Student Loans and Labor Supply Incentives

Gustavo Manso Alejandro Rivera Hui (Grace) Wang Han Xia^{*}

Abstract

We develop a dynamic household finance model showing that student loans – nondischargeable in the U.S. bankruptcy – alleviate the well-documented debt overhang in labor supply decisions. Non-dischargeability mutes opportunities for households to strategically reduce labor supply at the expense of creditors, thus correcting incentive distortions. This corrective effect, however, is partially undone by Income Driven Repayment (IDR) plans, which set student loan payments formulaically regardless of outstanding balance. IDR thus allows households to pseudo "discharge" student debt and re-activates debt overhang. We supplement our model with empirical analyses and uncover potentially unintended consequences of proposed reforms in student loans.

Keywords: Student loans, Household debt, Debt overhang, IDR plans.

JEL: D14, H52, H81, I22, I28.

^{*}We thank Ehsan Azarmsa (discussant), Asaf Bernstein, Philip Bond, Michael Boutros (discussant), Nuno Clara, Michael Dinerstein (discussant), Jason Donaldson, Brett Green, Menaka Hampole, Sasha Indarte, Urban Jermann, Paymon Khorrami, Elena Loutskina (discussant), Giorgia Piacentino, Marzena Rostek, Jacob Sagi, Tim de Silva, Vasiliki Skreta, Yufeng Wu, Liyan Yang, and Constantine Yannelis (discussant), as well as the audience at 2025 Education Finance Conference at the University of Chicago, 2025 UNC-Duke Corporate Finance Conference, 2024 Colorado Winter Conference, 2024 Texas Economic Theory Conference, 2024 MFA annual conference, Haas Business School, Southern Methodist University, University of Oklahoma, and UNC-Chapel Hill, for helpful comments. Manso is with Haas Business School at the University of California, Berkeley; manso@haas.berkeley.edu; Rivera and Xia are with the University of Texas at Dallas; axr150331@utdallas.edu; Han.Xia@utdallas.edu. Wang is with Bentley University; huiwang@bentley.edu.

1 Introduction

Student loans have become a predominant part of household debt over the past two decades. As documented in Looney and Yannelis (2015), student loan volume increased by almost tenfold from the early 1980s to the current years.¹ This drastic change has garnered attention from recent empirical literature studying the consequences of student debt on household decisions. A common theme of this literature is that student loans create excessive burdens and cause a negative impact on students' personal and professional life paths.²

In this paper, we propose one of the first theoretical frameworks in the literature to uncover an unexplored bright side of student loans – their ability to alleviate the debt overhang problem of household debt on labor decisions. We also study the nuances of this ability in consideration of various repayment plans – including *Income Driven Repayment* (IDR) plans and the designated student loan forgiveness, both of which have been under the spotlight of policymakers.

To do so, we build a dynamic model featuring a risk-averse household that borrows from the credit market and optimally chooses its labor supply. We first establish that household indebtedness induces a debt overhang problem – reaffirming this well-documented result from the prior literature (e.g., Donaldson, Piacentino, and Thakor 2018; Bernstein 2021). Debt overhang arises because households may not fully internalize the benefits of supplying labor, particularly when default is expected (e.g., He 2011; Diamond and He 2014). Specifically, households in the U.S. are protected by limited liability and often discharge their debt in bankruptcy. Therefore, any incremental wages generated by labor supply before default are partially used to fulfill debt obligations (via debt repayment), which postpones the discharge of debt and induces a wealth transfer from households onto lenders. In anticipation,

¹Outstanding student loan balance in the U.S. totals \$1.58 trillion. See https://www.newyorkfed.org /microeconomics/topics/student-debt. Although current default rates are below historic peaks (e.g., 1991), credible forecasters predict that nearly 40% of student borrowers will default over the 20-year horizon following their college entrance; Scott-Clayton (2018).

²The impact includes delayed homeownership and family formation (Gicheva 2011; Cooper and Wang 2014; Mezza, Ringo, Sherlund, and Sommer 2016; Goodman, Isen, and Yannelis. 2021), reduced small business formation (Amromin and McGranahan 2015), lower graduate school enrollment (Chakrabarti, Fos, Liberman, and Yannelis 2023), less entrepreneurship (Krishnan and Wang 2019), reduced stock market participation (Batkeyev, Krishnan, and Nandy 2017), and suboptimal labor market outcomes or human capital decisions (Minicozzi 2005; Rothstein and Rouse 2011; Weidner 2016; Lou and Mongey 2019; Ji 2021; Hampole 2023).

households choose to exert less effort ex ante. Put differently, the option to default – and the ability to discharge debt – creates an opportunity for the household to strategically reduce labor supply at the expense of lenders.

This opportunistic behavior, however, is curbed by the presence of student loans. Federal student loans – the dominant type on the student loan market – are distinct from other forms of consumer debt in that they are almost completely non-dischargeable in U.S. bankruptcy nowadays (Yannelis, 2020a). Upon default, student loans borrowers are subject to wage garnishment and are obligated to eventually pay off these loans (via garnished wages), even though their other non-student debt might be discharged during bankruptcy.³ In this case, opportunistic behavior – such as shunning away from the labor supply – cannot shield student borrowers from making up the shortfall in repayment. Since the creditors' debt is safe, any additional dollar generated through labor supply is not a wealth transfer from households onto lenders anymore. Instead, households are the residual claimants of additional labor income. Student loans effectively mute the shirking opportunity for households, correcting their disincentive to supply labor.

Such a *corrective* effect of student loans constitutes the first finding of our model. More specifically, we start by considering a benchmark case without student loans. We show that households' optimal labor supply decision exhibits a hump shape with respect to household indebtedness. As household indebtedness initially increases, a larger fraction of household income accrues to paying off debt, which lowers the households' overall consumption level and thus, increases the marginal utility of consumption. In this regime, households work harder (by increasing labor supply) to maintain the desired level of consumption. This incentive, however, is reversed after the indebtedness surpasses a threshold – when the debt overhang effect kicks in. In this regime, default becomes more probable and thus, households undertake opportunistic behavior to avoid transferring wealth (earnings from labor supply) onto lenders. Consequently, labor supply begins to decline as indebtedness further increases from this point, reflecting the aggravated debt overhang problem.

The presence of student loans alleviates debt overhang. When student loans take up a larger proportion of the household's total debt, the debt overhang does not kick in until a

 $^{^3 \}mathrm{See}$ Section 2.3.1 for institutional details including wage garnishment, bankruptcy filing, and automatic stay.

later stage: that is, the household's distorted incentive to cut back on labor supply does not show up until the indebtedness reaches a higher level than the case without student loans – indicating a "delayed" manifestation of debt overhang. Conditional on debt overhang, the decline in labor supply becomes less aggressive when student loans are present – indicating a "dampened" manifestation of debt overhang. Both the delayed and dampened manifestation of debt overhang point to the *corrective* effect of student loans.

In the next part of the model, we substantiate our baseline findings by considering student loan repayment options and whether the *corrective* effect of student loans varies with these options. The standard repayment plan for federal student loans follows the 10-year amortization schedule.⁴ An alternative repayment option that has received great regulatory attention pertains to the *Income Driven Repayment* (IDR) plans. Different from the standard amortization schedule, IDR sets the payment as a proportion of borrowers' discretionary income, and eventually forgives the remaining principal after 20-25 years of continuing payments.

Importantly, the proportion of income required as IDR payment is invariant with respect to borrowing amount and outstanding balance. This invariance creates a pseudo opportunity for borrowers to "discharge" student loans, thus undoing part of the *corrective* effect of student loans. To see this intuition, consider a student loan borrower making a decision on optimal labor supply. The borrower is presented with an option to later enter IDR. Once in IDR, the student's monthly payment becomes independent of how much he actually owes prior to IDR – as if this amount has been "discharged", in exchange for a fixed fraction of the borrower's future earnings.⁵

This pseudo dischargeability engenders a distortion in household incentives. Because student loans now resemble other types of consumer debt, the debt overhang problem resurfaces. That is, anticipating wealth transfers to lenders (due to debt "discharge"), households would again be discouraged from supplying labor ex ante. As such, IDR plans reactivate debt overhang and (partially) undo student loans' *corrective* effect as shown in the baseline

⁴The student loan interest rates are set by federal law and fixed for a cohort of students. Specifically, the interest rates are a function of loan type and disbursement date, and for the majority types, they are equivalent within a cohort regardless of borrowers' individual credit history or whether borrowers finance a degree from a four-year institution or a two-year degree from an online for-profit program.

⁵In Section 3.2, we discuss the possibility that the remaining balance forgiven may be taxable.

model. We label this reactivation of debt overhang as the *setback* effect of IDR. This finding highlights an important distinction between (1) wage garnishment upon student debt default and (2) the committed fraction of earnings as IDR repayment. In both instances, households forfeit part of income to lenders, but in the former, the (anticipation of) forfeited amount encourages household to supply labor – yielding the *corrective* effect, whereas in the latter, it discourages labor supply – yielding the *setback* effect. This "paradox" arises because in the former instance, the extent of garnished wage increases with the outstanding student loan balance – and by supplying labor ex ante, households can lower this balance and mitigate wage garnishment. In the latter instance, however, IDR monthly payment is set formulaically and invariant to the outstanding balance; borrowers are not better off entering IDR with a lower balance. Section 3.6 discusses this paradox in more detail.

Interestingly, we show the *setback* effect of IDR becomes aggravated when student loans make up a larger part of household debt. Households' decision to enter IDR faces a tradeoff. In addition to considerable non-monetary costs associated with IDR enrollment (Section 3.1), households commit a fraction of future earnings as IDR repayment (which they could avoid by waiting for a positive income shock and postponing IDR enrollment). Therefore, households are only willing to make such sacrifices when pseudo discharging student loans is worthwhile. This happens when they owe a large amount of student debt to begin with, in which case getting rid of the unaffordable payment (under standard amortization) is more valuable than incurring the costs of entering IDR and pledging a fraction of future earnings. Hence, a tension emerges. On the one hand, a larger proportion of student loans in household debt encourages them to supply more labor instead of shirking – the *corrective* effect. On the other hand, a larger amount of student loans makes IDR more attractive and more likely to be adopted – giving rise to a greater *setback* effect. Whether the *setback* effect can offset the *corrective* effect becomes an empirical question, and we shed light on this question by performing a calibration analysis.

Existing studies (e.g., U.S. Government Accountability Office (2015)) find that over the past decade, the IDR take-up rate is about 20% among students in need of repayment assistance. This low utilization indicates that in practice, households likely perceive the cost of filling for IDR substantial, relative to the benefits of doing so.⁶ In light of this

⁶The IDR take-up rate has grown over time. For instance, using the 2019 wave of the Survey of Consumer

observation, we calibrate the IDR take-up rate in our model – through the trade-off in households' decisions of whether to enter IDR – such that it matches the observed rate in practice. We solve our model under the calibrated parameters and find that in equilibrium, the *setback* effect introduced by IDR does not completely offset the *corrective* effect of student loans. Therefore, households with student loans in general are more resilient to debt overhang and are encouraged to supply labor.

We bring this prediction to observational data and perform supplementary empirical analyses. We employ data from the 1997 National Longitudinal Survey of Youth (hereafter, NLSY97), conducted by the U.S. Bureau of Labor Statistics. NLSY97 surveys a representative sample of Americans since their teen ages, and tracks various financial and professional information over their lives. The survey allows us to observe each household's itemized balance sheets (capturing household indebtedness), week-by-week labor records (capturing household labor supply), and the presence of student loans among household debt. Consistent with the model prediction, we find that given the *same* level of household indebtedness, a larger proportion of student loans among total debt is associated with less severe debt overhang. Specifically, households do not cut back their labor supply until a higher level of indebtedness – a delayed manifestation of debt overhang. Once debt overhang is present, the decline in labor supply is less aggressive among households with more student loans – a dampened manifestation of debt overhang. For identification, we employ location based variation in student loan borrowing to confirm our results.⁷

Overall, our theory, supplemented by empirical analyses, documents the significant role of student loans in mitigating household debt overhang, as well as its nuances with respect to IDR plans. Building on these analyses, we derive policy implications, speaking to the recent regulatory efforts promoting accessible and manageable student debt repayment.

The Biden administration announced the student loan relief program in 2022, including a few adjustments to the existing IDR plans. For instance, the program proposes to lower the IDR repayment to 5% of students' discretionary income, down from the current 10-

Finances (SCF), Catherine and Yannelis (2023) find that the rate has increased to over 30%.

⁷This empirical finding is in line with Daniels and Smythe (2019), who show that post-college income of individuals receiving student loans is 8-9% higher than the income of students without a loan. Such a difference is driven by more working hours, indicating that students borrowing loans are more active in supplying labor.

15%. Based on comparative statistics of our model, we show that a lowered proportion would make IDR more attractive to student borrowers and thus allure them to enter this plan at an earlier stage. Such anticipated heavier reliance on IDR, and in turn the greater (pseudo) discharge of student debt, discourages borrowers from providing effort *ex ante* – before they switch from the standard payment plan to IDR. Similarly, the program proposes to shorten the duration of IDR payment until loan forgiveness to 10 years, down from the current 20-year waiting period. Such an adjustment likewise increases the appeal of IDR and thus amplifies the IDR's *setback* effect, discouraging households from supplying labor ex ante. Overall, these implied effects may potentially undermine the government's initiative to relieve households from debt burden.

Literature Review

Our paper contributes to the growing literature on student loans by depicting their potential bright side in alleviating the debt overhang of household labor supply. Besides the long-term consequences of student debt (discussed at the beginning of the *Introduction*), this literature also investigates student loan credit risk factors including academic institution type and funding sources, students' non-cognitive traits, declining real estate values, tuition and student demographics, and family circumstances.⁸

Within this literature, a group of papers focuses on the IDR plans. Mueller and Yannelis (2022) find that aversion to administrative paperwork – part of the IDR application process – deters borrowers from adopting IDR plans. Catherine and Yannelis (2023) study the distribution consequences of student loan forgiveness, and compare them to an alternative form of forgiveness through expanding the current IDR plans. Herbst (2023) and Abraham, Filiz-Ozbay, Ozbay, and Turner (2018) study the impact of IDR on delinquencies, credit outcomes and career decisions. Cornaggia and Xia (2024) study the role of student loan servicers in influencing the decision of financially distressed borrowers on IDR take up.

In a recent study, Boutros, Clara, and Gomes (2024) build a dynamic model to examine

⁸See, e.g., Andolfatto and Gervais (2006), Lochner and Monge-Naranjo (2011), Shapiro (2014), Looney and Yannelis (2015), Fox, Bartholomae, Letkiewicz, and Montalto (2017), Mueller and Yannelis (2018), Eaton, Howell, and Yannelis (2020), Cornaggia, Cornaggia, and Xia (2021), Armona, Chakrabarti, and Lovenheim (2022), Cornaggia, Cornaggia, and Xia (2024), Gallagher, Billings, and Ricketts (2023). Other studies examine the effect of federal policies on student loan borrowing (Lucca, Nadauld, and Shen, 2019; Kargar and Mann, 2023).

households' choice between repaying student loans under the standard payment plan or IDR and delinquency. The authors calibrate the model and quantitatively analyze the welfare gains from alternative debt contracts – designed to defer payments until borrowers' later life stages. Our study complements theirs by examining how student loans, along with various repayment plans, shape household *labor supply incentives*. This aspect allows us to uncover a unique role of student loans in alleviating the well-documented debt overhang on labor supply.

In a related study by Karamcheva, Perry, and Yannelis (2020), the authors empirically show that IDR plans are subject to significant adverse selection: borrowers with low earnings and high loan balances are most likely to adopt IDR, and accordingly, loan forgiveness at the end of IDR plans is greatest for these borrowers. Even though our theoretical framework focuses on moral hazard in labor supply decisions, the implications embedded in our model – such as that IDR plans are most attractive to borrowers with larger student loan representation (Section 3.5) and low prospective income growth (Section 6.4) – are in line with the observations in Karamcheva, Perry, and Yannelis (2020).

As part of their analyses, Karamcheva, Perry, and Yannelis (2020) examine the introduction of new sub-programs under the system of IDR plans, which generously decrease the proportion of income borrowers commit as IDR payment. Using a discontinuity design, the authors find that the increase in plan generosity does not have significant impact on borrower earnings (and labor supply). Such an ex-post effect (or lack thereof) – i.e., effect conditioning on borrowers having entered IDR, is discussed in our model implications (Section 6.1). Relatedly, de Silva (2023) shows that under Australia's income-contingent repayment schedule (ICR), borrowers reduce labor supply to the level that bunches below the designated repayment threshold – the point at which borrowers start making payments. As the threshold (and payment stringency) moves, the bunching in labor supply adjusts accordingly. Because the ICR schedule is the only available contract in Australia, this observed labor supply response similarly represents an ex-post effect (without the possibility of switching from non-ICR to ICR).

Our main contribution to this literature is twofold. First, we document that student loans can correct the debt overhang in labor supply to begin with. This result, to our knowledge, is unexplored in the prior literature. Second, we document that such a corrective effect may be attenuated by the presence of IDR, and this attenuation works through an *exante* effect – i.e., labor supply decisions prior to IDR enrollment. That is, in anticipation of more attractive IDR plans, borrowers may be significantly discouraged from supplying labor when they are still making payments under the standard amortization plans. This ex-ante effect complements the ex-post investigations in Karamcheva, Perry, and Yannelis (2020) and de Silva (2023). In fact, in part of our theoretical analyses, we show that while the ex-ante effect of IDR on labor supply is unambiguous (i.e., IDR reactivates debt overhang), the ex-post effect is less straightforward and parameter dependent (Lemma 1).

