallocation of talent as some workers select into occupations for which their abilities are not optimally matched.

A key condition for credit constraints to give rise to a misallocation of talent is heterogeneity in occupational wages w_k . The variation in wages provides a trade-off in the

level and steepness of the lifecycle earnings profile when moving across occupations in a manner similar to the *within* occupation trade-offs offered by training on the job, s , reviewed in section 2.1. In particular, with heterogeneity in w_k , each occupation will have its own cutoff $\bar{\theta}_k$ such that all workers with talent $\theta_k > \bar{\theta}_k$ will be constrained. In other words, it is again the high ability individuals that will be most effected by the credit constraints. Intuitively, the cutoff $\bar{\theta}_k$ corresponds to the ability level at which individuals expect sufficiently rapid earnings growth during their career that they would want to borrow beyond the limit \bar{a} . Formally, the cutoff can be expressed implicitly by,

$$\frac{\beta}{1+\beta} [a_y + w_k(1 - s^*)] - \frac{1}{(1+\beta)(1+r)} w_k h(s^*, \bar{\theta}_k) = -\bar{a} \quad (2)$$

where s^* is optimal human capital investment on the job, as defined in equation (1), and the whole left hand side corresponds to optimal asset holdings a^* in the absence of constraints. The expression shows how the cutoffs depend on both initial household assets, a_y , and vary across occupations with w_k . Consistent with economic intuition, the cutoffs in all occupations are increasing in initial household assets a_y , so that it is highly talented individuals from poor households that are most effected by the constraints.

To illustrate the misallocation of talent, it is useful to consider a simple parametric case with log utility and human capital technology $h(s, \theta) = \theta^{1-\alpha} s^\alpha$, which satisfies the assumptions made above. Suppose households can choose between two occupations where, without loss of generality, occupation $k = 2$ offers higher wages, so $w_2 > w_1$. In the absence of binding credit constraints, households select the occupation which offers them the highest present discounted value of lifetime earnings. Given a realization of occupation specific abilities, the condition reduces to choosing occupation 1 if

$$\theta_1 > \frac{w_2 - w_1}{\kappa w_1} + \frac{w_2}{w_1} \cdot \theta_2 \quad (3)$$

where $\kappa = \left(\frac{1}{1+r}\right)^{\frac{1}{1-\alpha}} \left[\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}\right] > 0$. The condition in (3) is economically intuitive: it is optimal to choose the occupation offering a lower wage provided one has sufficiently high ability to generate greater overall lifetime earnings.

In contrast, we can consider an individual with high level of debt (or sufficiently low assets) such that they are constrained in both occupations. In this case, they value not

only discounted lifetime earnings, but also the timing of that income. Given a realization of occupation specific abilities, an individual chooses occupation 1 if

$$\theta_1 > \left(\frac{w_2 + a_y}{w_1 + a_y} \right)^{\frac{1+\alpha\beta}{(1-\alpha)\beta}} \cdot \frac{w_2}{w_1} \cdot \theta_2 \quad (4)$$

Comparing the occupational sorting rules in equation (3) to (4), one can see that there are individuals $\Theta = \{\theta_1, \theta_2\}$ who would choose occupation 1 in the absence of credit constraints, but choose occupation 2 if they were financially constrained by debt or insufficient assets.⁶ The extent to which these distortions to occupation choice effect aggregate outcomes depends on the joint distribution of assets, a_y , and talent, Θ , in the population. The more highly talented individuals are encumbered with large debts, the greater will be the effect on aggregate labor productivity.

Figure 2 illustrates the misallocation of talent by plotting how occupational sorting depends on the presence of credit constraints. For a given initial assets a_y , it shows how individuals sort into each occupation as a function of their talents. The cutoffs $\bar{\theta}_1$ and $\bar{\theta}_2$ are defined as in (2), indicating the regions where credit constraints bind in each occupation.⁷ The shaded region represents the population of workers who switch from occupation 1 to occupation 2 in the presence of credit constraints. The lower border of the region corresponds to the optimal occupation sorting rule in the absence of credit constraints expressed in (3). The low ability population, those with $\theta_1 < \bar{\theta}_1$ and $\theta_2 < \bar{\theta}_2$, continue to choose occupations according to this rule, since they are sufficiently untalented that credit constraints do not effect them. The high ability population, those with $\theta_1 > \bar{\theta}_1$ and $\theta_2 > \bar{\theta}_2$, are most affected by the constraints and sort according to equation (4). 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 hence switch to the latter.

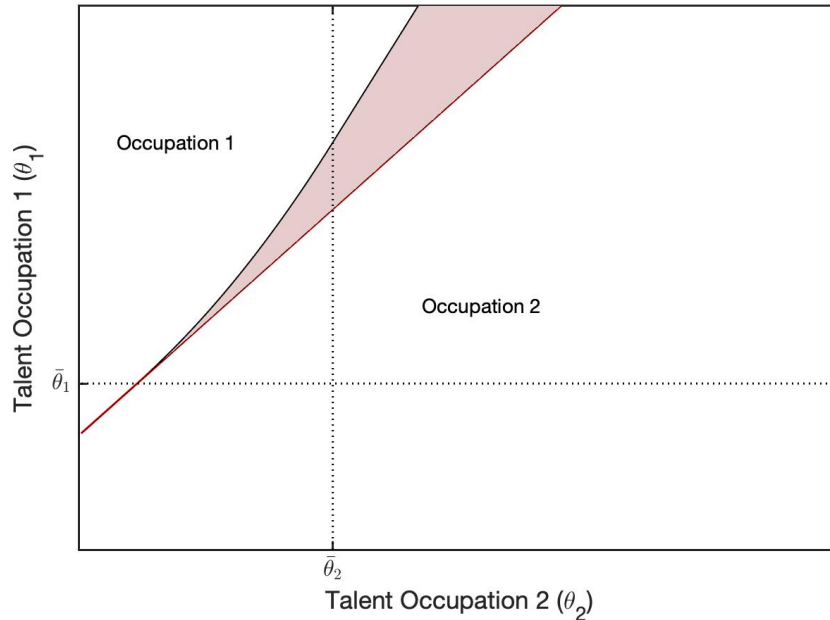
Moving to populations with higher assets, the cutoffs $\bar{\theta}$ increase and the upper bor-

⁶To see this more directly, note that the occupation sorting condition in (4) 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.

⁷Given the parametric example being considered, equation (2) can be solved explicitly so that

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

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, $w_2 > w_1$.

der of the misallocation region shifts inward toward the optimal sorting rule, shrinking the region where credit constraints distort occupation choices. Conversely, moving to households with lower assets leads the $\bar{\theta}$ cutoffs to shrink and the misallocation region to expand, resulting in greater distortions to occupation choice. As a result, the aggregate effect of the credit frictions will depend on the joint distribution of assets a and talents Θ in the population.

This joint distribution will also determine the marginal effect of policies like student debt forgiveness, since it will determine on the mass of workers near the threshold of misallocation region. As discussed in section 2.1, policies which alleviate credit constraints will boost human capital accumulation and alter the lifecycle earnings profile by reducing the shadow interest rate. For the population near the misallocation region, alleviating credit constraints can also induce occupation switching that leads to discrete adjustments in both productivity and lifecycle earnings. Figure 3 illustrates the effect, showing how initial earnings, returns to experience, and occupation choice of a constrained household responds to reductions in their debt. Initial reductions in household debt lead to increased human capital accumulation, reduce initial earnings and boost-

Figure 3: Household Assets and Occupation Switching



Notes: This figure illustrates the effects of reducing an individual household’s debt level on lifetime utility, initial earnings, and returns to experience. The simulation corresponds to an individual who has higher ability in occupation 1, $\theta_1 > \theta_2$, while occupation 2 offers the higher wages, $w_2 > w_1$.

ing returns to experience on the job. Eventually, household debt falls sufficiently far that the household passes out of the misallocation region and switches from occupation 2 to occupation 1, eschewing the higher initial earnings for greater lifetime earnings. The occupation switch leads to a large, discrete jump in earnings and returns to experience. After the switch, further debt reductions continue to effect human capital accumulation within the occupation, albeit with slightly larger effects in the more human capital intensive occupation.

