ASSETS AND JOB CHOICE:
STUDENT DEBT, WAGES AND AMENITIES

Mi Luo
Simon Mongey

ABSTRACT

If consumption and non-wage amenities of work enter utility, holding few assets may induce a trade-off between wages and amenities when searching for a job. We establish this in a model of search with asset accumulation, extended to accommodate amenities. We then provide empirical evidence of this trade-off in the context of student debt, finding that higher debt causes graduates to accept jobs with higher wages and lower job satisfaction. In a representative sample of college graduates, we infer causality by exploiting within-college, across cohort changes in financial aid. A quantitative extension of our theoretical framework that explicitly models student debt accounts well for our empirical results. Identifying the utility value of amenities through observed search behavior, we find that high satisfaction jobs are valued at 6 percent of lifetime consumption relative to low satisfaction jobs. This trade-off is economically significant; a policy maker using only wage data to assess the welfare effects of with an income-based repayment policy would mistakenly conclude that graduates prefer a fixed repayment policy.

Mi Luo
Emory University
1602 Fishburne Ln.
Rm 330
Atlanta, GA 30307
mluo26@emory.edu

Simon Mongey
Department of Economics
University of Chicago
1126 E. 59th Street
Chicago, IL 60637
and NBER
mongey@uchicago.edu

1 Introduction

How does the asset position of an individual affect the way that they choose between jobs? Does the substitutable role of wages and assets in providing consumption lead individuals with lower assets (or more debt) to accept jobs with higher pay but lower amenity value? A recent literature studying job choice in consumption-savings models has abstracted from non-wage amenities (Herkenhoff *et al.*, 2016; Lise, 2013). A number of recent papers have attributed an important role to non-wage amenities in job choice (Hall and Mueller, 2018; Sorkin, 2018; Mas and Pallais, 2017). Bringing these literatures together, we study the amenity-wage trade-off induced by assets in the context of student debt and job choice in the US. We ground this trade-off in theory, then show that it is both empirically observed and quantitatively relevant for the decisions and welfare of college graduates.

We focus on student debt for two reasons. First, a unique dataset positions us well to answer these questions. We merge restricted-use microdata from the National Center for Education Statistics' (NCES) *Baccalaureate and Beyond* surveys (BBS) with publicly available college financial data. These data allows us to (i) directly observe debt and earnings, as well as responses to a set of questions regarding job satisfaction which we view as reflecting amenities, (ii) exploit within-college across-time variation in institutional grants (i.e. those issued by the college) to construct instruments for individual assets (student debt), a crucial step in *assessing the causal effect* of assets on job choice in an empirical framework, (iii) exploit observed search behavior conditional on income and job satisfaction to identify the utility value of non-wage amenities, a crucial step in *quantifying the welfare effect* of assets on job choice in a quantitative model.

Second, student debt in the US has increased substantially over time (Figure 1A), with most the majority of this increase due to higher loans per borrower (Figure 1B). Understanding the effect of student debt on job choice is therefore of interest in and of itself.[1]

Theory. This paper proceeds in three steps. First, in a canonical job search model with consumption and savings, we establish that lower assets lead individuals to accept higher wage, lower amenity jobs. The partial equilibrium model combines a standard consumption and savings problem with three key ingredients: random search, random job offers that are heterogeneous in both

---

[1]As an example of recent empirical work, Alvaro Mezza and Sommer (2019) find that student debt has a negative effect on homeownership, with a $1,000 increase in student loan debt leading to an estimated postponement of homeownership by around 4 months.

Figure 1: Increases in total debt per borrowing student (in 2013, thousands of dollars)

A. Average loan per borrowing student

B. Decomposition of total loans
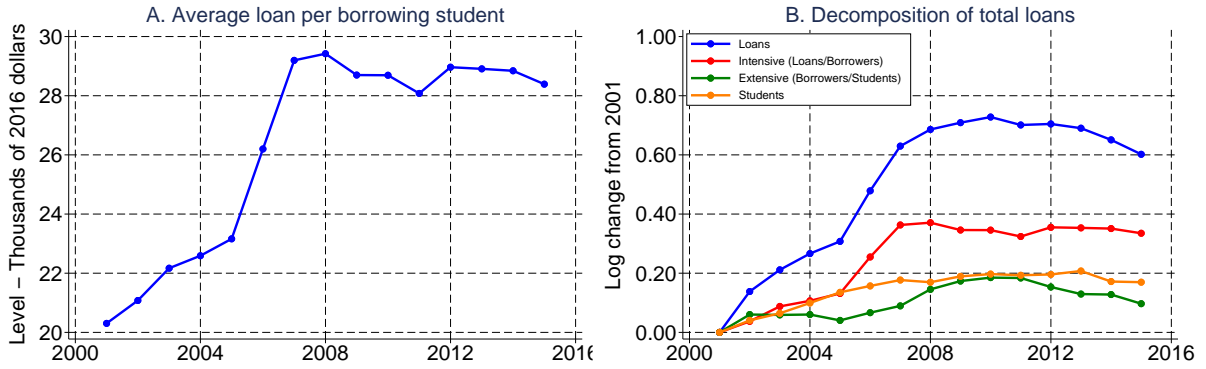
Notes (i) Data are from Integrated Post Secondary Data System (IPEDS) maintained by the National Center for Education Statistics, (ii) The sample is restricted to colleges that, according to the Carnegie classification, offer a highest qualification of at least a bachelor degree. We remove law, medical and business schools. (iii) The decomposition is exact in logs, $\log Loans_t = \log (Loans_t/Borrowers_t) + \log (Borrowers_t/Students_t) + \log (Students_t)$.

wage and non-wage features, and an assumption that amenities cannot be traded.

The theoretical model provides a framework that speaks to a broad literature. Kaplan and Schulhofer-Wohl (2017) find that survey measures of job satisfaction vary systematically across demographic groups in the cross-section and time-series. Mas and Pallais (2017) find evidence for sizeable compensating differentials at least for a tail of workers. Sorkin (2018) infers the existence of amenities through a structural model that accounts for the observed manner which workers systematically move from high to low wage firms. Estimated compensating differentials account for over half of the firm component in the variance of earnings. Hall and Mueller (2018) infer a dispersion of the value of the offered non-wage component of a job to be substantially larger than that of the offered wages. In their case, these non-wage components are inferred by workers accepting jobs below reported reservation wages. Bonhomme and Jolivet (2009) identify strong support for amenities but that these are not reflected in wages offered in the manner of compensating differentials. This will be consistent with our modelling of jobs as random draws of wage-amenity pairs. Meanwhile, other amenities, such as health insurance (Dey and Flinn, 2008) and job security (Jarosch, 2015), have also been studied in the context of search models. None of these papers study the role of assets in these decisions.

Empirical. Second, we contribute evidence for the theory by showing that assets induce just such a trade-off in the context of student debt: higher debt causes students to take jobs with higher wages but lower satisfaction.

We extend the one-college event studies of Rothstein and Rouse (2011) and Field (2009) to a

representative sample of college students.[2] Extending such analysis to a representative sample, we lack the quasi-experimental variation in student debt available to these authors. Still, we follow in the spirit of these papers, and exploit changes in the composition of grants and loans at the institution level—which differentially impact debt across cohorts within a college—to identify the causal effect of debt on job choice. A similar identification strategy has been used in contemporaneous work by Yannelis (2017) to understand the role of student debt in graduate school enrollment.[3]

Our main result is that higher student debt causes higher wages, lower job satisfaction, shorter unemployment durations out of college and lower rate of employment in low paying public interest jobs such as school teaching. Under our benchmark results, increasing student debt by $10,000 would lead to an increase of $2,110 in annual salary one year after graduation. The causal effect on wages is larger than the result in OLS, where we find higher debt is associated with no change or a negative change in wages. This suggests the IV is addressing negative selection into student debt, consistent with the findings of Yannelis (2015) over this period.[4]

This empirical result confirms our theory. A model without amenities, however, would imply the opposite: higher debt sharply reduces the value of unemployment relative to employment, reducing reservation wages, so reducing observed wages. The empirical observation that satisfaction falls therefore, through our model, rationalizes what would be an anomalous empirical observation.

The result contributes in three ways to the literature. First, we complement Herkenhoff *et al.* (2016) who find that increases in borrowing capacity encourage workers to move into self-employment. Second, comparison of our OLS and IV specifications clarifies why empirical studies of student debt and wages that do not control for the endogeneity of debt find that higher

---

[2] Both find that policies that substitute loans for grants reduce wages and increase uptake of public interest jobs. Rothstein and Rouse (2011) study a *"highly selective"* private college, while Field (2009) studies New York University Law School.

[3] Herbst (2019) uses the randomization of call-center operators managing borrower's phone calls as an instrument for take up of Income Based Repayment plans and finds that these plans increase consumption. The policy experiment in our quantitative model is consistent with these empirical results.

[4] We also make direct contributions to the study of the impact of student debt on graduating college students. We list examples for papers that have used the BBS data set . Lochner and Monge-Naranjo (2015) study the repayment and default patterns of the 1993 cohort ten years later in 2003, characterizing default along different demographics. Altonji *et al.* (2016) use the BBS and other data sets to study the impacts of recessions on labor market outcomes at labor market entry and long-run success. In an unpublished working paper, Chapman (2015) employs a different identification strategy that uses state-level merit aid to measure the causal effect of debts on earnings. She finds that exogenously increasing the loan burden of a college graduate by $1,000 increases her annual income by $400-$800 one year after graduation. Gervais and Ziebrath (2019) instead use a kink in parental contribution requirements, but do not find a statistically significant affect of debt on earnings, lacking power around the discontinuity in the BBS. They conclude that via this approach "student debt has a nonnegative effect on earnings".

debt is associated with *lower* wages (Ji, 2017). Third, that higher debt leads to lower job satisfaction provides explicit empirical support for the notion that amenities are important in job decisions, corroborating results from a growing literature (Hall and Mueller, 2018; Sorkin, 2018; Mas and Pallais, 2017; Delavande and Zafar, 2014).

Quantitative. Third, we extend our theoretical model to separately include student debt in the agent's portfolio, modelling its unique institutional features. A novel calibration strategy uses self-reported search behavior of the employed by income and job satisfaction to infer utility values of amenities. We invert this empirical relationship through the model to quantify the value of job satisfaction and its role in job choice. A free move to a high amenity job is—on average—valued at 6 percent of lifetime consumption. The value is lower when on-the-job search is less costly.

The model permits a policy experiment in which fixed repayment loans are replaced with income based repayment loans. This replicates the largest change in student debt policy in the US in recent years, as *Income Based Repayment* (IBR) plans, were introduced by legislation after our sample. Under an IBR, students only repay debt in a fixed proportion to disposable income when they are employed.[5]

We show that understanding the wage-amenity trade-off is important for evaluating the welfare consequences of this change. Keeping fixed the distributions of debt and available jobs, and adjusting taxes to maintain a balanced budget, we examine the effect of implementing the policy under a balanced budget with three main results. First, the policy has a positive effect on average graduate welfare, increasing average post-graduation expected utility by 1.3 percent. In the cross-section, 89 percent of students prefer the IBR, with low debt students preferring the fixed repayment system. Second, we use the model to decompose these gains. A large fraction of the welfare increase stems from shifting repayments into periods when marginal utility of consumption is lower. This is somewhat obvious, but we also find that 30 percent of welfare gains from graduates taking higher amenity jobs, while 7 percent is due to lower search costs. Third, we assess how policy analysis might be misguided if it does not account for amenities. A simple alternative metric would be discounted expected lifetime income, however this *decreases* under the IBR as students take jobs that pay less but they enjoy more. Understanding the role of amenities is crucial not only for the magnitude of the response, but also assessing its sign.

---

[5]This policy is the dominant repayment scheme in the United Kingdom, Australia, New Zealand and Singapore. The quantitative implications of the policy for college attendance are studied by Ionescu (2009), we study the implications for job choice.

Layout. The rest of the paper is organized as follows. Section 2 presents the theory. Section 3 presents our empirical results. Section 4 extends the theoretical model to a quantitative model of student debt. Sections 5 and 6 discuss estimation and model fit. Sections 7 and 8 we examine how student debt effects behavior in the model, compare our results to the data, and conduct our repayment scheme policy experiment. Section 10 concludes.

2 Theory

We present a labor-search model that motivates the following empirical study of the effect of student debt on job choice. Since it forms the core of the quantitative model that we will estimate in Section 4, we present the model in some detail.

Time and agents. Time is discrete. There is a continuum of infinitely lived workers. Each period a worker may be either employed or unemployed. Workers begin life unemployed.

Preferences. Workers maximize expected discounted lifetime utility which is time separable, and discount the future at constant rate $\beta \in (0,1)$. Period utility derives from consumption $c \in [0,\infty)$ and non-pecuniary amenity $x \in [0,\infty)$ which we call *job satisfaction* when employed, and has a fixed value $\overline{x} \in (0,\infty)$ when unemployed. Expected lifetime utility is therefore

$$\overline{U} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, x_t).$$

We assume that u has the following properties: twice continuously differentiable, with $u_c > 0$, $u_{cc} < 0$, $u_x > 0$, $u_{xx} < 0$, and that the Inada conditions hold for both arguments.[6] We assume that u is homogeneous of degree one in (c,x), which implies that $xu_{cx} + cu_{cc} = 0$, a sufficient condition for $u_{cx} > 0$.

Job. A *job* is a bundle of a wage and non-wage utility, or amenity, provided to the worker each period: $(w,x) \in [0,\infty) \times [0,\infty)$. Labor markets are frictional. At the beginning of each period unemployed workers draw one job with probability λ, from the distribution of jobs $F(w,x)$.[7] Employed workers are separated into unemployment at the beginning of each period with probability

---

[6]That is $\lim_{c\to 0} u_c(c,x) = \infty$, $\lim_{c\to\infty} u_c(c,x) = 0$, with the same being true for the limits in terms of x.

[7]We assume that $F(w,x)$ has continuous density $f(w,x)$ over bounded, connected supports $w \in \mathcal{W} = [\underline{w}, \overline{w}]$ and $x \in \mathcal{X} = [\underline{x}, \overline{x}]$.

δ. Unemployed workers receive unemployment benefits $\overline{w} \in (0, \infty)$. A separated worker does not receive a new job offer within the period in which they are separated.

Assets. Workers have access to a single liquid asset a that pays a riskless net return r. Holdings of the asset may be negative subject to the natural borrowing limit $\underline{a}$.[8]

Worker problem. Let $W(a, w, x)$ be the present discounted value of expected utility of a worker with assets a and employed in job (w, x), having survived the current period separation shock δ, but before payment of wage w. Let $U(a)$ be the present discounted value of expected utility of an unemployed worker with assets a after either (i) rejecting an offer, or (ii) being separated. We represent these values recursively as follows:

$$W\left(a, w, x\right) \;=\; \max_{a'} u\left(c, x\right) + \beta \left[\delta U\left(a'\right) + (1 - \delta)W\left(a', w, x\right)\right], \tag{1}$$
$$c \;=\; (1 + r)a + w - a' \quad, \quad a' \geq \underline{a}.$$

$$U\left(a\right) \;=\; \max_{a'} u\left(c, \overline{x}\right) + \beta \left[\lambda \int_{\underline{w}}^{\overline{w}} \max\left\{W\left(a', w', x'\right), U\left(a'\right)\right\} dF\left(w', x'\right) \right. \tag{2}$$
$$\left. + (1 - \lambda) U\left(a'\right)\right],$$
$$c \;=\; (1 + r)a + \overline{w} - a' \quad, \quad a' \geq \underline{a}.$$

Reservation policy. Given that $W(a, w, x)$ is monotonically increasing in each of w and x, the acceptance decision of the unemployed worker will be characterized by a *reservation job locus*. This locus $w^*(a, x)$ is defined implicitly:

$$W\left(a, w^*\left(a, x\right), x\right) = U\left(a\right). \tag{3}$$

---

[8]The natural borrowing limit is determined by the maximum amount that an individual may repay while receiving only unemployment benefits. This is given by $\underline{a} = \sum_{t=0}^{\infty} \overline{w}(1 + r)^{-t} = \overline{w}/r$.

A job with amenity value x is accepted if and only if its wage lies above the locus: $w \geq w^*(a, x)$. Using (3) we can establish that $w^*(a, x)$ is (i) downward sloping, and (ii) convex:[9]

$$w_x^* = -\frac{W_x^{(+)}}{W_w^{(+)}} < 0 \quad , \quad w_{xx}^* = -\frac{\left[W_{xx}^{(-)} + W_{xw}^{(+)}w_x^{*(-)}\right]^{(-)} W_w^{(+)} - \left[W_{xw}^{(+)} + W_{ww}^{(-)}w_x^{*(-)}\right]^{(+)} W_x^{(+)}}{W_w^2} > 0.$$

In the limit as x increases, the Inada conditions imply that $\lim_{x \to \infty} W_x = 0$. Meanwhile, the declining wage around $w^*(a, x)$ leads the marginal utility of consumption to increase: $\lim_{x \to \infty} W_w = \lim_{x \to \infty} u_c + \beta(1 - \delta)W_w' = \infty$. Therefore $\lim_{x \to \infty} w_x^* = 0$: at very large x and low w, a small increase in wages could keep utility constant despite a large decline in x. The converse is true as x decreases, with $\lim_{x \to 0} w_x^* = \infty$. Below we graph these and other properties of $w^*(a, x)$ in a numerical example (Figure 2A).