More broadly, our paper relates to the recent theoretical and empirical study documenting the debt overhang effect of household leverage. Using a labor-search model, Donaldson, Piacentino, and Thakor (2018) study labor supply decisions of indebted households. They show that a debt overhang problem makes households reluctant to work because they must use their wages to make debt repayments. This prediction is supported by Bernstein (2021), who empirically documents that underwater homes reduce household's incentive to supply labor. Similarly, Melzer (2017) shows that mortgage debt overhang reduces individuals' investment in home improvement.⁹ On the other hand, Zator (2020) shows that higher mortgage interest rates make household work harder in order to cover increased mortgage payments. Building a dynamic model, Manso, Rivera, Wang, and Xia (2023) find that labor supply exhibit a non-monotonic relation with household indebtedness; they contrast such a relation with that of households. Different from these studies, our paper focuses on one specific type of household debt – non-dischargeable student loans – to show their unique role in mitigating debt overhang, and how this role varies with repayment options.

⁹In addition, recent literature documents the negative effect of rising household debt in reducing labor income (Dobbie and Song 2015; Di Maggio, Kalda, and Yao 2019), labor mobility (Ferreira, Gyourko, and Tracy 2010, 2011; Bernstein and Struyven 2017; Brown and Matsa 2020; Gopalan, Hamilton, Kalda, and Sovich 2021), and inventors' innovation (Bernstein, McQuade, and Townsend 2021).

2 Model

2.1 Preferences and income process

We build a dynamic model to study households' optimal decision on labor supply in the presence of debt. We consider infinitely lived households that derive utility from consumption C_t , and disutility from supplying labor l_t . Households are assumed to be risk-averse. For tractability, we assume logarithmic consumption preferences and quadratic cost of supplying labor such that a household's per-period utility is given by:

$$u(C,l) = \log C - \theta \frac{l^2}{2}.$$
(1)

The household's life-time utility from consumption and labor supply streams, $\{C_t, l_t\}_{t \ge 0}$, is given by:

$$\mathbb{E}\left[\int_0^\infty e^{-\delta t} u(C_t, l_t) dt\right],\tag{2}$$

where $\delta > 0$ is the household's subjective discount rate.

Denote by $K_t \ge 0$ the household's hourly labor income (wage) per-period. The dynamics of K are given by the geometric Brownian motion (GBM) process:

$$dK_t = K_t(\mu dt + \sigma dB_t), \tag{3}$$

where B_t is a standard Brownian motion; $\sigma > 0$ is a proxy for labor income uncertainty, and μ captures labor income growth. It follows that the total wages derived by the household, W_t , are the product of hourly wage and the number of working hours (labor supply):

$$W_t = l_t K_t, \tag{4}$$

in which the household can freely supply labor at the prevailing hourly wage K_t .

2.2 Household saving, debt and student loans

Initially, households have complete access to credit markets and can freely borrow and save. The household savings S_t evolve according to:

$$dS_t = (r(S_t)S_t - C_t + W_t)dt \text{ if } t \le \tau_D,$$
(5)

$$S_t = 0 \text{ if } t > \tau_D, \tag{6}$$

where τ_D denotes the time at which the household reaches its credit limit, forcing it into default. We model the borrowing limit as a multiple <u>s</u> of the household's earnings upon default (see Appendix A.1). This modeling approach reflects an exogenous liquidity default akin to that in the dynamic corporate finance literature. See, e.g., Longstaff and Schwartz (1995) and Section V in Bolton, Chen, and Wang (2011). Savings can take either positive or negative values: positive savings denote deposits and negative savings indicate borrowing. We set the interest rate $r(S_t) = r_B$ when the household is borrowing (i.e., when $S_t < 0$) and $r(S_t) = r_S < r_B$ when the household is saving (i.e., when $S_t \ge 0$), reflecting the observation that interest rates for household savings are lower than those of household debt. Prior to default, Equation (5) states that total wages increase household savings, which are used to pay for household consumption.

The household has an exogenously given initial amount $S_0^{SL} < 0$ of student loans. This formulation reflects our expectation – supported by the empirical evidence (Section 5.4) – that among individuals attending colleges, the amount of student debt borrowing often results from plausibly exogenous factors, such as the tuition of local colleges available to a student. As such, our model and empirical analyses focus on how given student loan amounts affect households' labor supply decisions. We assume that these loans start out in the standard repayment plan (SRP) such that they accrue an interest rate $r_B > 0$. We then consider an alternative repayment schedule – the Income Driven Repayment (IDR) plans – which student borrowers can switch to from SRP.

The household's total borrowing (i.e., negative savings) is the sum of student debt and non-student debt: $S = S^{SL} + S^{NS} < 0$. For tractability reasons, we assume that the proportion of household borrowing corresponding to student loans Δ remains constant over time, so that:

$$S_t^{SL} = \Delta S_t, \qquad S_t^{NS} = (1 - \Delta)S_t. \tag{7}$$

Put differently, Equation (7) states that households pay down their student debt and nonstudent debt at the same proportion, so that their relative share is time-invariant.

2.3 Default

2.3.1 Default in practice

Before describing the modeling of household default, we provide practical details on the consequences of household default, and how they differ for student loans and other types of consumer debt.

Following default, creditors may use various means to recoup the value of the debt, and a common practice is wage garnishment. Wage garnishment occurs when an employer is obliged to withhold a proportion of household wages to repay outstanding debt. Private creditors are required to obtain a court order prior to the commencement of wage garnishment, and in certain states, private creditor garnishments are banned entirely. On the contrary, the federal government – creditor of federal student loans – can bypass the court order and begin administrative wage garnishment once a warning letter is sent to the borrower.¹⁰ Therefore, wage garnishment is more likely to occur when households default on student debt than on other debt. DeFusco, Enriquez, and Yellen (2024) find that the rise in student loan-related wage garnishment is the primary driver of the overall significant garnishments in the past few years.

After default, households that experience continuing financial distress may advance to bankruptcy proceedings. Unlike other forms of consumer debt, student loans are largely non-dischargeable during bankruptcy (Yannelis, 2020b). Borrowers are required to pay off these loans eventually, through garnished wages or forfeited social security, among other means. The non-dischargeability is thus akin to extending the wage garnishment imposed upon initial default – until the balance is paid off. By contrast, the wage garnishment ceases for other non-student debt once it is discharged during bankruptcy. In line with this distinction, DeFusco, Enriquez, and Yellen (2024) show that wage garnishment for student

 $^{^{10}\}mathrm{Federal}$ student loans constitute over 90% of the student debt market.

debt lasts longer than for non-student debt. Non-dischargeability adds a further penalty for households defaulting on student debt.

It is worth noting that upon bankruptcy filing, an *automatic stay* goes into effect, which prohibits almost all creditors from collecting the debts – including student loan lenders. Although this injunction temporarily pauses wage garnishment for student debt, it does not alter its non-dischargeability nature and thus, the fact that student debt is more punitive than other debt.¹¹

To model this punitive nature concisely, we assume that upon student loan default, wage garnishment starts immediately and continues until the debt is paid off, whereas wage garnishment for other types of debt ends with bankruptcy. In either case, we abstract from the possibility of automatic stay during bankruptcy. This treatment is without loss of generality and captures the greater penalty associated with student loans due to both the prompt wage garnishment and its non-dischargeability. Such greater penalty, as we document below, encourages households to work more diligently *ex ante*, in order to avoid entering default in the first place.¹² Importantly, we note that the fraction of income garnished after default is increasing in the student loan balance and this property plays a critical role in our ex-ante versus ex-post results as discussed in Section 3.6.

2.3.2 Modeling default

In our model, the household's student debt and non-student debt follow different trajectories in the case of default. The non-student debt is discharged and the household is hence shunned from credit markets, forcing its savings (and debts) to be zero as shown in Equation (6). This equation reflects that (i) a majority of households going through bankruptcy file Chapter 7 in the U.S. – in which case debtors discharge eligible (non-student) debts and (ii) default

¹¹The automatic stay last as long as the bankruptcy proceeding continues, and thus varies with the type of bankruptcy, Chapter 7 versus Chapter 13. According to Dobbie and Song (2015), about 80% of U.S households file for Chapter 7, which typically last around four months. After the conclusion of Chapter 7, eligible non-student debt is discharged (in exchange of borrowers' assets liquidation) – whereas the debt collection for student loan, including wage garnishment, resumes. The less common filing, Chapter 13, can last for several years. During this period, student loans are treated as other unsecured non-priority debt, and borrowers make payments according to the terms of plan they have made. Federal student loans remain non-dischargeable in this process.

 $^{^{12}}$ This ex-ante effect – the focus of our study – may differ from an ex-post effect after default occurs. As wage garnishment essentially imposes a tax on labor, households may instead opt to reduce the labor supply after default. See Lemma 1 in Section 3.4 for related discussions.

often damages debtors' credit worthiness, thereby limiting their borrowing capacity (e.g., Dobbie and Song (2015); Dobbie, Goldsmith-Pinkham, Mahoney, and Song (2020); Kleiner, Stoffman, and Yonker (2021)).¹³ In such a case, the household can no longer rely on credit markets to smooth its consumption. Instead, it will be consuming wages net of the amount required to make student loan payment.¹⁴

While payment to non-student debt terminates at this point, student debt requires continued repayments that match the outstanding amount as explained following Equation (9).

The value function and optimal policies after default can be computed in closed-form solutions:

$$H(K) = \frac{1}{\delta} \log K - \frac{\log(\theta)}{2\delta},\tag{8}$$

$$C(K) = \left(1 - \Delta |\underline{s}| \sqrt{\theta} (r - \mu)\right) K l(K), \qquad l(K) = \frac{1}{\sqrt{\theta}}.$$
(9)

Equation (9) follows a straightforward intuition. Because the household can no longer borrow or save following default, it becomes a hand-to-mouth household, whose consumption equals wages net of student loan repayment. The fraction of wages garnished to pay for student loans, $1 > \Delta |\underline{s}| \sqrt{\theta} (r_B - \mu) \ge 0$, is set to match the market value of the student loans at default – reflecting the non-dischargeability nature of student debt.¹⁵ Moreover, the optimal labor supply is inversely proportional to the square root of the parameter governing the cost of supplying labor, θ .

$$\mathbb{E}\left[\int_0^\infty e^{-r_B t} \left(\Delta |\underline{s}| \sqrt{\theta} (r_B - \mu)\right) K_t l(K_t) dt\right] = \Delta K_{\tau_D} |\underline{s}| = S_{\tau_D}^{SL}.$$
(10)

 $^{^{13}}$ Dobbie and Song (2015) report that 98.4% of Chapter 7 filings end with a discharge of debt. Almost all unsecured debts are eligible for discharge.

¹⁴Our main findings is robust to considering more borrower-friendly default outcomes, in which the household, for example, may still partially access credit markets following default.

¹⁵That is, for given household policies of labor supply l post default, the following equality holds:

2.4 Household's optimization problem

The household's problem consists of jointly choosing consumption and labor supply. We denote the household's value function as:

$$G(S,K) = \max_{C,l} \mathbb{E}\left[\int_0^{\tau_D} e^{-\delta t} u(C_t, l_t) dt + e^{-\delta \tau_D} H(K_{\tau_D})\right].$$
(11)

The first part of Equation (11) pertains to the value prior to default. It is a function of the consumption and labor supply decisions. The second part, H(K), is the value post default. In the main model, we assume that after default, the household's hourly wages remain intact even though it can no longer rely on credit markets to smooth its consumption. This assumption matches empirical findings by Dobbie, Goldsmith-Pinkham, Mahoney, and Song (2020), who show that personal bankruptcy information has an economically trivial impact on future earnings in the U.S. labor market.¹⁶ In an extension in Appendix A.6, we consider the case when household wages decline after default due to, e.g., more limited employment opportunities. Our main results remain.

The value function G(S, K) satisfies the dynamic programming equation:

$$\delta G(S,K) = \max_{C,l} \left\{ \log C - \theta \frac{l^2}{2} + G_S(S,K)(r(S)S - C + lK) + G_K(S,K)\mu K + \frac{1}{2}G_{KK}(S,K)K^2\sigma^2 \right\}.$$
(12)

The first two terms inside the brackets represent the household's instantaneous utility from consumption and labor supply. The third term captures the change in value for the household from changes in savings. The fourth and fifth terms are the change in value induced by the dynamics of hourly wages K.

Next, we need to pin down the boundary conditions. The value function satisfies (i) the value matching conditions at personal bankruptcy:

$$G(S, \underline{s}K) = H(K), \tag{13}$$

 $^{^{16}}$ As the authors explain, this finding is likely because bankruptcy contains little incremental value in predicting individuals' future job performance. However, the authors find modest effects of bankruptcy information on job-finding rates, consistent with Friedberg, Hynes, and Pattison (2022).

and (ii) a transversality condition specified in Appendix A.2. The value function and the resultant household optimal policies are fully characterized by the solution of Equation (12) subject to these two boundary conditions. The differential equation cannot be solved analytically. However, due to CRRA preferences and GBM dynamics for the hourly wage, the value function displays homogeneity of degree one. Hence, in Appendix A.1, we show how the two state variables K and S can be reduced to a single state variable $d_t \in [0, 1]$.

Importantly, as shown in Figure I, d_t predicts the probability of a household entering default. Therefore, we interpret the state variable d as household indebtedness. More precisely, we define z as the probability that the household eventually becomes default given its current level of indebtedness:

$$z(d) = \mathbb{P}(\tau_D < \infty | d_t = d).$$
(14)

 $z(\cdot)$ is increasing in the level of indebtedness, implying that the household is more likely to default when d is higher. As discussed in Section 5, our empirical counterpart for the state variable d will similarly predict household default, resembling the pattern of Figure I.

In the sequel, we characterize labor supply as a function of household indebtedness $l(d_t)$. That is,

$$l(d_t) = l(d(S_t, K_t)) = l(S_t, K_t).$$
(15)

2.5 Debt overhang and student loans

In this section, we characterize the labor supply policy in relation to household debt, and how this relation varies with the presence of student loans.

Figure II depicts the household's labor supply as a function of indebtedness l(d), for three different values of Δ – the ratio of student loans to total debt. The baseline parameter values used in the figure are discussed in Section 4 (and listed in Table 1). We normalize the indebtedness to be 1 at default and 0 when the household pays all its debts (i.e., when savings become non-negative). The solid line depicts the baseline case, in which the household does not have student loans on the balance sheet (i.e., $\Delta = 0$). Here we observe that household labor supply exhibits a hump-shaped relation with respect to indebtedness, consistent with



Figure I: Household default probability as a function of the state variable d. See Table 1 for baseline parameter values.

Manso, Rivera, Wang, and Xia (2023). The intuition is as follows. As household indebtedness initially increases, a larger share of income accrues to creditors (via debt repayment), which lowers the overall level of consumption and pushes up the marginal utility of an additional unit of consumption. As a result, the benefit of increasing effort to raise consumption is high, rendering labor supply an increasing function of indebtedness. This effect arises directly from the conventional diminishing marginal utility of consumption implied by risk-aversion.

However, when the indebtedness surpasses a certain threshold, the household optimally reduces its supply of labor due to debt overhang. At this point, the household is closer to default (and discharging its debts in bankruptcy), as shown by Figure I. Therefore, it no longer internalizes the full benefits of labor because any labor income will be partly used to repay the soon-to-be discharged debt, creating a wealth transfer from the household to creditors. Given that it bears the full cost of supplying labor while sharing the benefits with creditors, the household reduces effort. In this case, the option to default (and discharge debt) affords the household an opportunity to shirk at the expense of lenders.

Such opportunistic behavior, however, is curbed by the presence of student loans. Because the household has to repay student debt regardless of bankruptcy, any additional dollar generated through labor supply will entirely accrue to the household, instead of becoming a wealth transfer to lenders. As such, student loans mute the shirking opportunity for the



Figure II: Household debt overhang for different proportions of student loans. Labor supply as a function of household leverage l(d) is depicted for three values of Δ , the proportion of student loans among household debt. See Table 1 for baseline parameter values.

household, correcting its disincentive to supply labor.

The dashed and dotted lines demonstrate this *corrective* effect of student loans. The two lines depict labor supply when the proportion of student loans among total household debt (i.e., Δ) are 0.2 and 0.4, respectively. As household liabilities feature a higher proportion of student loans, labor supply becomes more resilient to debt overhang. Such resilience is shown in two manners. First, the decline in labor supply does not kick in until a higher level of indebtedness – a "delayed" manifestation of debt overhang. Second, the decline in labor supply is relatively flatter when student debt has the larger presence – a "dampened" manifestation of debt overhang. For instance, by the time the household indebtedness reaches the highest level, the drop of the dotted line from its peak – relative to the prior run-up (when indebtedness ranges between 0 and 0.3) – is less aggressive than the case of the solid line. Based on these observations, we state the first prediction of our model, which we later bring to the data:

Prediction 1. A larger proportion of student loans among household debt alleviates debt overhang in a household's labor supply.

