Borrow Now, Pay Even Later: A Quantitative Analysis of Student Debt Payment Plans

Michael Boutros   Nuno Clara   Francisco Gomes*

October 2023

Abstract

In the United States, student debt currently represents the second largest component of consumer debt, just after mortgage loans. Repayment of those loans reduces disposable income early in the borrower's lifecycle, when marginal utility is particularly high, and limits their ability to build a buffer stock of wealth to insure against background risks. In this paper, we study alternative student debt contracts that offer a 10-year deferral period. Borrowers defer either principal payments only ("Principal Payment Deferral", PPD) or all payments ("Full Payment Deferral", FPD) with the missed interest payments added to the value of the debt outstanding. We first calibrate an equilibrium with the current contracts, and then solve for counterfactual equilibria with the PPD or FPD contracts. We find that both alternatives generate economically large welfare gains, which are robust to different assumptions about the behavior of the lenders and borrower preferences. We decompose the gains into the percentages resulting from loan repricing and from the deferral of debt repayments. We compare these alternative contracts with the changes to Income Driven Repayment Plans being proposed by the current U.S. administration and show that they dominate such proposals. Crucially, the PPD and FPD contracts deliver similar welfare gains to the debt relief program considered by the administration, with no impact on the government budget constraint.

*Boutros: Bank of Canada (michael.boutros@bankofcanada.ca), Clara: Duke University, Fuqua School of Business (nuno.clara@duke.edu), Gomes: London Business School and CEPR (fgomes@london.edu). We thank Aleksander Andonov, John Campbell, Ruidi Huang, Tyler Hull, James J. Li, Jonathan Parker, Constantine Yannelis, and seminar participants at Duke University, the Southwestern Finance Association Conference, the Midwest Finance Association Conference, the Canadian Economic Association Conference, the North American Summer Meeting of the Econometric Society, the University of Oregon Summer Finance Conference for helpful comments. The presented views are those of the authors, not necessarily of the Bank of Canada.

1 Introduction

Student debt in the United States has more than tripled over the last 15 years, increasing from \$500 billion to almost \$1.8 trillion (Panel A of Figure 1). Student debt currently represents 7% of GDP, increasing from 4% of GDP in the early 2000s (Panel B), and is now the largest non-mortgage component of consumer debt. This increase is driven by both the extensive and intensive margins of student borrowing; in particular, the average student borrower now owes approximately \$37,000, up from \$18,000 in 2008 (Panels C and D).

The increasing volume of student debt and the repayment difficulties that many people now experience have opened a debate on whether public policy should further intervene in the student loan market. In this paper, we develop a quantitative dynamic model to evaluate the effects of alternative student debt repayment plans on consumption and savings over the lifecycle. We use the model to study both Standard Repayment Plans (SRP) and the increasingly popular Income-Driven Repayment Plans (IDRP) and to quantify the welfare gains from a set of modified plans that we propose. Our analysis centers on two standard channels through which student debt affects households: wealth and liquidity.

Student debt imposes a large wealth effect because, in addition to mechanically decreasing lifetime wealth, the presence of student debt on a household's balance sheet may preclude other leveraged purchases, such as real estate or vehicles, and this may further affect lifetime wealth. Young households (less than 35 years), who are at a critical juncture in their lives, hold almost 40% of student debt. For these households, massive levels of student debt have a first-order impact on behavior and can delay important discrete decisions such as becoming a homeowner, entering the stock market, or starting a family (Goodman et al., 2021; Folch and Mazzone, 2022; Hampole, 2022). The liquidity effects of student debt arise because payments begin shortly after graduation and draw upon liquidity, thereby depressing both immediate consumption and savings for future consumption, especially for lower-income households. These loans can be difficult to repay or renegotiate and are essentially nondischargeable in bankruptcy.

Recent proposals by the federal government, which issues more than 90% of all student debt in the United States, have centered around outright student debt forgiveness. These policies induce a large and positive wealth effect that is commensurately borne as a large fiscal cost. Ganong and Noel (2020) study a set of mortgage modification policies, enacted during the Great Recession, to isolate the wealth and liquidity effects associated with housing debt and that worked in fundamentally the same way as student debt does today. They find that reductions in principal payments, which increase wealth without affecting liquidity, have no impact on either consumption or delinquency. It follows that if student forgiveness is indeed effective, it will be because forgiving the debt outright will also have a positive liquidity effect. However, Catherine and Yannelis (Forthcoming) estimate that if student debt is forgiven for many high-income and likely very liquid households, this will mute the impact the additional liquidity will have on consumption.

In contrast, the policies we consider defer payments to later in the lifecycle, thereby offering concessions to borrowers only early in their lifecycles. This induces a large and positive liquidity effect early in these students' lives and, in some cases, a relatively small negative wealth effect. In principle, these deferrals are similar to maturity extensions, which Ganong and Noel (2020) find had a large and positive effect on consumption and delinquencies for mortgage borrowers during the Great Recession. These policies are also similar to the student debt payment pause program included in the 2020 CARES Act, which Dinerstein et al. (2023) document as having led to increased consumption and fewer delinquencies among holders of loans with paused payments. Importantly, these plans have near-zero fiscal cost because they are still fully repaid and, in fact, may generate more profits for the lenders over the lives of these loans.

The model serves to carefully quantify the welfare gains from increasing liquidity by deferring loan payments. Current student debt contracts impose a repayment schedule that starts (almost immediately) after graduation. At this stage of the lifecycle, agents have lower income and the marginal utility of consumption is particularly high. Furthermore, these agents have not yet been able to accumulate significant savings, and the requirement to make debt repayments further prevents them from accumulating much wealth early in life. In addition to reducing consumption overall, this debt repayment makes households less able to smooth consumption and much more vulnerable to income shocks. For these reasons, it makes sense to defer debt repayments to later in these agents' lifecycles, when they typically earn higher incomes and have accumulated significant savings.

We start by calibrating a life-cycle model of consumption and savings behavior where households are endowed with an initial level of student debt that must be repaid under the current rules: the Standard Repayment Plan (SRP) schedule with the option of enrolling in the Income-Driven Repayment Plan (IDRP). Under the IDRP, agents' total debt repayments are capped at a fraction of their disposable income, which can be particularly valuable for low-income individuals. Unsurprisingly this option is mostly taken early in life. We show that the model matches well the empirical fraction of individuals on an SRP, an IDRP or who are delinquent. It also replicates the empirical patterns of income, debt outstanding, debt to income and net wealth across these three scenarios.

We then use the calibrated model to solve for counter-factual equilibria, where we introduce alternative student debt contracts. Under the first alternative contract, ”Principal Payment Deferral” (PPD), agents only have to make interest payments during the first 10 years of the loan. Principal repayments only start in the 11th year. Under the second alternative contract, ”Full Payment Deferral” (FPD), agents are not required to make any payments at all during the first 10 years of the loan. During this period, interest payments are simply added to the value of the loan, which therefore grows every year at the rate of the interest on the loan.

For each of these alternative contracts we solve for a new equilibrium where lenders reprice the loans and households re-optimize subject to these new contracts. We evaluate these policies along several dimensions, including their impact on borrower consumption and welfare, but also delinquency rates and cash-flows to the lender(s). We find that under these alternative contracts individuals are better able to smooth consumption over time and insure against income shocks, leading to economically significant welfare gains: yearly certainty equivalent consumption gains of 1.35% for PPD contracts and 2.36% for FPD contracts.[1] We decompose these gains into the fraction that arises from re-pricing the loans and the fraction resulting from deferring the debt repayments. For the FPD contract, the gains are almost exclusively driven by the latter. By contrast, for the PPD contract, a significant component of the gains arises from a reduction in the loan interest rate.

Using our model, we show that these welfare gains are essentially identical to the ones implied by the debt-forgiveness proposals. This calculation is actually a conservative one, since our model ignores the potential costs arising from the additional taxes or the lower government spending that will be required to balance the federal budget.[2] Therefore, the actual level of debt forgiveness that would generate the same welfare gains is likely to be higher.

We also compare our certainty equivalent gains with those obtained under a simple 10-year contract extension. We find that while extending the maturity of the existing contracts is welfare improving, the corresponding gains are only 29% (51%) of those obtained under

[1]For comparison, there are similar welfare gains from stock market participation, computed in similar life-cycle models, e.g., Cocco et al. (2005).

[2]Alternatively, higher government debt or an increase in the money supply.

the FPD (PPD) modification. This confirms the intuition that the welfare gains are largely coming from reducing the debt burden early in life. Replicating these gains by using a simple maturity extension would require a much longer extension period than 10 years.

In addition to improving household welfare by postponing (most) debt repayments until later in the lifecycle, when agents have higher income and more wealth, these new contracts also imply substantially lower delinquency rates. Relative to the current equilibrium, delinquencies decrease by one-third in an economy with PPD contracts and by one-half in an economy with FPD contracts. In our analysis, we (conservatively) ignore the potential impact of these lower delinquency rates on the risk premia associated with the loan contracts. Any further reduction in the loan interest rates, due to this channel, would lead to even larger welfare gains.

For tractability, our baseline model abstracts from other features that could increase the welfare gains even further, such as housing decisions, family planning, job searches, and stock market participation. To the extent that individuals are forced to delay stock market participation or a housing purchase because they are required to repay their student debt early in life, the benefits of our proposed contracts would be even larger.[3] Similarly, the additional pressure to secure an income stream limits these individuals' job searches and forces them into suboptimal matches, as shown by Hampole (2022) and Folch and Mazzone (2022). Therefore, the welfare gains that we are measuring are likely a (fairly conservative) lower bound, relative to their full potential benefit for households. Student loans also have an impact on individuals' credit scores. Making regular payments on their student loans can help households build credit scores early in life. On the other hand, skipping a payment and becoming delinquent will trigger important negative credit events. Since both the FPD and the PPD contracts significantly reduce delinquency rates, this constitutes an additional source of welfare gains that is not captured in our analysis.[4]

Finally, we use our model to evaluate the current U.S. administration's proposals to change IDRP plans: (i) reduce the time for forgiveness in IDRPs from 25 years to 10 years and (ii) change the payment formula of IDRPs. We show that the first proposal has tiny welfare benefits for students, whereas the second proposal is similar in nature to our PPD/FPD

[3]An extension with an endogenous participation decision yields welfare gains that are around 20% higher than in the baseline model, even though we impose an exogenous portfolio allocation.

[4]Under FPD contracts, agents are not required to make loan payments early in life, so both the potential credit score benefits (from not becoming delinquent) and the costs (from becoming delinquent) are absent. But under the PPD contract, the payments still exist but they are only reduced early in life.

proposals and has welfare gains that are larger than those in our PPD proposal but smaller than those in the FPD.

Our paper contributes to the large literature studying student loans and repayment behavior, as surveyed in Lochner and Monge-Naranjo (2016).[5] Student debt is unique from other forms of leverage for two key reasons. First, it is unique to household balance sheets because student debt cannot be discharged in bankruptcy nor can human capital financed by student debt be seized during bankruptcy. Second, relative to how firms or governments finance their spending over long or infinite horizons, consumption and savings decisions depend heavily on the household's age, that is, its position in the lifecycle. The deferral policies we consider are welfare improving because of the upward sloping income profile and other age-related expenses early in the lifecycle.