Assets and job choice. Our main interest is in how the job acceptance policy varies over different levels of assets a. Differentiating (3) with respect to a and using envelope conditions for U_a, W_a:

$$w_a^* = \frac{U_a - W_a}{W_w} = \frac{1 + r}{W_w(a, w^*(a, x), x)} \left[u_c\left(c^U(a, \overline{x}), \overline{x}\right) - u_c\left(c^E(a, w^*(a, x), x), x\right)\right].$$

Let c^U and c^E be shorthand for the above consumption functions in unemployment and employment, respectively. We take a first order approximation of the marginal utility of consumption when unemployed around (c^E, x). Under the additional assumption that $u(c, x)$ is homogeneous of degree one in (c, x), and letting $\mathcal{R}(c, x)$ denote the coefficient of relative risk-aversion:[10]

$$\hat{w}_a(a, x) \approx \mathcal{R}(c, x)^{(+)} \frac{(1 + r)u_c^E\left(c^E(a, w^*(a, x), x), x\right)^{(+)}}{W_w(a, \hat{w}(a, x), x)^{(+)}} \left[\frac{\overline{x}}{x} - \frac{c^U(a)}{c^E(a, w^*(a, x), x)}\right].$$

A special case is useful. Let $\hat{w}_a^*(a, x)$ denote the approximation of $w^*(a, x)$ attained under an assumption that consumption and hence the marginal utility of consumption is constant in employment—as would be the case where the individual is saving at a constant rate—and let

---

[9]The recursive structure of the derivative properties of the value function with respect to x and w implies that they inherit the derivative properties of the utility function. For example, $W_{xw} = u_{cx} + \beta(1 - \delta)W_{xw}'$. Since it is always the case that $u_{cx} \geq 0$, then $W_{xw} \geq 0$.

[10]The approximation is as follows:

$$u_c\left(c^U, x_b\right) = u_c\left(c^E, x\right) + u_{cc}\left(c^E, x\right)\left(c^U - c^E\right) + u_{cx}\left(c^E, x\right)(\overline{x} - x) + \mathcal{O}_2\left(||x - \overline{x}||, ||c^E - c^U||\right).$$

Note that the residual terms in the approximation are second order with respect to the marginal utility of consumption, and so third order with respect to u.

relative risk aversion be given by a parameter $\mathcal{R}(c, x) = \overline{\mathcal{R}}$. In this case

$$\widehat{w}_a^*(a, x) \approx \overline{\mathcal{R}}^{(+)} [1 - \beta(1 - \delta)]^{(+)} \left[\frac{\overline{x}}{x} - \frac{c^U(a)}{c^E(a, w^*(a, x), x)} \right].$$

Working first with the special case, it is straight-forward to sign the derivative in the limits of x. Using the fact that $w^*(a, x)$ is decreasing in x and $c^E(a, w^*(a, x), x)$ is increasing in w:

$$\lim_{x \to 0} \widehat{w}_a^*(a, x) > 0 \quad , \quad \lim_{x \to \infty} \widehat{w}_a^*(a, x) < 0.$$

Note that we do not require that $c^E(a, w^*(a, x), x)$ is at any point *less* than $c^U(a)$, only that it is increasing in the wage. Since $w_a^*(a, x)$ is continuous, and monotone in x,[11] then there exists some $\widehat{x}(a)$ such that:

1. For all $x < \widehat{x}(a)$ an increase in assets leads to an *increase* in the reservation wage. With higher assets, workers *reject* more *high wage, low amenity* jobs, *accept* more *low wage, high amenity* jobs.

2. For all $x > \widehat{x}(a)$ an increase in assets leads to a *decrease* in the reservation wage. With lower assets, workers *reject* more *low wage, high amenity* jobs, *accept* more *high wage, low amenity* jobs.

Since the marginal value of wages is positive $(W_w(a, x) > 0)$, then inspection of (2) make it clear that this holds true for the general case as well. In fact, the result is more stark. Take a worker considering a high-w, low-x job. If the job were accepted, then the worker would be saving, with consumption increasing, and the marginal utility of consumption *falling* on-the-job.[12] Therefore $W_w < (1 + r) u_c^E / [1 - \beta(1 - \delta)]$, and so $w_a^*(a, x) > \widehat{w}_a^*(a, x)$. In the alternative case, at a low-w, high-x job, the marginal utility of consumption will fall on-the-job, such that $W_w > (1 + r) u_c^E / [1 - \beta(1 - \delta)]$, and $w_a^*(a, x) < \widehat{w}_a^*(a, x)$. Comparing the general and special cases, the derivative w_a^* is more positive (negative) than $\widehat{w}_a^*$ to the left (right) of $\widehat{x}(a)$.

Illustrative simulation. To visualize the above results we simulate a version of the model with simple functional forms and standard parameter values. We specify $F(w, x)$ as the product of independent log-normals with the same mean μ and standard deviation σ. For simplicity we set

---

[11]Note that $\widehat{w}_{ax}^* \propto \left[-\frac{\overline{x}}{x^2} + c^U(a) \cdot \frac{1}{(c^E(a))^2} \cdot [c_w^E \cdot w_x^*(a, x) + c_x^E] \right] < 0.$

[12]The Euler equation for consumption on-the-job is standard:

$$u_c^E = \beta(1 + r) \left[(1 - \delta) u_c^{E\prime} + \delta u_c^{U\prime} \right].$$

Table 1: Parameters for the illustrative monthly model depicted in Figure 2

β	R	λ	δ	μ	σ	ρ	$\overline{x} = \overline{w}$
$1.05^{-\frac{1}{12}}$	$\beta R = 0.99$	0.40	0.05	$\log\left(\frac{\$30,000}{12 \times \$1,000}\right)$	$\log\left(\frac{\$18,000}{12 \times \$1,000}\right)$	0.50	$\$1,000$

the unemployed benefits and amenity equal: $\overline{w} = \overline{x}$. The utility function $u(c,x)$ is CES with the same weights on each input and elasticity of substitution $\rho < 1$.[13]

In Figure 2A, a is held fixed and we plot the downward sloping, convex job acceptance policy $w^*(a,x)$. Note that the worker will accept high amenity jobs with wages below $\overline{w}$ and jobs with $x < \overline{x}$ but high wages. The baseline *low asset* case gives $w^*(a,x)$ under $a = 0$, the *high asset* case under twice the average wage draw. Under higher assets the reservation locus rotates clockwise around a stationary point. With higher assets the worker is liberated from the pains of high wage, low amenity work, and accepts more high amenity jobs at lower wages.

In solving the value function problem, third order polynomial splines are used to approximate value functions. In Panel B we use these differentiable splines used to determine the 'exact' derivative of $\widehat{w}_a = (U_a - W_a)/W_w$. The blue line gives the first order approximation used in the above argument which is shown to do a good job of characterizing the change in sign of $\widehat{w}_a$.

Panel C provides an alternative view of the change in sign of $\widehat{w}_a$. Recall that $\widehat{w}_a = (U_a - W_a)/W_w$. Let $\widehat{w}_a^1 = U_a/W_w$, and $\widehat{w}_a^2 = W_a/W_w$, such that $\widehat{w}_a = \widehat{w}_a^1 - \widehat{w}_a^2$. Since the risk-free asset allows the household to smooth consumption relatively well, the ratio of the marginal utility of wealth in employment to the marginal utility of wages is roughly constant around the reservation locus.[14] Meanwhile, with $U(a)$ fixed in this figure, higher amenities x are associated with lower wages w *around the job reservation locus* which delivers an *increasing W_w* and a steeply declining $\widehat{w}_a^2$.

Restrictions on utility. Homothetic utility implies that amenities and consumption are complements: $u_{cx} > 0$. However strict complementarity is not necessary. If $u_{cx} = 0$, then as x increases and $\widehat{w}(a,x)$ decreases, this causes $c^E(a, \widehat{w}(a,x), x)$ to decline such that the marginal utility of consumption in employment still increases. This is sufficient to cause W_w to increase and $\widehat{w}_a^2$ to decrease, which delivers the result. However if $u_{cx} < 0$, then the increase in x will have an off-setting effect on the marginal utility of consumption, leaving the change in W_w ambiguous. In the quantitative model we will consider separable utility between c and x—which allows us to be agnostic

---

[13]Formally $u(c,x) = [c^\rho + x^\rho]^{1/\rho}$.
[14]This ratio decreases slightly since at higher x, w falls around the job reservation locus, which increases W_w more than W_a.

Figure 2: Simulated example of job acceptance policy behaviour
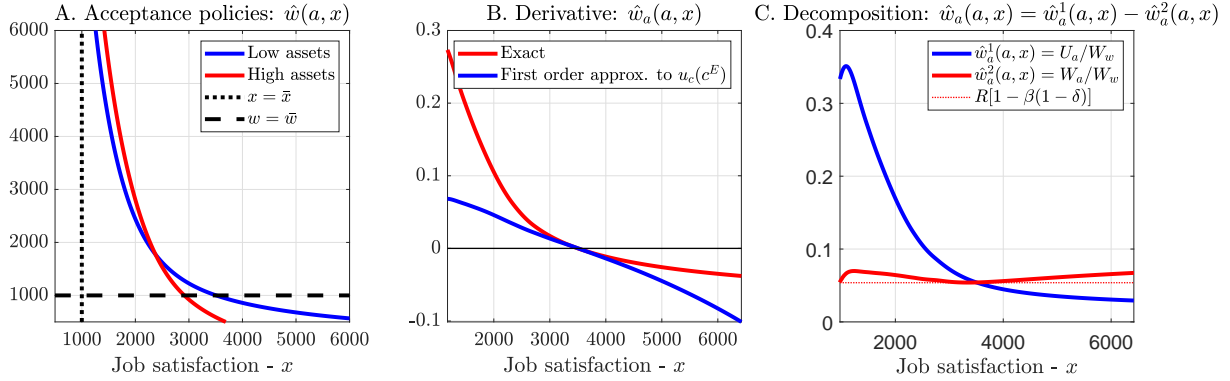
Notes: Panel A depicts the acceptance policies $w^*(a,x)$ for two asset levels (high and low). Jobs to the North-East of the locus are accepted, while others are rejected. As per the text, the locus is decreasing and convex in x for a fixed a, and as a increases rotates around a fixed point. Under lower assets, more high-x, low-w jobs are rejected (in the South-East), and more low-x, high-w jobs are accepted (in the North-West). Panel B depicts its derivative $\widehat{w}_a(a,x)$ with respect to a. Panel C decomposes the derivative $\widehat{w}_a(a,x)$ according to the equation $\widehat{w}_a = (U_a - W_a)/W_w$.

as to whether our data captures the flow or present discounted values of job satisfaction. Hence the fact that the results of this section go through with $u_{cx} = 0$ is not important for what follows.

Unemployment duration. Note that the model does not have a sharp prediction for the effect of assets on unemployment duration, only on the average wage and average job satisfaction of accepted jobs. As the reservation locus rotates, it may either lead to an increase or decrease in unemployment duration. In a model without amenities, a decrease in assets would cause U_a to decline by more than W_a, leading the reservation wage to fall and unemployment duration to shorten. In a model with amenities, the decrease in acceptance of low wage, high amenity jobs offsets some of this in unemployment duration. This is clear in Figure 2A, where it is not immediately clear whether the job acceptance region has expanded or contracted. This ambiguity will be borne out in our empirical results, where we do not find large effects of student debt on unemployment durations, but economically significant effects on job selection in w and x.

Summary. Our contribution in this section is to include amenities in a model of search with asset accumulation. The key insight is that when amenities are non-tradeable, and so do not enter the budget constraint of the worker, there exists a substitutability between wages and assets that does not carry over to job satisfaction and assets. Higher assets can be used to purchase consumption goods, which reduces the marginal value of a high wage job, leading the household to substitute

toward low wage, high amenity jobs.

Further questions.   This simple framework allows one to rethink existing studies with the added perspective of the interaction of assets and amenities. Over the business cycle, do recessions marked by sharp declines in asset values cause the cyclically unemployed to systematically choose higher wage, lower amenity jobs in the recovery? And if so, how would this change our view as to the true welfare losses of recessions? As an economy develops, does the increased provision of insurance through broader access to credit allow workers to take lower wage, higher amenity jobs? May this permit an interpretation of the flattening college wage premium as due to better insured college graduates taking jobs they like more? Does this help explain increasingly slow recoveries of wages in economic recoveries? Does the provision of insurance from parents to children, allow offspring of wealthy parents to take higher amenity, lower wage jobs such as low paying internships?

We think that all such questions would be quantitatively interesting provided that the utility values of job satisfaction are large. So in Section 4 we extend this model so that we can estimate the value of job satisfaction and understand how this may interact with policy. But first, we establish that this trade-off of amenities for wages when assets are lower is empirically observed, focusing on the special case of student debt.

3   Empirics

We first describe the data, sample selection and measurement. This determines both the sample for our empirics and for computing moments used in estimating the quantitative model in Section 5. We then detail our empirical strategy, results, how these results vary across the distribution of observables, and robustness to sample selection and measurement choices. Throughout we deflate all dollar amounts to 2009 according to the CPI.

3.1   Data

Individual.   At the individual level our main data source is the restricted-use National Center for Education Statistics' *Baccalaureate and Beyond Surveys* (BBS) microdata.[15] The data are available for three samples of undergraduates: those graduating in 1993, 2000 and 2008. We restrict attention

---

[15]Other papers to have used this data include Lochner and Monge-Naranjo (2015), Altonji *et al.* (2016), Gervais and Ziebarth (2017), and Weidner (2016).

to the latter two cohorts (BBS00, BBS09) since student debt is negligible in the first. Each sample is representative of that cohort, and consists of a mix of administrative data and self-reported data acquired in an interview one year after graduation. Administrative data include all responses on individuals' applications for free Federal funding (including parental income, parental education), college reported data (including GPA, major), and data from national administrative grant and loan records.[16] Self-reported data include various aspects of employment and job search between college and interview.

Institutional. Using harmonized college identifiers, we merge the BBS individual level data with institution level data from two sources. The first is the *Institutional Post-Secondary Database* (IPEDS), which includes annual data on grants and loans of all types (institutional, State, Federal, other). The second is the recently released *College Scorecard* data, which includes measures of SAT scores of entering college cohorts that we will use to control for time-varying college quality.

Sample selection. We keep the following observations. All references are with respect to characteristics of the individual one year after college when the BBS interview takes place. In terms of demographics, those under the age of 30. In terms of employment, those employed full-time in a single job, and that report earning above \$12,500 annually.[17] In terms of education, those that completed a four year undergraduate degree. In terms of debt, those with positive student loans from any source. We provide more explanation for these restrictions when considering robustness below.

Table 2 provides full-sample and restricted-sample summary statistics for both cohorts. We find that our sample selection procedure does not drastically change the general observable properties of the cohorts. The final two columns provides statistics for individuals under 30 with a college degree and employed full time in the CPS. A comparison of these statistics across samples shows the BBS to be broadly representative.[18]

---

[16]Respectively, these data sources combined in the BBS are as follows. The Free Application for Federal Student Aid (FAFSA) form is filled out by any student seeking Federal funding, which in practice is almost all students. The National Post-secondary Student Aid Study (NPSAS) is an administrative data set to which all colleges that receive any Federal funding report information on all students. The National Student Loan Data System (NSLDS) maintains records of all student loans.

[17]In 2009 dollars, one would earn \$12,445 (\$11,809) if working at the 2009 (2001) Federal Minimum Wage for 38 hours per week over a 50 week year. Full-time status is taken from an interview question regarding labor force participation. Ten categories are available, and we keep those responding *"One full-time job, not enrolled [in further education]"*. The majority of respondents (50.1 percent) give this response.

[18]Earnings are higher in the CPS because we cannot condition precisely on individuals having completed college one year ago. Hence the CPS statistics incorporate some income growth after college.

Table 2: Sample statistics

	Initial BBS sample		Final BBS sample		CPS	
	2001	2009	2001	2009	2001	2009
	(1)	(2)	(3)	(4)	(5)	(6)
A. Demographics						
Male	39.5	41.2	41.3	44.7	47.4	44.7
Age	25.2	24.4	22.8	22.8	26.8	26.5
White	77.9	72.4	79.6	75.2	80.2	80.4
Black	7.8	9.1	7.14	7.79	8.8	8.2
B. Outcomes						
$Earnings_{ijc}$ (\$)	30,522	26,831	30,788	36,418	38,833	44,025
$Satis_{ijc}$	0.854	0.736	0.858	0.758	-	-
$Search_{ijc}$	0.261	0.357	0.240	0.302	-	-
$Duration_{ijc}$	-	10.1	-	10.3	-	-
$Teach_{ijc}$	0.177	0.113	0.198	0.126	-	-
C. Other						
Observations	6,664	13,902	4,274	7,630	2,710	2,639
Colleges	594	953	543	953	-	-
Students/College	11	15	8	8	-	-

Notes: The table provides sample statistics for the main variables of interest in this paper. We provide the statistics for the two cohorts for the raw sample (columns 1 and 2) and the final sample for IV regressions (columns 3 and 4). We also report the statistics from the unrestricted Current Population Survey (CPS) for the two years, to show that our sample of the undergraduate graduates is close to the representative working population.

3.2 Specification

Estimating equation. To estimate the effect of student debt on job choices, we consider estimating the following equation:

$$y_{ijc} = \alpha_c + \beta d_{ijc} + \gamma_0 X_{ijc} + \varepsilon_{ijc}, \tag{4}$$

where y_{ijc} is an outcome measure for an individual i at college j in cohort $c \in \{2001, 2009\}$, d_{ijc} is a measure of student debt, X_{ijc} are a vector of individual level observables, and α_c is a cohort fixed effect. We seek an unbiased and consistent estimate of β. Two issues arise from such a specification.