We conclude this baseline analysis by formalizing the solution to the household's optimization problem, which allows us to analytically characterize the effect of student loans on the household's value function and its optimal labor supply (and consumption) policies.

Proposition 1. Suppose there exists a smooth solution for the Hamilton-Jacobi-Bellman (HJB) Equations (12) and (13), then this solution corresponds to the value function of the household optimization problem (11). In addition:

- The value function of the household is decreasing in the fraction of student loans Δ . That is, $\frac{\partial G(\cdot;\Delta)}{\partial \Delta} \leq 0$.
- The household's labor supply is increasing in the fraction of student loans Δ . That is, $\frac{\partial l(\cdot;\Delta)}{\partial \Delta} \geq 0.$
- The household's consumption is decreasing in the fraction of student loans Δ . That is, $\frac{\partial C(\cdot;\Delta)}{\partial \Delta} \leq 0.$

The intuition for this proposition is as follows. Compared to its counterpart with little student debt, a household with more student loans obtains the same payoff prior to default, but a lower value upon default since student loans cannot be discharged. Such a "non-dischargability penalty" renders its value function decreasing in Δ . Because default is more punitive with student loans, the household endeavors to avoid default by increasing its labor supply. This result is in line with the illustration of Prediction 1 depicted in Figure II, indicating the student loans' *corrective* effect in alleviating debt overhang.

Moreover, although not the focus of our study, the last point of Proposition 1 suggests an alternative venue for the household to avoid punitive default when facing more student loans: reducing consumption. As described in Section 5, our data do not contain clear information on household consumption, and value functions are unobservable by definition. Therefore, we empirically test Proposition 1 by focusing on households' labor supply activities and the subtle features contained in Prediction 1, such as the delayed and dampened manifestation of debt overhang.

3 Debt overhang and student loans under IDR

3.1 Institutional background

The federal government offers assistance programs to students unable to maintain the 10year amortization schedule to repay federal loans under SRP. One such program is the Income Driven Repayment (IDR) plans. IDR plans set payments as a proportion of student borrowers' discretionary income, and eventually forgive principal.¹⁷

There are four sub-programs under the IDR plans: Pay As You Earn Repayment Plan (PAYE), Revised Pay As You Earn Repayment Plan (REPAYE), Income-Based Repayment Plan (IBR), and Income-Contingent Repayment Plan (ICR). Depending on the programs, the percentage of income required as payment ranges from 10% to 20%. In case students' discretionary income is low enough, the federal government covers accrued interest that is not fulfilled by the IDR payment.¹⁸ After students make qualified payments for 20-25 years, IDR plans forgive the remaining principal, regardless of the amount of loans reimbursed by the cumulative IDR payments. The vast majority of federal student loans qualify for an IDR plan, as long as they are not in default.

In practice, while IDR plans provide insurance against unaffordable payment, the enrollment and maintenance process incurs significant non-monetary costs for student borrowers. First, the IDR application requires the completion of a lengthy 12-page form collecting extensive personal and financial information. Mueller and Yannelis (2022) show that this hurdle deters borrowers' IDR enrollment. Using a field experiment in which student borrowers received pre-populated IDR applications, the authors document a 34 percentage points increase in the enrollment.¹⁹

Second, to retain IDR eligibility, borrowers must re-certify their income annually and this process proves to be burdensome. Herbst (2023) shows that most borrowers fail to complete the re-certification and thus return to their pre-IDR repayment pattern after one

¹⁷Discretionary income is defined as gross income in excess of 150% of the applicable poverty line. See https://studentaid.gov/manage-loans/repayment/plans/income-driven

¹⁸The extent of this benefit varies across the sub-programs and the length of time in the IDR plan. See Cornaggia and Xia (2024).

¹⁹Other studies find that low IDR take-up rates can be attributable to the framing of IDR descriptions or parental support (e.g., Abraham, Filiz-Ozbay, Ozbay, and Turner 2020; Lochner, Stinebrickner, and Suleymanoglu 2021).

year of initial IDR enrollment. As a result, IDR's financial benefits (through reduced loan payment) only provide short-term incremental liquidity infusion to distressed borrowers.

Third, these procedural costs are compounded by the fact that borrowers are often in need of dealing with student loan servicers during IDR enrollment. The Department of Education (DoE) has reported that many borrowers lack awareness of IDR options or how they operate (U.S. Government Accountability Office (2015)).²⁰ Loan servicers are expected to provide assistance to borrowers and guide them through the IDR process. However, both anecdotal and academic evidence shows that loan servicers fall short of such expectations and are subject to conflicts of interest resulting in sub-optimal outcomes for students (Dynarski 2014; Cornaggia and Xia 2024). Dealing with servicers thus further raises reluctance from borrowers in taking up IDR.²¹

In our model below, we explicitly account for these costs and analyze how they affect the household's trade-off when deciding on IDR enrollment.

3.2 The household's value function while in IDR

As an important feature, the proportion of household income required as IDR monthly payment is formulaic; it does not vary with how much the household owes prior to IDR enrollment. By adopting IDR plans, it is as if the household would "discharge" the student debt in exchange for issuing an equity-type of claim to the creditor, in which it pledges a fraction of future labor income. In this section, we extend our baseline analysis provided in Section 2 by modeling such pseudo "dischargeability", and examining how it affects household labor supply and debt overhang.

Specifically, we incorporate the key institutional features of IDR discussed in Section 3.1 as follows. First, we display the pseudo "dischargeability" feature by subtracting S^{SL} (the initial student loan amount) from the household's total borrowing S once it enrolls in IDR – as if allowing the student loans to be removed from the household balance sheet.

²⁰Existing studies point out that the complexity of federal student aid program interferes borrowers' assessment and may obfuscate important benefits afforded by repayment plans such as IDR (e.g., Dynarski and Scott-Clayton 2006; Cox, Kreisman, and Dynarski 2020).

²¹Student loan servicers are designated by the DoE. The largest four are Great Lakes, Navient, FedLoan Servicing (a.k.a Pennsylvania Higher Education Assistance Agency), and Nelnet. See Cornaggia and Xia (2024) for the market shares of these servicers.

In exchange, a fraction ψ of the household's future wages is collected as IDR payment – representing the equity-type of claim issued to creditors. Thus, the household's net labor income after enrolling in IDR becomes $(1 - \psi)l_t K_t$.

It is worth noting that among the three most adopted IDR sub-programs, IBR and PAYE set a cap to the payment such that students pay no more than what they would have paid under the standard repayment plan, while REPAYE does not impose the cap. As of 2019, REPAYE accounts for approximately one third of all IDR enrollment.²² Relatedly, student borrowers may choose to transition out of IDR, effectively capping their payment to that set by the standard repayment plan. In our baseline model, we follow Abraham, Filiz-Ozbay, Ozbay, and Turner (2018) and abstract from such "payment cap" feature and assume that the IDR is irreversible. We do so for the tractability of the model. In Appendix B.1, we discuss that this treatment should not affects our main findings. There, we show through simulation that IDR is a highly persistent state – consistent with the finding in Boutros, Clara, and Gomes (2024) – making IDR unlikely to be reversed and the "repayment cap" unlikely to bind.

Second, to capture the non-monetary costs associated with IDR, we introduce a onetime utility cost u_I , corresponding to the up-front cost associated with the IDR application process, as well as the (present value of) ongoing cost of annual re-certification and the hassle of dealing with loan servicers.

Third, to model the fact that any remaining balance of student debt is forgiven after a period of continuing IDR payment (typically 20 years), we introduce the parameter $1/\lambda$, which denotes the number of expected waiting periods until the forgiveness is granted. The inverse, λ , captures the expected probability of student debt being forgiven in each period. This modeling device follows Leland (1998): it allows us to parsimoniously model stochastic debt forgiveness (rather than the less tractable deterministic debt) without loss of generality.²³

Lastly, we denote by τ_I the (endogenous) point in time at which the household chooses to switch from the standard repayment plans (SRP) to IDR.

With these elements in place, we solve the household's optimization problem recursively,

²²See more information here, https://studentaid.gov/data-center/student/portfolio#servicerportfolio-by-repayment-plan.

²³For example, a waiting period of 20 years before the designated forgiveness corresponds to $\lambda = 0.05$.

starting from the situation when the household is already in IDR, i.e., when $t > \tau_I$. The household's value function $\hat{F}(S, K)$ after exercising the option of IDR enrollment satisfies the dynamic programming equation:

$$\delta \hat{F}(S,K) = \max_{C,l} \left\{ \log C - \theta \frac{l^2}{2} + \hat{F}_S(S,K)(r(S)S - C + (1-\psi)lK) + \hat{F}_K(S,K)\mu K + \frac{1}{2}\hat{F}_{KK}(S,K)K^2\sigma^2 + \lambda(\hat{F}(S,K;\psi=0) - \hat{F}(S,K)) \right\}.$$
(16)

This value function resembles Equation (12), with two differences. From the third term of Equation (16), we see that a fraction ψ of the household income is taken away to fulfill the IDR payment. The last term of Equation (16) captures stochastic debt forgiveness. This equation is solved subject to a boundary condition similar to Equation (13).²⁴

As before, in Appendix A.1 we show that it is possible to numerically characterize this value function and optimal policies as a function of the state variable d_t , household indebtedness.

3.3 Household's value function prior to IDR

We now compute the value function for the household before it exercises the option of enrolling in IDR – that is, while it is still in the standard payment plan. As discussed in Section 3.2, enrolling in IDR offers the household the possibility of swapping its student debt (whose nominal value is ΔS_t) for an equity-like claim, entitling the creditors to a fraction ψ of household wages for the next $1/\lambda$ years. The value function before IDR, denoted by F(S, K), solves the household's optimization problem:

$$\max_{C,l,\tau_I} \mathbb{E}\left[\int_0^{\tau_D \wedge \tau_I} e^{-\delta t} u(C_t, l_t) dt + \mathbb{1}_{\tau_I \le \tau_D} e^{-\delta \tau_I} \hat{F}(S_{\tau_I}, K_{\tau_I}) + \mathbb{1}_{\tau_D < \tau_I} e^{-\delta \tau_D} H(K_{\tau_D})\right], \quad (17)$$

where the household in addition to choosing its optimal consumption and labor policies, now also chooses the optimal time to file for IDR, denoted by τ_I . This stochastic control problem can be solved by combining real-option techniques with standard dynamic programming. In

²⁴Such a boundary condition covers the case when the household defaults while in IDR.

the continuation region, it satisfies the dynamic programming equation:

$$\delta F(S,K) = \max_{C,l} \left\{ \log C - \theta \frac{l^2}{2} + F_S(S,K)(r(S)S - C + lK) + F_K(S,K)\mu K + \frac{1}{2}F_{KK}(S,K)K^2\sigma^2 \right\},$$
(18)

where the first order conditions are given by:

$$C(S,K) = \frac{1}{F_S(S,K)}, \quad \text{and} \quad l(S,K) = \frac{KF_S(S,K)}{\theta}.$$
(19)

Because the household can now optimally choose when to enroll in IDR, we need to solve for the endogenous exercise boundary denoted by $\hat{S}(K)$. Intuitively, this exercise boundary trades off the benefit of IDR – having student loans "discharged", against the costs of IDR – including (i) the utility cost associated with the enrollment and recertification process, u_I , and (ii) the cost of giving up the option of waiting, in which case a positive wage shock (σ) may significantly increase the household's income and thus, allows it to continue with the standard payment (without needing to pay u_I and commit the fraction ψ of future wages upton entering IDR).

The exercise boundary can be pinned down by the following conditions:

$$F(\hat{S}(K), K) = \hat{F}((1-\Delta)\hat{S}(K), K) - u_I \quad \text{and} \quad F_S(\hat{S}(K), K) = \hat{F}_S((1-\Delta)\hat{S}(K), K)(1-\Delta).$$
(20)

Here we use a "hat" symbol to denote quantities when the household is already in IDR. The condition on the left corresponds to the value matching condition, stating that the value to the household before and after exercising the IDR option is the same net of the utility cost u_I . The condition on the right is the smooth pasting condition stating that the household no longer has the incentive to wait to exercise the IDR option.

3.4 Student loans and debt overhang: revisited

Solving for the household's optimal policies following Sections 3.2 and 3.3, we plot labor supply as a function of the state variable, household indebtedness (d). Figure III considers three cases. The solid line represents the case in which student debt can only be paid under



Figure III: Labor supply in the absence of IDR (solid blue), before entering IDR (dashed red), and after entering IDR (dotted red). See Table 1 for baseline parameter values.

the SRP without the option of IDR. The dashed line represents household labor supply when IDR is available but before the option is exercised – i.e., when the household is still making payments under SRP. The dotted line represents labor supply after the household enrolls in IDR. As a first step, we take as given the proportion of student loans among household debt $(\Delta = 0.53)$ for all three cases. In Section 3.5, we vary this proportion to study the nuances of our results.

The solid line reiterates our baseline result in Section 2.5: it shows that with the IDR option, labor supply exhibits a hump shape with respect to household indebtedness, arising from the interplay of the diminishing marginal utility of consumption (when indebtedness is initially low) and debt overhang (when indebtedness surpasses a threshold).

The debt overhang problem is aggravated by the IDR option, as illustrated by the dashed line. Anticipating the opportunity to pseudo "discharge" student debt upon IDR enrollment, the household becomes discouraged from supplying labor ex ante, in order to avoid wealth transfers onto creditors. Such disincentive (partially) undoes the *corrective* effect of student loans as previously documented, uncovering a *setback* effect of IDR. Indeed, as shown by the dashed line, labor supply begins to drop (reflecting the onset of debt overhang) earlier than the solid line (around 0.3 versus 0.5 of indebtedness). The dashed line stops at around 0.6 of indebtedness. This is the point when the household chooses to exercise the option and switch to IDR – corresponding to the time point τ_I in the model.

The dotted red line takes over from this point. Two observations are worth noting. First, once in IDR, there is an immediate jump up in labor supply. This jump happens because, at this point, the household's student debt has been pseudo "discharged" – resulting in the *effective* indebtedness being lower than the *nominal* indebtedness. Such de-leveraging in turn revives part of the household's incentive to supply labor. Intuitively, the extent of this de-leveraging effect depends on how much student debt the household has borrowed and thus, how much has been pseudo "discharged" through IDR (i.e., Δ).

Lemma 1. The household's labor supply jumps up upon enrolling in IDR if and only if $\Delta > \psi$. Moreover, the magnitude of the jump is increasing in Δ and decreasing in ψ .

The proof is in Appendix A.4. Lemma 1 follows a straightforward intuition. The deleveraging effect due to IDR's pseudo "discharge" is countered by an opposite force. Upon entering IDR, the household pledges a fraction of future income as student loan payment, and such heightened "labor income taxes" lower the household's incentive to work, i.e., the well-known taxation effect. How much labor supply can be regained depends on how much student debt is "discharged" (Δ) versus how much additional "tax" is imposed (ψ). When the de-leveraging effect dominates the taxation effect, the household chooses to supply more labor after entering IDR than before. Otherwise, the household reduces labor supply even further. Accordingly, the larger the de-leveraging effect (Δ) is, the greater the regained labor supply.

Lemma 1 differentiates an *ex-ante* effect of IDR on labor supply (before the household enrolls in IDR) from an ex-post effect. It suggests that while the *ex-ante* effect is unambiguous – that is, IDR reactivates debt overhang and discourages labor supply – the *ex-post* effect is less clear. Depending on parameters, the ex-post labor supply can be either higher or lower than the case without IDR. In Section 6, we revisit these ex-ante versus ex-post effects, when we derive policy implications of our model in relation to the existing literature.

Another observation from Figure III is that the dotted line continues to exhibit a hump shape. This is because even after student debt is "discharged" by IDR, the household still holds other forms of debt (e.g., mortgages and credit cards) on the balance sheet. Therefore, debt overhang is still present. However, the decline of labor supply is less aggressive compared



Figure IV: Household debt overhang and IDR take-up for different values of Δ . Labor l(d) is depicted for two different values of Δ , the proportion of student loans in household debt. See Table 1 for baseline parameter values.

to that of the solid line, due to the de-leveraging effect and the lowered *effective* indebtedness (given the same level of the *nominal* indebtedness).

Overall, Figure III highlights the *setback* effect of IDR – which offers the household a pseudo opportunity to discharge student debt and create an incentive distortion *ex ante*. In the next section, we explore how this *setback* effect varies when student debt takes up different proportions in household debt.

3.5 The setback effect of IDR and student debt proportion

As discussed in Section 3.3, the household's decision to enroll in IDR faces a trade-off between the benefit of "discharging" student loans and the costs associated with entering IDR (i.e., the utility cost and the cost of giving up the option of waiting). The household is more willing to incur such costs if loan dischargability is perceived as worthwhile. This happens when the household has more student loans to begin with, in which case getting rid of the unaffordable payment under SRP is more valuable. As such, the proportion of student loans among household debt shapes the household's incentive to take up IDR and accordingly, the extent of IDR *setback* effect.

Figure IV explores this intuition by depicting household labor supply with respect to

indebtedness for different proportions of student loans among total debt (Δ). Panel A restates our baseline results without the IDR option (Section 2.5) and is presented here for ease of comparison. The blue (resp. red) line depicts the labor supply for a high (resp. low) value of Δ . A larger proportion of student loans ($\Delta = 0.50$ vs. $\Delta = 0.45$) alleviates debt overhang by both *delaying* and *dampening* the decline in labor supply as household indebtedness rises. This is the *corrective* effect of student loans.