From a market perspective, the demand for student debt has increased as both the returns to and the costs of education have increased over the last several decades, while the supply side of the market has responded with the expansion of government programs and a growing private lending sector (Sun and Yannelis, 2016; Ionescu and Simpson, 2016; Amromin et al., 2017; Lucca et al., 2019; Gallagher et al., 2022; Yannelis and Looney, 2022). For many, the decision to invest in higher education and accumulate human capital is closely linked to the ability to obtain student loans (Lochner and Monge-Naranjo, 2011; Gary-Bobo and Trannoy, 2015; Palacios, 2015; Abbott et al., 2019; Athreya et al., 2020). In this paper, we abstract from these larger forces that drive households to acquire student debt and we analyze the behavior of households whose members enter the workforce with student debt after having completed their educations.

Our main focus is on the specific payment plans that student borrowers use to pay down their debt. We quantitatively compare the standard repayment plan to income-driven repayment plans. Whereas there are a number of empirical papers looking into income-driven repayment plans, to the best of our knowledge we are the first to model the endogenous choice to enroll in income-driven repayment plans. Karamcheva et al. (2020) document that although the majority of borrowers enrol in the default standard repayment plan, income-driven plans have gained popularity in recent years. The evidence suggests that directly offering the income-driven plans to borrowers an alternative to the default plan increases enrolment (Abraham et al., 2020; Cox et al., 2020; Mueller and Yannelis, 2022) (in line with

[5]See also surveys by Avery and Turner (2012), Amromin and Eberly (2016), Bleemer et al. (2017), and Athreya et al. (2021).

similar behavior on interest-free student loans documented by Cadena and Keys (2013)), but there are important non-monetary costs which deter enrolment in income-driven plans (Lochner et al., 2021). Maggio et al. (2020) find large benefits from the discharge of student debt in a natural experiment. Herbst (Forthcoming) finds income-driven repayment plans reduce delinquencies, decrease outstanding balances, and have a positive effect on long-run measures of financial health. Goodman et al. (2021) show that increasing liquidity for student borrowers, which our modified plans do by adjusting the timing of payments to accumulate wealth for a period of time before paying down their student debt, has large welfare benefits for borrowers.

The rest of paper is organized as follows. Section 2 presents the model. Sections 3 and 4 discuss the calibration of the baseline economy and the baseline results, respectively. In Section 5, we introduce the alternative contracts, solve for the new equilibrium and compute our welfare measures. In Section 6, we compare our proposed student debt modifications to the changes proposed by the current U.S. administration. We conclude in Section 7.

2 Model

2.1 Environment

We consider an overlapping generations model with T generations of households and a single lender (the federal government) that provides student loans. Each period corresponds to one year. We only model the loan market and otherwise take a partial equilibrium approach.[6]

We model households that have just completed their educations and have begun their working lives with an initial endowment of student debt. In each period, in addition to plans for consuming and saving, each household makes a decision regarding its student debt payment. They can make a payment according to the default standard repayment plan (SRP), pay a cost to enroll in the income-driven repayment plan (IDRP), or become delinquent and pay a corresponding penalty. Each household lives for T periods. In the first K periods, the household receives stochastic labor income and faces borrowing constraints, and during the R retirement period, it receives pension income.

Student loans are issued at time zero by a single lender, the federal government. For

[6]We only consider households with student loans, so we are not modelling all consumers/savers.

our purposes, this would be equivalent to assuming a continuum of ex-ante identical and perfectly competitive lenders. In the baseline model, we calibrate the interest on student loans from the data. When considering alternative debt contracts, we compute the lender's net present value (NPV) associated with that particular interest rate and use it to re-price all of the other contracts. We discuss these calculations in detail in Section 5.

The default debt contract is the SRP, which is a constant-payment loan with a fixed payment schedule and maturity date. Alternatively, households can opt into the IDRP by paying a switching cost.[7] Under the IDRP, the payments are a function of their income; therefore, the maturity of the loan is variable. Payments in the IDRP are capped at those of the standard plan. Because of this cap, the maturity date of the SRP is a lower bound for the maturity in the IDRP but typically the debt is paid off over a longer period of time. The IDRP loans have a maximum maturity of 25 years and, at this time, any remaining debt is discharged without penalty.

We abstract from labor supply decisions and in particular how different contracts impact supply decisions. There is substantial empirical evidence that labor supply is not particularly responsive to the type of student debt payment plan. For example, Karamcheva et al. (2020) exploit a natural discontinuity to estimate that the labor supply effects of income-driven plans in the United States are negligible, and Britton and Gruber (2019) reach the same conclusion in the United Kingdom. With respect to the alternative policies we consider, Jacob et al. (2023), studying the federal Teacher Loan Forgiveness program using both a natural experiment and randomized control trial, find no evidence that outright debt forgiveness reduces borrowers' labor supply. Altogether, the supply of labor in each period appears independent of payment plan choice and modifications to the parameters of each payment plan. If, against existing evidence, our plans did alter labor supply decisions, this would manifest in different debt-to-income ratios, and we show in Section 5.5.3 that our conclusions are not sensitive to different levels of this ratio.

2.2 Debt Contracts

Student debt starts under the terms of the SRP, which is described below. However, households that qualify can apply for payments under an alternative plan, the IDRP. Under this plan, debt repayments are a function of income and are capped at the value of the SRP. There-

[7] In practice, this is largely a time cost associated with submitting the necessary paperwork; hence, we model it as a utility cost.

fore, low-income households have an incentive to switch plans. However, as a consequence of delaying their repayments, the maturity of the loan is extended.

The SRP structures payments using a standard amortization schedule that is common across many loan types, while the IDRP is designed to assist recent graduates as they enter the workforce and anticipate having increasing income profiles. Under most circumstances, students are given a six- or nine-month grace period between graduation and their first debt payment, so the student debt payments in our model begin in period two.

We denote the interest rate on the loans as $r_s \equiv r_f + \varphi^{Baseline}$, where $\varphi^{Baseline}$ is the student loan premium over the risk-free rate (r_f). The interest on the student loans is tax-deductible at the income tax rate τ.

2.2.1 Standard Repayment Plan

In the standard repayment plan, the loan is amortized over N_{SRP} periods (in the absence of delinquency). In each period, the principal (P_t^{SRP}) and interest payments (I_t^{SRP}) sum to a constant total payment, given by the following standard formula:

$$P_t^{SRP} + I_t^{SRP} = \left[\frac{1-(1+r_s)^{-N_{SRP}}}{r_s}\right]^{-1} S_0, \tag{1}$$

where S_0 is the initial balance of the loan and r_s is the interest rate on the loan. The fraction of the total payment allocated toward the principal and interest varies in each period according to the level of the outstanding debt:

$$I_t^{SRP} = r_s S_t, \tag{2}$$

$$P_t^{SRP} = \left[\frac{1-(1+r_s)^{-N_{SRP}}}{r_s}\right]^{-1} S_0 - r_s S_t. \tag{3}$$

As with all constant repayment loans, payments early in the amortization schedule are mainly applied to the interest, rather than the principal, with the pattern reversing as the loan reaches maturity.

2.2.2 Income-Driven Repayment Plan

Under the Income-Driven Repayment Plan (IDRP), the student loan payment in each period takes into account the household's income. Payments are reassessed annually, depending on changes in the household's tax filings, and correspond to the yearly frequency in the model.

A crucial feature of the IDRP is that the payment is capped at the alternative payment under the standard repayment plan. Specifically, payments in the IDRP are equal to the lesser between ω_{IDRP} of discretionary income or the standard payment:

$$P_t^{IDRP} + I_t^{IDRP} = \min\{P_t^{SRP} + I_t^{SRP}, \omega_{IDRP} \cdot \mathrm{DI}_t\} \tag{4}$$

Discretionary income, defined formally in the next section, is equal to net income minus a fraction, ω_{FPL}, of the Federal Poverty Level (FPL). The values of ω_{IDRP} and ω_{FPL} are parameters set by the policymaker

The IDRP has a maximum maturity of $N_{IDRP}(> N_{SRP})$ years, after which any remaining principal is discharged with no penalty.[8] As before, the interest is calculated using the outstanding level of debt and the principal payment is the remainder:

$$I_t^{IDRP} = r_s S_t, \tag{5}$$

$$P_t^{IDRP} = \min\{P_t^{SRP} + I_t^{SRP}, \omega_{IDRP} \cdot \mathrm{DI}_t\} - I_t^{IDRP}. \tag{6}$$

Note that since the total payment on the loan is capped at a fraction of discretionary income (equation (4)), it is possible that this would not be enough to cover the interest payments, that is, we might have

$$\omega_{IDRP} \cdot \mathrm{DI}_t < r_s S_t. \tag{7}$$

Under such a scenario, equation (6) implies that the amortization of the principal is actually negative and, therefore, the total principal increases between periods. We note that this occurs despite the fact that the household has neither become delinquent on the loan nor deviated from its payment schedule in any way.

2.2.3 Transitions Between IDRP and SRP

If an agent switches from the SRP to the IDRP, then they are subject to the rules described above. Likewise, if they later revert back to the SRP, then the loan terms are the standard SRP terms. However, since under the IDRP the loan amortization is lower (i.e., lower than the value implied by equation (3)), then the overall maturity of the loan will typically be higher than N_{SRP}. Making principal payments under equation (3) is not going to deliver a zero balance at $t = N_{SRP}$ because the current outstanding balance is higher than it would have been if the agent had remained in the SRP throughout the payment period.

[8] If the household's income is sufficiently large such that the payments are always equal to those under the standard plan, then the entire student loan is paid off in N_{SRP} periods. For intermediate levels of income and payments, the debt may be fully paid off between N_{SRP} and N_{IDRP} periods.

2.2.4 Delinquency

In practice, a student loan becomes delinquent immediately after missing a single payment. If the loan is delinquent for a period of time, typically 270 days, it goes into default. However, the government offers a number of plans to help distressed borrowers avoid default, and as a result, less than 0.5% of households (prior to the pandemic) were in default. The borrower can apply for temporary payment relief against a long list of eligibility criteria, e.g., economic hardship, medical expenses, and cancer treatment, in addition to case-by-case accommodations. All of these plans exist to ensure that borrowers receive the relief they require such that they eventually return to good standing and continuing making payments.

To model this, we allow the household to skip a payment on its student debt and become delinquent, but then return to good standing in the following period. Under delinquency, the household pays a utility penalty, ξ^D, in exchange for temporary liquidity relief. In keeping with reality, the debt is not discharged under delinquency, and the missed interest accrues to its existing balance:

$$P_t^D = -I_t^D, \tag{8}$$

and therefore the debt balance grows by the interest rate,

$$S_{t+1} = (1 + r_s)S_t. \tag{9}$$

Delinquency lasts for one period, and then household has access to the same menu of choices in the next period, including the option to become delinquent again.

2.3 Households

2.3.1 Budget Constraint

Households start each year with an initial endowment of wealth (W_t) and a stock of student debt (S_t), which could be zero if the debt had already been fully repaid. During the year, households receive labor income (or pension income if retired) and make their choices regarding how much to consume (C_t) and how much to pay on their debt (under the SRP, the IDRP, or to become delinquent). Household savings are invested at a riskless rate that earns a deterministic return, r_f. In the baseline model we do not allow agents to borrow from other credit markets (in addition to the student debt they have). We relax this assumption later on the paper.