First, there may be a correlation between the level of debt an individual takes on and the quality of the college attended. If a higher quality college costs more and produces better outcomes y_{ijc} conditional on observables X_{ijc}, then a positive $cov\left(\varepsilon_{ijc}, d_{ijc}\right)$ would bias β upward. To address this we consider only *within-college across-cohort* variation by adding a college fixed effect α_j, as

well as a continuous control for time-varying college quality Q_{jc}. Under $\varepsilon_{ijc} = \alpha_j + \gamma_1 Q_{jc} + \eta_{ijc}$, our estimate would be unbiased if $cov\left(\eta_{ijc}, d_{ijc}\right) = 0$.

However, even *within* a college there may be a correlation between the level of debt an individual takes on and unobservable quality or skill a_i which may be positively correlated with y_{ijc}. Two particular cases concern us. On the one hand, debt may be negatively selected. Low ability students may have low ability parents, who themselves have low wages, leading to additional borrowing. In this case a negative $cov(a_i, \eta_{ijc})$ would bias β downward. On the other hand, debt may be positively selected. High ability students forecast that they will receive high wages and so borrow to smooth consumption and purchase better quality college. In this case a positive $cov(a_i, \eta_{ijc})$ would bias β upward.

We address this in two ways. First, we add additional controls S_{ijc} for parent ability and student ability that are pre-determined at college entry, so that $\eta_{ijc} = \delta_2 S_{ijc} + v_{ijc}$. Second, we construct instruments for student debt that vary within-colleges, across-cohorts, are correlated with d_{ijc}, and uncorrelated with v_{ijc}. Motivated by Rothstein and Rouse (2011), we choose an instrument that depends on college policy. We take as our instrument Z_{jc} the ratio of the value of total institutional grants issued by the college $Grants_{jc}$ to the sum of $Grants_{jc}$ and total Federal loans $Loans_{jc}$. We later vary this instrument and show our results to be robust.

The final regression specification is therefore as follows:

$$\text{First stage:} \quad d_{ijc} = \alpha_c^1 + \alpha_j^1 + \psi Z_{jc} + \Pi'\left[\mathbf{X}_{ijc}, \mathbf{Q}_{jc}, \mathbf{S}_{ijc}\right] + \xi_{ijc}, \quad (5)$$

$$\text{Second stage:} \quad y_{ijc} = \alpha_c^2 + \alpha_j^2 + \beta \widehat{d}_{ijc} + \Gamma'\left[\mathbf{X}_{ijc}, \mathbf{Q}_{jc}, \mathbf{S}_{ijc}\right] + v_{ijc}.$$

We briefly describe the validity of our instrument:

1. **Relevance.** In Table 3 we show that Federal loans and Institutional grants account for over 70 percent of external funding.[19] Institutional grants account for more than half of $Grants_{jc} + Loans_{jc}$—i.e. $\mathbb{E}\left[Z_{jc}\right] \approx 0.5$. Changes in institutional grants have a large impact on a large amount of funding for a lot of students.

2. **Relevance.** In addition to being small in value, other forms of State and Federal funding, (i) show little across-cohort variation, (ii) are determined by formulae that are often functions

---

[19]In terms of percentage of total funding: Loans - Federal 26%, Other 6%; Grants - Federal 13%, State 10%, Institutional 45%. In terms of percentage of students receiving funding: Loans - Federal 54%, Other 7%; Grants - Federal 33%, State 34%, Institutional 55%. Federal loans and institutional grants therefore dominate funding (See Table 3).

of our existing controls S_{ijc}.[20]

3. **Variation.** There is significant within-college across-cohort variation in Z_{jc}. Figure B1 plots a histogram of $Z_{jc} - \overline{Z}_j$ using IPEDS data from 2001 to 2014, showing that there is significant variation in Z_{jc} within colleges, across time.

4. **Exogeneity.** One may be concerned that within-college across-time changes in Z_{jc} alter the composition of students beyond that accounted for by our rich set of observables and measures of college quality Q_{jc}. For example, suppose a student receives many offers from colleges $\mathcal{J}$ and then selects $j \in \mathcal{J}$ based on $Grants_{ijc}$—which is positively correlated with Z_{jc}, and will receive the same federal funding regardless of college choice. Comparing one college over time, if $Z_{jc'} < Z_{jc}$, then the lower $Z_{jc'}$ will lead fewer financially constrained (and thus potentially low ability) students in cohort c' to accept offers. Debt in c' will be higher and wages will also be higher due to selection on ability. Our estimate of β would be biased upward. However, we provide evidence that students do not in fact apply to many schools and then select on $Grants_{ijc}$. Nearly 80 percent, of Freshmen FAFSA applicants list only one college, and less than 4 percent list more than five (Figure B2). This suggests that college choice is primarily determined by idiosyncratic preferences and not individual selection on institutional grants, which are offered to the student *after* admission, which itself occurs *after* their FAFSA submission.[21]

We therefore summarize that the instrument is relevant for debt, fluctuates substantially across cohorts within institutions, and should be uncorrelated with unobserved student ability that is not already captured by our controls, given the limited scope of college applications.

---

[20]In an unpublished working paper Chapman (2015) studies a similar question to us, focussing only on the effect of student debt on wages. Using BBS09 and a regression-discontinuity design around SAT eligibility cut-offs for state-merit aid as an instrument for student debt, she finds a positive but statistically insignificant effect of debt on wages. Similar issues regarding power have affected researchers using Treasury data on the universe of student loans and parental income eligibility cut-offs for Federal need-based aid. This issue again occurs in the study of Gervais and Ziebrath (2019), who have insufficient power in the BBS around a kink in Federal grant issuance as a function of parental income. We suspect that much of the issue regarding statistical significance in both cases is that State and Federal grants make up small components of total external funding (Table 3) and—when not granted—can be made up with institutional grants if available. Hence, like Rothstein and Rouse (2011), we focus on within-college variation in the largest component of non-loan funding: institutional (or 'college') grants.

[21]Students can list up to ten colleges in their FAFSA submission. There is no monetary cost of listing additional schools. There is also no strategic cost: the number of colleges listed, and the order of colleges, is not disclosed to any school. Federal and State aid is administered using FAFSA based information, while most colleges also use the FAFSA to allocate institutional grants. For further information see FAFSA Help.

Table 3: External college funding at US colleges 2007-2014

	Fraction of total funding (1)	Fraction of students receiving funding (2)	Average value for a receiving student ($) (3)
A. Total loans and grants			
Loans	30.5	53.9	7,200
Grants	**69.5**	**72.7**	12,223
B. Disaggregated loans and grants			
Loans - Federal	24.7	53.1	5,932
Loans - Other	5.7	6.3	11,473
Grants - Federal	12.4	32.2	4,924
Grants - State	10.5	34.0	3,933
Grants - College (Institutional)	**46.7**	**56.0**	10,639

Notes: This table reports the decomposition of college funding at U.S. colleges from 2007 to 2014.

Measurement. The key independent variable d_{ijc} is total cumulative borrowing upon graduation.[22] Control variables X_{ijc} are age and dummies for sex (female omitted) and race (white omitted). Measures of student ability S_{ijc} are parental income as reported on the FAFSA form, which proxies for parental ability, and GPA which proxies for individual ability.

We consider separately the following dependent variables y_{ijc}. $Earnings_{ijc}$ is annual earnings in the current position, where we also use its logarithm for robustness. $Search_{ijc}$ and $Teach_{ijc}$ are binary variables capturing individual answers to the questions *"Are you currently looking for a job?"*, and *"Are you a teacher?"*, respectively. $Duration_{ijc}$ is the number of weeks between college and the beginning of first full-time job (where in this regression only, we condition our sample on students that did not have a job straight out of college).

A contribution of our analysis is the use of explicit measures of job satisfaction, which we briefly describe. Each individual is asked qualitative questions (depending on waves) regarding their job, each requiring a yes or no answer. We construct our baseline index $Satis_{ijc}$ as equaling one if the individual answers yes to the question *"Overall, would you say that you are satisfied with your job?"*. We also consider an alternative measure $\widetilde{Satis}_{ijc}$ which controls for systematic correlation of job satisfaction and earnings.[23]

---

[22]Since the codebook for the BBS data is not publicly available, we specify variable names from the BBS data and full descriptions in Table C1 in Appendix C.

[23]To construct this measure we estimate the following by OLS: $Satis_{ijc} = \theta_c^0 + \theta_c^1 Earnings_{ijc} + v_{ijc}$. We allow coefficients to vary by cohorts. We then set $\widetilde{Satis}_{ijc} \in \{0, 1\}$ according to whether the estimated residual is below or above 0.

Implementation. As shown in Table 2, the average number of students per college j is 8 in our final BBS data. With few students per college we are unable to precisely estimate college fixed effects α_j. We therefore group colleges into G groups $g_j \in \{1, \ldots, G\}$, and estimate fixed effects at the group level α_g. To isolate fixed effects for similar schools we group colleges by quantiles of average SAT scores of entering cohorts.[24] These are obtained from the publicly available College Scorecard data. We lag this measure by four years such that it aligns with the BBS cohorts. This serves the purpose of grouping together similar schools, and further controlling for college quality. In our baseline estimates we set $G = 10$, and cluster standard errors at the gc-level.

3.3   Results

Table 4 provides our main empirical results, where the instrument is the ratio of total grants to total financial aid (grants plus loans) for cohort c within college j.

Our benchmark specification shown in Column (1), implies that increasing a student's debt by $1,000 would lead to an increase of $211 for the initial wage, but a 0.22 percent decrease in our measure of the general satisfaction index. The empirical results establish the trade-off between a higher wage and a lower satisfaction level due to more student debts. The satisfaction measure conditioning out earnings, $\widetilde{Satis}_{ijc}$, responds more to debt, both economically and statistically. Furthermore, college graduates search more on-the-job, spend less time looking for their first job, and are less likely to take on a teaching job. As discussed above, shorter durations of unemployment would—in a model without amenities—necessarily imply a lower accepted wage. Instead we observe shorter durations and *higher* wages, rationalized by lower amenity. Table 5 reports coefficients on the additional controls for debt and satisfaction regressions.

OLS vs IV. Column (6) of Table 4 provides results from estimating equation (5) by OLS. Conditioning on our broad range of controls, the correlation of residual debt and wages is statistically insignificant, and the economic magnitude is much smaller than the IV coefficient. The point estimate of the coefficient in the log earnings case is even negative. This is consistent with debt being *negatively* selected, corroborating the interpretation of the data proposed by Yannelis (2015).

---

[24]The colleges are grouped (i.e. quantiles calculated) in the BBS data, which represents a subsample of all colleges, rather than being grouped in the College Scorecard data, which represents the universe of colleges, and then these groups being merged down onto our BBS data.

Table 4: Baseline empirical results: $Z_{jc} = \frac{Grants_{jc}}{Grants_{jc}+Loans_{jc}}$

	Baseline	Debt		Income		OLS	Follow-up
		Low	High	Low	High	2001-09	2012
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Earnings							
$Earnings_{ijc}$	0.2114***	0.9110	0.2430*	0.0253	0.1578*	0.0079	0.4548***
	(0.06804)	(0.61202)	(0.12884)	(0.03430)	(0.09256)	(0.01491)	(0.09058)
log $Earning_{ijc}$	0.0046***	0.0104	0.0073***	0.0006	0.0025*	-0.0000	0.0064***
	(0.00148)	(0.01314)	(0.00281)	(0.00138)	(0.00128)	(0.00032)	(0.00151)
B. Satisfaction							
$Satis_{ijc}$	-0.0022	0.0056	-0.0047	0.0004	-0.0056***	-0.0004	
	(0.00161)	(0.01458)	(0.00303)	(0.00265)	(0.00193)	(0.00028)	
$\widetilde{Satis}_{ijc}$	-0.0029*	-0.0029	-0.0048	0.0004	-0.0069***	-0.0004	
	(0.00162)	(0.01456)	(0.00304)	(0.00265)	(0.00198)	(0.00029)	
C. Other							
$Search_{ijc}$	0.0032*	0.0194	0.0025	0.0040	0.0040*	0.0009***	0.0017
	(0.00173)	(0.01579)	(0.00320)	(0.00280)	(0.00209)	(0.00032)	(0.00142)
$Duration_{ijc}$	-0.1351	0.2885	-0.2587	-0.2982	-0.0373		
	(0.09677)	(0.82459)	(0.16534)	(0.20041)	(0.10523)		
$Teach_{ijc}$	-0.0035***	-0.0162*	-0.0077***	-0.0032*	-0.0033***	0.0001	-0.0002
	(0.00114)	(0.00945)	(0.00208)	(0.00173)	(0.00122)	(0.00018)	(0.00121)

Notes: Results from estimating (5) by 2SLS, using the fraction of total grants to the sum of grants and loans within college j in cohort c as the instrument for individual loans of student i. Entries provide point estimates for β and standard errors in parentheses. $***$ indicates statistical significance at 1% level, $**$ means significance at 5% level, while $*$ for 10% level. Column (1) provides our baseline estimate. Columns (2) and (3) provide results splitting the sample into those with lower and higher than the median student debts level for robustness checks. Columns (4) and (5) provide results for the subsamples for those with lower and higher than the median earnings level. Column (6) provides results for the pooled OLS regression. The last column provides results using the follow-up survey in 2012 in place of the BBS09 in the baseline estimate, i.e. using the variation between BBS01 and BBS12.

Follow-up in 2012. Column (7) in Table 4 reports the IV results using the BBS01 cohort and a 2012 follow up survey for the BBS09 cohort. This new data were published only in 2017 and are helpful in understanding the long-term effects of student debt on labor market outcomes. The results on both earnings and log earnings are much larger and more precise.[25] A $1,000 increase in debts implies a $455 increase in the wage. The long-term effects of student debts on earnings are hence stronger than the short-run effects. While the 2012 follow-up does not contain questions regarding job satisfaction, the signs of the results for search and teaching are consistent with the baseline.

---

[25]We understand the precision to increase as the noisy process of sorting through the labor market settles down.

Heterogeneity. Next we provide results for estimation of (5) in subsamples under two different splits of the data. Columns (2) and (3) in Table 4 split the sample below and above median student debt level ($20,600). Columns (4) and (5) in Table 4 split the sample below and above median income ($36,200). Focussing on the results for earnings and satisfaction, we find that the magnitude and precision of the effects are largest for high debt and high income individuals. When we calibrate our model we will show that in fact earnings and satisfaction are correlated for individuals without college debt, implying a positive correlation in the underlying sampling distribution of jobs. The negative coefficient on satisfaction for high income individuals implies that job choice is undoing this positive correlation in the underlying sampling distribution of jobs, something our model will not target but will reproduce in Figure 9.

Robustness. We also consider two alternative specifications for our instrument Z_{jc} and find that our results are robust. Table A1 replicates Table 4 for estimation of (4) where Z_{ijc} is determined at the individual level, and given by total fees plus board as reported by the college $\left(StudentBudget_{jc}\right)$ minus institutional grants to the individual $\left(Grants_{ijc}\right)$. Table A2 reports results for a 3SLS approach which aims to isolate the increase in within college student debt correlated with within college grant policies, and use this as an instrument. In a new first stage we take college level average student debt from 2001-2014 $\left(\widehat{AveDebt}_{jc}\right)$ from IPEDS data and project this on year fixed effects, total number of students, our original instrument Z_{jc} and college fixed effects. We then use the predicted values, $\widehat{AveDebt}_{jc}$ as the instrument in our original two-stage least squares estimation. Under both alternative instruments, the coefficients on earnings and log earnings are positive and those on the two satisfaction measures are negative, consistent with the signs in the baseline estimation.

Summary. Our empirical results expand the results of Rothstein and Rouse (2011) to a representative sample and support the conclusion that higher student debt leads individuals to quickly take jobs that they like less but pay more and engage in more search once employed. Figure 3 summarizes the satisfaction component of this result, showing that, even outside of our regression framework, the fraction of individuals reporting high job satisfaction is decreasing in student debt, even within income groups.

If labor markets are frictional and search is costly then this trade-off may lead to persistently low job satisfaction and high search costs. But how is this valued in lifetime utility? And how

19

Figure 3: Student debt and job satisfaction
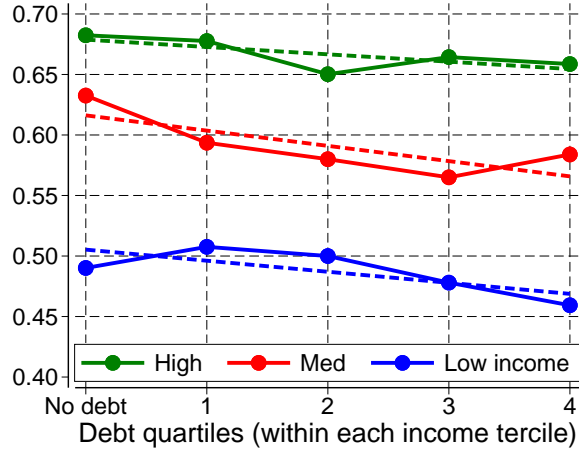
might this change our understanding of the welfare implications of debt repayment policy? These are the questions we now pursue, laying out a quantitative counterpart to the model in Section 2 that explicit accounts for student debt, and estimating the model using the empirical relationship between $Earnings_{ijc}$, $Satis_{ijc}$ and $Search_{ijc}$.

4   Quantitative model

Our theoretical model established an asset-induced trade-off between wages and job satisfaction, and we have provided empirical support for this trade-off in the context of student debt. We now extend the theory model, specializing it to the case of student debt so that we can put this trade-off in utility terms, answering: what is the utility value of job satisfaction to individuals? Through the model, the data informs us that these values are significant, which prompts us to consider how accounting properly for this trade-off affects the direction and magnitude of the response of welfare to a change in institutional repayment policies.