The illustration shows how endogenous occupation choice can mediate part of the impact of credit constraints on aggregate labor productivity and household earnings. Assessing the true aggregate contribution of these channels therefore requires measuring the joint distribution of assets and talent in the relevant population which will together determine the within and between occupation effects of credit frictions, as well as how they respond to policy interventions. We turn to addressing these challenges of measurement in the following sections, both empirically by exploiting plausibly exogenous variations in student debt as well as structurally through the calibration of a larger scale quantitative model which embeds these core mechanisms.

3 The Empirical Evidence

In this section, we provide reduced-form empirical evidence that individuals' earnings trajectories change in the presence of student debt, and that at least part of this effect is explained by occupational choice. We use panel data from the NLSY 1997 and an instrumental variables design to first estimate the impact of student debt on both initial earnings and returns to experience.

Our findings suggest that individuals with high student debt are forced to trade off between current and future income early in life – initial earnings are higher for those with student debt, while returns to experience are lower. Moreover, this trade-off appears to depend significantly on their *first* occupation and industry choice upon graduation. When we include occupation and industry fixed effects in the IV regression, they explain almost half of our estimated effects.

To further understand why occupational choice matters so much for our estimates, we conduct a supplementary analysis to test whether earnings trajectories vary considerably across occupations. We use the Current Population Survey to construct age-earnings profiles for each occupation in order to investigate how differential sorting into occupations potentially drives our earnings results. In the raw data, we indeed find a strong negative correlation between initial earnings and returns to experience across occupations. These empirical findings support the set-up of our theoretical model of occupational choice, student debt and human capital accumulation, and motivate our quantitative analysis.

3.1 Data

Our empirical analysis draws from several data sources. The primary dataset is the NLSY 1997, an individual-level panel dataset that contains information on higher education, student debt, and labor outcomes. It follows individuals from 1997 through 2015. Summary statistics are provided in Table 6 in the appendix. We restrict our analysis to individuals whose highest level of education is a bachelors degree.⁸

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

⁸We have re-run our analysis while also including individuals with an Associates degree. While the first and second stage results are largely consistent with those we find on the BA-only population, we refrain from making this our primary sample since BA and AA degree recipients make very different human capital investment decisions in college and have markedly different observed occupational choices post-degree.

labor market decisions and lifetime earnings trajectories. To construct our 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.

Supplementary analysis of occupation-specific age-earnings profiles uses the Current Population Survey. While the CPS does not contain information on student debt, it does contain comprehensive earnings information across the spectrum of occupations, industries, and ages.

3.2 Instrumental Variable Design

To estimate the effect of student debt on initial earnings and returns to experience, we consider the following equation:

$$\begin{aligned}
 y_{it} = & \underbrace{\alpha_0 + X_{it}\beta}_{\text{initial (log) earnings if no student debt}} + \underbrace{\alpha_1 \text{Exp}_{it}}_{\text{returns to experience if no student debt}} + \\
 & \underbrace{\alpha_2 \text{SD}_{it}}_{\text{effect of student debt on initial (log) earnings}} + \underbrace{\alpha_3 \text{SD}_{it} \times \text{Exp}_{it}}_{\text{effect of student debt on returns to experience}} + \epsilon_{it} \quad (5)
 \end{aligned}$$

where y_{it} is an outcome measure of individual i in year t for annual log earnings. The variable SD denotes the level of student debt. The variable Exp denotes the years of experience. The variable X_{it} includes gender, race and cohort fixed effects.

We seek an unbiased and consistent estimate of α_2 and α_3 . The effect of student debt on initial wages is measured by α_2 . The effect of student debt on the returns to experience is measured by α_3 . There are potential challenges to estimating equation 5 using OLS. For instance, we may be concerned that there is a correlation between the level of debt an individual takes on and the unobservable quality or ability of an individual. The bias can go either way. Individuals with high ability may expect to have higher future wage growth. They may 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 .⁹

⁹These identification challenges to identifying a causal impact of student debt on earnings are also

To address these identification challenges, we estimate the causal impact of student debt on earnings using an school-cohort-level instrumental variable. Our instrument follows that used in [Luo and Mongey \(2019\)](#) — 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, our instrument is defined as:

$$Z_{c(i),j} = \frac{\text{total grants}_{c(i),j}}{\text{total grants}_{c(i),j} + \text{total loans}_{c(i),j}} \quad (6)$$

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 \(2019\)](#), 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. To meet the exogeneity assumption of a valid instrument, we argue 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 10 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 compensated for *entirely* and *exclusively* by changes in student debt, not other sources of funding. Total funding for college remains constant in response to one standard deviation increase in the college grant share. And 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.

highlighted by ([Field 2009](#); [Rothstein and Rouse 2011](#); [Luo and Mongey 2019](#)). These papers uses variation in forgiveness of debt ([Field 2009](#)) and variation in grants ([Luo and Mongey 2019](#)) 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.

We also check whether variation in the college grant share changes other important educational outcomes which could confound our results, like the probability of college completion or the ability and wealth distribution of enrolled students. These results are shown in Table 11. Similar to [Luo and Mongey \(2019\)](#), we do not find a significant evidence of our instrument impacting enrollment or student selection on observables.

In our second stage, in which we regress our instrumented student debt variable on earnings, we include 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, our small sample size does not allow this – we have very few instances in which more than one student attended the same institution.

3.3 Estimated Impact of Student Debt on Initial Earnings and Returns to Experience

Using our instrument, we next investigate whether those who have subsequently higher debt choose jobs with significantly different earnings profiles. Table 1 shows the instrumented regression coefficients for α_2 and α_3 for log earnings. The coefficient on α_2 is positive, while the coefficient on α_3 is negative. These coefficients imply that an individual with more student debt has higher initial earnings upon graduation, but subsequently lower returns to experience.

To interpret the magnitudes, the coefficients in column (I) imply that a additional \$1K of student debt increases initial earnings by 3.14%. For our data sample, this equates to an additional \$508 annual earnings upon graduation, for every \$1K of additional student debt. While we use a different sample from the existing literature, we arrive at estimates on the effects for initial earnings that are consistent with existing estimates.¹⁰

Our new empirical evidence is shown in columns (II). Specifically, we find that total earnings grow by 1.37 ppts *slower* per year of experience respectively, for every \$1K of additional student debt. These effects are statistically significant. The magnitudes are also sizable, given the earnings of 25 – 30 year olds are estimated to grow at a rate of 7.75% on average each year ([Guvenen et al. 2021](#)).

We explore what factors explain the wage gap between those with and without student debt by including different controls. The evidence suggests that a large part of the wage gap between individuals that graduated with and without student debt is due to the se-

¹⁰For instance, [Luo and Mongey \(2019\)](#), [Rothstein and Rouse \(2011\)](#), [Chapman \(2015\)](#).

Table 1: IV coefficients of student debt on initial earnings and returns to experience, with and without first occupation fixed effects

Effect of student debt (\$000s) on:	Grants-based IV	
	(I)	(II)
(i) Log initial earnings (pvalue)	3.14% 0.08	1.56% 0.13
(ii) Mean returns to experience (pvalue)	-1.37% 0.08	-0.78% 0.01
Controls:		
Ability, age, college-type, race, gender	Yes	Yes
Occupation FE & Occupation x Exp		Yes
Source of variation:	Across-cohort, within college-type (among all students)	
Occupation FE & Occupation x Exp		
	No	Yes
Average initial earnings		50%
Average returns to experience explained		43%

Notes: The Table reports the instrumented estimates from regression 5 using the NLSY data. Our IV utilizes changes in the college-year grant share, which in turn impacts the amount of student debt taken out by individual students. Our dependent variable is log yearly earnings. See text for more details.

lection into different occupations upon graduation. The first occupation choice accounts for almost half of the gap in wage profiles between the two groups.