Household wealth therefore evolves according to the following dynamic budget constraint:

$$W_{t+1} = (1 + r_f)(W_t - C_t^j - P_t^j - I_t^j) + (1 - h_{t+1} - \tau)Y_{t+1}, \tag{10}$$

where h_t is the fraction of gross income on housing-related expenditures and τ is the income tax rate.[9] Net income is then given by $(1 - h_t - \tau)Y_t$. Switching costs (for the agents who enroll in the IDRP) or delinquency costs (for the agents who choose to skip a payment) are modelled as utility costs, so they do not enter the budget constraint.

Student debt, S_t, is measured at the beginning of the period. P_t^j and I_t^j denote the principal and interest payments, respectively, on student debt under each option $j \in \{\text{SRP}, \text{IDRP}, \text{D}\}$. Unless the household chooses delinquency, student debt evolves according to the principal payments:

$$S_{t+1} = S_t - P_t^j, \tag{11}$$

where P_t^j is given by equation (3) or equation (6), depending on the plan type. In the event of delinquency, S_{t+1} is given by equation (9).

2.3.2 Income Process

Income during the household's working life is modeled following Guvenen et al. (2021). In period t of household i's working life, income is given by

$$Y_t^i = (1 - \nu_t^i) \exp(g(t) + \alpha^i + z_t^i + \epsilon_t^i), \tag{12}$$

where $g(t)$ captures the age profile of the household's earnings and α is a household fixed effect calibrated to match average earnings. The unemployment shock, ν_t, generates a large decrease in income when the household is unemployed, while the stochastic processes, z_t and ϵ_t, capture, respectively, persistent and transitory income shocks for employed households.

The persistent income process, z_t^i, follows an AR(1),

$$z_t^i = \rho z_{t-1}^i + \eta_t^i, \tag{13}$$

with innovations drawn from a mixture of normal distributions. The persistent shock η_i^t is $\mathcal{N}(\mu_{\eta,1}, \sigma_{\eta,1})$ with probability p_z and $\mathcal{N}(\mu_{\eta,2}, \sigma_{\eta,2})$ otherwise.

[9] We do not model an explicit housing decision and instead incorporate housing expenditures in a reduced-form approach, following Gomes and Michaelides (2005).

The transitory shock, ϵ_t^i, is also a mixture of normal distributions drawn from $\mathcal{N}(\mu_{\epsilon,1}, \sigma_{\epsilon,1})$ with probability p_ϵ and $\mathcal{N}(\mu_{\epsilon,2}, \sigma_{\epsilon,2})$, otherwise. In both cases, the expected value of the mixed distribution is zero.

The unemployment shock, $1-\nu_t^i$, is given by

$$1-\nu_t^i = \begin{cases} 1 & \text{with prob. } 1-p_\nu(t, z_t^i), \\ \lambda & \text{with prob. } p_\nu(t, z_t^i), \end{cases} \tag{14}$$

where

$$p_\nu^i(t, z_t) = \frac{\exp(a_\nu + b_\nu t + c_\nu z_t^i + d_\nu z_t^i t)}{1+\exp(a_\nu + b_\nu t + c_\nu z_t^i + d_\nu z_t^i t)}. \tag{15}$$

This shock depends on the household's age and the persistent component of the income process. When the unemployment shock is realized, the household's income is scaled down by a constant fraction, λ.

As described in section 2.2, the debt repayments under the IDRP are a function of household discretionary income. In the model, for tractability, the measure of the discretionary income only includes the life-cycle component, the individual fixed effect, and the persistent component of household income:

$$\text{DI}_t = \exp(g(t) + \alpha + z_t) - \omega_{FPL} \times \text{FPL}. \tag{16}$$

Following Cocco et al. (2005), retired households receive a deterministic fraction, ω, of their income in the last period of their working lives. More precisely, for retired household i in period t, income is given by

$$Y_t^i = \omega \cdot \exp(g(K) + \alpha^i + z_K^i) \tag{17}$$

where K is the final working period.

2.3.3 Preferences and the Individual Optimization Problem

The individual optimization problem has three state variables: wealth, W_t, the level of student debt outstanding, S_t, and persistent labor income, z_t. We assume that households have Epstein-Zin preferences over consumption, as specified below.

In each period in which the household has student debt outstanding, it will decide between making payments under as SRP, an IDRP, or becoming delinquent on its student debt obligation. If household income is sufficiently low, then payments under the IDRP can be lower, but switching involves a non-monetary cost (ξ^{IRDP}).[10] Debt payments can be fully avoided via deliquency, but this is associated with a utility penalty, ξ^{D}. We can therefore write the household's value function as the maximum value associated with these three alternative choices:

$$V_t(W_t, S_t, z_t) = \max\{V_t^{SRP}(W_t, S_t, z_t), V_t^{IDRP}(W_t, S_t, z_t), V_t^{D}(W_t, S_t, z_t)\}, \quad (18)$$

where V_t^j denotes the auxiliary value function associated with each possible option $j \in \{\text{SRP}, \text{IDRP}, \text{D}\}$. The auxiliary value functions associated with each of three possible debt repayment decisions are given by

$$V_t^j(W_t, S_t, z_t) = \max_{C_t^j(\cdot)}\{(1-\beta)[C_t^j(W_t, S_t, z_t)^{1-1/\psi} - \xi^j] + \beta E_t[V_{t+1}(W_{t+1}, S_{t+1}, z_{t+1})^{1-\gamma}]^{\frac{1-1/\psi}{1-\gamma}}\}^{\frac{1}{1-1/\psi}}, \quad (19)$$

where ψ is the elasticity of intertemporal substitution, β is the subjective discount factor, and γ is the coefficient of the relative risk aversion. The utility cost of enrolling in each payment plan or of become delinquent is given by ξ^j.[11]

For all j, the continuation value on the right-hand side of equation (19) is the unconditional value function, V_{t+1}, since in the next period, the household can again choose between both plans or becoming delinquent.

2.3.4 Solution Method

This model cannot be solved analytically. The model has three continuous state variables (wealth W_t, level of student debt S_t and permanent income z_t), one continuous choice variable (C_t) and one discrete choice variable (whether to make the standard payment on the loan, enroll in IDRP, or become delinquent). In addition, we have to solve for the equilibrium loan premium under each alternative policy/economy (except the baseline for which this value is calibrated).

[10] This form of the IDRP cost is motivated by evidence in Lochner et al. (2021) and Mueller and Yannelis (2022) on the non-monetary costs associated with IDRP enrolment, primarily in the form of the attention costs associated with learning about the program and the cognitive costs of successfully completing the enrolment process.

[11] So ξ^{SRP}=0, since that is the initial plan, ξ^{IRDP} is the cost of transitioning to the IDRP plan, and ξ^{D} is the delinquency penalty.

We derive policy functions numerically using backward induction. We discretize the state space and follow Tauchen and Hussey (1991) to approximate the distribution of permanent income and idiosyncratic shocks. Conditional on a choice of repayment, for each point in our state space we maximize the utility function over consumption using a golden search method, and interpolate over the grids of state variables for values that do not lie on the grid. After solving for the optimal policy functions we solve a fixed-point problem to obtain the loan premia under different policies. We provide details on fixed point calculation in section 5.2.1 below. Each complete equilibrium takes between 40 to 100 hours to solve for, on a 12-core computer.

3 Calibration

3.1 Income Process

We calibrate the income process (equations (12) to (15)) using the estimates in Guvenen et al. (2021) and Cocco et al. (2005). We use the estimates of the Gaussian mixture parameters and the unemployment shock function from Guvenen et al. (2021). To ensure that the expected value of each mixture is zero, Guvenen et al. set the mean of the second component to zero (without loss) and estimate the mean of the first component. We report the full set of parameters in Table 1.

[INSERT TABLE 1 HERE]

We calibrate the parameter α_i to match the median income conditional on having student debt. Income and student debt data are taken from the 2019 wave of the Survey of Consumer Finances (SCF). The life-cycle component of earnings, $g(t)$, and the retirement income replacement rate are taken from estimates in Cocco et al. (2005) for college-educated households, and housing expenditures are taken from Gomes and Michaelides (2005).

3.2 Student Debt

Table 2 reports the specification of student debt payment plans. The payment period for the Standard Payment Plan, N_{SRP}, is 10 years, while maximum payment length of the Income-Driven Repayment Plan, N_{IDRP}, is 25 years. The federal government offers a number of

Income-Driven Repayment Plans that differ slightly in their construction, and we choose parameters that roughly resemble the average plan. Discretionary income is constructed using $\omega_{DI} = 150\%$ of the poverty level, and the IDRP payment is equal to $\omega_{IDRP} = 10\%$ of discretionary income.

We parameterize the remainder of the model and set calibration targets using data from the SCF.[12] The size of the student debt loan is set to the median initial amount borrowed by households, \$29,000. The loan premium for student debt, $\varphi^{Baseline}$, is the average rate, 3.5%. The calibration targets in our model are the fraction of households on each plan or delinquent. Since these three must sum to one, we focus on the fraction of households that are on the Income-Driven Repayment Plan, 31.3%, or are delinquent, 18.0%.[13] Naturally, to match these moments, we calibrate the utility costs of enrolling in IDRR, ξ^{IDRP}, and of becoming delinquent, ξ^{D}.

[INSERT TABLE 2 HERE]

3.3 Preferences and Other Parameters

In our baseline calibration, we set the coefficient of risk aversion, γ, to 2 and the elasticity of intertemporal substitution ψ to 0.5. Conditional on these, we then calibrate the discount factor, β, to 0.95 to match the median financial wealth of agents at the beginning of the life-cycle (age 25-34), conditional on having student debt. We use data from the 2019 wave of the SCF to compute financial wealth. Financial assets include transaction accounts, certificates of deposit, savings bonds, bonds, stocks, non-money market mutual funds, retirement accounts, cash value of life insurance, other managed assets and other financial assets. We set the risk-free rate to 1% and the income tax rate to 20%. In the robustness section, we report results for additional values of the preference parameters. They all deliver very similar conclusions for the welfare gains from introducing alternative student debt contracts.

[12]The SCF may undercount student debt due to the structure of the survey, since the unit of observation is the core economic unit of the household, which may omit adult children living with parents. For a more detailed discussion of the dataset and potential limitations, see Athreya et al. (2021) and Catherine and Yannelis (Forthcoming).

[13]Following Athreya et al. (2021), a loan is delinquent if the household reports in the SCF that they are not making payments on the loan either due to a lack of affordability or because the loan is in forbearance.

4 Baseline Results

In this section, we present results for the baseline model where we consider the current contract structure for student debt: the Standard Repayment Plan with an option to enroll in an Income-Driven Repayment Plan. The specific details of each of these and the implications of delinquency were presented in section 2.2. In the next section, we study the equilibrium and welfare implications of alternative debt contracts.