This pattern, however, is reversed in Panel B. The blue (resp. red) line again depicts the labor supply for a high (resp. low) value of Δ . The solid part corresponds to the labor supply before IDR enrollment, and the dotted part corresponds to the labor supply after. Contrary to Panel A, the solid blue line now exhibits more severe debt overhang than the red line, and the intuition is as follows. A larger proportion of student loans among total debt makes IDR more attractive because the pseudo "dischargability" can remove a greater burden of the loan payment under SRP. As a result, the household switches into IDR at an earlier stage – and this can be seen from the earlier stopping point of the blue solid line than the red solid line. Anticipating the quicker IDR adoption, debt overhang is reactivated earlier, giving rise to the earlier manifestation of debt overhang, that is, a greater *setback* effect.

Consistent with Figure III, labor supply jumps higher upon IDR enrollment for both values of Δ , as shown by the dotted lines. This jump reflects the revived labor supply incentives after "discharging" student loans – i.e., the de-leveraging effect. As expected, the revival of labor supply is greater when Δ is larger (the blue dotted line) due to a stronger de-leveraging effect (see Lemma 1). Based on Figure IV, we state the second prediction of our model as follows.

Prediction 2. Given the proportion of student loans among total debt is sufficiently high that a household considers filing for IDR, a higher proportion of student loans aggravates debt overhang in labor supply ex ante (i.e., a stronger setback effect).

Prediction 1 and Prediction 2 show that there is an asymmetric impact of student loans on labor supply incentives with and without IDR. Interestingly, this asymmetry also translates into the household's value function. The following lemma shows that, in contrast to Proposition 1 - in which the household is uniformly worse-off if it has a larger share of student loans – when IDR is available and upon filing for IDR, the household becomes better-off if it has more student loans.

Lemma 2. The household's value function at the time of filing for IDR (τ_I) is increasing in the fraction of student loans Δ . That is, $\frac{\partial F(\hat{S}(K;\Delta),K;\Delta)}{\partial \Delta} > 0$.

Intuitively, without the IDR option, more student loans always make the household worse off because they have to be necessarily repaid. When IDR is available – and when the option to pseudo "discharge" student loans is present, more student loans make this option more valuable, thus rendering the household's value function increasing in Δ .

3.6 IDR versus wage garnishment: a paradox

An important distinction regarding default versus entering IDR is worth discussing. In both instances, the households forfeit part of the income (as garnished wages or IDR monthly payments). Paradoxically, such garnishment in the former instance alleviates the debt overhang problem (the *corrective* effect). In contrast, in the latter instance it compounds the debt overhang problem (the *setback* effect). These opposite predictions pertain to household labor supply ex ante – i.e., before wage garnishment or IDR enrollment.

This paradox can be resolved as follows. In the case of default, the fraction of income garnished is increasing in the student loan balance upon default. Thus, supplying labor and generating income ex ante to lower the outstanding student loan balance is in the interest of the household.

By contrast, in the case of IDR, the fraction of monthly income garnished under incomedriven repayment programs is set formulaically and thus, *invariant* to the outstanding balance. Because borrowers are not better off entering IDR with a lower balance, they do not have incentives to generate income ex ante to lower student debt.²⁵

Not only do wage garnishment and IDR enrollment shape labor supply differently ex ante, they may also yield opposite implications ex post. After default occurs, wage garnishment effectively imposes a tax on labor, thereby discouraging households from continuing to supply

 $^{^{25}}$ Before 2021, the balance forgiven after the 20-25 year IDR payments was taxable. Hence, there was in practice a small benefit to entering IDR with a lower balance. However, starting in 2021 loan forgiveness became tax-exempt. See Appendix B.2 for a discussion of how these tax considerations may affect our results.

labor.²⁶ After IDR enrollment, although the committed monthly income as IDR payment similarly imposes a labor tax, a counteracting force is at play – the de-leveraging effect (due to IDR "pseudo" discharging student debt). This is the effect illustrated in Lemma 1, and there we show this counteracting force can encourage households' labor supply ex post.

The ex-post labor supply is not the focus of our study. Nevertheless, together with the ex-ante effect, it highlights different mechanisms underlying the paradox regarding wage garnishment and IDR enrollment – and thus the contribution of our model in delineating these differential effects.

4 Calibration

Prediction 2 provides a counter force to Prediction 1 in terms of how the presence of student loans may affect debt overhang. In this section, we assess which prediction is more likely to prevail in practice – guiding the subsequent empirical analyses. We start by describing the parameter choices for the calibration (and for our numerical solutions so far) in Section 4.1. In Section 4.2, we describe the calibration methods to match the observed IDR take-up rate in the existing literature.

4.1 Parameters

Our model's baseline parameters are summarized in Table 1 and described as follows. The subjective discount rate is set to $\delta = 0.03$, in line with the long-term subjective discount rate estimate in e.g., Laibson, Repetto, and Tobacman (2007). The savings interest rate (r_S) is based on the average federal funds rate during our empirical sample period (Section 5.2) - 0.015.²⁷ The borrowing interest rate $r_B = 0.08$ is a proxy for the average bank prime (mortgage) lending rate and credit card rate during the sample period. $\mu = 0.025$ is approximately the annual growth rate of hourly wage based on the private sector during the sample period, reported by the Bureau of Labor Statistics.²⁸ $\sigma = 0.35$ is an estimate based on the recent literature examining household earnings volatility (e.g., Shin and Solon

 $^{^{26}}$ Consistent with this intuition, Dobbie and Song (2015) find that employment and income increase after bankruptcy filing pauses wage garnishment.

²⁷https://www.newyorkfed.org/markets/reference-rates/effr.

²⁸https://beta.bls.gov/dataViewer/view/timeseries/CES0500000003.

2011; Dyna, Elmendorf, and Sichel 2013; Moffitt and Zhang 2018; Moffitt 2021).²⁹ We set the borrowing limit $\underline{s} = 7.6$ to match the 90th percentile of our sample distribution of debtto-income ratio when individuals file for bankruptcy. The cost of supplying labor $\theta = 2.5$ is set to ensure that the labor supply at the median value of $\Delta = 0.53$ fluctuates between 1,600 and 2,000 hours a year, in line with the median and the average labor supply in our sample. Lastly, the parameters $\psi = 0.15$ and $\lambda = 0.5$ are set directly to match the features of the standard IDR repayment programs (see Section 3.1).

The remaining parameter – utility cost of filing for IDR u_I – is unobservable in the data. As described in the next section, we calibrate this parameter to match the empirically observed IDR take-up rate, which in turn governs the prevalence of Prediction 1 and Prediction 2 in practice.

4.2 IDR take-up rate calibration

To set the context for our calibration of the IDR take-up rate, consider two households A and B. Household A does not borrow student loans ($\Delta = 0$), whereas Household B holds student loans among total debt ($\Delta > 0$). The *corrective* effect of student loans – pertaining to Prediction 1 – predicts that Household B is more resilient to debt overhang and thus, has stronger incentives to supply labor than Household A. On the other hand, the student debt held by Household B makes entering IDR, and thus pseudo "discharging" this debt, more appealing. Therefore, Household B may stop supplying labor sooner and become less resilient to debt overhang – the *setback* effect of IDR pertaining to Prediction 2. It is thus an empirical question whether the distortionary effect of household indebtedness on labor supply is mitigated or exacerbated for Household B, relative to Household A.

The key to unlocking this question lies in whether, and when, Household B exercises the IDR option. If this option is not considered by Household B at all – as if the option is non-existent, or if the option is not exercised by Household B until a very late stage – such that the disincentive of labor supply due to pseudo "discharging" plays a trivial role, then Prediction 1 should prevail in aggregate.

Whether and when Household B exercises the option in turn hinges on the trade-off

 $^{^{29}\}mathrm{We}$ confirm this estimate using our sample. See Section 5.1 for the data description.

governing IDR enrollment (see Section 3.5). Inspired by the insights of existing studies, the unobservable parameter u_I (the disutility of filing for IDR) plays a role in determining this trade-off. Mueller and Yannelis (2022) find that such costs play a significant role in explaining IDR take up, and using a field experiment, they show that once these costs are alleviated, the enrollment rate more than doubles. By calibrating u_I to match the observed IDR take-up rate, we can infer the household's trade-off in whether and when to enter IDR, thus shedding light on the prevalence of Predictions 1 versus 2 in practice. Based on the Congressional Budget Office (2020) report, IDR enrollment centers around 17% between 2010 and 2017 (the majority of our sample period). Hence, we calibrate u_I to approximately match a 17% IDR take-up rate.

Specifically, we take the following steps. Because IDR enrollment is irrelevant among households without student loans, our calibration focuses on those with student loans ($\Delta > 0$). For each of these households, we observe the outstanding volume of student loans as a proportion among the household's total debt Δ .³⁰ Since Δ varies in practice, we focus on five values – equal to the mid-point of each quintile of the Δ distribution: $\Delta =$ $\{0.05, 0.23, 0.35, 0.52, 0.76\}$. These values capture representative positions spanning the Δ spectrum and allow us to streamline the calibration without loss of generality.

For low values of Δ (0.05, 0.23, and 0.35), households never find it optimal to enroll in IDR. Their proportion of student loans is too low to make it worth the IDR costs. In such cases, the households are not subject to the IDR *setback* effect pertaining to Prediction 2. By contrast, for high values of Δ (0.52 and 0.76), households have sufficient student loans such that getting them "discharged" is worthwhile the cost of IDR. For these two upper quintiles, we run simulations starting at the state variable (d) that maps to the observed default probability of households with student loans (based on Figure I), and compute the fraction of households that eventually enroll in IDR.³¹ This is the fraction of households whose student debt burden turns out high enough to render IDR enrollment optimal. Matching this fraction to the empirically observed IDR take-up rate allows us to back out the utility cost u_I .

With the estimated u_I in place (along with other observable parameters discussed in

 $^{^{30}\}mathrm{See}$ Section 5 for the data description of the NLSY97 survey.

 $^{^{31}}$ This default probability is approximately 19%, derived from households' filing of bankruptcy or late payments on rents and mortgages. See Section 5.1 for detailed descriptions.



Figure V: Household debt overhang for households with and without student loans. Labor l(d) is depicted for household without student loans (solid blue line) and for the average of households with student loans (dashed red line) based on model calibration. See Table 1 for baseline parameter values.

Section 4.1), we then proceed to compute the labor supply of households for each value of $\Delta = \{0.05, 0.23, 0.35, 0.52, 0.76\}$. Lastly, we average out the labor supply across the five Δ values – and this average represents the labor supply policy of Household B in the example above. Figure V compares the labor supply of Household B ($\Delta > 0$) to that of Household A ($\Delta = 0$). It shows that even though Predictions 1 and 2 in theory generate opposite predictions, using the calibrated u_I that matches the empirical IDR take-up rate, households with student loans are more encouraged to supply labor than their counterparts without them. This result is consistent with Prediction 1 in that student loans provide resilience to debt overhang. In Section 5, we show empirical evidence supporting our calibrated results from this section. Readers more interested in the policy implications may skip ahead to Section 6.

5 Empirical evidence

5.1 Data and measurements

In this section, we provide empirical analyses guided by the findings from Section 4. Specifically, we test whether households' labor supply responds to a given level of indebtedness differently – depending on the proportion of student loans among total debt. This test pertains to the *corrective* effect of student loans in Prediction 1, which as shown in Section 4, is the dominating force in practice over the *setback* effect of IDR. We do not directly examine the *setback* effect, because such a test necessitates information on whether and when a borrower enrolls in IDR, and this information is not provided by the data available to us.

We employ the 1997 National Longitudinal Survey of Youth (NLSY97) for empirical analyses. NLSY97 is a program run by the U.S. Bureau of Labor Statistics (BLS). It surveys a nationally representative sample of individuals born between 1980 and 1984 and living in the United States at the time of the initial survey. The first round of interviews was conducted in 1997, when participants were 12 to 16 years old. Follow-up interviews were conducted annually until 2011, and biennially since then. The survey data are publicly available from Round 1 (1997-98) through Round 19 (2019-20). We obtain each respondent's geographical location through the license control of the BLS, which we employ to construct state fixed effects and to examine location based variation in student loan borrowing (Section 5.4).

The NLSY97 sample was selected to represent the civilian, non-institutional population of the United States. During the initial survey in 1997, 8,984 respondents were interviewed. Men accounted for 51% and women accounted for 49%. The survey included 51.9% non-Black/non-Hispanic, 26% Black non-Hispanic, 21.2% Hispanic or Latino, and 0.9% mixed raced respondents. The detailed description of the sample distribution can be found on the BLS website (https://www.bls.gov/nls/nlsy97.htm).

Three sets of information in the NLSY97 allow us to test the model predictions, including household labor supply, student loans, and balance sheets. We discuss each of them in detail.

Labor supply: The NLSY97 collects extensive information on respondents' labor market behavior. This information includes week-by-week records of a respondent's labor force status and associated job(s), the total number of hours the respondent works in each week at any job, and the hourly wage of the ongoing job during a survey interview. It allows us to observe a respondent's labor force activities, including working hours and the periods when he/she is unemployed or out of the labor force. We measure labor supply during a year as the respondent's total working hours in this year. In robustness analyses (Appendix C), we use annual earnings as an alternative measure of labor supply (e.g., Zator 2020; Bernstein 2021), and confirm our results.

Household balance sheets: NLSY97 collects comprehensive balance sheet information at four points in time, when a respondent is 20, 25, 30, and 35 years old. On the asset side, NLSY97 surveys each respondent's estimated market value of vehicles, the amount of savings and various financial assets (e.g., stocks and bonds), and the market value of residential property. On the debt side, NLSY97 surveys the amount of auto loans, credit card loans, student loans, mortgage loans, and money owed to other individuals or entities. To create a respondent-year panel (matching the annual labor supply panel), we expand the assets and debt information collected at these four ages through a linear interpolation (for years within a respondent's 20-35 ages) and extrapolation (for years beyond).³²

As we describe later, our empirical analyses use categorical variables to capture households' indebtedness (as opposed to a continuous measure), following Gopalan, Hamilton, Kalda, and Sovich (2021). This approach helps mitigate potential noise introduced by the interpolation and extrapolation.

Proxying for the state variable d: Our model in Section 2 characterizes the household's optimal policies as a function of the state variable d (indebtedness). As previously shown in Figure I, this state variable captures the probability of a household entering default. Therefore, to find an empirical counterpart of the state variable, we need a measure analogously capturing this default probability. To this end, we use household leverage, defined as the ratio of total debt to total assets. This measure is inspired by Melzer (2017), who documents that the ratio of household mortgage to property value – arguably the largest components of

³²For instance, a respondent's interpolated asset at age 26 equals: $Asset_{26} = Asset_{25} + \frac{(Asset_{30} - Asset_{25})}{5}$. Similarly, the respondent's extrapolated asset at age 36 equals: $Asset_{36} = Asset_{35} + \frac{(Asset_{30} - Asset_{30})}{5}$.

debt and assets for U.S. households – are highly indicative of default likelihood. In Figure 1, we verify that such an indication continues to hold using the broader measure of total debt to assets ratio.³³ We present all empirical results following the *References* section.

Specifically, in Figure 1, we plot the average default probability for each 10-point leverage bin from 0% to 130%. Default is identified by whether an individual has filed for bankruptcy as of a certain year. Overall, default risk increases with leverage. The increase remains relatively flat initially, and accelerates once leverage surpasses 60-70% – a pattern resembling the relation between the state variable d and default likelihood, as shown in Figure I of the model section.³⁴

We note that the magnitude of observed default is smaller than that indicated by the state variable d. This difference likely arises because the sample of NLSY97 consists of a younger generation (born between 1980 and 1984) and thus, the occurrence of default is limited. To overcome this limitation, in Appendix C Figure A4, we re-plot the relation between household leverage and default using a different cohort surveyed by the BLS – NLSY79. NLSY79 tracks a sample of individuals from their teens up to 60 years old, allowing us to observe default activities over a longer horizon. In Figure A4, we similarly observe that leverage strongly predicts default, as captured by bankruptcy (Panel A). Additionally, in Panel B, we identify default by whether an individual has missed any mortgage or rent payments as of a certain year.³⁵ The magnitude of default likelihood becomes significantly higher (up to 60%) in this longer horizon and is closer to that depicted by the state variable d in Figure I.

Taken together, this evidence verifies that household leverage provides a reasonable empirical counterpart of the state variable.

 $^{^{33}}$ An alternative measure to capture household default probability is the debt-to-income ratio, defined as the monthly total debt payment to net income – similar to the interest coverage ratio in corporate finance. This measure, however, is a less suitable proxy for our state variable. The existing literature often uses household income to capture labor supply (e.g., Zator 2020; Bernstein 2021). Therefore, this income component would appear in both the independent and the dependent variables.

³⁴Bankruptcy information is collected and mostly populated at three points in time, when a respondent is 25, 30, and 35 years old. Accordingly, the bankruptcy plot is restricted to the period when a respondent is between 25 and 35 years old.