To see this, Column (II) includes fixed effects for the first occupation that an individual chooses upon graduation. We also include the interaction of the fixed effects with the years of experience. The inclusion of these controls reduces the initial earnings gap by 50% (α_2 declines from 3.14 to 1.56). The marginal effect of student debt on returns to experience declines by 43% (α_3 declines from -1.37 to -0.78). These changes imply that initial occupational choices can explain much of the difference in the subsequent earnings profiles between those with and without student debt. These results imply

student debt impacts the selection into different occupations upon graduation.

3.4 The Role of Occupational Choice and Other Mechanisms

Our instrumental variables analysis finds that additional student debt generates age-earnings profiles with initially higher earnings and lower returns to experience. About half of this effect is explained by occupational fixed effects, meaning that individuals with student debt sort into professions that have a predictably flatter, front-loaded income trajectory. We provide some suggestive evidence in the data on potential sources for the heterogeneity in occupational earnings trajectories, and then test for sorting of individuals with student debt into occupations that are characterized by a flatter profile.

For this analysis, we use the Annual Social and Economic Supplement (ASEC) of the Current Population Survey (CPS) from 2010-2018 to construct occupation and industry specific age earnings profiles. These are cross-sectional profiles, meaning that the trajectory is constructed *across* individuals of different ages, rather than *within* person over time. While this technique is sensitive to changes in cohort composition over time, it allows us to look at longer trajectories and at more specific occupation groups. We pool multiple cross-sections from the CPS in order to control for some cohort effects.

We use a quadratic specification to estimate occupation-specific earning trajectories. In each 2-digit occupation category, we fit the regression specification $earnings_i = \alpha + \beta_1 age_i + \beta_2 age_i^2 + X_i + \epsilon_i$ to the cross-section. We restrict the sample to individuals who are between 21-60 years old and have exactly a bachelors degree; we control for household size, race, gender, state, and CPS cohort.

While we do not claim that these profiles are 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 our regression results. We find that there is considerable heterogeneity across occupations in the level of earnings received upon graduation and returns to experience. If we summarize the age-earnings trajectories by their intercept and slope, the trade-off that individuals face between front-loaded wages and lower growth becomes more apparent. As shown in Figure 4 there is a statistically significant negative relationship between occupations' initial log earnings after graduation (y-axis), and their average yearly growth rate in the first 15 years (x-axis).

An individual who is constrained at age 23 by student debt may choose a career in

Health Practice rather than in Law, since it provides higher initial wages that can be used to service the debt (see Figure 4). The debt constraint would thus induce them to choose a career with potentially lower expected net present value. For constrained individuals, the earnings trajectory becomes an imperfect, costly means of transferring consumption from the future to the present.

It is worth noting that these are equilibrium earnings trajectories, which encompass both endogenous labor supply decisions and occupation-specific technological differences. However, these differences in occupational earnings trajectories may stem from inherent technological differences. Performance, and consequently earnings, in some occupations may benefit more from on the job training. This will lead to a larger impact of human capital investment on returns to experience, and thus steeper earnings profiles, in particular occupations.


We next test whether individuals with higher student debt sort into occupations with a flatter, front-loaded earnings profile. We characterize occupation-specific earnings profiles using the annualized earnings growth rate between ages 23 and 48 from the CPS. We then use our instrumental variables regression from the NLSY to test whether individuals with more debt choose occupations with lower growth rates. There is a statistically significant negative relationship between instrumented student debt and occupation specific growth rate of earnings. This confirms our earlier IV results, that found instrumented student debt had a positive effect on earnings and a negative effect on returns to experience. A similar negative relationship exists for the OLS regression.

Using the OLS coefficients, we then predict the average debt level for individuals by occupation. Figure 4 plots the predicted student debt levels by occupation using differently sized bubbles. Large bubbles have relatively high predicted student debt levels compared to small bubbles. In line with our model predictions, high debt occupations also have higher initial earnings and flatter profiles – i.e. they are concentrated in the upper left hand portion of the plot. This graph not only documents the slope-intercept earnings trade-off across occupations, but also the sorting of constrained individuals into front-loaded options.

3.5 Summary of Empirical Evidence

In summary, our empirical estimates find that exogenous increases in student debt leads to higher initial earnings and lower returns to experience. Our evidence suggests that

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



Notes: This figure plots the predicted student debt levels by occupation using differently sized bubbles. Large bubbles have relatively high predicted student debt levels compared to small bubbles. In line with our model predictions, high debt occupations also have higher initial earnings and flatter profiles – i.e. they are concentrated in the upper left hand portion of the plot.

part of the effect comes from *within* occupation changes in earnings. The other part of the effect comes from changes in earnings resulting from the first-occupation choice upon graduation. These empirical findings support the intuition behind our illustrative model, which explicitly models the human capital investment and occupational choices of credit constrained individuals.

4 The Quantitative Model

In this section, we construct and calibrate a quantitative heterogeneous agent, incomplete markets model with occupation choice and on-the-job human capital investment.

Each period corresponds to five years. Life begins at age 18 when individuals are endowed with initial assets (a), realize their occupation specific talents (θ_k), and decide whether or not to go to college subject to a secondary school taste shock (ζ). Individuals

who decide to go to college may endogenously incur student debt d if they have insufficient assets to cover tuition. At age 23, graduates enter the labor market and choose an occupation $k \in \{1, 2, \dots, K\}$ to maximize their expected lifetime income.

$$k^* = \operatorname{argmax} \{ V_1, V_2, \dots, V_K \}$$

Working life continues until retirement at age 63; all households die at age 78. Households can be identified by their assets (a), human capital (h), student debt (d), occupation (k), employment status (z), and age (t). The problem of an employed working age household can be expressed recursively,

$$V_k(a, h, d, e, t) = \max_{c, s, a'} u(c) + \beta \mathbb{E} [V_k(a', h', d', z', t + 1)]$$

subject to

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

$$h'_k = \exp(z')H(s, h, \theta)$$

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

$$a' \geq -\bar{a} \quad , \quad s \in [0, 1]$$

where households have CRRA preferences $u(c) = \frac{c^{1-\rho}-1}{1-\rho}$. As in [Huggett, Ventura, and Yaron \(2011\)](#), human capital accumulation is risky and subject to log-normal idiosyncratic shocks, z' , with a [Ben-Porath \(1967\)](#) learning technology given by

$$H(s, h, \theta) = \theta_{ik}(sh_k)^\alpha + (1 - \delta)h_k$$

where h prevailing human capital, θ is individual talent, and s is the time invested in training. The functions $\mathbb{T}(w_k(1 - s)h_k)$ and $\phi(a, h, d, z, t)$ represent the tax system and student debt repayment rule, respectively, which we explain further below.

Student Debt and College Matriculation Households endogenously incur student debt when making the decision whether or not to go to college. Individuals make their matriculation decision at age 18—the beginning of life in the model. High school graduates solve the same lifecycle consumption-savings problem as college graduates above,

but make no occupation choice and have no student debt. We also abstract from earnings heterogeneity among high school graduates, allowing high school talent to be the same for all graduates and calibrating it to match the lifecycle income growth of high school graduates, as in [Lee and Seshadri \(2019\)](#). Households optimally enroll in college if the lifetime utility of doing so, given their chosen college occupation k^* , surpasses what they would get as a high school graduate. The college matriculation decision also includes a high school taste shock ζ to capture unobserved heterogeneity and smooth out the optimal enrollment policy.