4.1 Income, Debt and Wealth under SRP, IDRP, or Delinquency

Panel A of Table 3 shows the calibrated utility costs of delinquency and IDRP. To assess the model's performance in matching the targeted moments, Panel B panel shows the fraction of household/year observations where the household is making a student debt payment under the SRP, the IDRP, or is choosing delinquency. For each of these three options, Panel C reports four untargeted moments: median loan balance outstanding, income, debt-to-income, and net wealth. All of the statistics are conditional on households having debt outstanding.

[INSERT TABLE 3 HERE]

In the model, around 55% of the time, households make their scheduled payments according to the standard plan; 29% of the time, agents enroll in the income driven repayment plan; and 16% of the time, they become delinquent. These percentages are very close to their empirical counterparts, respectively, 51%, 31% and 18%.

Naturally agents tend to make payments under the IDRP when their income is lower, such that they can benefit from the corresponding payment reductions. Since there is a switching cost, they also have a greater incentive to use the IDRP when their debt balance is particularly higher. If income is very low and/or loan balances are particularly high, then households choose to become delinquent. In the model, income and loan balance outstanding are the only drivers of the delinquency decision, while in reality households might become delinquent for other reasons, or they might not make a fully rational delinquency decision. This explains why, although the results are qualitatively the same in both the model and in the data, the differences in median loan balance and median income in the delinquency states are more pronounced in the model. Finally, median net wealth is positive for households making repayments under the SRP but negative for those using the IDRP and even more negative for those who are delinquent. This is the case both in the model and in the data.

4.2 Life-Cycle Profiles

In table 4, we provide more detailed summary statistics from our baseline results and further include a breakdown by age group: 25-30, 31-35, 36-40 and 41-65.[14] Panel A reports statistics for households using the SRP, while the corresponding values for those using the IDRP or choosing to become delinquent are presented in panels B and C, respectively. All of the statistics (except for row 2) are conditional on households having debt outstanding.

[INSERT TABLE 4 HERE]

4.2.1 Income, Consumption and Leverage

The first row of table 4 shows the fraction of households enrolling in either a payment plan or becoming delinquent, conditional on having student debt outstanding (as in table 3). These numbers add to 1 across the three panels: households with debt outstanding must either make repayments under one of the two plans or become delinquent. In the second row, we report the percentage of households in each category, relative to the initial population, that is, those households that had student debt at age 25. This allows us to track down the fraction of households that fully repay their loans over time. For example, if we sum the percentages in the three categories for the age group 41-65, we obtain 6.2%, thus, indicating that 93.8% of households have fully repaid their student debt by age 41.

The next two rows of table 4 show the average income and average income growth. Income grows over the lifecycle, particularly early in life; hence, its growth rate is very high for those in the 25-30 age group and it decreases for the others. This pattern is not seen for households that choose to become delinquent (Panel C); instead, we observe growth rates that are negative or close to zero (for all age groups). This is because individuals are naturally more likely to skip a payment and become delinquent after experiencing a negative income shock. Despite the positive growth rates for all age groups (in panels A and B), from age 36, average income is actually flat or even already decreasing. This is because we are

[14] Since the maturity of the student loan is either 10 years (under the SRP) or a maximum of 25 years (under the IDRP), few households still have debt outstanding after age 40, which is why we consider them all as one group.

conditioning our findings on households that have student debt outstanding. Those with higher income growth are more likely to remain under the SRP throughout the payment period and to repay their loans more quickly. As a result, they do not appear in the next age group of this sample. Rows 5 and 6 of Table 4 report mean consumption and mean consumption growth, which follow the same patterns as their income counterparts due to the presence of borrowing constraints.

Younger agents, those under age 31, are more likely to use the IDRP (46.2% compared with 43.1% for the SRP). This is because income at the start of the lifecycle is low and the debt outstanding is still high, thus, making it more beneficial for households to pay the switching cost and enroll in the IDRP. As income grows and the debt is being gradually repaid, households are more likely to remain under the SRP: 65.4% and 69.3%, respectively, in age groups 31-35 and 36-40.

Although the percentage enrollment in the IDRP is lower for households in the 31-35 and 36-40 age groups (23.6% and 13.6%, respectively), it increases again for households in the age 41-65 cohorts. This increase is largely due to selection. As shown in row 2, while 53.7% of households still had student debt at age 36 (across all 3 groups), the number falls to 6.2% at age 41. Households that have not yet paid off their student loans when they reach their forties are likely to have low income and very high debt balances outstanding. Indeed, their income is even lower than for those in the 36-40 age group, even though their average income is growing, as confirmed by the growth rates (row 4). This selection is particularly evident when we look at the debt outstanding numbers (row 7), which are in fact substantially higher for the 41-65 age group than for the age 36-40 cohorts. Although the debt balance can increase due to the possibility of the negative amortization described in section 2.2, most households in the 36-40 age group are under the SRP (under which negative amortization is not possible).[15] These households are also more likely to have become delinquent earlier in the lifecycle. Recall that, in delinquency, interest accrues for missed payments, which further contributes to households' high debt balances.

4.2.2 Delinquency Rates and Debt Repayments

The previously discussed selection process also helps to explain the behavior related to delinquency rates over the lifecycle (panel C). Delinquency rates are initially very stable (comparing the age groups 25-30 and 31-35): early in life, debt balances are higher, while

[15] We describe the negative amortization events below.

both income and savings are lower, thus leading to high rates of delinquency. However, delinquency rates rise again later on and are particularly high for the 41-65 age group, reflecting the relatively lower income and higher debt balances of households that have not yet repaid their student loans by age 41. Those that have been particularly unfortunate in their income shock realizations likely also have low wealth (negative net wealth) and are therefore highly likely to become delinquent (again).[16]

The last four rows of Table 4 report average total debt payments, principal payments, interest payments and the ratio of debt payments to income.[17] As expected, total payments are typically lower under the IDRP, relative to the SRP. The 36-40 age group is an exception, but this is because the majority of agents in this group are close to paying off their debt and, as such, have very low outstanding balances. In fact, the ratio of total debt payments to debt outstanding is 74.5% for those using the SRP, but only 20.8% for those using the IDRP. This pattern of lower payments is particularly useful for agents earlier on in their lifecycle when their income tends to be lower. In fact, the ratio of total payments to income is actually higher for those using the IDRP, reflecting the lower income of households that opt for this payment plan.

A closer examination reveals that the lower total payments under the IDRP result from substantially lower loan amortization, between 30% to 50% lower than under the SRP. By contrast, interest expenses are in fact higher under the IDRP, as we would expect, since the loan balance is also larger, on average.

4.3 Understanding Debt Repayments

In this section, we study households' debt repayment decisions in more detail. Table 5 shows the different debt payments statistics for agents using the SRP (column 2), the IDRP or who are delinquent (column 5). It further separates IDRP repayments with positive or negative amortization (respectively, columns 3 and 4). Note that skipping a payment under delinquency leads to negative amortization, since the unpaid interest is added to the outstanding balance. The first row in Table 5 shows the probability of each event.[18]

[16] We discuss the persistence of delinquency later in the paper.

[17] When a borrow is delinquent, no payment occurs, as shown in rows 9 and 12. Since the debt accrues with interest, this is technically equivalent to a negative principal repayment and a positive interest payment of equal value, hence, the non-zero values in rows 10 and 11.

[18] The probabilities on Table 5 are slightly different than the probabilities shown on Table 3, since we are conditioning on a positive/negative amortization and we lose one period when doing so)

[INSERT TABLE 5 HERE]

Consistent with the previous results, agents choose to make payments under the SRP when they have low debt-to-income ratios. As this ratio increases, they are more likely to choose the IDRP or even to become delinquent. Comparing the actual payments with those that would have happened under the SRP plan (rows 4 and 5 versus rows 6 and 7), we see that switching to the IDRP provides substantial yearly savings. Even for those agents who have positive loan amortization (column 3), the total payment (principal plus interest) is on average 35% smaller.[19]

Negative amortizations under the IDRP tend to take place when debt to income is particularly high. Low individual income imposes a tighter cap on total payments, and a high level of debt implies a higher interest charge. As a result, there is a high probably that the interest expense will exceed the cap, leading to negative amortization. Interestingly, we see that these events are not very frequent. In only 1% of the cases are households achieving negative amortization within the IDRP, and, even as a fraction of the households that use this plan, this still represents only 3.4% of the total.

Table 6 shows the persistence in agents' decisions. The rows report the time t choices and the columns show the decisions at time $t + 1$.

[INSERT TABLE 6 HERE]

We see that these decisions are very persistent. If an agent is making payments under a standard repayment plan at time t with 94% probability, then they will again choose to make a payment under the standard plan at time $t+1$. With 5% probability, they will decide to use an income-driven plan and with 1% probability they will become delinquent. For a household that is enrolled in the IDRP, the degree of persistence is not as extreme but there is still a 62% probability that they will make the same decision the next period, an 28% probability that they will move to a standard repayment plan and an 11% probability that they will choose to be delinquent. Unsurprisingly, delinquency at $t + 1$ is more likely for agents who were previously in the IDRP versus those who remained in the SRP. But the difference in the conditional probabilities is quite large: 11% versus 1%.

Finally, there is also substantial persistence in delinquency rates. An agent who is delinquent at time t has a 76.2% probability of becoming delinquent again at time $t+1$. This helps

[19] This results exclusively from a reduction in the principal repayment, since the agent must always make the full interest payment, otherwise they are delinquent.

to explain why certain individuals still have very high debt balances later in life, with a significant portion of them still choosing delinquency at this stage of the lifecycle, as shown in Table 4: within the group of 41-60 year olds with positive student debt, 44.4% of those choose to be delinquent in any given year.

5 Modified Debt Contracts

In this section, we consider two payment plan modifications motivated by the patterns described in the previous section, especially as they relate to the timing of student debt repayment over the lifecycle. In the standard repayment plan with an option to select into the IDRP, households start repaying their loans early in their lifecycle. From a consumption-smoothing perspective this is highly suboptimal because this is when households' marginal utility of consumption is particularly high, due to the combination of an increasing income profile and borrowing constraints. Therefore, it would be optimal to (partially) defer these payments to a future date. Building on this intuition, we now consider alternative student debt contracts with payment plan modifications designed to deliver a shift in repayments over the lifecycle.

In our baseline calculations, we re-price the loans under each of the proposed modifications, such that the expected net present value (NPV) for the lender remains unchanged. However, we also report the results for two alternative scenarios. In one result, lenders require a higher NPV on the new loans, while in the other we keep the interest rate constant at the current value. These results will simultaneously allow us to decompose the sources of the welfare gains and provide robust evidence for our conclusions.

Although our model assumes a given initial level of student debt, we later show that our conclusions are robust to a scenario where households adjust their student loan size in response to the introduction of these policies.

5.1 Contract Terms

In this subsection, we present the terms of the two proposed student debt contracts. In the next subsection, we discuss the equilibrium pricing of these contracts, followed by a discussion of the new equilibrium and welfare implications.