We change the model in four ways. First, we include institutional features that distinguish student debt from other forms of debt, specifically its repayment requirements. Second, to map the model into the data in the previous section we specialize the utility function to two levels of job satisfaction x_L and x_H, while considering only log utility in consumption: $u(c, x) = \log c + x$. We normalize the value of job satisfaction in unemployment to zero. Third, higher student debt

Table 5: Baseline empirical results - Coefficients for additional covariates

	Baseline		High debt		High income	
	$Earnings_{ijc}$ (1)	$Satis_{ijc}$ (2)	$Earnings_{ijc}$ (3)	$Satis_{ijc}$ (4)	$Earnings_{ijc}$ (5)	$Satis_{ijc}$ (6)
Debt	0.2114***	-0.0022	0.2430*	-0.0047	0.1649*	-0.0056***
	(0.06804)	(0.00161)	(0.12884)	(0.00303)	(0.09294)	(0.00193)
Male	6.8748***	0.0222*	6.3125***	0.0218	4.5275***	-0.0024
	(0.50929)	(0.01203)	(0.74414)	(0.01739)	(0.74203)	(0.01539)
Age	-0.6859	-0.0353	-2.0485	-0.1130	-4.6126	-0.0166
	(2.20872)	(0.05265)	(3.19641)	(0.07604)	(3.01433)	(0.06358)
Black	2.0228**	0.0008	0.0577	0.0024	0.0989	0.0003
	(0.96772)	(0.00107)	(0.06472)	(0.00154)	(0.06079)	(0.00128)
Parental Income	0.4885**	0.0050	-0.1523	0.0121	-0.1971	0.0024
	(0.24249)	(0.00575)	(0.41793)	(0.00977)	(0.33822)	(0.00708)
GPA	12.0327***	0.1427***	12.2017***	0.1384**	8.1490***	-0.0460
	(1.82654)	(0.04331)	(2.53988)	(0.05964)	(2.60522)	(0.05438)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$ This table provides results from estimating (5) by 2SLS. Entries provide point estimates for β and standard errors. Columns (1) and (2) provide results for the baseline for outcomes of earnings and general satisfaction. Columns (3) and (4) provide robust results for the subsample with debts higher than the median level ($20,600). Columns (3) and (4) provide robust results for the subsample with income higher than the median level ($32,600).

may not affect job decisions if individuals have access to other forms of borrowing and saving, we therefore include liquid assets with realistic borrowing constraints. Fourth, in order to estimate the utility value of job satisfaction, we add costly on-the-job search. This allows us to leverage the empirical relationship between on-the-job search, wages and job satisfaction to estimate the magnitudes of x_L and x_H. Much of the structure of the simple model remains.

States. The state vector for an employed worker is $\mathbf{s}_E = (a, d, t, w, x)$ and contains liquid assets a, outstanding student debt d, a date variable for the loan t, wage w and job-satisfaction x. For an unemployed worker the state vector is $\mathbf{s}_U = (a, d, t)$. The initial state of the agent is unemployment with a_0, d_0, and $t_0 = 1$.

Repayment. We accurately model the repayment rules faced by student loan holders. To this end we introduce the following general objects which in later sections we specialize to replicate features of different repayment environments.

Definition. *A **repayment policy** $\mathcal{R}$ is a tuple of functions $\mathcal{R} = (\rho, \Delta, \tau)(\mathbf{s})$*

1. *The **repayment function** $\rho(\mathbf{s})$ specifies the full required repayment in the current period.*

2. *The **penalty function** $\Delta(\mathbf{s})$ specifies the evolution of debt $d' = \Delta(d, w, t)$ if the full repayment is not made. This includes penalties for late repayment.*

3. *The **deferral function** τ specifies the evolution of the date of the loan $t' = \tau(d, w, t)$ if the full repayment is not made. This captures renegotiation and other institutional features.*

Finally, we assume that if an individual's available resources, including credit up to their borrowing limit, exceed $\rho(\mathbf{s})$ plus a consumption floor $\underline{c}$, then the full repayment $\rho(\mathbf{s})$ must be made. If these resources cover less than $\rho(\mathbf{s})$, then a partial repayment is made and the penalty function is invoked. Since student debt can not be defaulted upon, we do not model default on the principle. The interest rate on student debt is given by r_d.

Assets. Since the standard repayment period for college debt is ten years, it would be counterfactual to impose $a' \geq 0$ until debt is repaid. When the worker is employed we limit credit to a multiple of income $a' \geq -\Gamma_E(\mathbf{s}_E) = -\gamma w$. An unemployed worker's credit cannot be further extended and so faces the borrowing constraint $a' \geq \Gamma_U(\mathbf{s}_U) = \min\{a, 0\}$.[26] We assume that the interest rate on student debt r_d is constant and allow the rate of return on liquid assets $r_a(a)$ to vary with assets.

Search. All unemployed workers search and face a probability λ_U of drawing from $F(w, x)$. For employed workers search is costly. At the convening of the labor market each period an employed worker draws an *iid* utility cost of search $\kappa \sim H(\kappa)$, $\kappa \in [\underline{\kappa}, \overline{\kappa}]$. If the cost is paid, the worker is deemed to be searching and an offer arrives with probability λ_E. The model therefore produces data on the fraction of workers searching across the observed states (w, x). The increase in income that is required to generate the same frequency of on-the-job search across low and high satisfaction jobs will therefore inform us about their utility values.[27]

Worker problem. Let $W(a, d, t, w, x)$ and $U(a, d, t)$ be the present discounted value of lifetime utility of an employed and unemployed worker (i) after the resolution of labor market risk, and

---

[26]This implies that on graduation the student cannot borrow until they find a job. The particular structure of borrowing limits is new and allows us to represent the problem recursively while also avoiding a common issue in Bewley style consumption savings models, which is that individuals that experience negative shocks are required to immediately delever.

[27]An additional benefit of the *iid* costs of search is to smooth expected value functions, allowing us to use sparse polynomial approximations when solving the model. With four continuous state variables (w, a, d, t) such polynomial approximations are crucial. The *iid* costs of search allow these to be implemented even in the presence of kinks induced by job acceptance decisions.

(ii) before the consumption and saving decision. These two values may be written recursively as follows:

1. Value of employment

$$
\begin{aligned}
W(a,d,t,w,x) \;=\; & \max_{c \geq \underline{c}} \log(c) + x + \beta \Bigg[\delta U(a',d',t') \ldots \\
& + (1-\delta) \int \max \Big\{ -\kappa' + W^S(a',d',t',w,x),\, W(a',d',t',w,x) \Big\} dH(\kappa') \Bigg] \\
W^S(a',d',t',w,x) \;=\; & \lambda_E \int \max \Big\{ W(a',d',t',w,x),\, W(a',d',t',w',x') \Big\} dF(w',x') \\
& + (1-\lambda_E) W(a',d',t',w,x)
\end{aligned}
$$

subject to

Case 1 - *Repayment* $(1+r_a(a))a + w + \Gamma_E(w) - \underline{c} \geq \rho(d,t)$

$$
a' \;=\; (1+r_a(a))a + w - c - \rho(d,t), \quad d' = (1+r_d)d - \rho(s_E), \quad t' = t+1
$$

Case 2 - *Delinquency* $(1+r_a(a))a + w + \Gamma_E(w) - \underline{c} < \rho(d,t)$

$$
a' \;=\; -\Gamma_E(w), \qquad d' = \Delta(a,d,t,w), \qquad t' = \tau(d,t)
$$

2. Value of unemployment

$$
\begin{aligned}
U(a,d,t) \;=\; & \max_{c \geq \underline{c}} \log(c) + \beta \Bigg[(1-\lambda_U)U(a',d',t') + \ldots \\
& \lambda_U \int \max \Big\{ W(a',d',t',w',x'),\, U(a',d',t') \Big\} dF(w',x') \Bigg]
\end{aligned}
$$

subject to

Case 1 - *Repayment* $(1+r_a(a))a + b + \Gamma_U(a) - \underline{c} \geq \rho(d,t)$

$$
\begin{aligned}
a' &= (1+r_a(a))a + b - c - \rho(d,t) \\
d' &= (1+r_d)d - \rho(d,t) \\
t' &= t+1
\end{aligned}
$$

Case 2 - *Delinquency* $(1 + r_a(a))a + b + \Gamma_U(a) - \underline{c} < \rho(d,t)$

$$a' \;=\; -\Gamma_U(a), \qquad d' = \Delta(a,d,t,b), \qquad t' = \tau(d,t)$$

Baseline policy.   In the baseline model we consider the repayment policy for Federal Stafford loans under the *Standard Repayment Plan*, which we denote $\mathcal{R}_S$:[28,29]

$$\rho_S(d,t) \;=\; \begin{cases} 0 & , \quad \text{if } t \leq T_G \\[2mm] \left[\frac{r_d}{1-(1+r_d)^{-(T-t+1)}}\right] d & , \quad \text{if } t \geq T_G + 1 \end{cases}$$

$$\Delta_S(a,d,t,y) \;=\; (1+r_d)d - \rho_p(a,y) + \phi_S\Big[\rho_S(d,t) - \rho_p(a,y)\Big]$$

$$\text{where} \quad \rho_p(a,y) = \max\Big\{ (1+r_a)a + y + \Gamma - \underline{c}, \, 0 \Big\}$$

$$\tau_S(d,t) \;=\; \begin{cases} T_G + 1 & , \quad \text{if } d > 0 \text{ and } t = T - 6 \\[2mm] t+1 & , \quad \text{otherwise} \end{cases}$$

The repayment ρ_S is calculated to amortize the balance d over $T - t$ periods. If available resources do not cover the full required repayment (i.e. $(1 + r_a(a))a + w - \Gamma(\mathbf{s}) \in (0, \rho_S(\mathbf{s}))$) then available resources are submitted in partial repayment $\rho_p(a,y)$. We call this *delinquency*. Under delinquency Δ_S states that the loan accrues interest and a penalty is added to the balance. This penalty is a fraction ϕ_S of the missed full repayment (ρ_S) net of the partial repayment (ρ_p).

The rules are then parameterized as follows. The penalty ϕ_S is set to 18.5 percent to approximate the effect of payment default.[30] Since any payment up to the borrowing constraint is

---

[28]Conditional on positive debt, 67 percent of the BBS sample receive only Federal funding and 30 percent receive a combination of Federal and private loans (BBS09 variable: *loansrc*). One hundred percent of Federal borrowers hold Stafford loans and 78 percent hold Stafford loans as their only form of Federal support (BBS09 variable: *fedlnpak*). The remainder receive a combination of Stafford and other Federal support (e.g. PLUS loans, Pell grants). 99 percent of Stafford borrowers held both *subsidized* and *unsubsidized* loans and of these 48 percent borrowed less than the total maximum for the program (BBS09 variable: *stafct1*). The *only* difference between subsidized and unsubsidized Stafford loans is the deferment of interest while studying at college under the latter, effectively changing the principle of the loan upon graduation for a given amount borrowed. Since $t = 0$ coincides with graduation we do not need to model the difference between these loans.

[29]For information regarding the *Standard Repayment Plan* the explaination at studentaid.ed.gov linked see here. Up until 2010 this was the only payment plan available to borrowers.

[30]Stafford loans are considered to be in *default* if they have remained *delinquent* for 270 days. A loan is *delinquent* if it is not completely up to date in its repayments. A defaulting borrower does not default in the traditional sense since student debt is essentially non-defaultable. Instead the loan is handed to a collection agency with a fee of 18.5 percent accrued to the principle. Since modeling this entirely would require additional state-variables we view the above as a reasonable approximation. As an additional reference, a major loan provider Sallie Mae issues a late fee of 6 percent of the repayment after a payment for each 15 days past its due date. Compounded over two periods to get to a month, this is 12.4 percent (see: http://lifehacker.com/how-one-late-student-loan-payment-affects-you-1326216867).

enforced, all delinquent borrowers end the period with assets $a' = \Gamma(\mathbf{s})$ and consume $c = \underline{c}$.[31] Graduates are initialized with $t = 1$. Payments are deferred for a six month grace period ($T_G = 6$) and the loan is amortized over 10 years ($T = 120$). We do not model renegotiation of the loan explicitly. As a stand-in τ_S implies that we reset the date of the loan to $t = T_G$ in the case that the individual gets close to T without it being fully amortized.[32]

5 Estimation

Given the baseline repayment policy the unknown parameters and functional forms are

$$\theta_1 = \left\{ \beta, \delta, r_d, r_a(a), \gamma, b, \underline{c} \right\} \quad , \quad \theta_2 = \left\{ x_L, x_H, H(\kappa), \lambda_U, \lambda_E, F(w, x) \right\}$$

Our strategy is as follows. We externally calibrate θ_1 and then make the following observation: all of the parameters in θ_2 enter the problem of a student graduating with $d = 0$. Given the complexity of the problem, our strategy is therefore to estimate θ_2 on the sub-sample of the BBS09 data *without* student debt, dropping two state variables. This also ties our hands. We are not using our main empirical results in the estimation of the model, which allows us to validate the model against them ex-post. The model is solved at a monthly frequency and estimated on BBS09 data.

5.1 Calibration - θ_1

We assume a monthly frequency, and choose β consistent with an 5 percent real interest rate in a representative agent model: $\beta = 0.95^{1/12}$. When discussing interest rates we compound monthly rates and express them at a yearly frequency. Following Kaplan and Violante (2014) (henceforth KV) we specify an annual real return on positive assets $r_a^+ = -2$ percent given the nominal return of zero after 2008 and two percent inflation. The interest rate on student debt is modeled on Stafford loans which are nominal, fixed rate loans. In academic years 2006-07 and 2007-08 the rate was 6.80 percent, so we set $r_d = 4.8$ percent.[33]

---

[31] In some cases the worker will not be able to finance $\underline{c}$ even when making no repayments, for example in the case that a worker has negative assets, is unemployed and $ra + b < \underline{c}$. Since unemployed the worker cannot extend borrowing. We simply assume that the individual remains at the borrowing constraint and that $\underline{c}$ is a subsidy from government to the household.

[32] Absent such a function, an individual may end up with a large balance d close to the amortization date. This will lead to required payments ρ_S that exceed available funds, and so the loan may never be repaid.

[33] Prior to 2006-07 rates were variable and fluctuated around this figure. From 2013 onwards - due to *The Bipartisan Student Loan Certainty Act 2013* - the rate will be equal to the minimum of the 10-year T-Bill rate and 8.25%. See: https://www.edvisors.com/college-loans/federal/stafford/interest-rates/

Figure 4: Median credit card limits and interest rates by age and education. Source: SCF01
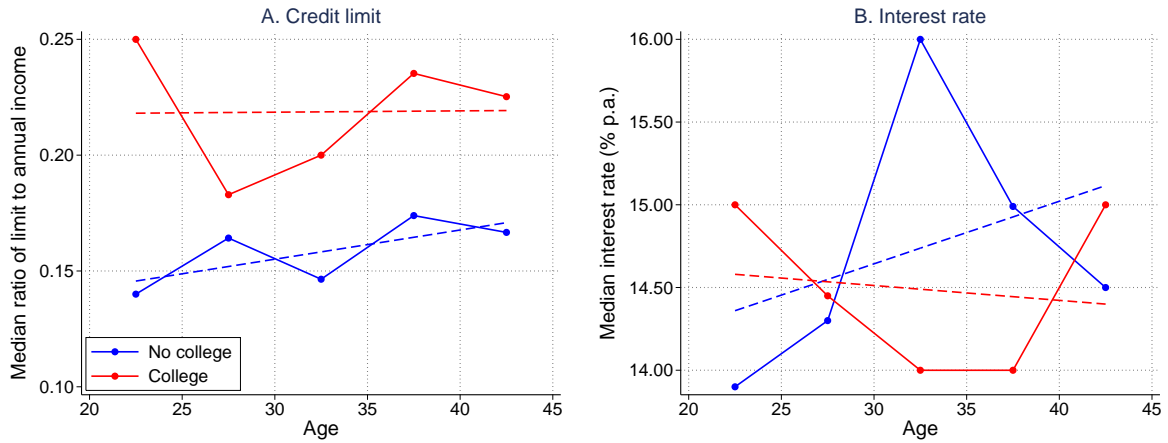
Notes: (i) Credit ratio is total available credit card limits divided by income, (ii) Income includes wage income, unemployment benefits child benefits, TANF, other, (iii) Values represent medians within age groups; approximately 1,000 observations per cell.

As borrowing of the liquid asset is uncollateralized, we model it around credit cards. Our convenient approach to the borrowing constraint is consistent with banks which do not force borrowers to pay down their debt in unemployment, but will no longer extend credit limits. The average nominal rate on consumer credit cards given by the Federal Reserve's Consumer Credit report was 14 percent in 2009 which is the same as we compute for the average rate paid by college graduates in the *Survey of Consumer Finance* 2001 (SCF01), so we set $r_a^- = 12$ percent.[34] Given the lack of information regarding credit limits in the BBS09 we use SCF01 to calibrate γ. The survey asks households to report their total credit limit. Using the same sample as KV, but now restricted to only college graduates aged 25 to 30, we find a median ratio of credit limits to annual labor income of 21 percent. This is higher than found in KV (18.5 percent) due to excluding those without a college degree, which have a median limit of only 15 percent of labor income (see Figure 4).[35]

The rate of separation from employment δ is taken from Lise (2013) which calculates a monthly rate for college graduates in the NLSY 1979 of 0.019. The Federal poverty threshold

---

[34]See: http://www.federalreserve.gov/releases/g19/: G19

[35]An interesting fact emerges when decomposing the data this way. When comparing college to non-college workers the *within* group increase in the credit limit with age is far less significant than the between group difference in average credit limits: the median credit limits for non-college students is 15.5 percent and 20.0 percent for college graduates. In fact the increase by age is statistically insignificant for college graduates while increasing from a median of 13.1 percent among 20 to 25 year old non-college graduates to 16.1 percent for ages 40 to 45.