If individuals decide to attend college (e.g. $V_{k^*} > V_{hs}\zeta$), then they will incur student debt if their initial assets are insufficient to cover the cost of college tuition net of any grants or family assistance they receive. Formally, letting τ denote tuition net of any grants or family assistance, household debt is given by

$$d = \begin{cases} 0 & \text{if } V_{k^*} < V_{hs}\zeta \\ \min\{0, a_0 - x \cdot \tau\} & \text{if } V_{k^*} > V_{hs}\zeta \end{cases}$$

where we note that the value of college $V_{k^*}(a, h, d, z, t)$, and hence the optimal college matriculation decision, implicitly depends on the amount of college debt households would take on. To capture the variety of individual circumstances determining access to college grants and family assistance, we assume that the out-of-pocket net tuition individuals must pay to attend college depends stochastically on family background. Parameter x captures extensive probability of having student debt; with probability $1-x$, an individual is able to attend college without incurring any out-of-pocket expenses. With probability x , individuals receive an out-of-pocket net-tuition cost τ of attending college. In calibrating the model, we allow probability x to depend on family assets, and conditional on facing out-of-pocket expenses, we assume net tuition costs are stochastic and jointly log-normal in the population such that,

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

We choose the parameters governing x and τ , which determine the net-tuition offers households receive, to match the *post-matriculation* realizations of student debt by family background. The correlation between family assets, a , and the extensive margin of debt,

x , is implemented via a copula. The flexible parameterization allows us to capture the endogeneity of student debt and patterns of selection into college that we observe in the data. Finally, we note that even after receiving grants and financial aid, some households may still have insufficient assets to pay for college and so are unable to matriculate even if very talented. These households highlight the key credit market friction that motivates government education loans and grants in the first place: individuals cannot borrow against their future human capital.

Student Debt Repayment Households begin repaying student loans after they graduate college. The payments a household makes to service its student loan debt depend on its outstanding balance d , time to maturity $\bar{T} - t$, household assets, potential earnings, and employment status z . The repayment function $\psi(a, h, d, z, t)$ summarizes how repayments depend on individual circumstances.

In normal circumstances, a household with sufficient financial resources will make payments $\rho(d, t)$ to amortize its student loan over a repayment period $t < \bar{T}$, as in [Luo and Mongey \(2019\)](#), given by

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

Due to the stochastic human capital risk, it is possible that households find themselves unable to make their student loan payments. This may happen if a household has experienced long sequences of negative shocks which result in human capital depreciation. Consistent with “undue hardship” provisions of student debt repayment programs, we capture the effect of these circumstances on debt repayment through the function $\xi(z)$ such that the repayment function is given by,

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

so that households are never forced to make student debt payments in excess of $\xi(z)$ of their households net worth (e.g. income and assets). In other words, households become delinquent on their student loan debt if their annual payments are in excess of $\xi(z)$ fraction of their total net worth.

Finally, while economic hardship can allow households to defer debt repayments and accrue interest on their outstanding balance, we do not allow debt deferral to continue

beyond the maturity ceiling \bar{T} . In the final period of debt maturity households must pay off their total outstanding balance or default. Households do not make any student debt payments beyond the maturity period, e.g. $\phi(a, h, d, z, t) = 0$ for all $t > \bar{T}$.

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

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

where $\tau(y)$ represent the effective marginal tax rates for the tax bracket of individuals with income y . The brackets and marginal rates $\tau(y)$ are chosen to match the effective tax rates estimated by the Congressional Budget Office (CBO) displayed in Figure 12. Accurately modelling the effective marginal tax rates is quantitatively important since these will influence household incentives to raise their income by attending college and investing in human capital (Saez, Slemrod, and Giertz 2012; Jones 2019). Furthermore, several of our policy counterfactuals are redistributive (e.g. debt forgiveness) and we use tax system $\mathbb{T}(y)$ to quantitatively match the progressivity of tax incidence required to finance these programs.

Retirement and Retirement Benefits. Households retire deterministically at age 63 and continue to make consumption and savings decisions until mortality at age 78. Retired households fund consumption out of savings and retirement benefits $\pi(\hat{y})$, which they begin receiving after retiring. Formally, retired households solve

$$V_R(a, \pi) = \max_{c, a'} u(c) + \beta V_R(a', \pi)$$

subject to

$$c + a' = e + (1 + r)a$$

$$a' \geq -\bar{a}$$

Following Huggett, Ventura, and Yaron (2011), we assume that the social security system pays everyone a common retirement transfer, denoted as e . We set this number to

\$20,447.78 per year, which is the average annual social security benefit as documented in the US Social Security Administration in 2023.

Production Technology and Equilibrium Wages. Given the focus on occupation choice, we endogenize occupational wage rates with an aggregate production technology that employs both high school and differentiated college labor. In particular, we employ a nested CES production technology which is common in other applied work (see for instance, [Jones \(2014\)](#)). The technology is given by,

$$Y = A [H_{hs}^\sigma + A_c H_c^\sigma]^\frac{1}{\sigma}$$

where σ parameterized the substitutability of high school and college labor, H_{hs} is total high school effective labor, A_c captures college augmenting TFP – calibrated to match the college wage premium—and H_c is the total labor input of college graduates. The aggregate labor input of college graduates itself depends on the composition of employment and occupation specific human capital such that,

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

where $H_{c,k}$ is total effective labor of college graduates in occupation k , a_k is occupation specific TFP chosen to match relative wages amongst college occupations, and η parameterizes the elasticity of substitution amongst college graduates.

4.1 Internal Calibration and Model Fit

The household risk preference parameter, the structural parameters of the human capital technology, and the interest rates on debt are set exogenously. The CRRA preference parameters ρ is set equal to 2, consistent with standard values employed in the literature. While the population distribution of occupation specific talents θ_{ik} are calibrated internally, the structural parameters governing the human capital technology are set exogenously to be consistent with the literature. Following [Huggett, Ventura, and Yaron \(2011\)](#), α is set to 0.7 and the annual depreciation rate δ is set to 0.05 to match the in-sample end of working life decline in earnings. The annual exogenous interest rate on household savings r equals 0.04 and the annual interest rate on student debt r_d is set to 0.042, consistent with the average federal student loan rate observed in the data.

Table 2: Internal Calibration Targets: Assets, Debt, and College Matriculation

	Data	Model
Mean of log initial assets (population)	10.50	10.22
Variance of log initial assets (population)	2.47	2.16
Mean of log initial assets (college grads only)	11.10	10.61
Variance of log initial assets (college grads only)	2.19	2.18
Mean student debts (\$)	17,784	17,468
Standard deviation of student debts (\$)	24,889	24,173
Corr(Asset, Student Debt)	-0.16	-0.16
Fraction of sample with student debts (%)	67.0	62.1
Corr(Asset, Having Student Debt)	-0.25	-0.22
College completion rate by family background (%)		
- first asset quintile (Q1)	18.2	24.4
- second asset quintile (Q2)	22.8	26.2
- third asset quintile (Q3)	26.7	29.5
- fourth asset quintile (Q4)	37.7	35.4
- fifth asset quintile (Q5)	50.7	51.2
Variance of log earnings in the first 10 years	0.39	0.39
Ratio of assets, old-to-young workers	5.76	5.79

The remaining parameters are simultaneously internally calibrated so that the model endogenously generates a joint distribution of assets and student debt (as in Figure 1), college matriculation patterns, and occupational heterogeneity in earnings, returns to experience, and worker sorting (as in Figure 4) that are consistent with the data.