5.1.1 Principal Payment Deferra (PPD)

The first alternative contract, Principal Payment Deferral (PPD), shifts the original amortization schedule of the loan forward by N_{PPD} periods. In these N_{PPD} periods, the household is still required to make interest payments,which are simply the student loan interest rate multiplied by the initial loan amount. Since interest is being paid, the balance of the loan does not increase over time (in the absence of delinquency). For $t \leq N_{PPD}$ we have

$$I_t^{PPD} = r_s^{PPD} S_t \tag{20}$$

$$P_t^{PPD} = 0, \tag{21}$$

where r_s^{PPD} is the interest rate on student loans under the PPD contract.

After the initial interest-only periods, the household can choose to make principal and interest payments under the standard repayment plan (equations (1) to (3)) or to pay the switching cost (ξ^{IRDP}) and enroll in the income-driven plan (equations (4) to (6)); that is, the contract reverts back to the current one. As before, in the case of delinquency, the loan balance increases by the value of the missed interest payments (equation (9)). In our analysis below, we set $N_{PPD} = 10$.

5.1.2 Full Payment Deferral (FPD)

The second alternative contract, Full Payment Deferral (FPD), defers both principal and interest payments for N_{FDP} periods. The initial interest payments are not forgiven, they are just deferred. In these periods, the household is still charged an interest payment as in the Principal Payment Deferral contract. However, these are not actually paid, but instead added to the principal of the loan. For $t \leq N_{FPD}$ we have

$$I_t^{FPD} = r_s^{FPD} S_t \tag{22}$$

$$P_t^{FPD} = -I_t^{FPD}, \tag{23}$$

where r_s^{FPD} is the interest rate on student loans under the FPD contract.

As a result, after the initial N_{FPD} deferred-payment periods, the new loan balance becomes the initial balance multiplied by $(1 + r_s)^{N_{FPD}}$:

$$S_{N_{FPD}}^{FPD} = (1 + r_s^{FPD})^{N_{FPD}} S_0 \tag{24}$$

At this point, the payment schedule is re-calculated by using the new loan balance and the debt contract reverts back to the baseline set-up: principal and interest payments according to the Standard Repayment Plan (equations (1) to (3)), with the option to pay the switching cost (ξ^{IRDP}) and enroll in the Income Driven Repayment Plan (equations (4) to (6)). If the agent chooses to become delinquent, then the loan balance increases by the value of the missed interest payments (equation (9)). In our analysis below, we set $N_{FPD} = 10$.

5.2 Equilibrium Loan Premia

In this section we describe how we compute the loan premium for each of the modified contracts: the PPD, φ^{PPD}, and the FPD, φ^{FPD}.

We assume that the lender is risk-neutral, or can fully diversify the cash-flow risk (namely delinquency risk) associated with the different repayment schedules. Therefore, we discount all cash flows at the riskless rate. We make this assumption because the lender is typically the federal government, which is much better able to diversify this risk than private lenders. If we incorporate a risk premium in our calculations, then the welfare gain from the introduction of the new contracts would be even higher, since they imply lower delinquency rates, as discussed below.

5.2.1 Baseline Case

In our baseline case, we price the new debt contracts such that they deliver the same NPV as the SRP/IDRP contract. We take this NPV as the normal level of revenue that the lender, typically the federal government, requires to cover the costs of originating and administrating these loans. Therefore, we impose that the contracts must generate the same level of revenue. Under this assumption, we compute the equilibrium loan rates for each new contract using the following fixed point algorithm:

(i) Compute the implied average NPV on those loans ($NPV^{Baseline}$) by simulating the model under the SRP/IDRP plan and discounting cash-flows at the risk-free rate.

(ii) For each of the two new equilibria—the equilibrium with the PPD and that with the FPD contracts—simulate the economy by using the same premia as in the baseline economy (φ^{FPD})) and compute the implied NPV for the lender: $NPV^i(\varphi^{Baseline})$, for $i \in \{\text{PPD}, \text{FPD}\}$

(iii) If the resulting NPV ($NPV^i(\varphi^{Baseline})$) is lower (higher) than the target ($NPV^{Baseline}$), construct a sparse 10-point loan premia grid with higher (lower) loan premia.

(iv) Compute the NPV of the loans for each of the 10 new values for the loan premia and pick the premia that delivers the NPV closest to the target.

(v) Repeat (iii) and (iv) until convergence.

Table 7 shows the equilibrium loan premia for each of the two new contracts, the PPD and the FPD.

[INSERT TABLE 7 HERE]

For both policies under consideration, the equilibrium premia are lower than under the baseline contract. This is because these loans are outstanding for longer and, as a result, they accrue higher total interest payments. The reduction in interest rates is more significant under the FPD contract because, under this scenario, during the deferral period the loan balances are increasing at a rate (the loan premium) that is higher than the discount rate. As a result of this premium, extending the maturity leads to a higher NPV. Therefore, the same NPV can be obtained with a lower loan interest rate.

The decreases in the interest rate are economically large, ranging from 1.77% for the PPD contract to 2.00% for the FPD, but it is important to highlight that the welfare gains reported below are only partially driven by this repricing of the loans. In fact, if we fix the interest rate at the baseline level, the majority of the utility gains remains.

5.2.2 Alternative Cases

As discussed above, in our baseline calculations we assume that the lender has a target level of revenue that corresponds to the one obtained under the baseline contract ($NPV^{Baseline}$) and, therefore, all of the alternative contracts must deliver the same expected discounted cash flow. However, there are three potential considerations that might imply a different assumption.

First, there might be differences in cash-flow risk and, in particular, delinquency and missed payment risk, in the economies with the different contracts. If the lender cannot fully diversify this risk, then we should not match the NPVs that are discounted at the riskless rate. However, as we show below, the new contracts actually deliver lower delinquency rates, so, if we were to include a risk premium in the discount rate, this would deliver a higher NPV than the baseline case. Therefore, the fixed point algorithm would imply an even larger difference between the equilibrium loan rate for the new contracts and the baseline rate, leading to even larger welfare gains.

Second, under the new contracts, the loans remain in existence for a longer period and this might increase the administrative costs for the lender.

Third, since the loans are (on average) paid off over a longer period, the discount rate should increase to reflect an additional term premium.[20]

To take the last two considerations into account, we consider a case where the required NPV on the new contracts is 10% higher than the corresponding NPV in the baseline economy ($NPV^{Baseline}$).

Finally, we also report the results for an extreme case where we keep the interest rate on the loans constant at baseline level ($r_f+\varphi^{Baseline}$). This will serve as a maximum conservative assumption on the pricing of the new loans and will allow us to decompose the sources of the welfare gains.

5.3 Results

We now present the results in the two alternative economies: the one with the principal payment deferred (PPD) debt contracts (hereafter PPD-economy), and the one with the full payment deferred (FPD) debt contracts (hereafter FPD-economy). Under PPD the loan amortization schedule is shifted for $N_{PPD} = 10$ years. During those first 10 years, households are only required to pay interest. In the FPD scheme, agents defer both loan interest and principal payments for $N_{FPD} = 10$ years. During this period, debt outstanding increases over time as the "missing" interest payments are being added.

Table 8 shows statistics conditional on agents having student debt outstanding, as in the previous sections. We present results for the full lifecycle later on. To facilitate the comparison across the different scenarios, panel A shows the corresponding statistics for the baseline case. Panel B reports the results for the economy with the PPD contracts, while panel C presents the results for the economy with the FPD contracts. Within each panel,

[20]It is important to note that, although a 10-year extension is significant, the loan maturities in the baseline contract are already between 10 (under the SRP) and 25 years (under the IDRP). So, with the 10-year deferral, this is now being extended to between 20 and 35 years. The term structure of the interest rates at these very high maturities is relatively flat, as shown in Augustin et al. (2021), for example.

we report the average across all ages and results for three different age groups that capture important stages of the lifecycle in the different economies. The first group covers agents aged 26-35 years in the period during which agents in the PPD and FPD economies enjoy the (partial) deferral of debt payments. In the second group, those aged 36-40, agents are making substantial debt repayments in all economies. Finally, after age 40, most individuals in the baseline economy have already repaid their student loans, but those in the PPD and FPD economies have not yet done so, due to the initial deferral period.

[INSERT TABLE 8 HERE]

5.3.1 Debt Repayments

Under the baseline economy, most households have repaid their debt by age 41, as previously discussed. By contrast, under PPD and FPD contracts, younger agents are spared from making (large) repayments; therefore, we have a much larger fraction of the population still with debt outstanding after age 40. Usage of the IDRP option is much less common under the alternative contracts: it falls from 28.5% in the baseline case to 10.0% and 10.9% in the PPD and FPD economies, respectively. This is due to the fact that households are no longer required to make large debt payments early in life when their incomes are still low. We saw in Table 4 that, in the baseline economy, the IDRP option is mostly used by those in the 26-30 age group (46.2%) and much less after that. Under the alternative contracts, these agents are enjoying the deferral period and, as such, only have to make either (much smaller) interest payments (PPD economy) or no debt payments at all (FPD economy). During the first 10 years of the lifecycle, under the PPD debt contracts, agents are only making interest payments and, therefore, never choose to enroll in the IDRP.[21]

5.3.2 Net Wealth

Since (the main) debt repayments in the PPD and FPD economies take place later in life, they happen when individuals are earning higher incomes and after they have had time to accumulate a more-substantial level of wealth. Not only is agents' wealth accumulation typically higher late in life but, since they have not been forced to make (large) debt repayments early on, they were also able to save more in the two alternative economies. As a result, net wealth is substantially higher for agents in age groups 36-40 and 41-65 under the PPD and

[21]Under the FPD contracts, agents make no payments before age 35.

FPD economies, even though their debt outstanding is also higher (as expected, since it is only being repaid now).

5.3.3 Delinquency Rates

[INSERT FIGURE 2 HERE]

The combination of higher income and higher wealth when the principal repayments are due leads to substantial lower delinquency rates with the alternative contracts. Figure 2 shows that delinquencies fall by about one-third in the PPD economy and by about one-half in the FPD economy. The large reductions in the delinquency rates represent important benefits, for both borrowers and lenders, from the two alternative contracts.[22] These lower delinquency rates also suggest that the equilibrium loan premium on these contracts might even be lower than what we have assumed by imposing the same NPV as for the baseline contract.

5.3.4 Debt Outstanding

Under the baseline contract, leverage falls quite rapidly early on as agents repay their loans and it increases again for the last age group (41-65) because of the previously discussed sample selection: the few individuals who still have debt after age 40 are mostly those who became delinquent (or negatively amortized) in the past. By contrast, under the PPD contracts, leverage will remain largely constant until age 35, only increasing slightly due to the occasional delinquency events.[23]. After age 35, the repayments start and we observe a gradual reduction of the loan amount outstanding until retirement. In the economy with FPD contracts, leverage is actually increasing until age 35 as the accrued interest is being added to the principal.[24]. As a result, the loan amount outstanding is still higher for agents between ages 36 and 40 and only decreases after that.

As previously discussed, in the baseline economy, the average loan amounts for individuals with debt outstanding actually increase substantially for the age group 41-65. Interestingly, this does not happen in either the PPD or FPD economies. Likewise, although

[22]We report explicit welfare gains later in the paper.

[23]The principal payments are slightly negative, reflecting the few delinquencies that occur during this period.