³⁵Missed payments are those at least 60 days past due. The NLSY79 surveys whether an individual has missed payments in the past five years. See Manso, Rivera, Wang, and Xia (2023) for a detailed description of the data. As the authors note, the NLSY79 data are not suitable for studying the effects of student loans on household decisions. This is because in this older cohort of individuals, student loans have not become a prominent part of household debt, consistent with the evolution of student debt documented in Looney and Yannelis (2015).
Student loans: NLSY97 collects information on educational experiences and allows us to extract student loan borrowing information, along with the type of student loans. Similarly to the debt (assets) information, NLSY97 provides the amount of outstanding student debt when a respondent is 20, 25, 30, and 35 years old. A respondent is considered to have borrowed student loans if, at any of the four ages, his/her outstanding student debt is positive. Otherwise, the respondent is considered without student loans. In our empirical analyses, we examine whether the two groups of individuals exhibit different labor supply in relation to household indebtedness, as predicted by Figure V.

Two points are worth noting. First, our categorization captures the cross-sectional variation in whether a household has student debt during the sample period. This way is therefore unlikely confounded by the time-series (life-cycle) variation.³⁶ Second, the *corrective* effect of student loans proposed in our model hinges on non-dischargeablility in bankruptcy. Therefore, we only include federal student loans for our empirical analyses. Private student debt, on the other hand, may be eligible for discharge.

5.2 Sample and summary statistics

Our sample includes individuals who have attended college in the sample period. This restriction ensures that differences between households with and without student loans ($\Delta > 0$ versus $\Delta = 0$) are unlikely attributable to education attainment. To construct the sample, we start with the NLSY97 surveys from 1997 to 2019. We keep respondents who are 25 or older to ensure that they are likely in the labor force and hence, labor supply decisions are relevant. In this case, most of our sample (approximately 80%) falls between 2010 and 2019. As of the latest survey interview, the oldest respondents turn 39 years old. We obtain 38,022 respondent-year observations representing 4,878 individuals.

Next, we construct a host of control variables. *Male* and *White* indicate a respondent's gender and ethnicity. *MaritalStatus* captures a respondent's marital status in a given year. *GraduateDegree* indicates whether the respondent has obtained a graduate degree in a given year. To control for factors related to the life cycles of households, we include the respondent's age and its quadratic form, Age and Age^2 . We winsorize all continuous

 $^{^{36}}$ Boutros, Clara, and Gomes (2024) show that student debt balance decreases with age as households pay it off gradually.

variables at the 1th and 99th percentiles to eliminate undue effects of outliers. Lastly, we obtain the geographic location of each respondent from the restricted-use NLSY97 Geocode files supplementing the main NLSY97 survey. The Geocode files track each respondent's residential location at the time of an interview. We obtain a license to use this information from the BLS. In our analysis, we control for respondents' state (by year) fixed effects to absorb confounding macro-economic factors that may affect labor supply decisions. After requiring the availability of control variables and geographic location information, we obtain a sample consisting of 35,781 respondent-years. This sample constitutes the basis for our analysis.

Table 2 reports summary statistics. LaborSupply (in hours) has a mean of 1,832, a median of 2,080, and a standard deviation of 1,030. The median labor supply corresponds to approximately 40 working hours a week, and 52 weeks per year. As described in Section 5.3, our main variables of interest are the indicators for each category of household leverage following Gopalan, Hamilton, Kalda, and Sovich (2021), in order to mitigate potential data noise. The indicator functions $1_{Lev \in (l_k, h_k]}$ equal one if the total debt to total asset ratio is between l_k and h_k . We divide leverage into six non-overlapping categories: (0,25%], (25%, 50%], (50%, 75%], (75%, 100%], (100%, 125%], and (> 125%). The statistics of these indicators are reported in Table 2. They depict the distribution of household leverage in our sample.

In terms of student loans, the indicator SL equals one if a household has outstanding student debt in the sample period – corresponding to the case of $\Delta > 0$ in the model; it equals zero if $\Delta = 0$. This variable shows that more than half of households have borrowed student loans. Conditional on having student loans (SL = 1), the debt amount is slightly below \$20,000, accounting for 49% of total household debt. These statistics are in line with the existing studies (e.g., Looney and Yannelis 2015; Cornaggia, Cornaggia, and Xia 2021).

5.3 Baseline empirical results

Guided by the prediction of Figure V in Section 4, our analyses aim to compare households that have the *same* level of leverage but differ in its composition – in terms of the proportion of student loans among total debt. That is, we test whether households' labor supply re-

sponds to a given level of leverage differently, depending on whether there are student loans on the balance sheets (SL = 1 versus SL = 0).

Figure 2 plots (the logarithm of one plus) labor supply for different leverage groups. The dashed line represents households with student loans (SL = 1) and the solid line represents households without student loans (SL = 0). Household leverage is categorized into six groups; for instance "50-75" denotes the group whose leverage ratio is between 50% and 75%. Household leverage is the total debt divided by total asset, as defined and discussed in Section 5.1.

Two observations are consistent with Figure V. First, labor supply exhibits a hump-shape with respect to household leverage. Individuals are initially more likely to supply labor as leverage rises, but once leverage surpasses a certain threshold – i.e., when debt overhang kicks in, they become less likely to do so.

Second and importantly, households with student loans are more resilient to debt overhang. Specifically, the decline in labor supply for the dashed line does not emerge until leverage reaches the 100-125% range, compared to 75-100% for the solid line – a *delayed* manifestation of debt overhang. The decline in labor supply is also less aggressive for the dashed line: even at the highest leverage (> 125%), labor supply remains sizable; but in contrast, the solid line has dropped significantly by that time – to a level lower than the initial point when leverage is at 0-25%. This pattern indicates a *dampened* manifestation of debt overhang. Collectively, the delayed and dampened manifestation results in a widening gap between the two lines as leverage rises. This graphical analysis provides the first step evidence supporting the *corrective* effect of student loans.

We next estimate an OLS regression to formalize the graphical evidence. To capture the non-linear relation between labor supply and household leverage without imposing functional form restrictions, we follow the specifications in Gopalan, Hamilton, Kalda, and Sovich (2021). That is, we include indicators for categories of household leverage as the main independent variables of interest:

$$LaborSupply_{i,t} = \alpha + \sum_{k} \beta_k \mathbb{1}_{Lev_{i,t-1} \in (l_k,h_k]} + \gamma_1 Z_{i,t-1} + \gamma_2 X_i + FE + \epsilon_{i,t}.$$
(21)

Labor Supply is defined as the number of hours a household has worked (in logarithm) from

the last year (t-1) to the current year (t). In Appendix C Table A1, we alternatively use household earnings to capture labor supply and confirm our results. The indicator functions $1_{Lev_{i,t-1}\in(l_k,h_k]}$ equal one when the total debt to total asset ratio of the household *i* in year t-1 is between l_k and h_k . Leverage is divided into six categories: (0, 25%], (25%, 50%], (50%, 75%], (75%, 100%], (100%, 125%], and (> 125%).³⁷ The indicator for category (50%, 75%] is omitted as the base case. The coefficient β_k captures the average labor supply of households with leverage in each category relative to the base case.

Here we choose the median category (50%, 75%] as the base case for an easier empirical interpretation. The goal of our analyses is to depict the hump-shaped relation between labor supply and leverage, and importantly, the differences in this relation among households varying by student loans. With (50%, 75%] as the base case, negative β_k coefficients for the leverage categories on both sides of the base would indicate a hump shape – and the comparison of these coefficients across households with and without student loans would demonstrate any differences in this shape.

The vectors Z and X include time-varying and time-invariant respondent characteristics. Time-varying characteristics include respondent age, marital status, and the concurrent degree status. Time-invariant characteristics include gender and race. We include state×year fixed effects to control for macro level economic conditions, which might affect both household leverage and labor supply. Standard errors are clustered at the state by year level.³⁸

We start by estimating Equation (21) separately for households with and without student loans. Table 3 reports the results. Columns (1) to (3) include households with student loans, and columns (4) and (6) include those without student loans. To visualize these results, we plot the coefficients of each leverage category (i.e., β_k) in Figure 3 Panel A based on estimates in columns (3) and (6).

The patterns in Panel A largely resemble that of Figure 2. First, both groups of households exhibit a humped-shape relation between labor supply and leverage. This shape is statistically significant, as shown by the negative coefficients of the leverage categories on both sides of the base case – e.g., (0, 25%] and [> 125%]. These coefficients are significant at

³⁷We remove households with zero leverage from the analyses to ensure that households without student loans (SL = 0) are on an equal footing with those having student loans (SL = 1), which by definition have non-zero leverage.

³⁸Our main findings are similar when we cluster standard errors at the individual level.

the 5% level or better. Second, the hump shape exhibits differences between households with and without student debt. The switching point is higher for households with student loans (following leverage 100-125%), compared to that of households without student loans (75-100%). This result is consistent with the *delayed* manifestation of debt overhang. Once debt overhang kicks in, labor supply declines less aggressively when student debt is present, and it remains at a relatively high level than the other group – indicating a *dampened* manifestation of debt overhang.

The statistical significance of the delayed and dampened manifestation of debt overhang can be inferred from column (7) of Table 3. This column reports regression results testing the differences in β_k of each leverage category between column (3) and column (6). Specifically, we estimate an augmented Equation (21), including the interaction between each leverage category and SL (the indicator for the presence of student loans). That is,

$$LaborSupply_{i,t} = \alpha + \sum_{k} \omega_k \mathbb{1}_{Lev_{i,t-1} \in (l_k,h_k]} \times SL_i + \sum_{k} \beta_k \mathbb{1}_{Lev_{i,t-1} \in (l_k,h_k]} + \eta SL_i \qquad (22)$$
$$+ \gamma_1 Z_{i,t-1} + \gamma_2 X_i + FE + \epsilon_{i,t}.$$

The coefficients ω_k indicate the differences in each β_k between the two groups – and the differences in the hump shape between the two groups. In column (7), we observe that when leverage is relatively low (smaller than 100%), the labor supply of households with student loans stays at a similar level as households without student loans – indicated by the largely insignificant coefficients ω_k .³⁹ However, this pattern begins to diverge as leverage reaches the range of (100%, 125%): the labor supply of households with student loans stay significantly higher than the other group; the two coefficients, 0.317 and 0.712, become statistically significant at the 1% level. As visualized in Figure 3 Panel A, such a divergence arises because debt overhang (1) kicks in at a later stage when student loans are present (i.e., the delayed manifestation of debt overhang) and (2) is less aggressive in reducing labor supply than the case without student debt (i.e., the dampened manifestation). In Panel B of Figure 3, we plot the coefficients ω_k in column (7), which further illustrates the onset of labor supply divergence between cases with and without student debt over the spectrum of

³⁹The coefficients ω_k are not identical to the differences between β_k in column (3) and column (6). This is because a few singleton observations are dropped in the estimation of Equation (22).

household leverage.

The economic significance of this hump-shape divergence is sizable. Based on column (6), when the leverage of households without student debt initially increases from 0 to 100%, labor supply increases by approximately 140 hours – about 8% of the sample mean (Table 2). This increase is followed by a 452-hour decline, as debt overhang kicks in after the leverage level of 100%. In contrast, based on column (3) examining households with student debt, the initial increases in labor supply amounts to 118 hours, followed by only a 207-hour decline, and this decline does not emerge until leverage exceeds the 125% level.

Overall, the graphical and regression analyses support our model Prediction 1 and Figure V, and they illustrate the *corrective* effect of student loans in mitigating debt overhang on household labor supply.

5.4 Location based variation in student loan borrowing

Our empirical analyses contrast households that have the same level of leverage but differ in their student loan representation. One potential concern is that individuals with and without student debt may exhibit different characteristics, which simultaneously affect labor supply decisions. For example, individuals from financially abundant families may afford college expenses out of wealth instead of borrowing, and such family backgrounds may nurture these individuals to be less resilient to the impact of household debt. In this case, it would appear that individuals without student loans are more susceptible to debt overhang. In this section, we address this concern by identifying plausibly exogenous variation in student loan borrowing.

To illustrate the intuition, consider students A and B entering colleges in the same year (cohort) but originating from different states, X and Y. To the extent that students typically attend local colleges (due to proximity to family and/or lower costs), the average tuition in their origination states – particularly that of public universities – likely represent the education expenses these students are subject to.⁴⁰ If, given an enrollment year (cohort), colleges in state X charge a higher average tuition than those in state Y, then student A is more

⁴⁰This assumption is in the same spirit of Armona, Chakrabarti, and Lovenheim (2022), who employ local supply of for-profit colleges to identify variation in students' enrollment in those colleges and in turn, student loan borrowing.

likely to take out student loans to afford college expenses than student B. Because a student's origination state, along with the average college tuition in that state, is unlikely correlated with an individual resident's labor supply decision post college, this supply-side variation in education expenses offers a plausibly exogenous variation in student loan borrowing.

To operate this intuition, we collect university tuition information from the Integrated Postsecondary Education Data System (IPEDS) of the National Center for Education Statistics, and calculate the average in-state tuition of all public universities in each state-year (AvgTuition). We identify an individual's origination state using the location reported in the latest interview prior to college enrollment. In Table 4 Panel A, we verify that the average tuition of individuals' origination states significantly predicts student loan borrowing. Specifically, we perform a Probit model regressing the indicator SL (equal to one if the individual has borrowed student debt) on AvgTuition, along with cohort fixed effects. These fixed effects ensure that we compare students entering college at the same time; they help mute the potential confounding effects of economic or labor market conditions, and isolate the variation stemming from student locations. Column (1) of Table 4 Panel A includes AvgTuition as the main variable of interest and column (2) includes other controls used in the baseline analyses in Table $3.^{41}$ In both specifications, we observe a positive and significant correlation between the average tuition of an origination state and the likelihood of student loan borrowing.

Based on column (2), we then generate a predicted propensity of each individual borrowing student loans. We categorize individuals with a propensity in the top tercile of the sample distribution as "*Most likely to have student loans*", and those with a propensity in the bottom tercile as "*Least likely to have student loans*". We repeat our main analyses using these two groups.

Figure 4 presents the univariate graphical analysis following the same manner as Figure 2. The interpretation is similar: individuals that are most likely to have student loans (the dashed line) exhibit both delayed and dampened manifestations of debt overhang, compared to those least likely to have student loans (the solid line). This interpretation is further

 $^{^{41}}$ Because this analysis is based on cross-sectional (rather than panel) data, we collapse time-varying variables to form cross-sectional measures. For example, *MaritalStatus* now indicates an individual's marital status in the year when he/she first enrolls in colleges. *GraduateDegree* now indicates whether an individual has ever obtained a graduate degree in the sample period.

confirmed by Table 4 Panel B, in which we present regression analyses as those in Table 3.

Column (7) estimates Equation (22) and tests the differences in the coefficients of each leverage category between columns (3) and (6). We again observe that when leverage is lower, the labor supply of individuals most likely to have student loans stays at a similar level as that of individuals least likely to have student loans. However, as leverage grows beyond 100%, labor supply of the former group rises significantly above the latter. As discussed in Table 3, this observation suggests that in the presence of student debt, household labor supply does not decline until a later stage and the decline emerges in a less aggressive manner – both of which keep the labor supply at a higher level. These observations again confirm that student loans delay and dampen the manifestation of debt overhang.

In Appendix C Table A2, we supplement these tests with a two-stage least square (2SLS) instrumental variable (IV) analysis. This analysis is performed on Equation (22), in which the original variable of interest SL (the indicator for having student loans) – along with its interaction with the indicator for each leverage category – is instrumented by AvgTuition and the respective interactions with leverage indicators (except the base category omitted from regressions). The first stage regressions therefore contain six equations. The second stage re-estimates Equation (22) using the instrumented SL and its five interactions.

Column (1) of Table A2 reports the first stage. To conserve space, we only present the equation pertaining to SL (omitting other equations pertaining to its interactions). Here we observe that the average tuition of borrowers' origination states significantly predicts student loan borrowing, as previously shown. The Sanderson-Windmeijer F statistic is 30.36, significant at the 1% level.⁴² Column (2) reports the second stage, and shows a similar pattern as that in column (7) of Table 3 and Table 4. That is, when leverage is lower, the labor supply of individuals with a higher (instrumented) likelihood of having student loans stay at a similar level as, or somewhat lower than, that of individuals with a lower (instrumented) likelihood. This ranking reverses as leverage grows, and eventually flips when leverage reaches the highest level – indicating that student loans delay and dampen the effect of debt overhang on labor supply.⁴³

 $^{^{42}{\}rm The}\ F$ -statistics of the other five first-stage equations are between 30.38 and 71.34, significant at the 1% level.

⁴³The coefficient estimates in column (2) are generally larger than those in the previous tables because the instrumented SL is a continuous (instead of binary) measure.

6 Policy implications

In this section, we derive policy implications that speak to the recent regulatory effort promoting accessible and manageable IDR plans. The Biden administration announced the student loan relief program in 2022, proposing various adjustments to the existing IDR plans. We perform comparative statistic to our model, and provide predictions on how these changes would impact the household's labor supply incentives.

Our analysis in this section complements several studies focusing on an ex-post effect, that is, how changes in IDR programs affect households' labor supply conditioning on them having entered IDR (e.g., Karamcheva, Perry, and Yannelis 2020; de Silva 2023). Our contribution is to assess the *ex-ante* effect – i.e., the households' labor supply responses prior to enrolling in IDR, when they are still under the standard repayment. These ex-ante effects help us uncover several unintended consequences of the loan relief program.