The parameters $\{\beta, \mu_a, \sigma_a, \mu_\tau, \sigma_\tau, \rho_{a\tau}, \mu_x, \rho_{ax}, \epsilon, \sigma_\theta\}$ primarily determine the joint distribution of household assets, student debt, college matriculation, and earnings inequality. Table 2 summarizes the 16 targets most closely associated with these 9 parameters. These parameters are fit internally since student debt will be incurred endogenously based on selection into college, which itself will depend on the parameters governing expected payoffs in the labor market (see Table 3). Parameters μ_a and σ_a pin down the distribution of initial assets in the population. Together with $\mu_\tau, \sigma_\tau, \rho_{a\tau}$, these parameters determine the joint distribution of assets and student debt among the indebted, conditional on selection into college. Similarly, parameters μ_x and ρ_{ax} will determine

Table 3: Internal Calibration Targets: Occupational Heterogeneity

Occupation Group	Mean Log Earnings (\$)		Returns to Experience (%)		Employment Share (%)	
	Data	Model	Data	Model	Data	Model
Executive Admin Management	9.95	9.93	12.06	11.83	7.35	7.14
Math and Computer Science	9.93	9.91	13.28	12.85	8.75	8.88
Architects and Engineers	10.14	10.12	11.34	11.17	4.56	4.26
Counselors	10.09	10.07	11.20	11.07	6.12	5.93
Teachers	9.69	9.66	10.44	10.65	5.69	5.95
Education	9.83	9.82	8.41	8.47	14.17	14.16
Sports	9.55	9.56	4.23	3.89	1.61	1.73
Media	9.37	9.35	18.09	18.49	3.44	2.99
Health Practitioners	9.60	9.56	11.60	11.68	2.90	3.03
Health Support	10.29	10.27	4.28	4.49	4.03	4.06
Food Services	9.74	9.74	8.23	8.09	2.20	2.20
Cleaning	9.76	9.74	2.81	2.87	5.31	5.34
Service Workers	9.44	9.46	10.85	9.84	1.29	1.29
Sales	9.38	9.39	7.80	7.43	3.17	3.10
Office & Admin	9.73	9.75	13.44	12.77	11.11	12.29
Maintenance	9.69	9.69	9.49	9.43	15.30	14.71
Transportation	10.32	10.33	3.49	3.09	1.07	1.06
	9.20	9.16	18.25	18.83	1.93	1.88

Notes: This table summarizes the 54 calibrated parameters $\{w_k, \mu_k, \nu_k\}_{k=1, \dots, 18}$ and their data counterparts. Employment share is calculated among college graduates, so they add up to 100%.

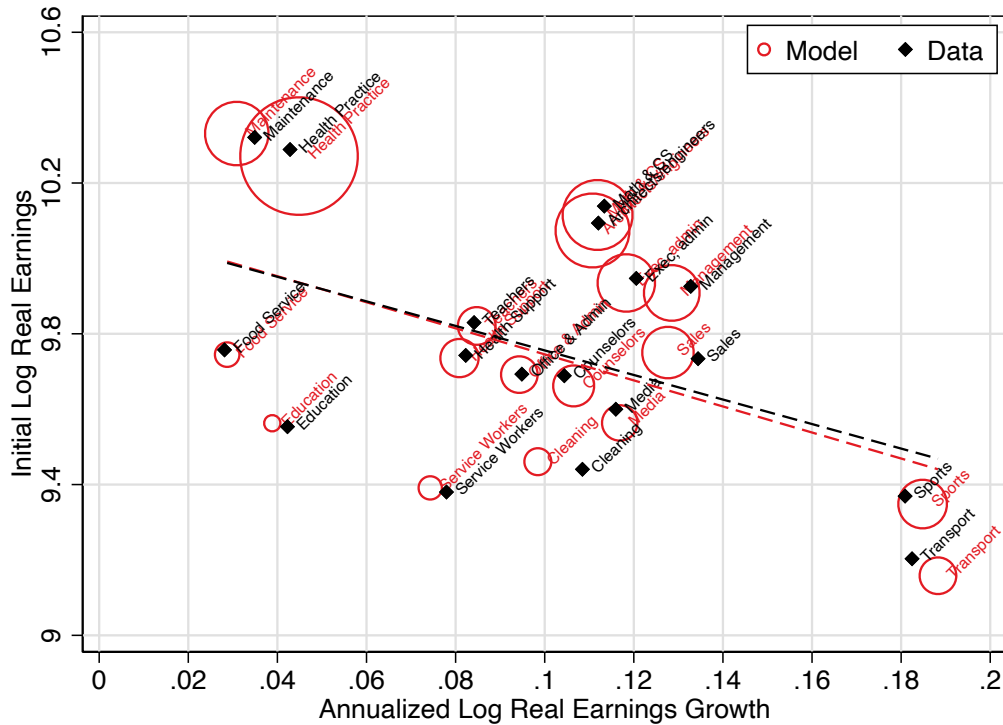
the extensive margin of college debt among graduates. Combined with the distribution of assets and student debt, parameter ϵ governing the high school taste shock will determine college matriculation patterns by family background. The common variance of student ability, σ_θ , is set to match the variance of log earnings we observe in the data. Finally, given lifecycle risk and the human capital technologies, the discount fact β determines the ratio of assets held by old-to-young workers.¹¹ Table 7 summarizes the

¹¹We used the ratio of liquid assets (excluding real estate) between under age 40 workers and age 55-69 workers in 2022 to calculate the data and model moments. Source data from the Distributional Financial Accounts, Federal Reserve Board Governors.

calibrated parameters.

The parameters $\{w_k, \mu_k, \nu_k\}_{k=1, \dots, 18}$ primarily determine that occupational heterogeneity in mean earnings, returns to experience, and worker sorting across jobs. Table 3 displays the 54 data targets primarily associated with these parameters and the model fit. The occupation-specific wage rates w_k mainly control occupation k 's average earnings; the mean population talent μ_k in occupation k most effects the returns to experience of workers who select into occupation k ; and the occupation-specific amenity ν_k are chosen to match the job choice probabilities, given the equilibrium occupation earnings profiles. These parameters must be fit simultaneously with those in Table 2 since the model targets are interdependent. Labor market returns to a college education will determine the willingness of households matriculate and take on student debt; at the same time, the distribution of debt among graduates will influence their subsequent investments in on-the-job training and occupational choice.

Figure 5: Model Fit: Occupational Heterogeneity and Student Debt



Notes: This figure plots the mean log earnings and returns to experience for each occupation from the calibrated model along with their data counterparts. The size of the bubbles represent the average amount of student debts among individual in that occupation.

Figure 5 plots the calibrated model's fit of the relationship between student debt and

occupational heterogeneity displayed originally in Figure 4. While the occupational heterogeneity in mean earnings and returns to experience are targets of the calibration process, the sorting of students with different levels of student debt across the occupations is not. Nevertheless, Figure 5 shows the model does a good job at replicating these patterns of sorting, consistent with the mechanism we study. The calibrated model also does relatively well at replicating the response of the aggregate earnings profile to changes in student debt, as summarized by the IV results in section 3. In particular, simulating the effect of an exogenous increase in student debt among the population of college graduates within the model predicts that a thousand dollar increase in debt leads to a 2.48 percentage point rise of initial earnings and -0.28 percentage point drop in returns to experience. Though untargetted, the model's predictions are close to our IV results and within range presented in the literature.

5 Quantitative Results

GOING THROUGH MAJOR REVISION

This section uses the calibrated model to compute the aggregate welfare and productivity consequences of federal policies which aim to alleviate student debt burdens. The analysis focuses in particular on two classes of policies: extended repayment programs and student debt forgiveness.

5.1 Extended Repayment Programs

Table 4 reports the aggregate welfare and productivity effects of extending student loan repayment periods. The current Federal Standard Repayment Plan requires students repay loans in fixed monthly installments over the course of ten years.¹² The results in Table 4 report the impact of extending the standard repayment period on federal loans by 2 and 5 years. In each case, the interest rate on outstanding loans is adjusted so that the net present value of individual liabilities do not change. In other words, the program has no net cost to the government and does not require any cross-household redistribution; rather, the program redistributes costs *within* each individual's lifecycle. The model captures the extensive participation margin by allowing households to select out of extended repayment programs and stay with the standard repayment plan.