[24]The principal payments are negative and equal, in absolute value, to the interest payments

the delinquency rates for this age group increase in all three scenarios, they are 57% and 65% lower with the PPD and FPD contracts, respectively. In these two economies, we do not have the same strong selection as in the baseline case where only agents with past delinquencies and/or low income reach age 41 with student debt still outstanding. For the same reason, in the alternative economies, we have a much smaller fraction of agents using the IDRP contracts at this stage of the lifecycle.

5.3.5 Consumption

In both the PPD and FPD economies, agents have much higher average consumption while their student debt outstanding is 23% and 32% higher than consumption in the PPD and FPD economies, respectively. The qualitative pattern was expected since earlier in the lifecycle agents either make no payments at all (deferred payments) or make only interest payments (interest only). Also, since their loans are extended for longer periods, the statistics for the PPD and FPD economies are also capturing households later in life, when their income is on average higher. Even when we condition on age, the differences in consumption remain very large, but we are only considering agents with positive student debt here. In the next section, we provide a more direct comparison of the full consumption profiles in the different economies.

5.4 Welfare Analysis

5.4.1 Consumption Over the Lifecycle

The PPD and FPD contracts benefit individuals early in life by deferring (most of) their student debt repayments. Relative to the baseline contract, this will allow them to increase consumption early in life but this should be reflected in lower consumption late in life when debt is finally being amortized. In the previous section, we reported statistics for households with positive debt balances only. In this section, we track all households through their lifecycle to capture this important trade-off. More precisely, Table 9 reports consumption, net wealth and debt outstanding over the lifecycle, without conditioning on positive debt balances.

[INSERT TABLE 9 & FIGURE 3 HERE]

Table 9 confirms that, under the PPD and FPD contracts, households are better able to smooth consumption over the lifecycle by increasing it early in life at expense of lower consumption late in life. Figure 3 plots the percentage difference of consumption and savings over the lifecycle for PPD and FPD contracts, relative to the baseline model. The gains early in life are quite substantial under both alternative contracts. Under the FPD (PPD) contract, average consumption at ages 26-30 is 6.1% (4.2%) higher and average bond saving increases by 6.5% (8.7%). At ages 31-35, consumption gains are again quite sizeable, 2.2% (1.5%), while bond savings skyrocket by 36.1% (29.9%). Average consumption relative to the baseline turns negative for those aged 36-40, while average bond saving remains very elevated relative to the baseline. These patterns demonstrate that the household's immediate concern is to increase consumption and then increase wealth buffers. This is exactly why liquidity early in the lifecycle is so important: first, to immediately increase consumption, and second, to insure against negative income shocks by building up wealth.

By comparison, the reductions in average consumption later in life are much smaller. From age 36 to 65, the differences in consumption relative to the baseline case are between 0.2% and 1.0% under the PPD contracts and between 0.4% and 1.2% with the FPD contracts. Crucially, these decreases in consumption late in life are less important in marginal utility terms, since average consumption is now much higher than it was before age 35.

Differences in wealth accumulation for retirement are also quite small. Relative to the baseline economy, we see that average wealth in the 61-65 age group is only 1.7% lower in the PPD economy and only 2.6% lower in the FPD economy. The implied differences in consumption at retirement are even smaller, 0.04% and 0.06% lower, respectively, since it is also being financed by social security payments.

In addition to ignoring the differences in the marginal utility of consumption at different ages, these comparisons also ignore risk. By reducing (or even eliminating) debt payments early in life, the PPD and FPD contracts allow agents to better smooth income shocks exactly at the stage of their lifecycle during which they are more vulnerable to such shocks, since they have not yet had the opportunity to build a significant buffer stock of wealth. In the next section, we explicitly take these into account by measuring the certainty equivalents associated with each debt contract.

5.4.2 Welfare Gains

As shown in Table 9, the PPD and FPD contracts allow agents to increase consumption early in life, when marginal utility is highest, at the expense of lower consumption late in life, when marginal utility is lower. In addition, they allow for better consumption smoothing against income shocks, since younger agents have less wealth and are therefore more vulnerable to these shocks.

In this section, we formally quantify the utility gains of the alternative debt contracts by measuring the certainty equivalent consumption level associated with each type of contract and comparing these gains with the certainty equivalent consumption level obtained under the baseline economy. The corresponding percentage gains are reported in Table 10. To shed light on the source of the welfare gains, we report the results with the contracts priced under the equilibrium loan premium in each economy (see Table 7) and the results with the loan premium fixed at the value in the baseline economy. Furthermore, we also report the welfare gains resulting from a simple 10-year maturity extension of the existing contract.

[INSERT TABLE 10 HERE]

The welfare gains from the proposed contract modifications are economically large. The certainty equivalent consumption gain from moving to an economy with PPD debt contracts is 0.62% per year. The gain from moving to an economy with FPD contracts is an even larger 1.78% per year. For comparison, these values are similar to the welfare gains obtained for stock market participation in the context of similar life-cycle models (see Cocco et al. (2005)).

As shown in Table 7, the PPD and FPD economies are characterized by lower interest rate premia on student loans. It is therefore important to understand how much of the welfare gains result from having these lower interest rates versus the deferral of repayments. To answer this question, the first column of Table 10 reports the welfare gains for households in a counterfactual economy where the interest rate is kept at the (higher) baseline value.[25]

Without the interest rate reduction, the gains in the PPD economy are smaller: 0.63% versus 1.35%. On the other hand, the welfare gains from switching to the FPD contracts are barely affected: 1.96% versus 2.35%. These results highlight the importance of shifting debt repayments from the early stage of the lifecycle to later in life, when agents are earning higher incomes and have accumulated more wealth. Under the FPD contracts, all payments

[25] The higher interest rate in this counterfactual economy is reflected in larger profits for lenders, which is why we emphasize that we are now comparing household welfare only.

are deferred. As a result, the total welfare gains are three times larger and the reduction in the loan interest rate is largely irrelevant. With the PPD contracts, the total gains are still sizeable but, since the agent is still making interest payments early in life, the reduction in the loan interest rate plays a much more important role.

Finally, we can also compare the welfare gains of the PPD and FPD contracts to those obtained when the contract maturity is extended by 10 years, but we keep all other features unchanged.[26] This maturity extension allows individuals to delay full debt repayments until 35 years later, as in our proposed modifications. However, repayment of both principal and interest start in year one as under the current contract. As we can see in Table 10, although extending the maturity of the loans is welfare improving, the gains are only 51% (29%) of those obtained under the PPD (FPD). If we wish to replicate the certainty equivalent gains of these two contracts, we would have to provide individuals with a much more substantial maturity extension.

5.5 Robustness

In this section, we show that our conclusions regarding the welfare gains from the PPD and FPD contracts are robust to alternative assumptions about the required NPV of the lenders or the preference parameters of the borrowers. We also show that the gains are larger if we augment the model to include an endogenous stock market participation decision, and that the welfare gains are not very sensitive for allowing agents to borrow from private sources of debt.

5.5.1 Relaxing the Net Present Value Assumptions

In our previous analysis, we computed the equilibrium loan premium for each proposed modification (φ^{PPD}) and (φ^{FPD}), by imposing the condition that the NPV of the loan remains the same as under the original contract formulation ($NPV^{Baseline}$). However, as discussed in section 5.2.2, it is possible that the new debt contracts might be associated with a different equilibrium NPV.

The positive NPV on the loans is presumably compensation for the costs associated with the loan provision and subsequent maintenance.[27] The costs of loan origination should re-

[26] We also reprice these loans and the corresponding equilibrium interest rate is 2.02%.

[27] This could also result from market power if these are provided by private lenders and not the federal

main unchanged but the maintenance costs might increase, given that we are extending the duration of the loan. In addition, it is possible that lenders might require a higher term premium as compensation for the higher loan maturity. Both of these arguments would suggest higher NPV requirements for the PPD and FPD loans, hence a higher loan premia. On the other hand, as shown in section 5.3, the new contracts are associated with significantly lower delinquency rates implied by the new contracts. This in turn would imply lower equilibrium loan premia.

As a conservative robustness exercise, we consider the case in which the impact of the increased costs dominates the reduction in risk and, therefore, lenders require a higher NPV to originate the loans. More precisely, we assume that the new NPV must be 10% higher than the one in the baseline economy. With this assumption, we recompute the calculation of the equilibrium loan premium for the PPD and FPD economies by using the algorithm described in 5.2.1 but with the target in step iii) now set at $1.1 * NPV^{Baseline}$. Panel A from Table 11 reports the corresponding endogenous loan premia in each of the two economies.

[INSERT TABLE 11 HERE]

With the new equilibrium loan premia, we can again study the outcomes in the corresponding PPD and FPD economies. In particular, we are interested in the corresponding welfare gains, which we report in panel B of Table 11. As before, we can decompose the welfare gains in two sources: (i) a welfare gain coming from the deferral of the payments and (ii) a welfare gain coming from the endogenous change in the interest rate.

Compared to the values reported in Table 10, the overall welfare gains are naturally smaller but the differences are minimal. For the PPD contract, the yearly certainty equivalent consumption gain falls from 1.35% to 1.30%. For the FPD contract, the welfare gain falls slightly less, from 2.36% to 2.33%. This is consistent with the previous result showing that the welfare gains of this contract are almost exclusively due to the deferral of payments rather than from the reduction in the loan interest rate.

5.5.2 Alternative Preference Parameters

We have shown that the welfare gains from introducing the new contracts are largely driven by the ability to shift payments over the lifecycle. Therefore, these gains should vary across

government.

agents with different discount factors, or different elasticities of intertemporal substitution. Table 12 shows the welfare gains for different values of these two parameters.

[INSERT TABLE 12 HERE]

In all cases that we consider, the welfare gains from both the PPD and, especially, the FPD contracts remain economically large.

Compared to the baseline calibration, the welfare gains of the policies increase (decrease) when households are more (less) impatient, that is, when they have a lower (higher) subjective discount factor. More-impatient households value consumption today relatively more than consumption tomorrow. Therefore, policies such as those with PPDs and FPDs that allow indebted students to have higher consumption earlier in the lifecycle bring larger welfare gains.

When agents have a lower EIS, they are less willing to substitute consumption across time, so policies in this scenario have higher welfare gains compared to the baseline. In contrast, with a higher EIS, households are happier with having lower consumption today (and lower utility) to guarantee higher consumption tomorrow. In this scenario, policies that allow agents to defer consumption are less beneficial.

5.5.3 Different Levels of Initial Student Debt

Even though we do not model the households' endogenous education choice, our model can still help in understanding the effects of our policies for different values of the loan-to-income ratio. These different initial values of wealth proxy for households' different major choices, different lifetime earnings, or differential levels of financial help from their families. This analysis therefore serves two purposes. First it documents which groups of households will benefit more/less from these alternative contracts. Second, it shows the robustness of the welfare gains in a scenario where households might change their initial level of debt in response to the introduction of such contracts.

[INSERT TABLE 13 HERE]

We solve our model for different levels of initial debt while keeping all other parameters unchanged. In particular, we solve the model for a 25% lower and 25% higher initial level of debt. Table 13 reports the welfare gains of the two policies we consider (PPD and FPD) for the different levels of initial debt. The welfare gains are monotonically increasing with the initial level of debt. Households with higher initial balances have greater benefits from having the option of deferring principal payments or all the payments, whereas households with lower initial balances have the least benefits. Therefore, these types of policies are particularly helpful for students with no financial help from their families, who are likely to have higher balances and who enroll in majors with lower lifetime earnings and, thus, start their working life with higher loan-to-income ratios.