6.1 Fraction of income ψ required for IDR payment

The loan relief program proposes to lower the IDR payment to 5% of students' discretionary income, down from the current 10-15%. Figure VI depicts comparative statics of labor supply policies prior to IDR enrollment with respect to ψ – the proportion of income required as IDR payment. The dashed (resp. solid) line depicts labor supply for a low (resp. high) value of ψ . The end of each line represents the point when the household switches to IDR from SRP.

As shown by the dashed line, a lower ψ aggravates debt overhang ex ante: the household's disincentive in supplying labor kicks in at an earlier stage than the solid line (0.3 versus 0.4 of indebtedness). This observation is intuitive. A lower proportion of income required as payment makes IDR more appealing. As a result, the household chooses to exercise this option earlier due to the anticipated pseudo "discharge", activating debt overhang.

However, the prediction regarding ex-post differences in labor supply (after filing for IDR) between a high and low ψ is ambiguous – as shown in Figure A2 of Appendix A.7. Two opposing forces are at play. On the one hand, the proportion of income required as IDR payment constitutes effectively an additional income tax, and a lowered ψ encourages the household from supplying labor. On the other hand, a lowered ψ increases the household



Figure VI: Household debt overhang and IDR uptake for different values of ψ . Labor l(d) is depicted for two different values of ψ , the fraction of income pledged as IDR payment. See Table 1 for baseline parameter values.

wealth, reduces the marginal utility of additional consumption, and discourages labor supply. In the presence of these opposing forces, Figure A2 shows that there is no clear ranking between the dashed (representing a lower ψ) and solid lines (representing a higher ψ).

This observation is in line with Karamcheva, Perry, and Yannelis (2020), who empirically find that a more generous IDR plan – which lowers the required proportion of household income as payment – does not significantly affect the income (labor supply) of borrowers already enrolled in IDR. Our model complements this finding by showing that before IDR enrollment, a reduction in the required IDR payment can generate greater distortion in borrowers' incentive to supply labor.

6.2 Debt forgiveness waiting period $1/\lambda$

The loan relief program also proposes to shorten the duration of IDR payment until loan forgiveness, down from the current 20-year waiting period. Figure VII depicts comparative statics of household labor supply with respect to $1/\lambda$ – the (expected) waiting period.

The dashed (resp. solid) line depicts the labor supply for a shorter (resp. longer) waiting period (5 versus 20 years). The observation in Figure VII follows a similar intuition as that in Figure VI. Shortening the waiting period makes IDR more attractive, thus activating debt



Figure VII: Household debt overhang and IDR uptake for different values of λ . Labor l(d) is depicted for two different values of λ , the parameter governing student debt forgiveness upon entering IDR. See Table 1 for baseline parameter values.

overhang in labor supply at an earlier stage.

Overall, the results in Sections 6.1 and 6.2 suggest that increasing the appeal of IDR to households could distort their labor supply decisions ex ante. The resultant disincentive may partially undermine the government's initiative to relieve households from student debt burden. To this extent, our model provides additional considerations regarding the pros and cons of the proposed IDR adjustments by the loan relief program.

6.3 Cost of filing for IDR u_I

Figure VIII depicts comparative statics of household labor supply with respect to u_I , the utility cost of filing for IDR. This analysis is inspired by Mueller and Yannelis (2022), who show that after a field experiment aiming to relieve households from this cost, IDR adoption increased substantially.

The dashed (resp. solid) line of Figure VIII depicts household labor supply for a low (resp. high) value of u_I . We observe that a lower filing cost again makes IDR more appealing, leading to an earlier manifestation of debt overhang ex ante. This result speaks to regulatory efforts targeting an easier and less time consuming IDR enrollment process for households. Two such efforts are the possibility of reducing the paperwork required to enroll in the



Figure VIII: Household debt overhang and IDR uptake for different values of u_I . Labor l(d) is depicted for two different values of u_I , the parameter governing how costly it is to file for IDR. See Table 1 for baseline parameter values.

program (e.g., by pre-filling some of the responses) and encouraging student loan servicers to provide better guidance on how to enroll in the program.

6.4 Income growth rate μ

Lastly, Figure IX depicts comparative statics of household labor supply with respect to μ , the parameter governing the growth rate of hourly wages (in Equation (3)). The dashed (resp. solid) line depicts the labor supply for a low (resp. high) value of μ . A higher growth rate makes IDR less appealing. This is because the household may forfeit a larger amount as IDR payment when their prospective income rises. This prospect in turn delays households' IDR enrollment – as shown by the later stopping point of the solid line. In turn, the labor supply collapses less sharply than the dashed line.

This finding is in the same spirit of Karamcheva, Perry, and Yannelis (2020), who empirically document that IDR plans are subject to adverse selection: borrowers with low earnings (or high loan balances) are most likely to adopt IDR – and vice versa. Here we complement this observation by considering households' prospective earnings, as well as how it affects their ex-ante IDR enrollment and ultimately, the labor supply.

From policymakers' perspective, this result implies that any effort in increasing households' prospective wages can delay the occurrence of debt overhang in labor supply – through



Figure IX: Household debt overhang and IDR uptake for different values of μ . Labor l(d) is depicted for two different values of μ , the parameter governing the growth rate of hourly wages. See Table 1 for baseline parameter values.

the channel of deterred household enrollment in student loan IDR. To the best of our knowledge, this is a benefit not yet documented in the prior literature.

7 Conclusion

In this paper, we develop a dynamic model to document an unexplored effect of student loans on household decisions – their ability to alleviate the debt overhang on household labor supply. This effect arises because non-dischargeable student debt mitigates households' strategic incentive to cut back labor supply when default becomes probable at the expense of creditors. Indeed, we show that as student loans take up a larger proportion of total household debt, labor supply becomes more resilient to debt overhang: the debt overhang does not emerge until a higher level of household indebtedness (a delayed manifestation of debt overhang), and the decline in labor supply exhibits a less aggressive pattern (a dampened manifestation of debt overhang).

Interestingly, when we consider various student loan payment options, we find that such a *corrective* effect of student loans is partially undone by the IDR plans, an alternative program recently under the spotlight of policymakers. Because IDR plans formulaically set the repayment as a proportion of borrowers' discretionary income – regardless of outstanding loan balance – IDR allows households to pseudo "discharge" the student debt in exchange for a fixed fraction of future earnings. This "dischargability" in turn re-activates debt overhang, giving rise to a *setback* effect of IDR. We supplement our model with a calibration and empirical analyses based on data obtained from the 1997 National Longitudinal Surveys of Youth. We find support for our theoretical predictions.

Our findings generate important policy implications, particularly speaking to the recently proposed reforms in IDR plans and student debt forgiveness. The student loan relief program announced by the Biden administration in 2022 proposes a few adjustments to make the IDR plans more accessible and manageable. We find that more appealing IDR plans may induce households to enroll in these plans at an earlier stage, thereby amplifying the setback effect of IDR and reactivating debt overhang to a larger extent. As such, these adjustments may ex ante worsen households' incentive to supply labor, potentially undermining the government's initiative to relieve households from debt burden.

References

- Abraham, K. G., E. Filiz-Ozbay, E. Y. Ozbay, L. J. Turner, 2018. Behavioral effects of student loan repayment plan options on borrowers' career decisions: Theory and experimental evidence. Working paper.
- ——, 2020. Framing effects, earnings expectations, and the design of student loan repayment schemes. Journal of Public Economics, 104067.
- Amromin, G., L. McGranahan, 2015. The Great Recession and Credit Trends across Income Groups. American Economic Review: Papers and Proceedings 105(5), 147–153.
- Andolfatto, D., M. Gervais, 2006. Human capital investment and debt constraints. Review of Economic dynamics 9(1), 52–67.
- Armona, L., R. Chakrabarti, M. Lovenheim, 2022. Student debt and default: The role of for-profit colleges. Journal of Financial Economics 144(1), 67–92.
- Batkeyev, B., K. Krishnan, D. K. Nandy, 2017. Student debt and stock market participation. Working paper.
- Bensoussan, A., J. L. Lions, 1978. Applications des inéquations variationnelles en contrôle stochastique. Dunod, Paris, France.
- Bernstein, A., 2021. Negative home equity and household labor supply. Journal of Finance 76(6), 2963–2995.
- Bernstein, A., D. Struyven, 2017. Housing lock: Dutch evidence on the impact of negative home equity on household mobility. Available at SSRN 3090675.
- Bernstein, S., T. McQuade, R. R. Townsend, 2021. Do household wealth shocks affect productivity? evidence from innovative workers during the great recession. Journal of Finance 76(1), 57–111.
- Bolton, P., H. Chen, N. Wang, 2011. A unified theory of Tobin's q, corporate investment, financing, and risk management. Journal of Finance 66(5), 1545–1578.
- Boutros, M., N. Clara, F. Gomes, 2024. Borrow now, pay even later: A quantitative analysis of student debt payment plans. Journal of Financial Economics 159, 103898.
- Brown, J., D. A. Matsa, 2020. Locked in by leverage: Job search during the housing crisis. Journal of Financial Economics 136(3), 623–648.
- Catherine, S., C. Yannelis, 2023. The distributional effects of student loan forgiveness. Journal of Financial Economics 147(2), 297–316.
- Chakrabarti, R., V. Fos, A. Liberman, C. Yannelis, 2023. Tuition, Debt, and Human Capital. Review of Financial Studies 36(4), 1667–1702.

- Congressional Budget Office, 2020. Income-Driven Repayment Plans for Student Loans: Budgetary Costs and Policy Options. .
- Cooper, D., J. C. Wang, 2014. Student loan debt and economic outcomes. Federal Reserve Bank of Boston working paper.
- Cornaggia, J., K. Cornaggia, H. Xia, 2021. College student behavior and student loan default. Working paper.

——, 2024. Natural disaster, financial shocks, and human capital. Working paper.

- Cornaggia, K., H. Xia, 2024. Who mismanages student loans, and why?. Review of Financial Studies, forthcoming.
- Cox, J., D. Kreisman, S. Dynarski, 2020. Designed to fail:Effects of the default option and information complexity on student loan repayment. Journal of Public Economics 192, 104298.
- Daniels, G. E., A. Smythe, 2019. Student debt and labor market outcomes. AEA Papers and Proceedings 109, 171–175.
- de Silva, T., 2023. Insurance versus moral hazard in income-contingent student loan repayment. Working paper.
- DeFusco, A. A., B. Enriquez, M. Yellen, 2024. Wage garnishment in the United States: New facts from administrative payroll records. American Economic Review: Insights 6(1), 38–54.
- Di Maggio, M., A. Kalda, V. Yao, 2019. Second chance: Life without student debt. Unpublished working paper. National Bureau of Economic Research.
- Diamond, D. W., Z. He, 2014. A theory of debt maturity: the long and short of debt overhang. The Journal of Finance 69(2), 719–762.
- Dobbie, W., P. Goldsmith-Pinkham, N. Mahoney, J. Song, 2020. Bad credit, no problem? Credit and labor market consequences of bad credit reports. Journal of Finance 75(5), 2377–2419.
- Dobbie, W., J. Song, 2015. Debt relief and debtor outcomes: Measuring the effects of consumer bankruptcy protection. American economic review 105(3), 1272–1311.
- Donaldson, J. R., G. Piacentino, A. V. Thakor, 2018. Household debt and unemployment. Journal of Finance, Forthcoming.
- Dyna, K., D. Elmendorf, D. Sichel, 2013. The evolution of household income volatility. Brookings.
- Dynarski, S. M., 2014. An economist's perspective on student loans in the United States. Economic Studies, 1–28.

- Dynarski, S. M., J. E. Scott-Clayton, 2006. The cost of complexity in federal student aid: Lessons from optimal tax theory and behavioral economics. National Tax Journal 59(2), 189–402.
- Eaton, C., S. T. Howell, C. Yannelis, 2020. When investor incentives and consumer interests diverge: Private equity in higher education. Review of Financial Studies 33(9), 4024–4060.
- Ferreira, F., J. Gyourko, J. Tracy, 2010. Housing busts and household mobility. Journal of urban Economics 68(1), 34–45.
- Fox, J. J., S. Bartholomae, J. Letkiewicz, C. Montalto, 2017. College student debt and anticipated repayment difficulty. Journal of Student Financial Aid 47(2), 1–26.
- Friedberg, L., R. M. Hynes, N. Pattison, 2022. Who benefits from bans on employer credit checks. Journal of Law and Economics, Forthcoming.
- Gallagher, E. A., S. B. Billings, L. R. Ricketts, 2023. Human capital investment after the storm. Review of Financial Studies 36(7), 2651–2684.
- Gicheva, D., 2011. Does the student-loan burden weigh into the decision to start a family?. Working paper.
- Goodman, S., A. Isen, C. Yannelis., 2021. A Day Late and a Dollar Short: Liquidity and Household Formation among Student Borrowers. Journal of Financial Economics 142(3), 1301–1323.
- Gopalan, R., B. H. Hamilton, A. Kalda, D. Sovich, 2021. Home equity and labor income: The role of constrained mobility. Review of Financial Studies 34(10), 4619–4662.
- Hampole, M. V., 2023. Financial Frictions and Human Capital Investments. Working paper.
- He, Z., 2011. A model of dynamic compensation and capital structure. Journal of Financial Economics 100(2), 351–366.
- Herbst, D., 2023. The Impact of Income-Driven Repayment on Student Borrower Outcomes. American Economic Journal: Applied Economics 15(1), 1–25.
- Ji, Y., 2021. Job search under debt: Aggregate implications of student loans. Journal of Monetary Economics 117, 741–759.
- Karamcheva, N., J. Perry, C. Yannelis, 2020. Income-driven repayment plans for student loans: Budgetary costs and policy options. Working paper.
- Kargar, M., W. Mann, 2023. The incidence of student loan subsidies: evidence from the PLUS program. Review of Financial Studies 36(4), 1621–1666.

- Kleiner, K., N. Stoffman, S. E. Yonker, 2021. Friends with bankruptcy protection benefits. Journal of Financial Economics 139(2), 578–605.
- Krishnan, K., P. Wang, 2019. The cost of financing education: Can student debt hinder entrepreneurship?. Management Science 65(10), 4522–4554.
- Laibson, D., A. Repetto, J. Tobacman, 2007. Estimating discount functions with consumption choices over the lifecycle.
- Leland, H. E., 1998. Agency costs, risk management, and capital structure. Journal of Finance 53(4), 1213–1243.
- Lochner, L., T. Stinebrickner, U. Suleymanoglu, 2021. Parental support, savings, and student loan repayment. American Economic Journal: Economic Policy 13(1), 329–371.
- Lochner, L. J., A. Monge-Naranjo, 2011. The nature of credit constraints and human capital. American economic review 101(6), 2487–2529.
- Longstaff, F. A., E. S. Schwartz, 1995. A simple approach to valuing risky fixed and floating rate debt. Journal of Finance 50(3), 789–819.
- Looney, A., C. Yannelis, 2015. A Crisis in Student Loans? How Changes in the Characteristics of Borrowers and in the Institutions They Attended Contributed to Rising Loan Defaults. Brookings Papers on Economic Activity (Fall), 1–89.
- Lou, M., S. Mongey, 2019. Assets and Job Choice: Student Debt, Wages, and Job Satisfaction. Working paper.
- Lucca, D. O., T. Nadauld, K. Shen, 2019. Credit supply and the rise in college tuition: Evidence from the expansion in federal student aid programs. Review of Financial Studies 32(2), 423–466.
- Manso, G., A. Rivera, H. Wang, H. Xia, 2023. Household Debt Overhang and Human Capital Investment. Available at SSRN 4190550.
- Melzer, B. T., 2017. Mortgage debt overhang: Reduced investment by homeowners at risk of default. Journal of Finance 72(2), 575–612.
- Mezza, A. A., D. R. Ringo, S. M. Sherlund, K. Sommer, 2016. On the effect of student loans on access to homeownership. Federal Reserve Board Finance and Economics Discussion Series.
- Minicozzi, A., 2005. The short term effect of educational debt on job decisions. Economics of Education Review 24(4), 417–430.
- Moffitt, R., 2021. Trends in earnings volatility among US men. The Reporter.
- Moffitt, R., S. Zhang, 2018. Income volatility and the PSID: Past research and new results. AEA Papers and Proceedings 108, 277–280.

- Mueller, H., C. Yannelis, 2018. The rise in student loan defaults in the great recession. Journal of Financial Economics 131, 1–19.
- Mueller, H. M., C. Yannelis, 2022. Increasing Enrollment in Income-Driven Student Loan Repayment Plans: Evidence from the Navient Field Experiment. Journal of Finance 77(1), 367–402.
- Rothstein, J., C. E. Rouse, 2011. Constrained after college: Student loans and early-career occupational choices. Journal of Public Economics 95(1), 149–163.
- Scott-Clayton, J., 2018. The looming student loan default crisis is worse than we thought. Economic Studies at Brookings, Evidence Speaks Reports 2(34).
- Shapiro, R. J., 2014. The flawed reasoning and evidence for the Department of Education's Gainful Employment Regulation of private, for-profit colleges and universities. Working paper.
- Shin, D., G. Solon, 2011. Trends in men's earnings volatility: What does the Panel Study of Income Dynamics show?. Journal of Public Economics 95, 973–982.
- U.S. Government Accountability Office, 2015. Federal student loans: Education could do more to help ensure borrowers are aware of repayment and forgiveness options. U.S. Government Accountability Office 15-633.
- Weidner, J., 2016. Does student debt reduce earnings?. Working paper.