¹²While 10-years is standard repayment period on federal loans, in practice repayment periods often vary with individual circumstances. For instance, students with Direct Consolidation Loans can face repayment periods between 10 to 30 years.

Table 4: Welfare and Productivity Effects of Extended Repayment Programs

	Δ Welfare			Δ TFPR
	Total	Income	Amenity	
2 Year Elongated Repayment				
Everyone	0.25%	0.45%	-0.20%	0.41%
Switchers	1.80%	45.72%	-43.92%	58.55%
Stayers	0.25%	0.25%		0.22%
% of job switchers	0.38%			
% choosing longer repayment	73.49%			
5 Year Elongated Repayment				
Everyone	0.45%	0.66%	-0.21%	0.57%
Switchers	1.98%	43.35%	-41.37%	60.87%
Stayers	0.45%	0.45%		0.37%
% of job switchers	0.41%			
% choosing longer repayment	84.68%			

Notes: This table report the impacts of extending the standard 10-year repayment period to 12 and 15 years on welfare and aggregate TFPR. % of job switchers report the fraction of job switchers induced by the policy, and % choosing longer repayment is the fraction of college graduates who prefer/choose the longer repayment duration. Change in TFPR is defined as average of change in discounted lifetime income across households

The results show that the extended repayment programs modestly increase welfare and labor productivity, and that the effect is monotonically increasing in the duration of the extensions analyzed. Furthermore, while the number of workers induced to switch occupations was small, ranging from 0.38% - 0.41% of the population, they experience by far the largest welfare and productivity gains from the program. While welfare increased by 0.25 - 0.45% for those who did not switch jobs, it grew by 1.80 - 1.98% for occupation switchers. Despite these large effects, the aggregate welfare effects of the program still closely mirror those of job stayers. In part, this is due to the fact that policy induced job-switchers remain a small fraction of the overall population. It is also due to the *direction* of job switches induced by policy. As the welfare decompositions in Table 4 show, the pattern of occupation switching was predominantly individuals leaving high amenity jobs for occupations which offered more scope for human capital accumulation and hence higher lifetime income. Due to the offsetting effects of income and amenities,

the net welfare consequences of this reallocation of workers remains modest despite large changes in the underlying sources of welfare.


While the occupation switching induced by the policy has a minimal effect on *aggregate* household welfare, it is responsible for nearly half of the gains in aggregate labor productivity. By extending repayment periods, the policy change reduces the shadow cost of borrowing which stimulates investment in human capital accumulation. Households switch from high amenity jobs to those with more scope for investment on the job and hence greater returns to experience and higher lifetime earnings. Consistent with theory in the preceding sections, high talented individuals will be most effected by the constraints, making the productivity gains among induced switchers especially large. Despite only representing a small share of the population, the productivity gains among switchers nearly doubles the *aggregate* productivity gains from the policy relative to those experienced by job stayers, from 0.22% - 0.37% to 0.41% -.57%. The results suggest that the impact of student debt and credit constraints on the occupation choice of college graduates can have substantial economic effects on aggregate labor productivity.

Figure 6 accounts for the sources of aggregate gains in labor productivity across the occupations using a shift-share decomposition. It reports the total contribution of each occupation to aggregate productivity, as well as its within-occupation and between-occupation components. Nearly one-third of the aggregate rise in productivity is due to the reallocation of workers across occupations; the total between-occupation contribution is 0.17% while the within-occupation contribution is 0.40%.¹³ The decomposition also shows that nearly all of the productivity gains are due to a small number of occupations: Sales, Management, Exec. Admin, Architects, Engineering, Math and Computer Science. As repayment elongation policies alleviate credit constraints, many workers switch to these occupations because they offer greater scope for human capital accumulation and hence higher lifetime earnings. The effect is evident in both the large between and within contributions of these occupations to aggregate productivity. Not only do repayment elongation policies lead more people to choose these occupations, they also increase the intensity of human capital investments on the job within these professions.

Finally, even though the repayment extension policies are budget neutral and do not require cross-household redistribution, the policies nevertheless have distributions consequences. The distributional consequences are the result of the uneven way that student


¹³In the case of the two-year debt elongation policy, the total within-occupation contribution is 0.24% and the between-occupation contribution is 0.17%.

Figure 6: Repayment Elongation (5yr): Shift-Share Decomposition of Aggregate TFPR



Notes: This figures shows the shift-share decomposition of aggregate gains in labor productivity across occupations, for the five-year debt elongation policy. Black bars represent the total gains in the labor productivity. Gray and red bars represent the across- and within-occupation shares of the labor productivity gains, respectively.

Figure 7: Five Year Extended Repayment



Notes: This figures shows the distribution of welfare gains under the five-year debt elongation policy across the household net asset distribution. Net asset is defined as initial asset minus total amount of student debt at the initial period. Black bars represent the welfare gains among individual who had positive student debts. Gray bars represent the welfare gains among everyone, including individuals without student debts.

debt is distributed across households (recall Figure 1). Figure 7 plots the distribution of welfare gains under the five year extended repayment program across the household wealth distribution. It shows that the welfare gains at the bottom of the wealth distribution are 3 to 4 times as large as the average. Consequently, alongside welfare and productivity gains, extended repayment policies can also reduce economic inequality even without explicit redistribution.

5.2 Debt Forgiveness Programs

Table 5 reports the aggregate welfare and productivity consequences of student debt forgiveness policies. It reports the effect of forgiving outstanding student debt up to a cap of 10k, 50k, and 100k; the final cap of 100k is essentially equivalent to complete student debt forgiveness. In contrast to extended repayment policies like those analyzed in section 5.1, student debt forgiveness requires requires new net expenditures by the government and hence redistribution across households. To be consistent with the pre-

vailing tax system, we assume that the costs associated with any debt forgiveness are distributed across households as lump-sum tax liabilities with the same progressivity as the existing income system.

The results in Table 5 show that the effect of student debt forgiveness on household welfare and productivity is non-linear in the size of the program. The smaller (10k) and larger (100k) programs actually reduce aggregate household welfare, while the smaller program also leads to a reduction in aggregate labor productivity. In contrast, the middle sized 50k program delivers modest welfare and labor productivity gains. The mixed effects reflect the distortionary consequences of redistribution, since those who bear the costs of taxation are not the same as those who benefit from student debt forgiveness. In particular, higher income households bear most of the cost while the benefits accrue to lower income households (see Appendix Figures 14 and 15).

The net impact of each program is determined by the countervailing effects of relaxing credit constraints for those who receive debt forgiveness and the added distortions from taxation required to fund the programs. The small 10k program leads to both modest welfare and productivity losses because the amount of debt forgiven is too small to have a substantial effect on the human capital investments of the heavily indebted population, but it still results in substantial additional tax burdens for higher income earners. In contrast, the median sized 50k program boosts aggregate welfare and productivity since it is sufficiently large to alleviate credit constraints for much of the indebted population. The large 100k debt forgiveness programs again yields negative household welfare. The larger program helps a smaller population of constrained graduates at the margin, and leads to far greater tax burdens for higher income households, discouraging the pursuit of higher income through human capital accumulation.