Figure 5: On-the-job search, job satisfaction, income and student debt
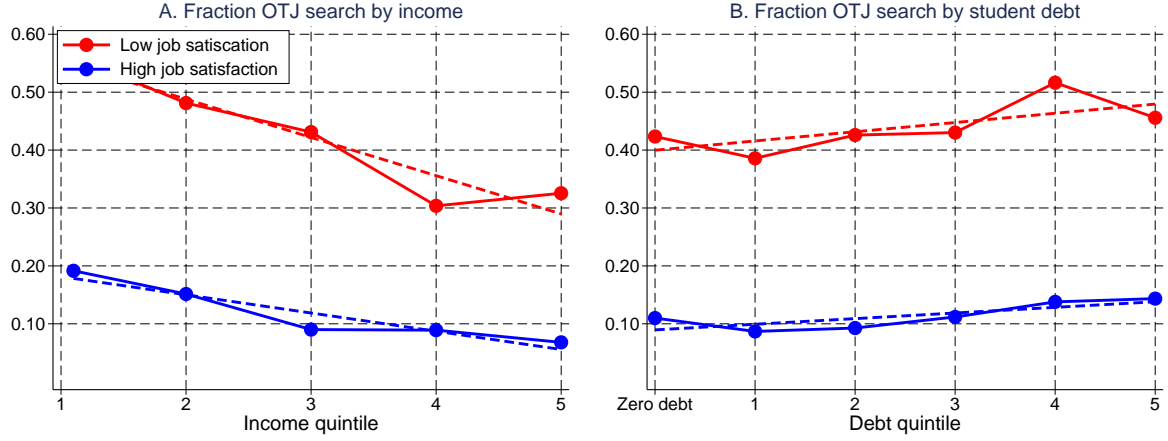
Notes: (i) The question asked is "Are you currently looking for a job?" (ii) Zero on the *x*-axis corresponds to students with zero debt (23% of the sample) (iii) Part-time workers are those workers that are employed and report average hours of less than 35 per week (iv) Income is measured as annual income and debt is cumulative student borrowing upon graduation (v) Since all unemployed workers search by definition we limit the figure to employed workers only.

for an individual living alone in 2008 was \$991 per month. We therefore set $\underline{c}$ to \$991.[36] In our estimation sample of 1,439 workers only 20 have a monthly income less than this amount. We assume that unemployment benefits b are sufficient to cover half of $\underline{c}$. Computing average benefits for unemployed workers in the SCF01 gives approximately this result.

5.2 Indirect inference - θ_2

The parameters of $\theta_2 = \{x_L, x_H, H(\kappa), \lambda_U, \lambda_E, F(w, x)\}$ are jointly estimated by indirect inference given functional forms that we now specify. Search costs are assumed to be distributed uniformly with mean κ and upper bound $\overline{\kappa}$.

Given our assumption on two values of amenity we can fully characterize joint distribution $F(w, x)$ using two conditional densities $F_L(w)$ and $F_H(w)$ and a probability p_H of drawing x_H. These conditional densities are assumed to be log-normal:

$$
x = \begin{cases} x_L & \text{w.p.} \quad (1 - p_H) \\ x_H & \text{w.p.} \quad p_H \end{cases}
$$
$$
\log w | x_k \sim N\left(\mu_k, \sigma_k^2\right)
$$

---

[36]See: https://www.census.gov/hhes/www/poverty/data/threshld/thresh08.html

27

The *population distribution* $F(w, x)$ from which agents draw offers is distinct from what we call the *sample distribution* $\tilde{F}(w, x)$ which is the distribution of workers over (w, x) one year after graduating college which we observe in the data.

We impose the same parametric form on the sample distribution to give estimates of observed $\tilde{p}_H$ and $\{\tilde{\mu}_k, \tilde{\sigma}_k\}_{k \in \{L, H\}}$ which we treat as moments in the estimation. To form these moments we assign workers in our sample to low and high satisfaction groups in the same way as we did for our earlier empirics.[37] Having split individuals into high and low satisfaction groups we fit a log-normal distribution $\log w_k \sim \mathcal{N}(\tilde{\mu}_k, \tilde{\sigma}_k)$ to each of the conditional wage densities. The sampling proportion of high satisfaction jobs is $\tilde{p}_H = 0.374$. Figure 6A shows the log-normal to be a good approximation of these conditional densities.

Under these assumptions we have eleven parameters to estimate:

$$\theta_2 = \left\{ x_L, x_H, \kappa, \overline{\kappa}, \lambda_U, \lambda_E, p_H, \mu_L, \mu_H, \sigma_L, \sigma_H \right\}.$$

Identification. A key hurdle in the quantitative study of amenities such as ours is identifying the utility values x_L and x_H. We have structured the model to leverage observables in the BBS in this regard. Our approach is to use observed search behavior over job satisfaction and income groups to infer these values. Figure 5A shows that, conditional on job satisfaction, the fraction of individuals searching on the job is decreasing in income. It is also sharply decreasing in job satisfaction. With a random cost of search, the model generates exactly this data, with some fraction of individuals at each income and job satisfaction level choosing to search. This relationship effectively allows us—through the model—to put a utility value on job satisfaction. The two parameters associated with the cost of search $(\kappa, \overline{\kappa})$, and x_L and x_H are chosen to jointly target the average rate of search, the difference across satisfaction groups, the slope with respect to income and the amount of search not explained by income and debt. Panel B is the non-regression counterpart of our empirical results, showing that search is also increasing in student debt.

Data. The estimation sample consists of $n = 940$ students from the BBS09 that satisfy the following conditions: (i) unemployed upon graduation, (ii) without student debt. The data is $X_n^{Data} = \{E_i, j_i, w_i, d_i, x_i, S_i\}_{i=1}^{n}$ which consists of observations on employment status $E_i \in \{0, 1\}$, number of jobs since graduation j_i, monthly income w_i, duration of search after graduation d_i, our constructed measure of job satisfaction x_i, and an indicator of active job search $S_i \in \{0, 1\}$.

---

[37] Again we refer the reader to Appendix C for the exact questions asked.

Figure 6: Wage and asset distributions
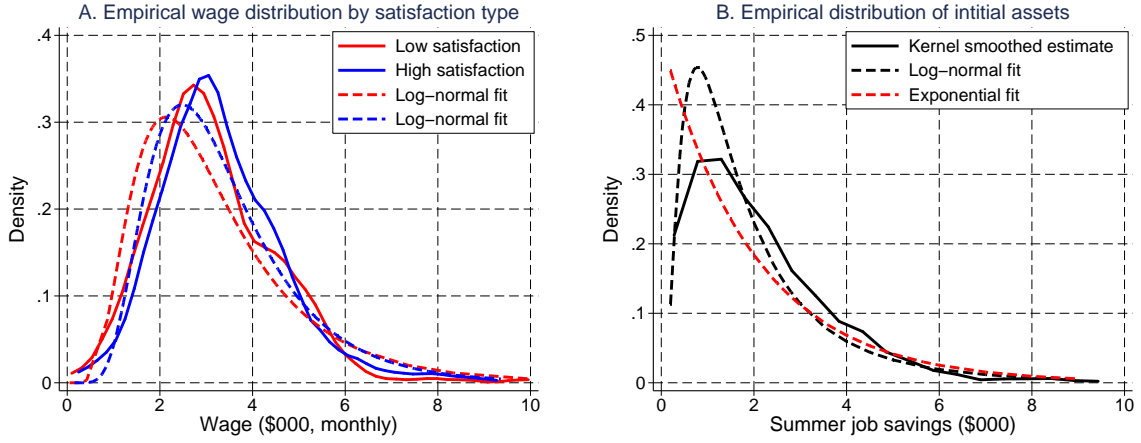
Notes: (i) Panel (A) plots kernel smoothed densities of monthly wages (in thousands of dollars) in solid and the log-normal fit of these distributions in dashed lines. (ii) Panel (B) plots a kernel smoothed density of summer savings from the year before college (BBS09) *jobsave* (in thousands of dollars) together with an exponential and log-normal fit. (iii) For both panels the data are from BBS09 corresponding to our estimation sample for workers without student debt.

Simulation. Given a vector of parameters θ we compute moments from the model as follows. Policies are solved under θ and then used to simulate $s = 1, \ldots, S$ samples of size n for 12 months to derive a dataset X_n^{Model}. Moments are computed for each of the s samples and averaged across samples to compute an expectation of the moments. We initialize simulations at $t = 0$ by specifying that each worker is unemployed and endowed with asset $a_{i,0}$ which are drawn from $\log a_{i,0} \sim \mathcal{N}(\mu_a, \sigma_a)$ with probability p_a and set to zero with probability $(1 - p_a)$. BBS09 does not provide data on savings, though it does contain data on savings due to work over the previous summer. We use this to estimate $p_a = 0.35$, $\mu_a = 0.376$ and $\sigma_a^2 = 0.817$ (see Figure 6B).[38]

Moments. The 12 moments used in our estimation of the 11 parameters are as follows. Let $n_E = \sum_{i=1}^n \mathbf{1}_{[E_i=1]}$.

1. Maximum likelihood estimates of the parameter ν of a simple hazard model of unemployment $p(u) = \nu e^{-\nu t}$ so that $\widehat{\nu} = \overline{D}_n = \frac{1}{n_E} \sum_{i=1}^n \mathbf{1}_{[E_i=1]} d_i$

---

[38]Two other ways to proceed may be as follows. First, we could assume that $\log a_{i,0} \sim N(\mu_a, \sigma_a)$ and estimate these parameters along with p_a, adding three parameters to the joint estimation. Alternatively we could assume that $a_{i,0} = \omega \times \tilde{a}_{i,0}$ where $\tilde{a}_{i,0}$ is summer savings as distributed in the data and estimate only ω, i.e. assume that initial wealth is perfectly correlated with summer savings. To estimate these new parameters would require us to add moments regarding assets to the model, which we do not have in the BBS data. We view our approach as a trading off the assurance that we are using data from a single unified source, against using better asset data from an alternative source such as the NLSY which has a different sampling scheme.

2-4. Average number of jobs since graduation $\bar{J}_n$, fraction of employed workers searching $\bar{S}_n$, and fraction of workers unemployed $\bar{U}_n$

$$\bar{J}_n = \frac{1}{n_E}\sum_{i=1}^{n}\mathbf{1}_{[E_i=1]}j_i, \qquad \bar{S}_n = \frac{1}{n_E}\sum_{i=1}^{n}\mathbf{1}_{[E_i=1]}s_i, \qquad \bar{U}_n = \frac{1}{n}\sum_{i=1}^{n}\mathbf{1}_{[E_i=0]}$$

5-7. The coefficients of a linear probability model estimated on employed workers $(\hat{\beta}_w, \hat{\beta}_x)$, and the implied standard deviation of residuals $\hat{\sigma}_e$:

$$S_i = \hat{\beta}_0 + \hat{\beta}_w\frac{w_i}{1,000} + \hat{\beta}_x\mathbf{1}_{[x_i=x_H]} + \hat{e}_i \qquad , \qquad \hat{\sigma}_e = \frac{1}{n}\sum_{i=1}^{n}\hat{e}_i^2.$$

8-12. Maximum likelihood estimates of the parameters of the *sample distribution* of (w_i, x_i) under the same parametric specification as the *population distribution*

$$\widehat{\mu}_k = \frac{1}{n_{E,k}}\sum_{i=1}^{n}\mathbf{1}_{[E_i=1,\,x_i=x_k]}\log w_i, \quad k \in \{L, H\}$$

$$\widehat{\sigma}_k^2 = \frac{1}{n_{E,k}}\sum_{i=1}^{n}\mathbf{1}_{[E_i=1,\,x_i=x_k]}\left(\log w_i - \widehat{\mu}_k\right)^2, \quad k \in \{L, H\}$$

$$\widehat{p}_H = \frac{1}{n_E}\sum_{i=1}^{n}\mathbf{1}_{[E_i=1,\,x_i=x_H]}$$

Estimation. Estimation of the parameters is achieved using a minimum distance estimator (MDE) based on the set of 12 moments $\widehat{m}_n$ described above. We define the following criterion function

$$\mathcal{Q}_n(\theta) = -\frac{n}{2}(\widehat{m}_n - m(\theta))' W_n (\widehat{m}_n - m(\theta))$$

and the associated MDE $\widehat{\theta}_n^{MDE} = \arg\min_{\theta \in \Theta} \mathcal{Q}_n(\theta)$. The weighting matrix W_n is constructed from a consistent estimate of the asymptotic variance covariance matrix of the moments Σ which satisfies $\sqrt{n}(\widehat{m}_n - m(\theta_0)) \xrightarrow{d} \mathcal{N}(0, \Sigma)$. This is found by boot-strapping from the data to obtain $\widehat{\Sigma}_n$ and then taking $W_n = \left(\text{diag}\left[\widehat{\Sigma}_n\right]\right)^{-1}$.[39] Recall that $\widehat{m}_n$ is the average across boot-strapped samples of the data to remove small sample bias.

Table 6: Parameter estimates - θ_2

Parameter		Value
A. Search costs		
Mean	κ	0.828
Upper bound	$\overline{\kappa}$	1.617
B. Disutility of labor, relative to $x_{Unemployment} = 0$		
Low satisfaction	x_L	-2.344
High satisfaction	x_H	-1.027
C. Job offer arrival rates		
Unemployed	λ_U	0.649
Employed conditional on paying κ	λ_E	0.081
D. Population distribution parameters		
Probability of high job satisfaction draw x_H	p_H	0.611
Mean of log w for x_L	μ_L	0.837
Mean of log w for x_H	μ_H	0.993
Variance of log w for x_L	σ_L^2	0.140
Variance of log w for x_H	σ_H^2	0.124

Notes: Table provides estimates of parameters of quantitative model from Section 4. The model is estimated on the 2009 cohort of the BBS data, and estimated only on students reporting zero student debt.

6 Model fit

Parameter estimates are given in Table 6 and the values of target moments in Table 7. The overall fit of the model is good, with the moments closely matching those found in the data. Recall that the amenity value of unemployment was set to zero, so we read the estimates of x_L and x_H as suggesting that work still has a lower amenity value than unemployment, but a low satisfaction job has more than twice the level of disamenity.

Table 7 column (3) provides the standard deviation of the moments from the data, constructed by taking the standard deviation of each moment across boot-strapped samples from the data. Column (4) provides similar statistics from the model, constructed by taking the standard deviation of the moment across repeated simulations of the model. These columns do not provide a formal assessment of fit, and are not statistics we have seen in similar tables in other papers where models are estimated by SMM. However we find these statistics useful. They reassure us that even given the very limited degree of ex-ante heterogeneity, the model generates similar sized variation in these moments in small samples, suggesting that the statistical properties of the data generated

---

[39]In taking the diagonal we ignore the correlations between moments which may be imprecisely measured in small samples given that some are fourth order.

Table 7: Target moments

Moments		Mean Data (1)	Mean Model (2)	Std. deviation Data (3)	Std. deviation Model (4)
A. Means					
Duration	$\mathbf{E}[h_i]$	2.500	2.943	0.108	0.052
Number of jobs	$\mathbf{E}[j_i]$	1.524	1.298	0.031	0.018
Search	$\mathbf{E}[s_i]$	0.187	0.180	0.014	0.013
Probability of x_H	$\mathbf{E}[\mathbf{1}[x_i = x_h]]$	0.716	0.732	0.017	0.014
B. Wage distribution					
Mean log w for x_l	$\mathbf{E}[\log w_i | x_l]$	1.097	1.003	0.029	0.039
Mean log w for x_h	$\mathbf{E}[\log w_i | x_h]$	1.148	1.098	0.016	0.012
Variance log w for x_l	$\mathbf{V}[\log w_i | x_l]$	0.175	0.138	0.023	0.019
Variance log w for x_h	$\mathbf{V}[\log w_i | x_h]$	0.143	0.103	0.010	0.006
C. Regression coefficients					
Wage ($000) coefficient	$\hat{\beta}_w$	-0.031	-0.033	0.012	0.010
High satisfaction coefficient	$\hat{\beta}_x$	-0.283	-0.308	0.036	0.042
Std. dev. residuals	$\hat{\sigma}_e$	0.364	0.313	0.011	0.011

Notes: (i) Column (1) gives the mean of moments from a bootstrap of the data with 10,000 re-samples (these are associated with the diagonal of the weighting matrix W_n used in estimation) (ii) Column (2) gives the mean of the moments from the $S = 1,000$ simulations of the model used to compute $\mathcal{Q}_n(\theta)$ (iii) Column (3) gives the standard deviation of the moments from the same bootstrap exercise used for column (1) (iv) Column (4) gives the standard deviations of the moments from 10,000 samples from the model.