Yannelis, C., 2020a. Strategic default on student loans. Working paper.

- ———, 2020b. Strategic default on student loans. Unpublished working paper. Working paper.
- Zator, M., 2020. Working More to Pay the Mortgage: Interest Rates and Labor Supply. Working paper.



Figure 1: Default probability and household leverage

This figure plots the default probability – the percentage of households entering default – for each 10-point leverage bin. Default is identified by whether a household has filed bankruptcy as of a given year. Bankruptcy information, including the date when a respondent first filed bankruptcy, is collected by the NLSY97 survey when a respondent is 20, 25, 30, and 35 years old.



Figure 2: Labor supply and household leverage

This figure plots the relation between labor supply and household leverage. It reports the average number of hours that a household has worked (in logarithm) since the previous year, across household leverage bins. The bin of 0-25 consists of households whose leverage is between 0-25%. The bin of 25-50 consists of respondents whose total leverage is between 25-50%, and so forth. The vertical lines surrounding each point represent the two-sided 90% confidence intervals. Household leverage is the ratio of total debt to total asset, as defined in Section 5.1.



Figure 3: Regression coefficients and the difference

Panel A plots the regression coefficients β_k of each leverage category indicator in Equation (21) separately for households with student loans (SL = 1) and without student loans (SL = 0). The coefficients correspond to column (3) and column (6) in Table 3, respectively. The vertical lines surrounding each point represent the two-sided 90% confidence intervals. Panel B plots the differences in these regression coefficients, corresponding to column (7) in Table 3.



Figure 4: Labor supply and household leverage, using location based variation in student borrowing

This figure plots the relation between labor supply and household leverage in a similar manner as Figure 2, using location based variation in student loan borrowing to categorize the presence of student loans among total debt. The variation stems from the average tuition of public universities in students' origination states during their college enrollment year. The detailed procedure of this categorization is in Section 5.4. The figure reports the average number of hours that a household has worked (in logarithm) since the previous year, across household leverage bins. Household leverage is the ratio of total debt to total asset, as defined in Section 5.1. The vertical lines surrounding each point represent the two-sided 90% confidence intervals.

Table 1: Calibration parameters

Parameter	Value	Source/Moment
Subjective discount rate, δ	0.03	Laibson, Repetto, and Tobacman (2007)
Savings interest rate, r_S	0.015	Average Federal funds rate
Borrowing interest rate, r_B	0.08	Average mortgage and credit card rates
Volatility, σ	0.35	Shin and Solon (2011) among others
Growth rate, μ	0.025	Bureau of Labor Statistics
Borrowing limit, \underline{s}	7.6	90th percentile of debt-to-income upon bankruptcy
Labor supply cost, θ	2.5	Average hours worked
Fraction of income as IDR payment, ψ	0.15	Current IDR program
For giveness intensity for IDR, λ	0.5	20-year waiting period until forgiveness
Cost of filing for IDR, u_I	8.5	Matching observable IDR take-up rate
		(Section 4.2)

This table summarizes the baseline calibration parameters discussed in Section 4.1.

Table 2: Summary Statistics

This table reports summary statistics of the sample. Labor supply is the total number of hours a household has worked since the previous year (t - 1) to the current year (t). $1_{Lev \in (0,25\%]}$ is an indicator for whether the ratio of total debt to total asset, measured at the previous year (t - 1), is between 0% and 25%. $1_{Lev \in (25,50\%]}$ is an indicator for whether the ratio of total debt to total asset is between 25% and 50%, and so forth. SL is an indicator for whether the household has outstanding student debt in the sample period. The detailed categorization is in Section 5.1. SL amount and SL proportion are the amount of student loans and the ratio of student loans to total debt, respectively, conditional on the household has student loans (SL = 1). Age is a respondent's age in the current interview. Male and White are indicators of a respondent's gender and ethnicity. GraduateDegree is an indicator for whether a respondent has obtained a graduate degree as of the previous year. MaritalStatus is an indicator for whether a respondent is married as of the previous year.

Variable	Ν	Mean	S.D.	p5	p50	p95
Labor supply (hrs)	35,781	1,832.420	1,029.780	0	2,080	3,190
$1_{Lev \in (0,25\%]}$	35,781	0.204	0.403	0	0	1
$1_{Lev \in (25, 50\%]}$	35,781	0.213	0.410	0	0	1
$1_{Lev \in (50,75\%]}$	35,781	0.210	0.407	0	0	1
$1_{Lev \in (75, 100\%]}$	35,781	0.131	0.338	0	0	1
$1_{Lev \in (100, 125\%]}$	35,781	0.005	0.228	0	0	1
$1_{Lev>125\%}$	35,781	0.186	0.389	0	0	1
SL	35,781	0.525	0.499	0	1	1
SL amount (Conditional on $SL=1$)	18,776	19,937	$18,\!477$	$1,\!167$	14,000	$65,\!667$
SL proportion (Conditional on $SL=1$)	$18,\!060$	0.492	0.576	0.017	0.308	1.604
Age	35 781	31 326	3 /128	26	31	37
Male	35,781	0.459	0.498	0	0	1
White	35,781	0.100	0.190	0	1	1
GraduateDearee	35,781	0.020	0.400	0	0	1
MaritalStatus	35,781	0.100 0.425	0.300	0	0	1
1/1 ///////////////////////////////////	55,761	0.420	0.434	0	0	1

Table 3: Baseline regression of household indebtedness and labor supply

This table presents regression analyses examining the effect of household leverage on labor supply. Columns (1) to (6) estimate Equation (21). Columns (1) and (3) include households with student loans on the balance sheets (SL = 1), and columns (4) to (6) include households without student loans (SL = 0). Labor Supply is the logarithm of one plus the number of hours that a household has worked since the previous year. The main independent variables are the indicator functions $1_{Lev_{i,t-1} \in (l_k,h_k]}$, which equal one when the household's total debt to total asset ratio (i.e., leverage) is between l_k and h_k . Leverage is divided into six nonoverlapping categories: (0%, 25%], (25%, 50%], (50%, 75%], (75%, 100%], (100%, 125%], and(> 125%). The indicator for category (50%, 75%) is omitted as the base case. Column (7) estimates Equation (22), testing the differences in each β_k between the two groups SL = 1and SL = 0. The control variables include a respondent's age, gender, ethnicity, graduate degree attainment, and marital status. The definitions of these control variables are in Table 2. State FE are indicators for the respondent's residential state reported in each interview. Year FE are indicators for calendar years. Each regression includes a separate intercept. Standard errors are clustered at the state-year level and reported in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Dep. Var.				Labor supply	y		
	Wit	th student lo	Dans	With	Difference		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$1_{Lev \in (0,25\%]}$	-0.285***	-0.274***	-0.281***	-0.408***	-0.382***	-0.429***	0.102
	(0.064)	(0.064)	(0.063)	(0.053)	(0.054)	(0.055)	(0.084)
$1_{Lev \in (25, 50\%]}$	-0.135**	-0.129**	-0.128**	-0.042	-0.028	-0.059	-0.105
	(0.054)	(0.054)	(0.054)	(0.053)	(0.053)	(0.053)	(0.076)
$1_{Lev \in (75, 100\%]}$	0.039	0.022	0.020	-0.087	-0.057	-0.019	0.055
	(0.057)	(0.057)	(0.057)	(0.068)	(0.072)	(0.071)	(0.091)
$1_{Lev \in (100, 125\%]}$	-0.003	-0.012	-0.036	-0.460***	-0.417***	-0.368***	0.317^{**}
	(0.058)	(0.058)	(0.059)	(0.133)	(0.136)	(0.136)	(0.142)
$1_{Lev>125\%}$	-0.341***	-0.347***	-0.389***	-1.111***	-1.059^{***}	-1.080***	0.712^{***}
	(0.048)	(0.049)	(0.050)	(0.110)	(0.110)	(0.112)	(0.121)
Male			0.301***			0.535***	
			(0.034)			(0.046)	
White			0.118***			0.185***	
			(0.040)			(0.047)	
GraduateDegree			0.627***			0.734***	
0			(0.052)			(0.062)	
MaritalStatus			-0.223***			-0.324***	
			(0.040)			(0.042)	
A g e			-0.191			0.002	
			(0.133)			(0.158)	
Aae^2			0.003			-0.001	
			(0.002)			(0.002)	
State \times Year FE	NO	YES	YES	NO	YES	YES	
Observations	18,787	18,749	18,749	17,010	16,953	16,953	
R-squared	0.004	0.044	0.056	0.011	0.056	0.075	

Table 4: Location based variation in student loan borrowing

Panel A presents Probit regression analyses examining the relation between student loan borrowing and college tuition. The dependent variable is the indicator SL, which equals one if an individual has borrowed student debt, and zero otherwise. The main independent variable is AvqTuition, defined as the average in-state tuition of public universities in an individual's origination state. Origination state is identified as the state reported in the latest interview prior to the year when the individual enrolls in college. Cohort fixed effects are indicators for the year of college enrollment. Panel B presents regression analyses examining the effect of household leverage on Labor Supply, using location based variation in student loan borrowing. Columns (1) to (6) estimate Equation (21). Columns (1) to (3) include households that are most likely to have borrowed student loans, and columns (4) to (6)include households that are least likely to have borrowed student loans. The likelihood of student loan borrowing is predicted by column (2) of Panel A, using the average in-state tuition of public universities in an individual's origination state during the college enrollment year. The categorization of households that are most (least) likely to have student loans is described in Section 5.4. Labor Supply is the logarithm of one plus the number of hours that an individual has worked since the previous year. The main independent variables are the indicator functions $1_{Lev_{i,t-1} \in (l_k,h_k]}$, which equal one when the household's total debt to total asset ratio (i.e., leverage) is between l_k and h_k , as defined in Table 3. Column (7) estimates Equation (22), testing the differences in each β_k between the two groups of households. The definitions of control variables are in Table 2. State FE are indicators for the respondent's residential state reported in each interview. Year FE are indicators for calendar years. Each regression includes a separate intercept. Standard errors are clustered at the state-year level and reported in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Dep. Var.	SL			
	(1)	(2)		
AvgTuition	$\begin{array}{c} 0.154^{***} \\ (0.029) \end{array}$	0.161^{***} (0.028)		
Controls Cohort FE	NO YES	YES YES		
Observations R-squared	$4,962 \\ 0.017$	$4,962 \\ 0.039$		

Panel A: Relation between student loan borrowing and tuition

Dep. Var.	Labor supply							
	Most likely	ely to have student loans		Least likel	Difference			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
$1_{Lev \in (0.25\%]}$	-0.513***	-0.452***	-0.429***	-0.262***	-0.245***	-0.297***	-0.080	
	(0.080)	(0.080)	(0.082)	(0.072)	(0.074)	(0.074)	(0.109)	
$1_{Lev \in (25, 50\%]}$	-0.138**	-0.093	-0.115*	0.035	0.089	0.049	-0.134	
	(0.065)	(0.068)	(0.067)	(0.063)	(0.065)	(0.065)	(0.093)	
$1_{Lev \in (75, 100\%]}$	-0.018	-0.010	0.014	-0.003	0.018	0.032	0.033	
	(0.072)	(0.073)	(0.073)	(0.072)	(0.075)	(0.077)	(0.105)	
$1_{Lev \in (100, 125\%]}$	0.118	0.121	0.094	-0.567***	-0.565***	-0.619***	0.762***	
	(0.092)	(0.094)	(0.094)	(0.131)	(0.135)	(0.133)	(0.166)	
$1_{Lev>125\%}$	-0.294***	-0.264***	-0.322***	-0.794***	-0.825***	-0.821***	0.573***	
	(0.064)	(0.068)	(0.070)	(0.077)	(0.079)	(0.078)	(0.101)	
Controls	NO	NO	YES	NO	NO	YES		
State \times Year FE	NO	YES	YES	NO	YES	YES		
Observations	11.018	10.958	10.958	10,757	10.695	10.695		
R-squared	0.007	0.064	0.082	0.013	0.071	0.095		

Panel B: Regression analyses using location based variation in student loan borrowing

A Online Appendix

A.1 Dimensional reduction

In this Appendix, we provide details on the reduction of the dimensionality of the household's optimization problem from the original two state variables (S, K) to indebtedness (d) as the single state variable. d is defined as follows:

$$d_t = \frac{S_t}{l_{\tau_D}\underline{s}K_t} = \frac{S_t\sqrt{\theta}}{\underline{s}K_t},\tag{23}$$

where the equality follows from substituting the labor supply at default from Equation (9). Therefore, $|\underline{s}|$ represents the highest earnings multiple the household is allowed to borrow, since $d_t \leq 1$ implies that $S_t \leq \underline{s} l_{\tau_D} K_t$, for all t.

We illustrate the dimensional reduction by solving the value function before the household enrolls in IDR (as in Section 3.3). The other cases can be solved analogously.

Recall that we aim to find a function F(S, K) that solves Equation (18) subject to Equation (20). Due to CRRA preferences and GBM, we conjecture the following functional form for the value function:

$$F(S,K) = f(S/K) + \frac{1}{\delta}\log K.$$
(24)

Upon substituting into Equation (18), we observe that letting s = S/K, the equation becomes a second order ODE in $f(\cdot)$ given by:

$$2\delta + 2\delta \log (f'(s)) + 2\delta^2 f(s) + \sigma^2 = 2\delta s f'(s) \left(-\mu + r(s) + \sigma^2\right) + \delta s^2 \sigma^2 f''(s) + \frac{\delta f'(s)^2}{\theta} + 2\mu.$$
(25)

The value matching and smooth pasting conditions given by Equation (20) become:

$$f(\hat{s}) = \hat{f}((1-\Delta)\hat{s}) - u_I$$
 and $f'(\hat{s}) = \hat{f}'((1-\Delta)\hat{s})(1-\Delta),$ (26)

where \hat{f} denotes the scaled value function while in IDR that satisfies the same functional form as Equation (24) with respect to $\hat{F}(S, K)$. Finally, we need a transversality condition, derived below in Section A.2. These three boundary conditions allow us to solve using a standard shooting method for the exercise boundary \hat{s} and the scaled value function before enrolling in IDR f(s) where $s \in [\hat{s}, \infty)$.

Finally, since s and d are linearly related, it is straightforward to go from the scaled savings s to the variable of interest – indebtedness d – for all the quantities of interest.

A.2 Transversality condition

For the value function while in IDR (Section 3.2), we also need a transversality type of condition that must apply when savings become large. This condition is obtained by noting that the limiting case – when the household has no labor income (i.e., when wages are zero) – implies that the household consumes fraction δ of his savings due to logarithmic preferences. That is,

$$\lim_{K \to 0} C(S, K) = \delta S \iff \lim_{s \to \infty} s f'(s) = \frac{1}{\delta}.$$
 (27)

A.3 Proof of Proposition 1:

We recall that the value function in the baseline model is denoted by G(S, K). The verification argument given a smooth solution follows from Bensoussan and Lions (1978). Next, making an identical conjecture as in Equation (24), we denote by g(s) the scaled value function. Upon substituting this conjecture, the HJB becomes:

$$2\delta + 2\delta \log (g'(s)) + 2\delta^2 g(s) + \sigma^2 = \delta s^2 \sigma^2 g''(s) + 2\delta s g'(s) \left(-\mu + r(s) + \sigma^2\right) + \frac{\delta g'(s)^2}{\theta} + 2\mu$$
(28)

and optimal consumption is given by K/g'(s) and labor supply by $g'(s)/\theta$.

We study the behavior of the value function and the policies as a function of Δ . Let $0 \leq \Delta_1 < \Delta_2$. Then, we note that $g(s; \Delta_1) \geq g(s; \Delta_2)$. That is, ceteris paribus, more student loans make households worse off – since for any given policy, the household with a lower student loan balance will get a higher value upon filing for bankruptcy. Next, we prove the following Lemma:

Lemma 3. The derivatives of the value function for $\Delta_1 < \Delta_2$ satisfy:

$$g'(s;\Delta_1) < g'(s;\Delta_2) \tag{29}$$

for $s \in [\bar{s}, 0)$. That is, the value function of a household with higher student loans is uniformly steeper in the range $[\bar{s}, 0)$.

Proof of Lemma 3:

Evaluating Equation (28) at s = 0 reduces it to a first order ODE, which implies that there is only one value of g(0) that satisfies both the ODE and the transversality condition (27). Since the ODE and the transversality condition are independent of Δ , then we must have that $g(0; \Delta_1) = g(0; \Delta_2)$. Moreover, since we already know that $g(s; \Delta_1) > g(s; \Delta_2)$, it must be the case that $g(0; \Delta_1) \leq g(0; \Delta_2)$.