Furthermore, the gains remain economically large in all cases, ranging from 1.13% to 1.44% for the PPD and from 1.83% to 2.64% for the FPD. This confirms that even if households respond to these policies by taking on different initial levels of debt, they will still benefit substantially from their introduction.

5.5.4 Additional Sources of Unsecured Debt

In reality, households with student debt may also choose to take on additional unsecured debt in the form of credit cards or lines of credit. If households can borrow to increase disposable income, this may detract from the welfare gains of the proposed plans. To study this, we relax the borrowing constraint in the model by allowing our agents to borrow at the risk-free rate plus a premium ($r_f + \varphi^{PrivateLoan}$), up to a borrowing limit $\bar{\kappa}$. We calibrate the borrowing premium over the risk-free rate to 8% and the borrowing limit to be 30% of average income over the life-cycle, and note that results we present here are not very sensitive to the choice of these two parameters.

Relative to the benchmark model, we find that allowing agents the ability to borrow in unsecured markets increases their welfare by 1.12% in consumption equivalent units per year. On top of relaxing borrowing constraints, allowing them to either defer interest payments (PPD) or all payments (FPD) for 10-years, yields additional welfare gains of 1.31% and 2.21%. Thus even when we allow for other sources of debt, our policies deliver similar welfare gains to the 1.35% (PPD) and 2.36% (FPD) in the baseline scenario.

5.5.5 Stock Market Participation

High student debt payments early in the lifecycle of agents may delay some important decisions such as the decision to buy a house or enter the stock market. Insofar as our proposed plans defer payments and increase liquidity early in the lifecycle, they will have an impact on the timing of decisions that are made early in the lifecycle. In this section, we extend the model to consider an endogenous stock market participation decision and measure the welfare benefits of our proposed plans.

As in Gomes and Michaelides (2008), we assume that households must pay an entry cost to invest in stocks for the first time. Following their calibration, we set the cost to 5% of average annual income. To avoid introducing one additional choice variable in the model, we abstract from the optimal portfolio choice and endow agents with a $60\%/40\%$ stock/bond portfolio throughout their working life. Naturally, if we allowed agents to choose their optimal portfolio every year, the benefits from stock market participation would be even larger. [28].

The risky asset has a gross return of r_t^{stock} and its excess return is given by

$$r_{t+1}^{stock} - r_f = \mu^{stock} + \sigma^{stock}\epsilon_{t+1}, \tag{25}$$

where r_f is the risk-free rate, μ^{stock} is the excess average stock market return, which we calibrate to be 6%, σ^{stock} is the volatility of the stock market return, with we calibrate to be 15.7%, and ϵ_{t+1} is the period $t+1$ innovation to excess returns, which we assume to be independently and identically distributed (i.i.d.) over time and distributed as $N(0,1)$.

[INSERT TABLE 14 HERE]

Table 14 reports the results. Under the PPD/FPD contracts, households start participating in the stock market almost a year earlier. Their lower commitments to debt repayments allow them to accumulate more wealth early in life and, therefore, to pay the entry cost sooner. They also enter the stock market at higher levels of leverage and lower levels of income and net wealth. Finally, the welfare gains when we allow for stock market participation are even higher than in the base model. The PPD contract yields a welfare gain of 1.65% (vs 1.35% under the base model) and the FPD policy contract provides a welfare gain of 2.36% (vs 2.71% under the base model).

6 Comparison to the Biden Administration's Proposals

Our model setup can also help us understand the welfare benefits of the Biden administration's proposed modifications to student loan repayment plans. The proposals center

[28] During retirement, we assume agents only hold bonds. This is also a conservative assumption, since it reduces the benefits of stock market participation.

around two major features: outright debt forgiveness and modifications to the IDRP plans.

6.1 Outright Debt Forgiveness

The first part of the Biden administration's proposal is a one-time debt relief by the U.S. Department of Education that forgives between $10,000 and $20,000 of debt for borrowers who meet certain income requirements.[29] The estimated fiscal cost of this proposal is $400 billion. We use our model to evaluate the level of debt forgiveness that makes borrowers indifferent between debt forgiveness and the PPD and FPD modifications. In other words, we calculate the level of debt discharged that yields the same welfare as under the PPD and FPD modifications. Since our model does not take into account the potential adverse impacts from a reduction in government revenue associated with the official policy proposal, our calculation provides a conservative estimate of the debt-forgiveness equivalent for each of the two policies.[30]

[INSERT TABLE 15 HERE]

Panel A of Table 15 reports our findings. Households in our model derive the same welfare from the PPD proposal as $11,300 in debt forgiveness and the same welfare from the FPD proposal as $17,500 in debt forgiveness. Given these figures and the number of households with outstanding student debt, a back-of-the-envelope calculation yields that the fiscal cost of obtaining the same welfare as the PPD or FPD proposals is between $450 billion and $700 billion. This is slightly higher than the estimated cost of the Biden proposal because, in our model, the PPD and FPD options are presented to all borrowers and not just those beneath an income threshold. This highlights that our policies are able to deliver the same magnitude of welfare gains as the debt forgiveness proposals currently being debated while also remaining budget neutral for the government.

[29]Details of the current proposal can be found here: https://studentaid.gov/manage-loans/forgiveness-cancellation/debt-relief-info.

[30]By reducing government revenues, debt forgiveness will imply a combination of an increase in taxes, a reduction in government spending, and additional government debt or an increase in the monetary base. To the extent that these adjustments will reduce household welfare, our calculation overestimates the benefits of debt reduction and, therefore, underestimates the debt forgiveness equivalent for each of the proposed contracts.

6.2 Modifications to IDRP Plans

The Biden administration has proposed two main changes to the income-driven repayment plans. First, decreasing the repayment period before which outstanding debt is forgiven to 10 years, which in our model maps to changing N_{IDRP} to 10 years, instead of the current 25 years. Second, decreasing payments uniformly by making two changes to their construction: (a) increasing the means-testing threshold of the Federal Poverty Level to 225%, and (b) setting payments equal to only 5% of discretionary income. In the model, (a) and (b) correspond to changing ω_{FPL} and ω_{IDRP}, respectively. We analyze each of these proposed changes in turn, which allows us to understand their welfare benefits while keeping the federal government budget unchanged (i.e., repricing the loans to deliver the same NPVs).

Panel B of Table 15 reports the welfare gains for each modification relative to the baseline. The first modification decreases the number of years to forgiveness from 25 to 10 years, meaning that after 10 years of payments, all debt is discharged at no cost. This change has a significant impact on the profitability of the loans, thus, the interest rate on student loans must increase to 9.7% to ensure that the loans have the same profitability. This change induces a welfare gain of 0.31%, much lower than the welfare gains of our proposed PPD and FPD policies of 1.35% and 2.36%, respectively. The second proposal significantly decreases payments made under the IDRP, which yields larger welfare gains that are closer to those from our proposed policies. After repricing the loans, this proposal yields a welfare gain of 1.94%, which lies between the welfare gains of our PPD and FPD proposals.

7 Conclusions

We build a quantitative life-cycle model of consumption and saving to study the impact of student debt repayment plans. We calibrate the model to generate behavior consistent with observed patterns on enrolment in the standard and income-driven repayment plans that are currently offered. We consider two modifications to each plan. The first, "Principal Payment Deferral" (PPD), defers principal payments for 10 years. The second, "Full Payment Deferral" (FPD), defers all payments for 10 years and, during this time, the deferred interest is added to the principal.

These alternative plans lead to significant welfare gains of 1.30% in yearly certainty equivalent consumption for the PPD contracts and 2.33% for the FPD contracts. By comparison, these are on par with the welfare gains from currently proposed debt forgiveness plans, even ignoring the potential costs from the additional taxes or lower government spending that will be required to balance the federal budget under the official proposals. These gains come primarily from postponing payments early in the lifecycle, when margin utility is high, to later in the lifecycle, when the household has had the opportunity to accumulate wealth. Under the current plans, households make large payments early in life instead of accumulating wealth. Not only does this reduce consumption in each period but it also reduces households' wealth accumulation and the ability to smooth consumption across periods.

Although the FPD plan yields the largest gains, it is possible that the PPD contract is more appealing from a behavioral and/or political perspective, by avoiding a large jump in household liabilities at mid-life. The welfare gains in our analysis likely understate the true gains, since we abstract from job searches and other financial decisions such as housing and family planning. Since student debt is difficult to renegotiate and is essentially nondischargeable in bankruptcy, student debt repayments crowd out these other financial decisions. Future work that better studies the adjustments along these additional dimensions will reveal the true benefits of the policies we propose.

References

Abbott, Brant, Giovanni Gallipoli, Costas Meghir, and Giovanni L. Violante, 2019, Education policy and intergenerational transfers in equilibrium, *Journal of Political Economy* 127.

Abraham, Katharine, Emel Filiz-Ozbay, Erkut Ozbay, , and Lesley Turner, 2020, Designed to fail: Effects of the default option and information complexity on student loan repayment, *Journal of Public Economics* 192.

Amromin, Gene, and Janice Eberly, 2016, Education financing and student lending, *Annual Review of Financial Economics* 8, 289–315.

Amromin, Gene, Janice Eberly, and John Mondragon, 2017, The housing crisis and the rise in student loans, Working paper.

Athreya, Kartik, Christopher Herrington, Felicia Ionescu, and Urvi Neelakantan, 2021, *Student Loan Borrowing and Repayment Decisions: Risks and contingencies 1*, 458–511 (Routledge).

Athreya, Kartik, Felicia Ionescu, Urvi Neelakantan, Jessie Romero, and Ivan Vidangos, 2020, Who values access to college?, Economic Brief 20-03, Federal Reserve Bank of Richmond (U.S.).

Augustin, Patrick, Mikhail Chernov, Lukas Schmid, and Dongho Song, 2021, Benchmark interest rates when the government is risky, *Journal of Financial Economics* 140, 74–100.

Avery, Christopher, and Sarah Turner, 2012, Student loans: Do college students borrow too much—or not enough?, *Journal of Economic Perspectives* 26, 165–192.

Bleemer, Zachary, Meta Brown, Donghoon Lee, Katherine Strair, and Wilbert van der Klaauw, 2017, Echoes of rising tuition in students' borrowing, education attainment, and homeownership in post-recession america, Staff Reports 280, Federal Reserve Bank of New York.

Britton, Jack W., and Jonathan Gruber, 2019, Do income contingent student loan programs distort earnings? evidence from the uk, Working paper 25822, National Bureau of Economic Research.

Cadena, Brian, and Benjamin Keys, 2013, Can self-control explain avoiding free money? evidence from interest-free student loans, *Review of Economics and Statistics* 95, 1117–1129.

Catherine, Sylvain, and Constantine Yannelis, Forthcoming, The distributional effects of student loan forgiveness, *Journal of Financial Economics* .

Cocco, Joao F, Francisco J Gomes, and Pascal J Maenhout, 2005, Consumption and portfolio choice over the life cycle, *The Review of Financial Studies* 18, 491–533.