Table 8: Additional moments

Moments		Mean Data (1)	Mean Model (2)	Std. Deviation Data (3)	Std. Deviation Model (4)
A. Unemployment, consumption and assets					
Unemployment	$\mathbf{E}[u_i]$	0.079	0.095	0.010	0.010
Mean log consumption	$\mathbf{E}[\log c_i]$	0.685	0.516	-	0.014
Consumption - Income ratio	$\mathbf{E}[c_i]/\mathbf{E}[w_i]$	0.750	0.602	-	0.006
Mean log assets	$\mathbf{E}[\log a_i]$	2.321	2.316	-	0.021
B. Other regression moments					
Regression R-squared	R^2	0.124	0.346	0.027	0.043

Notes: (i) Column (1) gives the mean of moments from a bootstrap of the data with 10,000 re-samples, (ii) Column (2) gives the mean of the moments from the $S = 1,000$ simulations of the model used to compute $\mathcal{Q}_n(\theta)$, (iii) Column (3) gives the standard deviation of the moments from the same bootstrap exercise used for column (1) (iv) Column (4) gives the boot-strap standard deviations of the moments from 10,000 samples from the model.

by the model are close to the data not only in terms of means of moments.[40]

In Table 8 we produce additional over-identifying moments. The model overstates unemployment only a little at 12 months over graduation in the BBS sample. We also take data from the SCF01. We find that the consumption to income ratio is roughly consistent, as is the mean of log consumption and log assets. The R^2 of the linear probability model for on-the-job search is too high in the model relative to the data, indicating that there is additional heterogeneity in the data not captured in the model. That being said, the fact that fit is not perfect ($R^2 = 0.346$) implies that the *iid* search costs and asset dispersion are capturing a lot of this heterogeneity.
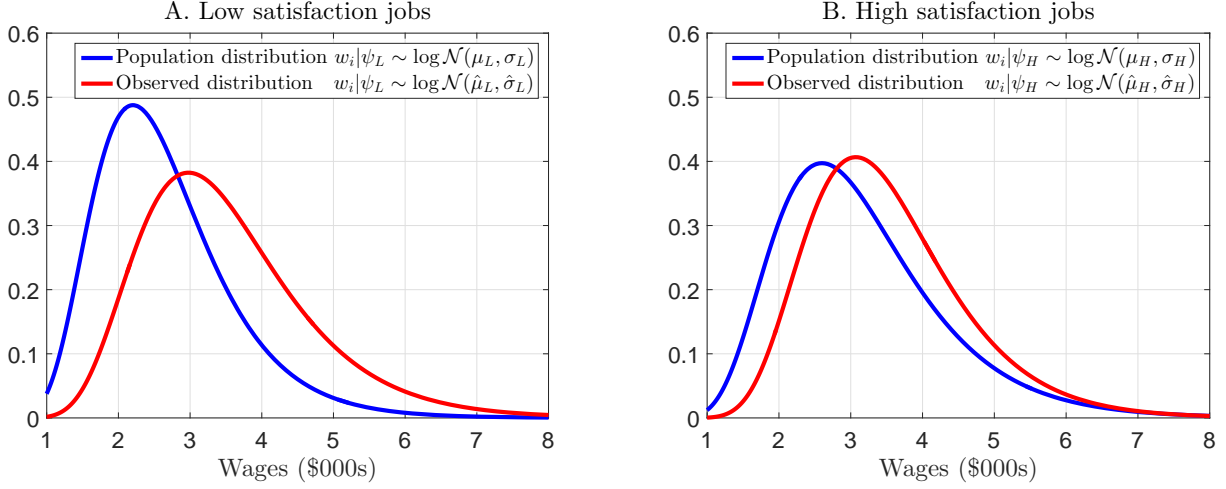
Sampling and population distributions. In Figure 7 we plot the population wage distribution (from which agents draw offers) and the resulting sample distribution which is the model counterpart to the distribution measured in the data one year after graduation. Recall from Table 7 that we match the sample distribution well. The model has a wage and amenity ladder, as workers first accepting low wage or low amenity jobs may pay to search for better wages or amenities. The fact that the mean wage of high satisfaction jobs is higher than the mean wage of low satisfaction jobs in the *sample distribution* therefore may be driven by dynamics up the job ladder rather than a higher mean in the population distribution. We find that to match the data we still need an underlying positive correlation between wages and job satisfaction in the population distribution: $\mu_L < \mu_H$. From the point of view of the model, this correlation is an innate technological feature of the jobs that workers can draw.

The amenity ladder causes workers with initially low amenities to have higher incentives to search—a relationship targeted in our calibration. Their wages grow if in their search for a high amenity job when they draw another low amenity job with a higher wage. Workers with low amenities therefore drift up the wage ladder more quickly. This leads to a wider gap between population and sample distributions for low amenity jobs.

Low vs. high asset graduates. In Figure 8 we show the evolution of cohort means of variables in the model for the one year period following entry into the labor market. We compare the paths for a cohort in which all workers have zero assets and one in which all workers have high assets which are drawn the upper quartile of the initial distribution of initial assets $\log a_{i,0} \sim \mathcal{N}(\mu_a, \sigma_a^2)$.

---

[40]To see this point consider a case in which the model was generating the moment relating to unemployment duration by a deterministic mechanism that delivers an acceptable match at exactly this frequency. The model would match this moment, but across many simulations of the panel of data from the model there would be no variance in this moment.

Figure 7: Sampling and population distributions: $\log w | x_k \sim \mathcal{N}\left(\mu_k, \sigma_k^2\right)$

Notes: (i) The figure compares the sampling and population distributions implied by the model, (ii) The population distributions are log-normal with parameters $\left\{\mu_k, \sigma_k^2\right\}_{k \in \{L,H\}}$ given in Table 6, (iii) The sampling distributions are the log-normal fits to data generated by the model on low and high satisfaction workers 12 months after graduation with parameters given by $\left\{\widehat{\mu}_k, \widehat{\sigma}_k^2\right\}_{k \in \{L,H\}}$ in column 2 of Table 7.

We discuss the wage-amenity trade-off below in the context of student debt, but here note that the model describes sensible dynamics. Workers in a higher initial asset position when unemployed are more selective in the labor market, rejecting more offers leading to higher initial wages and job satisfaction when employed. Since low asset workers quickly accept less satisfactory offers when unemployed they also incur search costs and search more when employed. They then continue to accept more offers when employed than high asset workers.

7   Results I - Role of student debt

In this section we compare the predictions of the model with student debt against the behavior detailed in Section 3, specifically the influence of student debt on the wage-amenity trade-off and search. Since we did not target moments related to student debt in our estimation of the model we view this as a clear out of sample test of the model.

High vs. low student debt.   We next consider the evolution of average labor market outcomes for cohorts by the level of student debt. We simulate three cohorts, one with no debt, one with medium debt ($d_{i,0} = \$40,000$) and one with high debt ($d_{i,0} = \$85,000$). For simplicity we set

Figure 8: Comparison of low and high debt cohorts
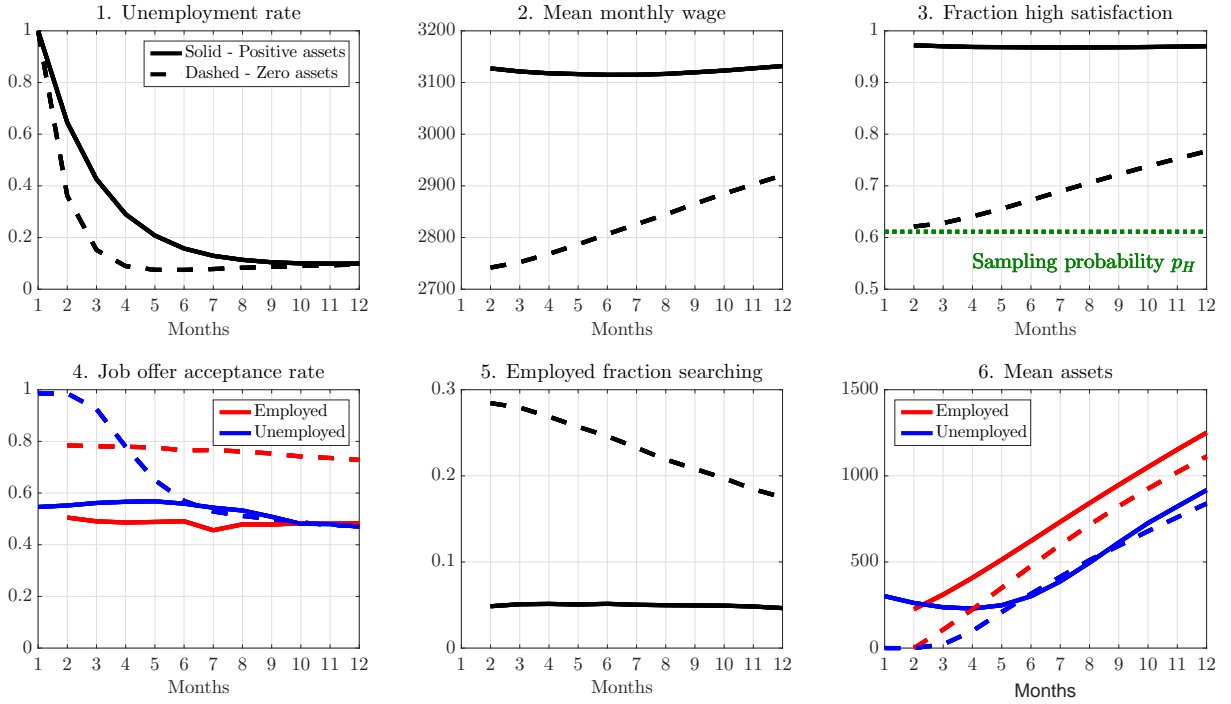
Notes: (i) Simulated cohorts have 100,000 workers, (ii) All workers in low asset cohort are initialized with $a_{i,0} = 0$, (iii) All workers in the high asset cohort draw $a_{i,0}$ from the upper quartile of the initial distribution of assets $\log a_{i,0} \sim \mathcal{N}\left(\mu_a, \sigma_a^2\right)$

initial assets to zero for all individuals.

The behaviour shown in Figure 9 is consistent with our theory and empirics. In the first period all workers are unemployed. Those with higher debts desire higher wage jobs to meet the repayments that will become due after the 6 month grace period ends. Given that on-the-job search is costly, high debt graduates use the grace period to accept fewer initial wage offers than medium debt graduates who in turn accept slightly fewer than those with no debts. This job selection leads to high debt graduates having initial jobs with higher wages.

Since wages are correlated with job satisfaction in the population distribution they also have higher job satisfaction. When employed, high debt graduates both search more *and* are more selective, seeking higher wage jobs. In particular, those in high satisfaction jobs are prepared to trade off their amenities for higher wages. This is found in their lower acceptance rate and the fact that although their wages remain higher, their job satisfaction falls relative to lower or no debt graduates. All of these patterns are maintained when comparing the medium and no debt

Figure 9: Comparison of zero, low and high debt cohorts
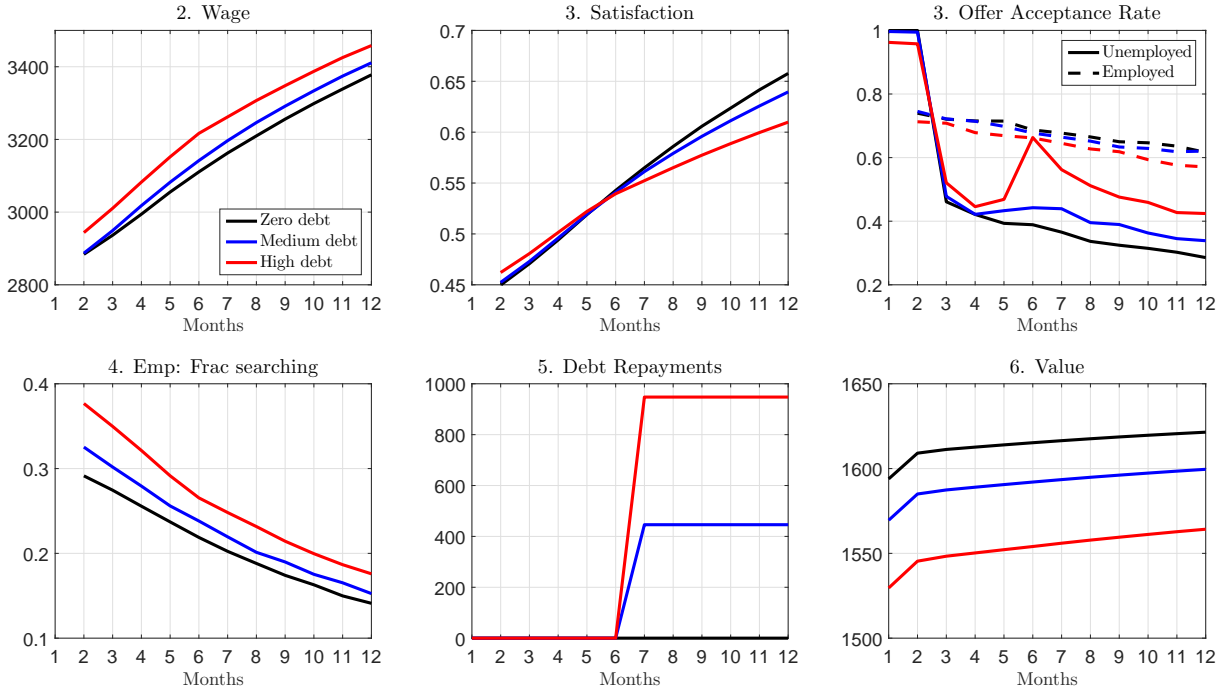
Notes: (i) Simulated cohorts have 100,000 workers each, (ii) All workers are initialized in unemployment with $a_{i,0} = 0$, (iii) High debt workers are initialized with $d_{i,0} = 85,000$, medium debt workers are initialized with $d_{i,0} = 40,000$.

graduates. At 12 months, despite a positive correlation in the population distribution, higher debt students have higher wages and lower amenities. We view this as validation of our empirical exercise.

First repayment.   Finally, we note the effect that the timing of the first repayment has on job acceptance for high debt students. Seven months after graduating first repayments on student loans become due and the acceptance rate for unemployed high debt graduates jumps up, diverging from medium and no debt and causing satisfaction to further diverge from other students. We also see a kink in the profile of high debt wages as the unemployed accept anything to meet repayments.

Summary.   Overall we find strong qualitative support of the effect of student debt on search, wages and job satisfaction and this appears to be of potential quantitative importance. In the following two sections we put numbers around this by first studying the importance of the amenity channel for assessing the welfare consequences of a change in repayment policy and second quan-

tifying amenity values in lifetime utility terms.

8 Results II - Evaluating repayment policies

External validity. First we wish to be perfectly clear regarding the external validity of the following results given our sample and estimation. Consider computing the model under policies $\mathcal{R}$ and $\mathcal{R}'$ and computing moments m and m' from a simulated cohort of graduates under each policy. Then $m' - m$ is the model's prediction of the counterfactual change in that moment under the assumption that $\mathcal{R}'$ is an unanticipated policy change enacted upon the student's date of graduation.

Welfare. We use the following measure of welfare throughout. Given a joint initial distribution of assets and student debt $F(a,d)$ upon graduation, and a repayment policy $\mathcal{R}$, *total welfare of student debt holders* $\mathcal{W}$ is

$$\mathcal{W}(\mathcal{R}) = \int U_{\mathcal{R}}(a,d,0)dF(a,d)$$

Consumption welfare $\mathcal{W}_c(\mathcal{R})$ and non-pecuniary welfare $\mathcal{W}_x(\mathcal{R})$ are computed similarly but only count the utility flows from consumption and job satisfaction, respectively.

Income based repayment. Our main policy experiment considers the effect of a change from the standard repayment plan $\mathcal{R}_S$ under which we have so far proceeded to an *Income Based Repayment Plan* (IBR) $\mathcal{R}_I$. In the United States, IBR plans were introduced in 2009, after our sample.[41] Therefore a sudden change to $\mathcal{R}_I$ seems like a valid experiment. The repayment policy $\mathcal{R}_I = (\rho_I, \Delta_I, \tau_I)$ is as follows, noting that the condition for delinquency remains the same: $(1 + r(a))a + w - \Gamma(\mathbf{s}) < \rho_I(\mathbf{s}) + \underline{c}$:

$$\begin{aligned}
\rho_I(\mathbf{s}) &= \max\left\{0.15 \times (w - 1.5\underline{c}), 0\right\} \\
\Delta_I(\mathbf{s}) &= \Delta_S(\mathbf{s}) \\
\tau_I(\mathbf{s}) &= \tau_S(\mathbf{s})
\end{aligned}$$

---

[41] In 2008 IBR plans were approved by congress. In 2009 these were implemented on a small scale. In 2010 executive action by President Obama lowered the rate of repayment for loans issued after Academic Year 2013. Since then the administration has passed actions requiring matriculating college students to be informed of IBR plans. This is very much an active policy area.

Following institutional data, under $\mathcal{R}_I$ repayments are 15 percent of disposable income, where disposable income is defined as wages minus 150 percent of the Federal poverty level $\underline{c}$ which in 2008 was \$933/month.[42] If wages are less than 150 percent of the Federal poverty level then repayments are zero.[43]

Under the US policy three further features were added that we do not consider in our analysis: (i) debt would be forgiven after 25 years, (ii) in the case that repayments $\rho_I(\mathbf{s})$ are less than interest accrued $r^d d$, the government would pay the remaining interest to avoid *negative amortization*, (iii) the actual required repayment would be the minimum of $\rho_I(\mathbf{s})$ as defined here and the payment required under a Standard Repayment policy $\rho_S(\mathbf{s})$. For now we ignore these changes and assess welfare under only the changes to ρ. This is a cleaner experiment and is comparable to the implementation of income based repayment schemes in other countries such as Australia and the United Kingdom.

Government budget. So that the policy is not providing a free-lunch to graduates, we modify the fiscal environment of the model such that a government taxes labor income at a constant rate τ to finance the cash-flows associated with loan repayment. We view the government as being able to borrow freely at an exogenous interest rate r_g and constrained by a lifetime budget constraint for each cohort. We require this budget constraint to hold in both repayment schemes upon graduation ($t = 0$).