Now suppose for a contradiction that Equation (29) is not satisfied for some $s \in [\underline{s}, 0)$. This means that there exists \tilde{s} such that $g'(\tilde{s}; \Delta_1) = g'(\tilde{s}; \Delta_2)$. Moreover, let \tilde{s} be the largest value (i.e., closest to 0) such that the two derivatives are equal. Since it is the largest, then we must have that $g''(\tilde{s}; \Delta_1) < g''(\tilde{s}; \Delta_2)$. Next, we evaluate Equation (28) at \tilde{s} for both values of Δ . We then subtract these two expressions to obtain:

$$2\delta^2(g(\tilde{s};\Delta_1) - g(\tilde{s};\Delta_2)) - \tilde{s}^2\delta\sigma^2(g''(\tilde{s};\Delta_1) - g''(\tilde{s};\Delta_2)) = 0.$$
(30)

But this is a contradiction since each of the terms in the LHS are strictly positive. Therefore, we conclude that no such \tilde{s} exists and that $g'(s; \Delta_1) < g'(s; \Delta_2)$ for all $s \in [\bar{s}, 0)$, which completes the proof of Lemma 3.

Finally, having characterized the slope of the scaled value as a function of Δ , it follows that a household with more student loans will consume less (since consumption is inversely proportional to g'(s)) and work more (since labor is directly proportional to g'(s)), which completes the proof of Proposition 1.

A.4 Proof of Lemma 1:

The first order condition with respect to labor supply (l) before entering IDR given by (19) upon substituting (24) becomes $l(s) = \frac{f'(s)}{\theta}$. Similarly, labor supply while in IDR is given by $\hat{l}(s) = \frac{(1-\psi)\hat{f}'(s)}{\theta}$.

Evaluating the expressions above for labor supply at the IDR exercise threshold \hat{s} combined with the smooth pasting conditions (26) yields:

$$l(\hat{s}) = \frac{f'(\hat{s})}{\theta} = \frac{\hat{f}'((1-\Delta)\hat{s})(1-\Delta)}{\theta} < \frac{(1-\psi)\hat{f}'((1-\Delta)\hat{s})}{\theta} = \hat{l}(\hat{s}(1-\Delta)),$$
$$\longleftrightarrow$$
$$\Delta > \psi.$$

Finally, the ratio of the labor supply immediately after filing for IDR to that immediately before filing for IDR is given by:

$$\frac{\hat{l}(\hat{s}(1-\Delta))}{l(\hat{s})} = \frac{1-\psi}{1-\Delta},\tag{31}$$

which is increasing in Δ and decreasing in ψ .

A.5 Proof of Lemma 2:

By (1) making the dependence of the value function and the IDR exercise boundary on Δ explicit, (2) differentiating the value matching condition in (26) with respect to Δ , and (3) then cancelling some terms using the smooth pasting condition in equation (26), we obtain:

$$\frac{\partial f(\hat{s}(\Delta);\Delta)}{\partial \Delta} = \underbrace{f'((1-\Delta)\hat{s}(\Delta);\Delta)}_{\geq 0} \underbrace{[-\Delta\hat{s}(\Delta)]}_{\geq 0} + \underbrace{\frac{\partial \hat{f}((1-\Delta)\hat{s}(\Delta);\Delta)}{\partial \Delta}}_{=0} \geq 0.$$
(32)

The first inequality follows from the monotonicity of the value function on scaled savings (i.e., more savings generates higher utility). The second inequality follows from the fact that the household only files for IDR when scaled savings are negative (i.e., when it has debts). The last equality holds because the terms of IDR are invariant to the student loan amount. Thus, the household's value function after enrolling in IDR is independent of Δ (i.e., $\frac{\partial \hat{f}(\cdot;\Delta)}{\partial \Delta} = 0$).

A.6 Wage reduction post default

In our baseline model, we assume that a household's hourly wage remains intact after default, as documented by Dobbie, Goldsmith-Pinkham, Mahoney, and Song (2020). We now consider the possibility that hourly wages decline moderately after default. Such decline may arise because of resistance from employers to the household's unfavorable credit history – resulting in reduced employment, or because of wage garnishment until the household's debts are repaid – which effectively lowers the hourly wage. In this Appendix, we show that our key result (that student loans reduce debt overhang) is robust to a post-default decline in hourly wages. To this end, we extend the model to incorporate a parameter $\phi > 0$ that captures the post-default fraction of hourly wages retained by the household upon default. That is, the value function post default for the household now becomes $H(\phi K)$, where $1 - \phi > 0$ captures the hourly wage decline after default.



Figure A1: Robustness with respect to ϕ . See Table 1 for baseline parameter values.

Figure A1 shows the baseline cases in solid lines in which there is no post-default reduction in hourly wages (i.e., $\phi = 1$) for two values of Δ , and in dashed the cases in which there is a 25% reduction in hourly wages post default (i.e., $\phi = 0.75$). We observe that even with the wage reduction, student loans continue to mitigate debt overhang, that is, the solid red line continues to show a less aggressive debt overhang pattern than the red dashed line.

A.7 Household labor supply while in IDR with different ψ

Figure A2 depicts comparative statics of the ex-post household labor supply – after the household enrolls in IDR – with respect to ψ . ψ is the proportion of income required as IDR payment. The dashed (resp. solid) line depicts the labor supply for a low (resp. high) value of ψ . As previously discussed, the net effect of an increase in ψ on the labor supply is



Figure A2: Labor supply while in IDR for different values of ψ . See Table 1 for baseline parameter values.

ambiguous due to the opposing effect of a higher effective income tax (discouraging labor) and a wealth effect (encouraging labor). Here we observe that the discouraging effect prevails in a larger portion of the state space (i.e., where the solid blue line is below the dashed red line).

B Discussion of supplementary features regarding IDR

B.1 Payment cap in IDR repayment

In our main analyses, we abstract from the "payment cap" feature incorporated in certain IDR sub-programs (such as IBR and PAYE) for tractability. Under these sub-programs, borrowers will not pay more than what they would have paid under the standard repayment plan, should their future income grow substantially. We now account for this feature.

We start by discussing conceptually how this feature affects our core findings regarding IDR. Recall from Section 3.3 that the household's decision to enroll in IDR trades off the benefit against the costs – including the possibility that a positive wage shock may significantly increase the household's income and result in the IDR payment exceeding the standard payment (which could have been avoided if the household did not give up the option of waiting). This "overpay" situation will be eliminated by the repayment cap. To this end, IDR cap lowers the cost of IDR enrollment and makes IDR more appealing to households. Consequently, the household is expected to exercise the IDR option even earlier, reactivating debt overhang. Put differently, the repayment cap of certain IDR sub-programs likely strengthens the *setback* effect, which is at the core of our paper's contribution in depicting the nuances of how student loans affect labor supply.

We then consider to what extent the payment cap feature is salient to our model based on both existing empirical evidence and a simulation analysis. In a sample of financially distressed student borrowers, Cornaggia and Xia (2024) estimate that in order for an average student's IDR payment to surpass the 10-year standard payment (thus making the repayment cap binding), her prospective income needs to more than triple of the current level. Such a "super growth" is unlikely (even in the span of 20 years when loan forgiveness takes effect at the end of IDR), given that the average annual growth rate of U.S. median household earnings is less than 5%. Consistent with this implication, Karamcheva, Perry, and Yannelis (2020) find that the typical borrower in IDR is negatively amortizing. This happens when borrowers are not making payments high enough to pay down interest while in IDR, resulting in the accrued interest capitalized and the balance growing over time (even exceeding the original amount). This finding therefore suggests that IDR is likely an absorbent state and the repayment cap is unlikely to bind.

We next consider to what extent the payment cap feature may alter the main findings our model based on a simulation analysis. Our goal is to assess the likelihood that the IDR payment cap becomes binding. To do so, we simulate household income paths and estimate the percentage of scenarios in which the household's income grows substantially – resulting in the IDR payment surpassing the standard payment. Specifically, for each household, we first compute the amount of student loans at the time point immediately prior to the IDR enrollment. We then calculate the standard payment for this amount under the 10-year amortization plan. This payment is set as the IDR repayment cap. Next, we simulate the household's income paths over the next 20 years, and compute the required IDR payment over this time interval. The IDR payment is computed as a 15% of discretionary income, where discretionary income is defined as the gross income in excess of 150% of the poverty



Figure A3: Simulations of IDR payment paths. See Table 1 for baseline parameter values.

line.44

For the ease of exposition, Figure A3 illustrates 100 simulations of the paths of IDR repayment. The dashed horizontal line denotes the IDR payment cap. In this illustrative simulation, only 5 paths go above the cap during the entire 20-year span.⁴⁵ In fact, when we run 10,000 simulations, we find that 5.2% of the simulated paths go over this cap. This observation is in line with recent studies suggesting that borrowers enrolled in IDR remain financially distressed, and are unlikely to experience substantive income growth to push their IDR payment above the standard payment.⁴⁶ We thus conclude that the repayment cap feature of IDR abstracted from our model is unlikely to affect our results.

⁴⁴Based on Cornaggia and Xia (2024), borrowers enrolled in IDR have gross income that is approximately 1.6 times the poverty line (based on a family size of two), which we use as input for our simulations. We run 10,000 simulations.

⁴⁵More specifically, the initial IDR repayment is normalized to 1, and in order for the IDR repayment cap to bind, the households' income would need to increase by almost 3 times from its initial level.

⁴⁶In a sample of financially distressed student borrowers, Cornaggia and Xia (2024) estimate that in order for an average student's IDR payment to surpass the 10-year standard payment (thus making the repayment cap binding), her prospective income needs to more than triple of the current level. Such a "super growth" is unlikely. Karamcheva, Perry, and Yannelis (2020) find that the typical borrower in IDR is negatively amortizing. This happens when borrowers are not making payments high enough to pay down interest while in IDR, resulting in the accrued interest capitalized and the balance growing over time (even exceeding the original amount).
B.2 Taxable forgiven balance

Another feature associated with IDR is that in earlier years, the remaining balance forgiven after 20-25 years of IDR payments was taxable. However, since 2021, the government has eliminated this tax burden.

Considering taxes on the forgiven balance is relevant because it may alter one condition generating the IDR's *setback* effect. That is, the IDR payment needs to be invariant to the borrower's outstanding loan balance – in which case additional labor supply does not benefit the borrower, discouraging her from supplying labor ex ante. If IDR payments become decreasing on the outstanding balance, then the borrower would be incentivized to work harder ex ante, reduce the balance before IDR, and enjoy a lower monthly payment during IDR. This is the key intuition discussed in Section 3.6.

Taxes to be paid on the IDR forgiven balance (prior to 2021) can be regarded as an *addition* to IDR month payments. This addition does depend on the outstanding loan balance – that is, the higher the balance upon IDR enrollment, the larger amount will likely be forgiven in 20-25 years, and the more taxes will be incurred at that time. As such, IDR payments are not entirely invariant to the loan balance – offsetting the IDR *setback* effect.

We believe this consideration is unlikely to change our main finding in a material way. Ultimately, the amount of taxes to be paid on the forgiven balance is determined by the borrower's average tax rate, and given that borrowers enrolled in IDR are often in financial distress, this rate is likely low – resulting in a relatively small amount of taxes to be paid when the loan balance is eventually forgiven.⁴⁷ Further spreading this small amount over 20-25 years of IDR repayments (before the forgiveness takes place) diminishes the impact of these taxes. Put differently, the "balance-dependent-part" of the IDR repayments (due to taxes on forgiven balance) is small relative to the "balance-invariant-part" of the IDR repayments (the fixed 10%-15% of monthly income). As such, IDR payments are approximately invariant to the loan balance, thereby preserving the *setback* effect.

⁴⁷For instance, Cornaggia and Xia (2024) estimate that among financially distressed students needing federal assistance such as IDR, the average income is \$28,300. Based on the federal income tax rate for single filers in 2005, their average tax rate is about 13.7%.

C Additional figures and tables



Figure A4: Default probability and household leverage (NLSY79 Cohort)

This figure uses survey data from the NLSY79 to plot the default probability for each 10-point leverage bin. In Panel A, default is identified by whether a household has filed for bankruptcy as of a given year. In Panel B, default is identified by whether a household has missed mortgage or rent payment as of a given year. Missed payments are those at least 60 days past due. The NLSY79 surveys whether an individual has missed payments in the past five years. Detailed descriptions of the NLSY79 data and sample construction are in Manso, Rivera, Wang, and Xia (2023).

Table A1: Robustness – Labor income as an alternative measure of labor supply

This table presents regression analyses examining the effect of household leverage on labor income as an alternative measure of labor supply. Columns (1) to (6) estimate Equation (21). Columns (1) and (3) include households with student loans on the balance sheets (SL = 1), and columns (4) to (6) include households without student loans (SL = 0). Labor Income is the logarithm of one plus the total wage and salary that a household has earned since the previous year. The main independent variables are the indicator functions $1_{Lev_{i,t-1} \in (l_k,h_k]}$, which equal one when the household's total debt to total asset ratio (i.e., leverage) is between l_k and h_k . Leverage is divided into six non-overlapping categories: (0%, 25%], (25%, 50%], (50%, 75%], (75%, 100%], (100%, 125%], and (> 125%). The indicator for category (50%, 75%) is omitted as the base case. Column (7) estimates Equation (22), testing the differences in each β_k between the two groups SL = 1 and SL = 0. The control variables include a respondent's age, gender, ethnicity, educational attainment, and marital status. The definitions of these control variables are in Table 2. State FE are indicators for the respondent's residential state reported in each interview. Year FE are indicators for calendar years. Each regression includes a separate intercept. Standard errors are clustered at the state-year level and reported in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Dep. Var.	Labor Income							
	With student loans				Without student loans			Difference
	(1)	(2)	(3)		(4)	(5)	(6)	(7)
$1_{Lev \in (0,25\%]}$	-0.630^{***} (0.102)	-0.544^{***} (0.101)	-0.453^{***} (0.098)		-0.694^{***} (0.081)	-0.606^{***} (0.082)	-0.538^{***} (0.082)	0.079 (0.120)
$1_{Lev \in (25, 50\%]}$	0.034 (0.092)	0.057 (0.090)	0.085 (0.089)		0.176** (0.077)	0.189^{*} (0.078)	0.140^{*} (0.076)	-0.058 (0.120)
$1_{Lev \in (75, 100\%]}$	0.179^{**} (0.089)	0.238^{***} (0.089)	0.249^{***} (0.087)		-0.252^{**} (0.099)	-0.129 (0.101)	-0.018 (0.096)	0.307^{**}
$1_{Lev \in (100, 125\%]}$	-0.210^{**}	-0.122	-0.130 (0.096)		-0.758^{***}	-0.667*** (0.200)	-0.492^{**}	0.386^{*}
$1_{Lev>125\%}$	(0.050) -0.150^{***} (0.079)	(0.033) -0.933^{***} (0.077)	(0.030) -0.867^{***} (0.076)		(0.131) -2.070^{***} (0.142)	(0.200) -1.900*** (0.142)	(0.135) -1.765*** (0.142)	(0.200) 0.993^{***} (0.156)
$\begin{array}{l} \text{Controls} \\ \text{State} \times \text{Year FE} \end{array}$	NO NO	NO YES	YES YES		NO NO	NO YES	YES YES	
Observations R-squared	18,787 0.020	$18,749 \\ 0.091$	$18,749 \\ 0.114$		$17,010 \\ 0.022$	$16,953 \\ 0.072$	$16,953 \\ 0.114$	

Table A2: Instrumental variable analyses

This table presents a two-stage least square (2SLS) instrumental variable (IV) analyses on Equation (22). The first stage regresses SL (the indicator for having student loans), along with its interactions with each leverage category, on AvgTuition and the respective interactions with leverage categories (except the omitted base category). The first stage therefore contains six regressions. These equations also include all control variables in the second stage regressions. The second stage re-estimates Equation (22) using the instrumented SL and its interactions. Column (1) reports the first stage equation pertaining to SL only. Column (2) reports the second stage equation. The definitions of variables are in Table 2. Cohort FE are indicators for the year of college enrollment. State FE are indicators for the respondent's residential state reported in each interview. Year FE are indicators for calendar years. Standard errors are clustered at the state-year level and reported in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	First stage	Second stage
Dep. Var.	SL	Labor supply
	(1)	(2)
	0 00	
AvgTuition	0.035^{***}	
Ano Truition V 1	(0.008)	
Avg $Iullion \times 1_{Lev \in (0,25\%]}$	(0.009)	
Ana Truition × 1.	(0.010)	
$A \ bg \ 1 \ a \ b \ b \ c \ c \ c \ c \ c \ c \ c \ c$	(0.010)	
AnaThitian $\times 1_{L_{2}} = (75, 100\%)$	0.029***	
$Lev \in [75, 100\%]$	(0.025)	
AvaTuition $\times 1_{L_{mc}(100,125\%)}$	0.027**	
$100 1000000 \times 1200 (100,125\%)$	(0.014)	
AvaTuition $\times 1_{Lev \in (>125\%]}$	0.007	
	(0.009)	
Instrumented SL		0.271
		(1.290)
Instrumented SL× $1_{Lev \in (0.25\%]}$		-2.751*
		(1.435)
Instrumented SL× $1_{Lev \in (25,50\%]}$		0.116
		(1.450)
Instrumented SL× $1_{Lev \in (75,100\%]}$		0.149
		(1.444)
Instrumented SL× $1_{Lev \in (100, 125\%]}$		2.444
		(1.661)
Instrumented SL× $1_{Lev>125\%}$		5.256^{**}
		(2.092)
First Stage Sanderson-Windmeijer F stat	30.36	
Controls	VES	VES
State \times Year FE	YES	YES
Cohort FE	YES	YES
	1 100	
Observations	32,478	32,478