As in the case of elongated repayment policies, the population of switchers play an important role in driving aggregate outcomes, especially in terms of productivity. The share of the population switching occupations following debt forgiveness is also much larger than under repayment elongation policies. In part, this reflects the fact that forgiving debt provides a more drastic alleviation of credit constraints compared with extending the duration of repayments. In addition, debt forgiveness induces some occupation switching among higher income households as the increases in taxation necessary to fund debt forgiveness discourages human capital investments at the top of the income distribution (Saez, Slemrod, and Giertz 2012; Jones 2019). Figure 8 shows how the incidence of taxation to fund debt forgiveness effects the composition of occupation

Table 5: Welfare and Productivity Effects of Student Debt Forgiveness

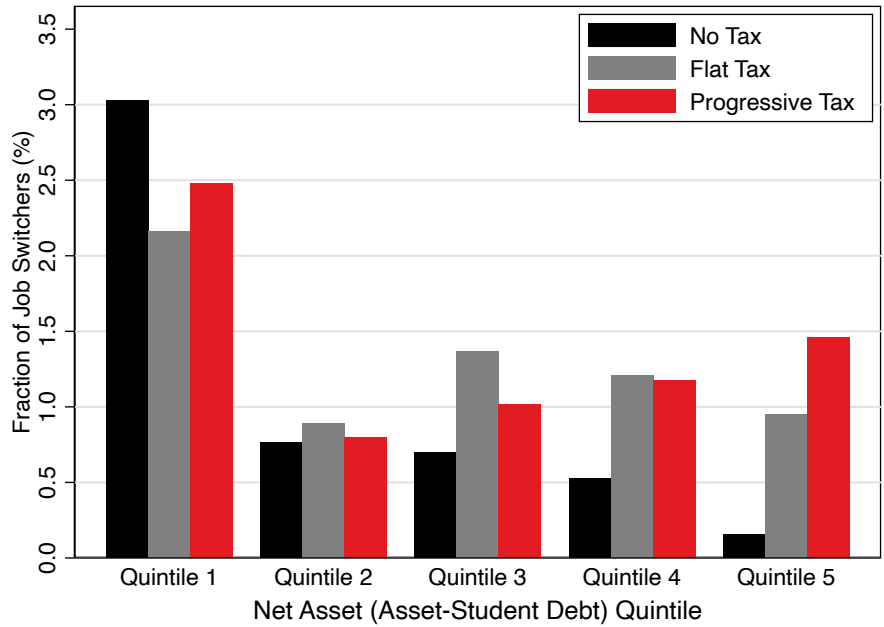
	Δ Welfare			Δ TFPR
	Total	Income	Amenity	
10K cap student debt forgiveness				
Everyone	-1.05%	-0.70%	-0.35%	-0.17%
Switchers (4.36% of treated)	-3.30%	34.00%	-37.30%	50.65%
Stayers	-1.03%	-1.03%		-0.49%
50K cap student debt forgiveness				
Everyone	0.02%	0.20%	-0.18%	0.20%
Switchers (6.61% of treated)	0.43%	10.08%	-9.65%	8.89%
Stayers	0.01%	0.01%		0.07%
100K cap student debt forgiveness				
Everyone	-0.18%	0.09%	-0.27%	0.20%
Switchers. (7.48% of treated)	1.47%	13.84%	-12.37%	14.15%
Stayers	-0.21%	-0.21%		-0.04%

Notes: This table report the effects of student debt forgiveness on welfare and aggregate TFPR. Fraction of job switchers induced by the policy reported in parentheses. Change in TFPR is defined as average of change in discounted lifetime income across households within each group.

switchers. Without any new taxes, the population of occupation switchers would decline monotonically with household net debt, as in the case of debt elongation policies. Using flat taxes to fund the debt forgiveness would lead to many more occupations from the middle of the income distribution. Funding debt forgiveness through progressive taxation, as in the benchmark results of Table 5, discourages human capital accumulation at the top of the income distribution and leads to occupational downgrading among wealthier households.

Figure 9 displays a shift-share decomposition of the contribution of each occupation to aggregate labor productivity following a 50k-cap student debt forgiveness policy. Noteworthy is the fact that the relative contribution of each occupation has changed relative to the decomposition of the debt elongation policies in Figure 6. Whereas under the repayment elongation programs most switchers moved to higher earning occupations with greater scope for human capital accumulation such as Management, Engineering, Math, and Computer Science (recall Figure 6), most switching following debt forgive-

Figure 8: Occupation Switching by Net Assets under 30K Debt Forgiveness




Notes: This figure plots the fraction of job switchers within each of the five net asset quintiles, under 30K-cap student debt forgiveness policy with different ways of financing the policy.

ness is toward middle income occupations, such as Teachers, Office Administration, Media, and Counseling. As discussed above, these differences in the patterns of occupation switching reflect the countervailing forces of debt forgiveness and increased taxation, which together predominantly effect households in the tails of the aggregate wealth distribution and incentives them to move to the middle.

Partly as a consequence of these differential switching patterns, the composition of within-occupation and between-occupation contributions to aggregate productivity are also meaningfully different under the two policies. Under the 50k-capped program, within-occupation effects raise aggregate labor productivity by 0.21% while between-occupation effects lead to a -0.01% decline. Moreover, the composition of these effects varies non-linearly across program sizes as the relative impact of debt forgiveness and increased taxation varies, leading to different patterns of occupation switching. Under the small 10k-capped forgiveness program, within-occupation effects reduce aggregate productivity by -0.43% while between occupation effects raise it by 0.26%. Under the 100k-capped program, which effectively amounts to forgiving nearly all outstand-

Figure 9: Debt Forgiveness (50k cap): Shift-Share Decomposition of Aggregate TFPR



Notes: This figures shows the shift-share decomposition of aggregate gains in labor productivity across occupations, for the 50K-cap debt forgiveness policy. Black bars represent the total gains in the labor productivity. Gray and red bars represent the across- and within-occupation shares of the labor productivity gains, respectively.

ing student debt, the within-occupation contribution to aggregate productivity is 0.12% while the between-occupation contribution is 0.08%. While the source of productivity gains varies student debt forgiveness programs, their cumulative contribution to aggregate labor productivity always appears less than repayment elongation policies which similarly alleviate credit constraints without incurring the distortionary effects of redistributive taxation.

6 Conclusion

We empirically and theoretically examine how household assets affect individual career development and aggregate labor market outcomes. Using panel microdata on the early career development of recent college graduates, we document the relationship between assets, debt, occupation choice, and the earnings lifecycle. Exploiting exogenous variation in student debt burdens following changes in the generosity of university tuition grants, we find that those with more initial debt chose careers with higher initial earnings but lower returns to experience over the next 10-15 years. Initial occupation choice mediates a substantial part of the measured effect of debt on the earnings lifecycle.

To understand the data and its implications, we develop a model in which credit constraints interact with human capital decisions. High debt burdens lead workers to distort labor market choices toward careers which offer more front-loaded compensation. The adjustment process occurs both on the intensive margin, by reducing on-the-job investment, and through an extensive margin adjustment in occupation choice. Calibrating the model to replicate key features of the microdata, we investigate the consequences of extended repayment and student debt forgiveness programs on lifecycle earnings, occupation choice, welfare, and aggregate productivity.

Our counterfactual analysis suggests that both policies increase welfare and labor productivity by allowing households with large amounts of student debt to choose occupations better matched to their abilities. We find that although the fraction of household who are induced to switch occupation by the policies is small, they drive the majority of welfare and productivity gains. This means that the occupational choice channel is important to account for. While the repayment extension policy is universally welfare improving, the distributional effects of debt forgiveness policies is only welfare improving under certain parameters. This is because debt forgiveness requires the government to introduce distortionary taxation to fund the program.