The government's lifetime budget constraint for a cohort is as follows:

$$\sum_{t=0}^{\infty} \sum_{i=1}^{I} \frac{\tau \mathbb{1}_{it}^E w_{it} + \rho_{it}}{(1 + r_g)^t} = \sum_{t=0}^{\infty} \sum_{i=1}^{I} \frac{\left(1 - \mathbb{1}_{it}^E\right) b + r_g d_{it}}{(1 + r_g)^t}. \tag{6}$$

On the revenue side, the government receives taxes at a constant rate τ on all wages of employed workers $\left(\mathbb{1}_{it}^E = 1\right)$, and all loan repayments. On the expenditure side, the government finances unemployment benefits and must meet interest payments on the stock of debt.

When comparing welfare under $\mathcal{R}_S$ and $\mathcal{R}_I$ we solve for the tax rate τ that balances the budget. We set $r_g = 3$ percent. Since the IBR will reduce the repayments close to graduation, revenue will be lower and interest payments on debt higher, requiring higher taxes.[44] Quantitatively, these

---

[42]See: https://www.census.gov/hhes/www/poverty/data/threshld/thresh08.html

[43]All plans are computed annually and then effective for the following year however if large income fluctuations occur within a year then individuals can have their repayment re-assessed. Therefore the monthly determination model seems fine (see: StudentAid.gov).

[44]The government can also levy a lump-sum tax on all the wage earners. We also calculate the welfare comparison

A. Student welfare under $\mathcal{R}_S$ and $\mathcal{R}_I$

B. Distribution of student debt $d_{i,0}$
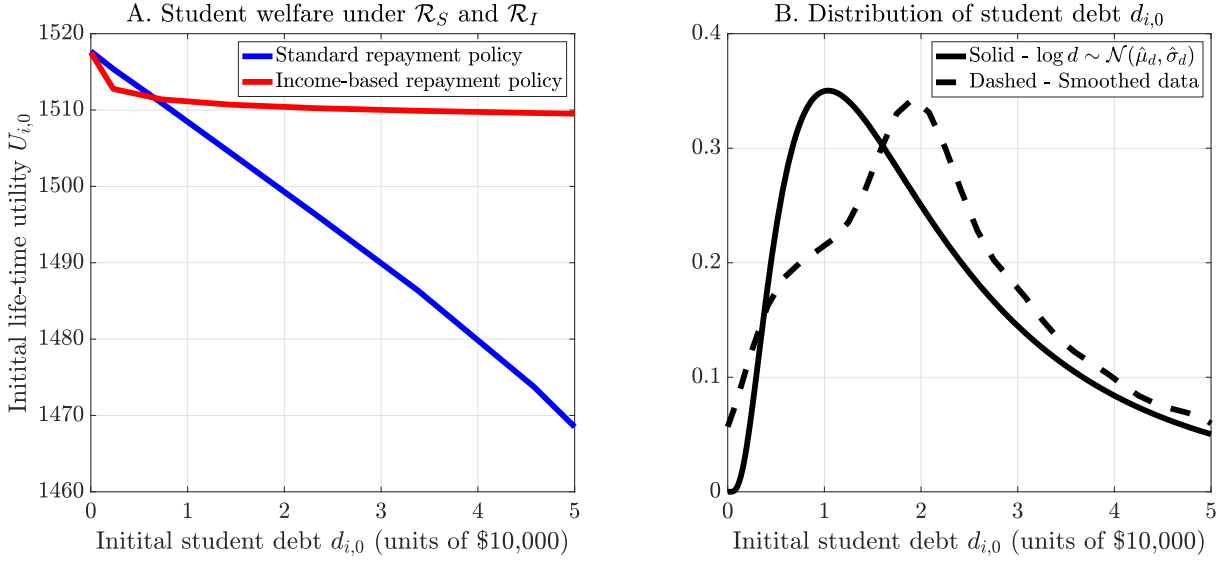
Notes: The figure plots the values taken on by the value function for unemployment $U(a, d, t)$ under the range of values for d given on the x-axis, and values of $a = 0$ and $t = 1$. These are plotted for the model solved under the baseline set of parameters under $\mathcal{R}_I$ and $\mathcal{R}_S$.

differences turn out to be minor: τ_S under the Stafford system is 1.13 percent, and τ_I under the IBR is 1.14 percent.

Welfare and debt. Figure 10 plots welfare under each of the repayment plans. At very low debt levels individuals prefer the standard repayment policy, while at moderately higher debt levels the income based repayment program is preferred. At low debt levels individuals prefer to 'get on with it' and pay down their loans quickly rather than bearing the additional interest burden that comes with the IBR. At high debt levels the standard repayment policy requires large repayments that cripple the borrower. A borrower with \$40,000 worth of debt has an initial repayment of \$500 per month at the end of the grace period, while under *IBR* the average repayment is only \$200. This has the consequence of delivering higher consumption to the worker early on in their career when wages are lower, and the marginal utility of consumption is high. For these reasons welfare under the *IBR* is flat in debt since the repayments are determined only by the wage, the slight downward slope is due to a smaller present discounted utility of consumption due to a higher future repayments due to higher accrued interest.

under this specification and results are qualitatively the same. This is most likely driven by our abstraction from an intensive margin of labor supply.

Average welfare. We can compute the average welfare under both plans by integrating initial values across the observed distribution of initial assets $a_{i,0}$ and debt $d_{i,0}$. We treat the two as independent and model their marginal densities using log-normal approximations to the data (Figure 10B). We find that the $\mathcal{W}(\mathcal{R}_S) = 1,499$ and $\mathcal{W}(\mathcal{R}_I) = 1,520$, such that the policy delivers a 1.3 percent welfare gain. We find that 89 percent of individuals prefer $\mathcal{R}_I$.

Decomposing welfare. We can use the model to decompose the gains in total welfare into that due to (i) utility of consumption, (ii) disutility of labor, (iii) search costs.

We derive three main findings, reported in Table 9. First, 63 percent of the increase in welfare under the income-based repayment plan is due to the expected utility value of consumption flows. This stems from two sources: consumption is delivered when marginal utility of consumption is higher, and the present discounted value of consumption itself is higher by 2.45 percent (final row). Second, the welfare value of lower disutility of labor (higher job satisfaction) under $\mathcal{R}_I$ is significant, accounting for 30 percent of the total difference. Third, a computation of life-time utility based on the present discounted value of wages would suggest that the Stafford scheme is preferred to the income based scheme. This is an important result and stems from the fact that wages are endogenously higher under the Stafford plan as graduates sacrifice amenities. Note that this is the type of calculation that one would employ in the standard on-the-job search model with linear utility, no savings and student debt. A general message of this exercise is that understanding the role of amenities in the utility of workers can deliver a different prediction for not only the magnitude but also the *sign* of the welfare affects of policy.

9 Results III - Evaluating job satisfaction

An interesting quantitative exercise is to use the model to measure the value of job satisfaction. Given the solution of the model under the baseline set of parameters we can ask how the worker values a transition from a low to a high satisfaction job in terms of (i) life-time consumption, (ii) willingness to pay in terms of wages, and (iii) how search behavior changes over satisfaction levels. We abstract from student debt, setting $d_i = 0$ for all workers.

(i) Consumption compensation First we determine the certain reduction in consumption in all future states that an individual in state $\mathbf{s} = (a, w, x_L)$ would be willing to incur in order to transition to a high satisfaction job today, keeping the same wage w. That is, we compute the function

Table 9: Decomposing welfare across repayment plans

		$\mathcal{R}_S$ (1)	$\mathcal{R}_I$ (2)	$\mathcal{R}_I - \mathcal{R}_S$ (3)	Fraction of diff. (4)
Total welfare	$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u(c_{i,t}) - x_{i,t} - \mathbf{1}_{S_{i,t}} \kappa_{i,t} \right]$	1499.4	1519.6	+1.35%	
Consumption	$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_{i,t})$	1714.8	1727.6	+0.75%	0.63
Work disutility	$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t x_{i,t}$	210.4	204.3	-2.90%	0.30
Search costs	$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \mathbf{1}_{S_{i,t}} \kappa_{i,t}$	4.9	3.6	-26.53%	0.07
Wages ($000's)	$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t w_{i,t}$	552.9	549.9	-0.56%	
Consumption ($000's)	$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t c_{i,t}$	477.7	489.3	+2.45%	

Notes: The table provides a decomposition for the total welfare under two different repayment regimes: $\mathcal{R}_S$ and $\mathcal{R}_I$, and compares the difference between the two. The three components in the welfare are: consumption, work disutility, and disutility from search costs. Column (4) gives the percentage of the difference between welfare under the repayment policies that is accounted for by each margin.

$\Omega(a, w)$ that satisfies

$$W^{\Omega}(a, w, x_H) := \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t u\big((1 - \Omega(a, w))c_t\big) - x_t \,\middle|\, x_0 = x_H, \, w_0 = w \right] = W(a, w, x_L).$$
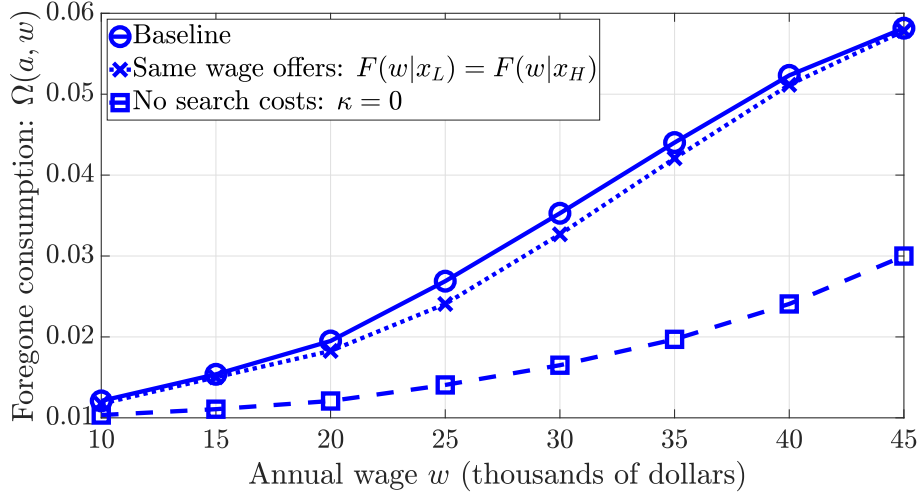
We repeat the exercise for three cases: (i) under the baseline parameters, (ii) setting the conditional wage distribution for low satisfaction draws equal to high satisfaction draws. (iii) under free on-the-job search ($\underline{\kappa} = \overline{\kappa} = 0$). We plot results in Figure 11.[45]

We first consider the baseline case. A worker in a low satisfaction job at date t would be willing forego 2 to 6 percent of *life-time* consumption from date t onwards to transition to a high satisfaction job at the same wage *today*, even when the average duration of the job is $1/\delta = 5.5$ years. This varies by wage. Higher wages in the low satisfaction job imply higher consumption and a lower marginal utility of consumption so the worker is prepared to give up more of this consumption to access a better amenity job. Eliminating the positive correlation between wages and job satisfaction in the offer distribution $F(w, x)$ has little impact, demonstrating that it is not the underlying correlation of wages and amenities that is driving the baseline results.

In the case with zero search costs, the value of a free transition from x_L to x_H— which this experiment measures—is less valuable since if not taken, the worker can still search for free in subsequent periods which increases the value of $W(a, w, x_L)$. This leads to $\Omega_{[\kappa=0]}(a, w) < \Omega(a, w)$ for all (a, w). Search frictions are important for understanding the role of amenities in worker

---

[45]Under log utility we have we have the following analytical expression $\Omega(a, w) = 1 - \exp\left\{(1 - \beta)\left[W(a, w, x_L) - W(a, w, x_H)\right]\right\}$.

Figure 11: Life-time consumption equivalent value of high job-satisfaction

Notes: The figure depicts the life-time consumption equivalent value of high job-satisfaction relative to low job-satisfaction for six cases: low and high assets, respectively, for (i) under the baseline parameters, (ii) free search ($\kappa = 0$), and (iii) setting the conditional wage distribution equal for low and high satisfaction draws. For this figure a is set to average assets in the steady state of the model.

decisions.

(ii) Wages vs. job satisfaction. Next we consider the *wage* that would make an employed worker in state $\mathbf{s} = (a, w, x_L)$ indifferent between that job and a job with higher job-satisfaction: $\mathbf{s'} = (a, w', x_H)$. That is, we compute the function $\widetilde{w}(a, w)$ that satisfies

$$W(a, \widetilde{w}(a, w), x_H) = W(a, w, x_L)$$

We compute $\widetilde{w}(a, w)$ as function of w for $a = \$0$ and $a = \$50,000$ and plot these as solid lines in Figure 12. Comparing these wages can be thought of as analogous to the indifference relationships derived in the theoretical model in Section 2 where we had a continuous support for x.

The first take-away from Figure 12 is that the function $\widetilde{w}(a, w)$ lies below the 45°-line: a worker with a low satisfaction job, zero assets and an annual wage of $20,000 will accept a high satisfaction job with an annual wage of around $12,500. In this sense the high satisfaction job is valued at $7,500 in *current* annual wages to the worker.

Second, as the wage of the low satisfaction worker increases, so does the pay-cut that the worker is prepared to take in moving to a high satisfaction job. At higher wages the worker is already able to build up savings to insure against job-loss, and is consuming more so has a

Figure 12: Wage offer indifference across assets and wages
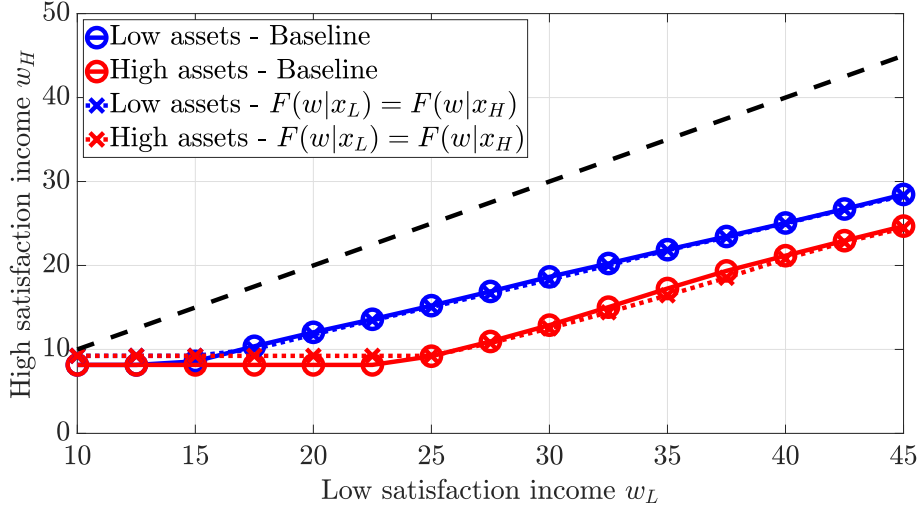
Notes: The figure shows the wage offer indifference across assets and wages for the six cases: low and high assets, respectively, for (i) under the baseline parameters, (ii) free search ($\kappa = 0$), and (iii) setting the conditional wage distribution equal for low and high satisfaction draws.

lower marginal utility of consumption. In this position the worker cares less about monetary compensation and is prepared to take a larger pay-cut to move into a more satisfying job.

Third, for a given wage a high asset worker is prepared to take a larger pay-cut than a low asset worker. This replicates the theoretical result from Section 2: higher assets tilt the worker's job acceptance policy away from wages and towards job satisfaction.

Finally, the gap between the pay-cuts acceptable to low and high asset workers narrow as wages increase. The probability of job separation is low and so a high wage worker with low assets will quickly increase their assets. Their job acceptance behavior therefore comes to approximate that of a high asset worker. In this sense wages, through the budget constraint, substitute for assets. Again, the underlying correlation in the population distribution is not important for these results.

Overall these conclusions are consistent with those in Section 2 and provide evidence for the quantitative value of job-satisfaction; in many cases the worker is prepared to take a pay-cut of around 50 percent.

(iii) Search decisions   We quantify how the satisfaction level of a job directly effects the search behavior of a worker. Recall that the conditional mean effects of wages and satisfaction on search are fitted in our indirect inference estimation - matching the parameters of the linear probability

Figure 13: Reservation wage $\underline{w}(a,x)$ for on-the-job search at κ by assets and job satisfaction
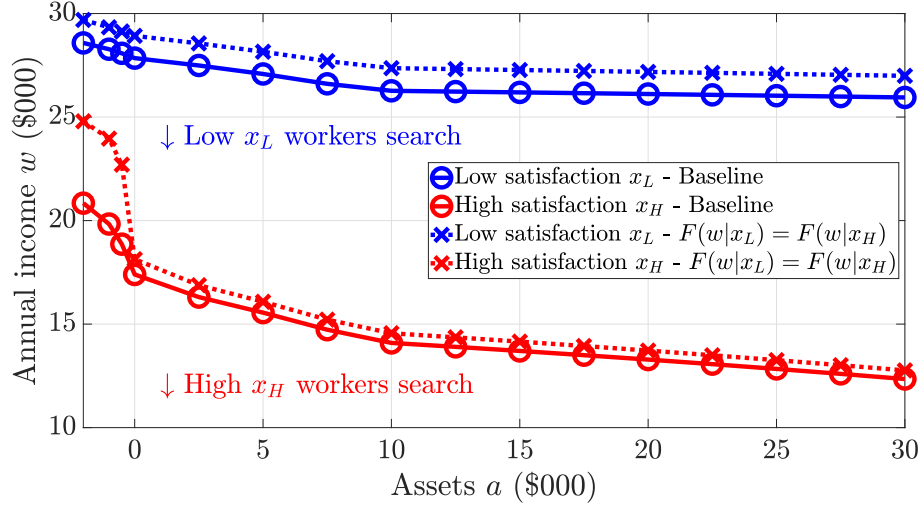
Notes: The figure depicts the on-the-job search decisions $\underline{w}(a,x_L)$ and $\underline{w}(a,x_H)$ under the mean search cost κ.

model $S_i = \alpha + \beta_w w_i + \beta_x x_i$. We now use the model to extend the analysis to include the effect of assets, which were not observed in our data. For a given value of κ and states (a,x) we can determine the threshold *reservation wage for search* $\underline{w}(a,x)$ such that for all $w < \underline{w}(a,x)$ the individual searches. That is, $\underline{w}(a,x)$ equates the marginal benefit of search to the marginal cost, satisfying

$$W(a,\underline{w}(a,x),x) - \int \max\left\{W(a,\underline{w}(a,x),x), W(a,w',x')\right\} dF(w',x') = \frac{\kappa}{\lambda_E}$$

Figure 13 shows $\underline{w}(a,x_L)$ and $\underline{w}(a,x_H)$ for the mean value κ. For both levels of job satisfaction the reservation wage for search is declining in assets. For asset values over \$60,000, no high satisfaction worker searches at the mean κ since they can no longer find higher satisfaction jobs and the value of search in terms of wage outcomes does not out-weigh the cost. As the wage and asset levels of the worker decrease, workers search more. The main difference in search behavior, however is due to the satisfaction of the worker in their current job. Consider workers with assets of \$20,000: only high satisfaction workers in jobs paying less than around \$14,000 search, while the cut-off wage for low satisfaction workers is around \$27,000.

10   Conclusion

In this paper we established that the level of student debt held by a graduating college student has a statistically significant effect on early labor market behavior and outcomes. Specifically, higher

levels of debt cause workers to end up in jobs with higher wages and lower job satisfaction, and to search more on-the-job. We showed that these outcomes can be neatly rationalized by a simple extension of the Lise (2013) model of search with asset accumulation to accommodate amenities. Since wages and asset levels are linked through the budget constraint whereas job satisfaction is not, higher levels of debt cause workers to *substitute* higher wages for lower job satisfaction in their reservation policies.

We extended this simple model to a quantitative framework which we estimate using the novel observables provided by the NCES BBS09 data. We found a quantitatively important role for job satisfaction in shaping the labor market behavior and strong evidence for its interaction with assets. Modelling the institutional framework of US Federal student loans in 2008 we showed that the model's out-of-sample predictions for students with student debt fits well with the data, recommending the model for policy analysis.

We considered a simple policy experiment of a transition to a *income based repayment scheme* (IBR). The IBR is strictly preferred by students with higher debt burdens as it allows student to intertemporally shift large repayments to periods when the marginal utility of consumption is lower. This consumption effect accounts for 63 percent of the welfare gains. The remainder is dominated by the higher job satisfaction achieved by graduates under the less pressing repayment requirements. Importantly, the large trade-off of job satisfaction and wages found under the standard repayment system would lead one that considers only wages to mistakenly infer that the standard repayment system is preferred.

This paper can be extended along a number of dimensions. In particular, the quantitative model could be extended to accommodate college choice. Such a model would treat the continuation values in the existing model as data, and utilize the new *College Scorecard* data to estimate production functions for colleges. This setting would allow policy experiments of the type: *"Suppose the repayment policy changed from $\mathcal{R}$ to $\mathcal{R}'$, what is the total effect on welfare, including debt take up and college choice?"* rather than being qualified by our statements at the beginning of Section 8. We establish that the role of amenities, and their ability to rationalize the otherwise puzzling evidence we provide that debt can cause wages to *increase*, will be key to any such exercise.

References

ALTONJI, J. G., KAHN, L. B. and SPEER, J. D. (2016). Cahiser or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success. *Journal of Labor Economics*, **34** (S1), S361–S401.

ALVARO MEZZA, S. S., DANIEL RINGO and SOMMER, K. (2019). Student loans and homeownership. *Journal of Labor Economics (forthcoming)*.

BONHOMME, S. and JOLIVET, G. (2009). The pervasive absence of compensating differentials. *Journal of Applied Econometrics*, **24** (5), 12–25.

CHAPMAN, S. (2015). Student loans and the labor market: Evidence from merit aid programs.

DELAVANDE, A. and ZAFAR, B. (2014). University choice: The role of expected earnings, non-pecuniary outcomes and financial constraints.

DEY, M. and FLINN, C. (2008). Household search and health insurance coverage. *Journal of Econometrics*, **145** (1-2), 43–63.

FIELD, E. (2009). Educational debt burden and career choice: Evidence from a financial aid experiment at nyu law school. *American Economic Journal: Applied Economics*.

GERVAIS, M. and ZIEBARTH, N. (2017). Life after debt: Post-graduation consequences of federal student loans.

— and ZIEBRATH, N. (2019). Life after debt: Postgraduation consequences of federal student loans. *Economic Inquiry (forthcoming)*, **57**.

HALL, R. E. and MUELLER, A. I. (2018). Wage dispersion and search behavior: The importance of nonwage job values. *Journal of Political Economy*, **126** (4), 1594–1637.

HERBST, D. (2019). Liquidity and insurance in student loan contracts: The effects of income-driven repayment on borrower outcomes (working paper).

HERKENHOFF, K., PHILLIPS, G. and COHEN-COLE, E. (2016). The Impact of Consumer Credit Access onEmployment, Earnings and Entrepreneurship.

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

JAROSCH, G. (2015). Searching for job security and the consequences of job loss.

JI, Y. (2017). Job search under debt: Aggregate implications of student loans.

KAPLAN, G. and SCHULHOFER-WOHL, S. (2017). Understanding the Long-run Decline in Interstate Migration. *International Economic Review*, **58** (1), 57–94.

— and VIOLANTE, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, **82** (4), 1199–1239.

LISE, J. (2013). On-the-Job Search and Precautionary Savings. *Review of Economic Studies*, **80** (3), 1086–1113.

LOCHNER, L. and MONGE-NARANJO, A. (2015). Default and repayment among baccalaureate degree earners. In B. Hershbein and K. Hollenbeck (eds.), *Student Loans and the Dynamics of Debt*, W.E. Upjohn Institute for Employment Research, vol. 8.

MAS, A. and PALLAIS, A. (2017). Valuing Alternative Work Arrangements. *American Economic Review*, **107** (12), 3722–3759.

ROTHSTEIN, J. and ROUSE, C. E. (2011). Constrained after college: Student loans and early-career occupational choices. *Journal of Public Economics*, **95** (1), 149–163.

SORKIN, I. (2018). Ranking firms using revealed preference. *Quarterly Journal of Economics*, **133** (3), 1331–1393.

WEIDNER, J. (2016). Does student debt reduce earnings?

YANNELIS, C. (2015). Asymmetric Information in Student Loans.

— (2017). Debt and human capital: Evidence from student loans (working paper).

Appendix For Online Publication

This Appendix is organized as follows. Section A includes additional tables referenced in the text. Section B includes additional figures referenced in the text. Section C includes additional details regarding the data.

A Additional tables

Table A1: Alternative instrument 1: $Z_{ijc} = Budget_{ijc} - Grants_{ijc}$

	Baseline	Debt		Income		OLS
		Low	High	Low	High	2001-09
	(1)	(2)	(3)	(4)	(5)	(6)
A. Earnings						
$Earnings_{ijc}$	0.0506	0.0470	0.1044	-0.0014	0.0402	0.0036
	(0.04916)	(0.53181)	(0.09569)	(0.02301)	(0.07251)	(0.01263)
$\log Earning_{ijc}$	0.0010	-0.0012	0.0026	-0.0003	0.0005	-0.0002
	(0.00107)	(0.01172)	(0.00206)	(0.00093)	(0.00100)	(0.00030)
B. Satisfaction						
$Satis_{ijc}$	-0.0003	0.0070	-0.0010	-0.0003	-0.0006	-0.0005*
	(0.00118)	(0.01299)	(0.00226)	(0.00179)	(0.00149)	(0.00030)
$\widetilde{Satis}_{ijc}$	-0.0005	0.0046	-0.0012	-0.0003	-0.0010	-0.0005
	(0.00119)	(0.01299)	(0.00227)	(0.00179)	(0.00151)	(0.00031)
C. Other						
$Search_{ijc}$	0.0010	-0.0295	-0.0049	0.0027	-0.0010	0.0009**
	(0.00128)	(0.06082)	(0.01399)	(0.00188)	(0.00151)	(0.00033)
$Duration_{ijc}$	-0.1978**	-0.7700	-1.0546	-0.3485**	-0.0557	
	(0.08750)	(2.17102)	(0.89504)	(0.14602)	(0.10361)	
$Teach_{ijc}$	-0.0012*	-0.0021	-0.0044***	-0.0005	-0.0018**	0.0004**
	(0.00074)	(0.00792)	(0.00147)	(0.00114)	(0.00092)	(0.00014)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$ This table replicates Table 4, but for an alternative specification of the instrument. Here the instrument considered is the individual student budget (reported by the university) minus institutional grants for individual i in college j for cohort c. For further details of these variables see Table C1.

Table A2: Alternative instrument 2: $Z_{ijc} = \widehat{AveDebt}_{jc}$

	Baseline	Debt		Income		OLS
		Low	High	Low	High	2001-09
	(1)	(2)	(3)	(4)	(5)	(6)
A. Earnings						
$Earnings_{ijc}$	0.1955***	0.4090	0.2569**	0.0132	0.1591*	0.0081
	(0.06709)	(0.47158)	(0.12772)	(0.03336)	(0.09236)	(0.01497)
$\log Earning_{ijc}$	0.0042***	0.0008	0.0077***	0.0002	0.0025*	-0.0000
	(0.00145)	(0.01037)	(0.00279)	(0.00134)	(0.00127)	(0.00033)
B. Satisfaction						
$Satis_{ijc}$	-0.0019	0.0052	-0.0046	0.0005	-0.0053***	-0.0004
	(0.00159)	(0.01152)	(0.00300)	(0.00259)	(0.00191)	(0.00028)
$\widetilde{Satis}_{ijc}$	-0.0026	0.0001	-0.0048	0.0005	-0.0065***	-0.0004
	(0.00159)	(0.01151)	(0.00301)	(0.00259)	(0.00196)	(0.00029)
C. Other						
$Search_{ijc}$	0.0029*	0.0032	0.0023	0.0038	0.0037*	0.0009**
	(0.00171)	(0.01230)	(0.00317)	(0.00273)	(0.00208)	(0.00032)
$Duration_{ijc}$	-0.2120**	0.0984	-0.2493	-0.4317**	-0.0682	
	(0.10563)	(0.65243)	(0.16428)	(0.21970)	(0.11295)	
$Teach_{ijc}$	-0.0034***	-0.0131*	-0.0075***	-0.0028*	-0.0035***	0.0003*
	(0.00102)	(0.00750)	(0.00204)	(0.00169)	(0.00122)	(0.00016)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$ This table replicates Table 4, but for an alternative specification of the instrument. Here the instrument considered is the predicted value from a college level regression of $AveDebt_{jt}$ on year and college fixed effects as well as controls for the number of students, and $Grants_{jt}/(Grants_{jt} + Loans_{jt})$. For further details of these variables see Table C1.

B Additional figures

Figure B1: Variation of the instrument across cohorts
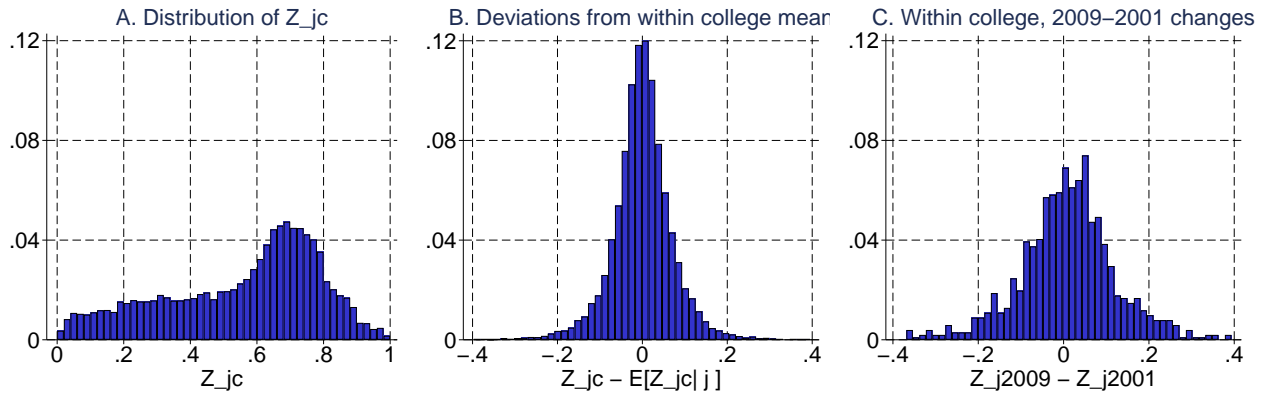
Notes: These figures plot distributions of statistics constructed from our instrument $Z_{jc} = Grants_{jc}/(Grants_{jc} + Loans_{jc})$. Data are from IPEDS 2001-2015, with the sample of colleges restricted to those in Carnegie categories discussed in the main text, and for which we have information from the College Score Card data regarding entering cohort median SAT scores. The median number of colleges each year is $1,202$. Panel A plots the distribution of Z_{jc} in the pooled sample of colleges and cohorts. Panel B plots the distribution of deviations of Z_{jc} from within college means $Z_{jc} - \overline{Z}_j$. Panel C plots the distribution of changes in Z_{jc} between the two cohorts in our analysis: 2001 and 2009.

Figure B2: Student debt and job satisfaction
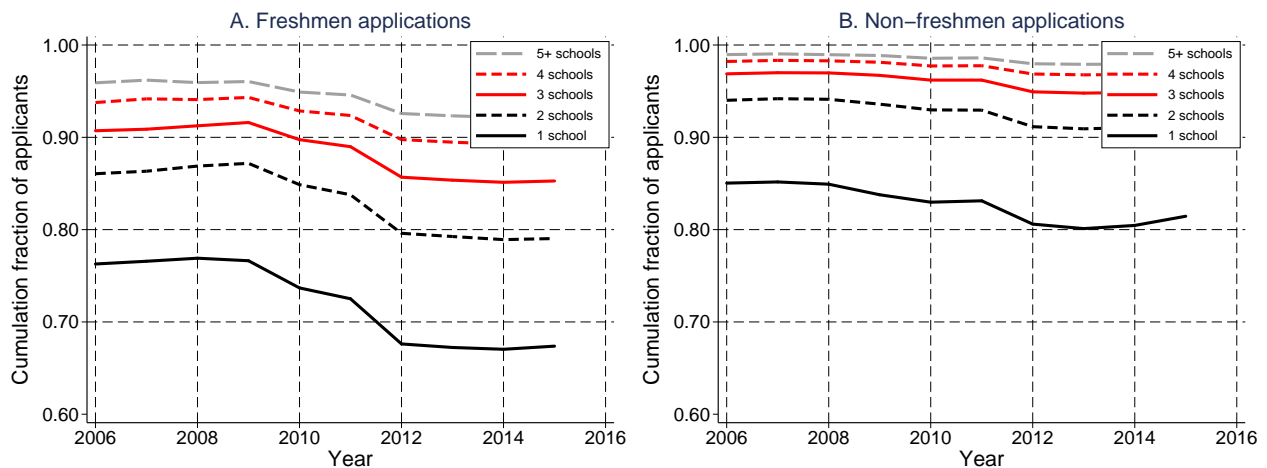
Notes: Data from StudentAid.gov (link below). Graphs plot the cumulative fraction of FAFSA applications for Federal Student Aid that list the specified number of colleges. Data are for individuals filing original applications for student aid. The average annual number of Freshman applications for student aid over 2006-2015 was 6,640,518. The average annual number of Non-freshman applications for student aid over 2006-2015 was 12,453,944. The latter is larger since individuals also file for college aid once in college. https://studentaid.ed.gov/FAFSA-data

C Data

Table C1 provides the precise mapping between the variables we consider and the questions asked in the BB03 and BBS09 surveys. In all cases these questions are consistent over time.

Table C1: BB data variable definitions

Variable in paper		Code	Description
A. Outcome variables			
Income	$Earnings_{ijc}$	b1ansal	Annual income from all sources
Teaching	$Teach_{ijc}$	b1occ33	Are you currently employed as a $K-12$ teacher?
On the job search	$Search_{ijc}$	b1search	Are you currently looking for a different job? (Only asked if employed)
Duration of unemployment	$Duration_{ijc}$	b1timoff	As of the 2009 interview, the number of weeks after the respondent began applying for job that the respondent received the offer for the job currently held.
Job satisfaction	$Satis_{ijc}$	b1jbover	Overall, would you say that you are satisfied with your job?
B. Control variables			
Parental income		cincome	Indicates the total 2006 income of parents of dependent respondents (Used in Federal need analysis to determine AY 07-08 eligibility)
GPA measure		gpa2	Cumulative undergraduate grade point average (between 1.00 and 4.00)
SAT measure		tesatcre	The sum of reported SAT verbal and math scores

Notes: The codebook for the NCES BBS surveys is itself restricted use since it discloses statistics along with each question, and so is available only with a data license. Here we have reproduce the text from the codebook for the variables we have considered in this paper. More information regarding these variables is available upon request.