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.
- 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.
- 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.
- 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.
- 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. 2019, May. "Assets and Job Choice: Student Debt, Wages and Amenities." Working paper 25801, National Bureau of Economic Research.
- Mincer, Jacob. 1974. *Schooling, Experience, and Earnings*. National Bureau of Economic Research, Inc.
- 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

Table 6: Summary Statistics for NLSY Sample

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

Notes: This table provides summary statistics for the NLSY97 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 10: 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 11: 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 7: Parameter Values

Category	Notation	Parameter value
Preferences	$\beta, u(c) = \frac{c^{1-\rho}-1}{1-\rho}$	$(\beta, \rho) = (0.985^5, 2)$
Human capital technology	$H(s, h, \theta) = \theta(sh)^\alpha + (1 - \delta)h$	$(\delta, \alpha) = (1 - (1 - 0.05)^5, 0.7)$
Human capital shocks	$z \sim \text{LogNormal}(\mu_z, \sigma_z)$	$\mu_z = 0.0, \sigma_z = \sqrt{(0.151^2) * 5}$
Interest rates (annual)	(r, r_d)	$(r, r_d) = (0.04, 0.042)$
Initial conditions	$\begin{pmatrix} a_0 \\ \tau \end{pmatrix} \sim \ln N \left[\begin{pmatrix} \mu_a \\ \mu_\tau \end{pmatrix}, \begin{pmatrix} \sigma_a^2 & \rho_{a\tau} \\ \rho_{a\tau} & \sigma_\tau^2 \end{pmatrix} \right]$	$(\mu_a, \mu_\tau, \sigma_a, \sigma_\tau, \rho_{a\tau}) = (3.84, 3.60, 1.86, 1.73, -2.50)$
Extensive margin of student debt	$x \sim \text{Bernoulli}(p)$	$p = 0.8$
Student debt repayment	\bar{T}	$\bar{T} = 2$
High school taste	$\zeta \sim \text{Frechet}(\epsilon)$	$\epsilon = 10$
Occupational heterogeneity	(K, σ_θ) $\{w_k, \mu_k, \nu_k\}_{k=1, \dots, 18}$	$(K, \sigma_\theta) = (18, 0.52)$ See Table 8

Table 8: Occupation-Specific Parameter Values

Occupation	k	w_k	μ_k	ν_k
Executive Admin	1	5.397	-1.529	0.0385
Management	2	5.387	-1.429	0.0376
Math and Computer Science	3	5.984	-1.523	0.0354
Architects and Engineers	4	5.848	-1.516	0.0365
Counselors	5	4.343	-1.714	0.0449
Teachers	6	4.332	-1.593	0.0449
Education	7	2.922	-1.897	0.0522
Sports	8	3.828	-1.419	0.0396
Media	9	4.277	-1.792	0.0448
Health Practitioners	10	5.192	-1.773	0.0404
Health Support	11	4.145	-1.880	0.0459
Food Services	12	3.127	-1.793	0.0514
Cleaning	13	3.748	-1.923	0.0478
Service Workers	14	3.078	-1.814	0.0516
Sales	15	4.897	-1.483	0.0410
Office & Admin	16	4.119	-1.585	0.0461
Maintenance	17	5.152	-1.970	0.0406
Transportation	18	3.519	-1.530	0.0419

Figure 12: 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.

Figure 13: Validation: Aggregate Lifecycle Profile

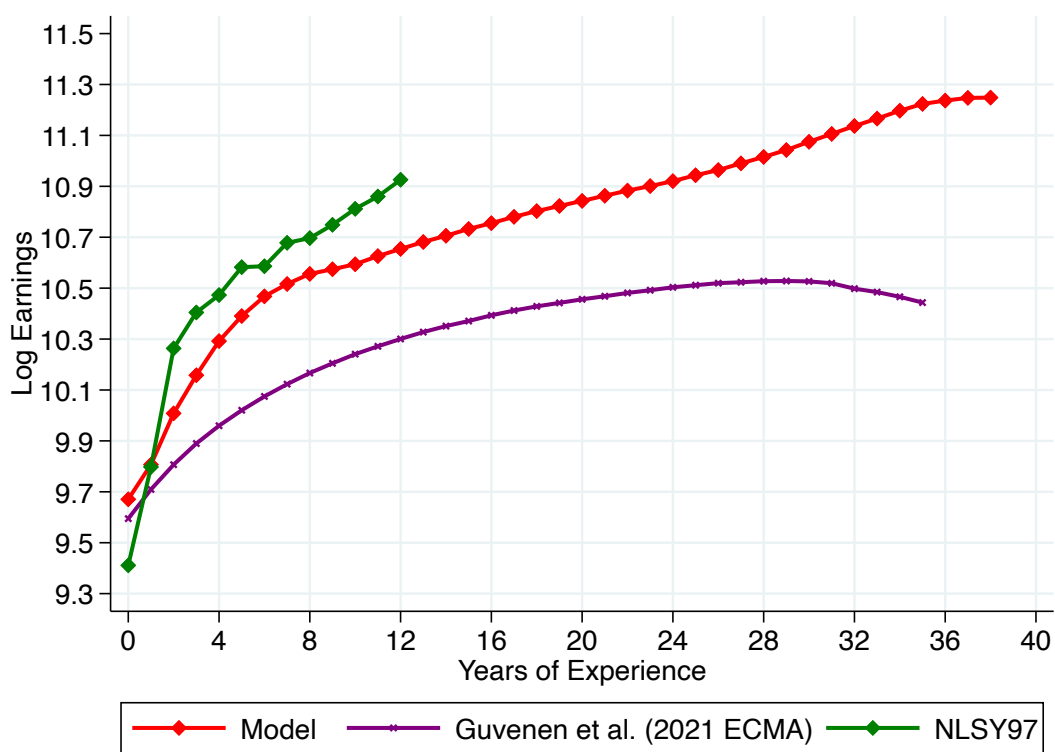

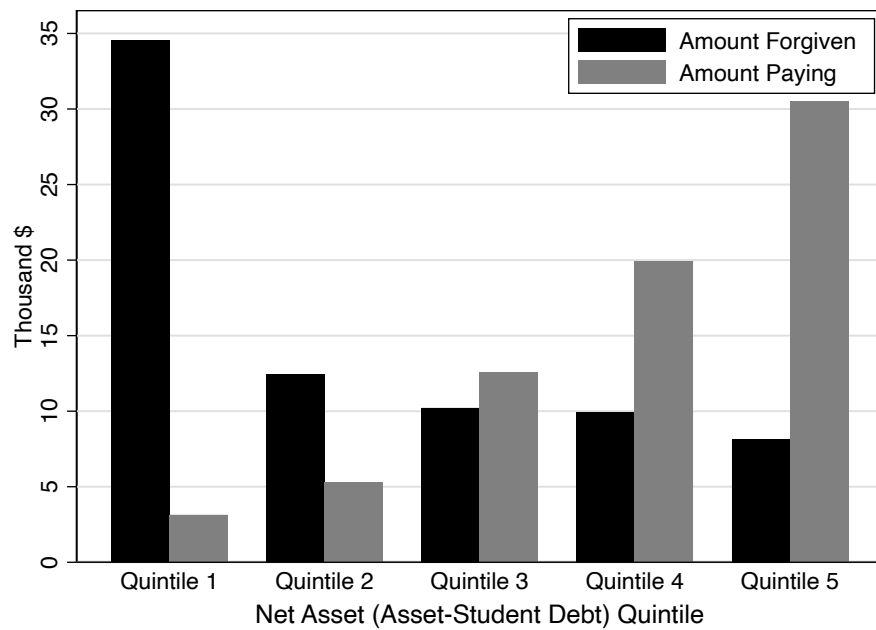


Figure 14: Welfare by Net Assets under 50k Debt Forgiveness



Notes: This figure shows the distribution of welfare gains under the 50K-cap debt forgiveness policy across the household net asset distribution. Net asset is defined as initial asset minus total amount of student debt at the initial period. Black bars represent the welfare gains among individuals who had positive student debts. Gray bars represent the welfare gains among everyone, including individuals without student debts.

Figure 15: Distribution of Costs and Benefits, 50k Debt Forgiveness



Notes: This figures shows the average amount of student debts forgiven and average amount of lump-sum taxes paid by the individuals within each net asset quintile. Net asset is defined as initial asset minus total amount of student debt at the initial period.