Cox, James, Daniel Kreisman, and Susan Dynarski, 2020, Framing effects, earnings expectations, and the design of student loan repayment schemes, *Journal of Public Economics* 183.

Dinerstein, Michael, Constantine Yannelis, and Ching-Tse Chen, 2023, Debt moratoria: Evidence from student loan forbearance, Working paper.

Folch, Marc, and Luca Mazzone, 2022, Go big or buy a home: The impact of student debt on career and housing choices, *Available at SSRN 3805220* .

Gallagher, Emily A., Stephen B. Billings, and Lowell Ricketts, 2022, Human capital investment after the storm, Working paper.

Ganong, Peter, and Pascal Noel, 2020, Liquidity versus wealth in household debt obligations: Evidence from housing policy in the great recession, *American Economic Review* 110, 3100–3138.

Gary-Bobo, Robert, and Alain Trannoy, 2015, Optimal student loans and graduate tax under moral hazard and adverse selection, *The RAND Journal of Economics* 46, 546–576.

Gomes, Francisco, and Alexander Michaelides, 2005, Optimal life-cycle asset allocation: Understanding the empirical evidence, *The Journal of Finance* 60, 869–904.

Gomes, Francisco, and Alexander Michaelides, 2008, Asset pricing with limited risk sharing and heterogeneous agents, *The Review of Financial Studies* 21, 415–448.

Goodman, Sarena, Adam Isen, and Constantine Yannelis, 2021, A day late and a dollar short: Liquidity and household formation among student borrowers, *Journal of Financial Economics* 142, 1301–1323.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, 2021, What do data on millions of US workers reveal about lifecycle earnings dynamics?, *Econometrica* 89, 2303–2339.

Hampole, Menaka V, 2022, Student debt and labor market trajectories .

Herbst, Daniel, Forthcoming, The impact of income-driven repayment on student borrower outcomes, *American Economic Journal: Applied Economics* .

Ionescu, Felicia, and Nicole Simpson, 2016, Default risk and private student loans: Implications for higher education policies, *Journal of Economic Dynamics and Control* 64, 119–147.

Jacob, Brian, Damon Jones, and Benjamin J. Keys, 2023, The value of student debt relief and the role of administrative barriers: Evidence from the teacher loan forgiveness program, Technical Report 31359, NBER Working Paper Series.

Karamcheva, Nadia, Jeffrey Perry, and Constantine Yannelis, 2020, Income-driven repayment plans for student loans, Working paper 2020-02, Congressional Budget Office.

Lochner, Lance, and Alexander Monge-Naranjo, 2011, The nature of credit constraints and human capital, *American Economic Review* 101, 2487–2529.

Lochner, Lance, and Alexander Monge-Naranjo, 2016, *Student Loans and Repayment: Theory, Evidence and Policy*, volume 5, 397–478 (Elsevier).

Lochner, Lance, Todd Stinebrickner, and Utku Suleymanoglu, 2021, Parental support, savings, and student loan repayment, *American Economic Journal: Economic Policy* 13, 329–371.

Lucca, David O., Taylor Nadauld, and Karen Shen, 2019, Credit supply and the rise in college tuition: Evidence from the expansion in federal student aid programs, *The Review of Financial Studies* 32, 426–4666.

Maggio, Marco Di, Ankit Kalda, and Vincent Yao, 2020, Second chance: Life without student debt, Technical Report 25810, NBER Working Paper Series.

Mueller, Holger, and Constantine Yannelis, 2022, Increasing enrollment in income-driven student loan repayment plans, *Journal of Finance* 77, 367–402.

Palacios, Miguel, 2015, Human capital as an asset class: Implications from a general equilibrium model, *The Review of Financial Studies* 28, 978–1023.

Sun, Stephen, and Constantine Yannelis, 2016, Credit constraints and demand for higher education: Evidence from financial deregulation, *Review of Economics and Statistics* 98, 12–24.

Tauchen, George, and Robert Hussey, 1991, Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models, *Econometrica: Journal of the Econometric Society* 371–396.

Yannelis, Constantine, and Adam Looney, 2022, The consequences of student loan credit expansions: Evidence from three decades of default cycles, *Journal of Financial Economics* 143, 771–793.

A Figures

Figure 1: Student Debt in the United States

Panel A shows the total student debt outstanding in the U.S. between Q1 2006 and Q4 2021. Panel B plots the ratio of student debt outstanding to nominal GDP. The student debt and GDP data are from the Federal Reserve Economic Data (FRED). Panel C plots the number of people with student debt outstanding and Panel D reports the average outstanding balance per recipient. The data comes from the Office of Federal Student Aid (FSA).

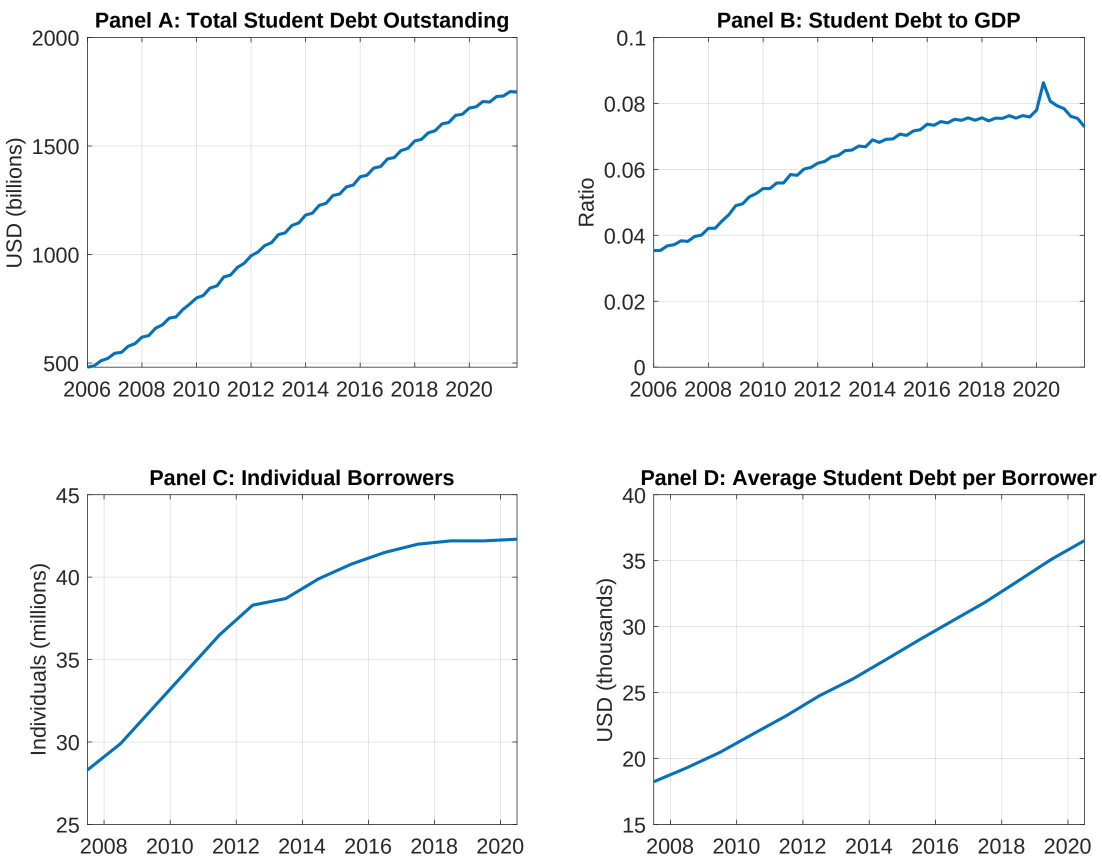

Figure 2: Repayment Choices Under Baseline and Alternate Plans

This figure shows the proportion of all households with student debt that choose to make payments according to the standard repayment plan (SRP), make payments under the income-driven repayment plan (IDRP), or skip the payment and become delinquent. Each bar represents one version of the model: baseline, Partial Payment Deferral (PPD), and Full Payment Deferral (FPD).

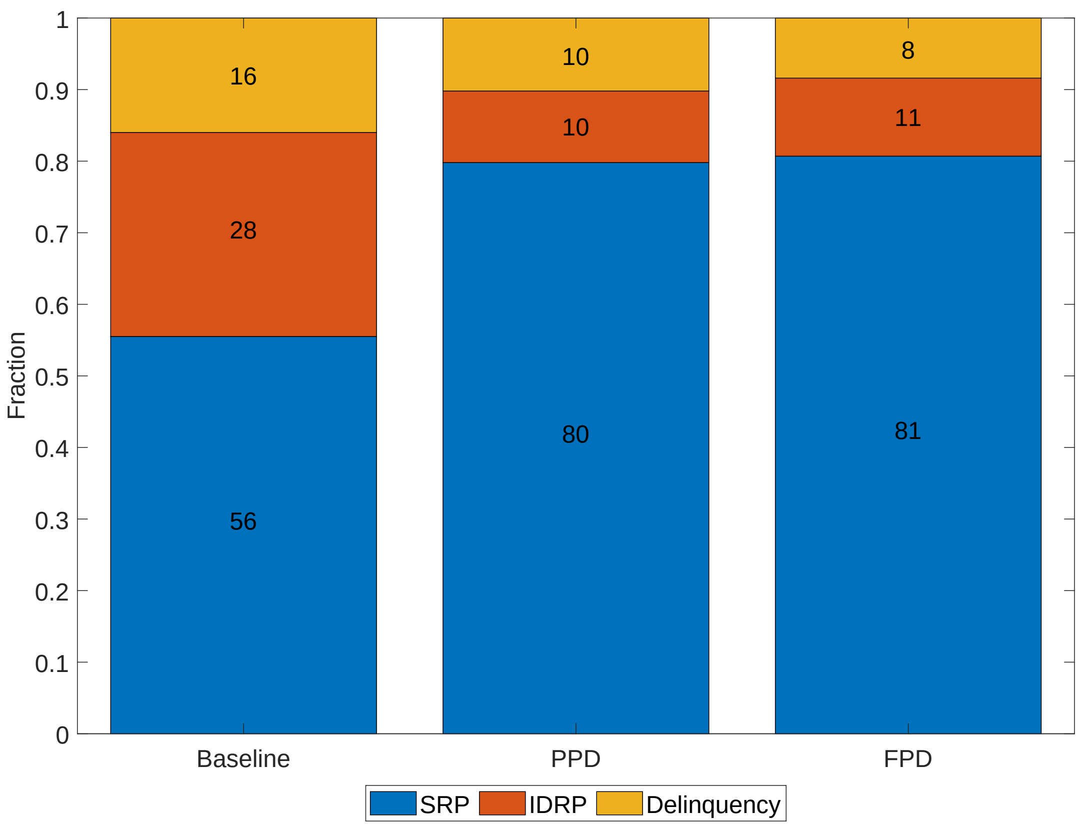

Figure 3: Consumption and Saving Over the Lifecycle in Alternate Repayment Plans

Panel A (B) shows the percentage difference between consumption (bond savings) in the modified contracts, Partial Payment Deferral (PPD) or Full Payment Deferral (FPD), relative to the baseline model for households in each group.

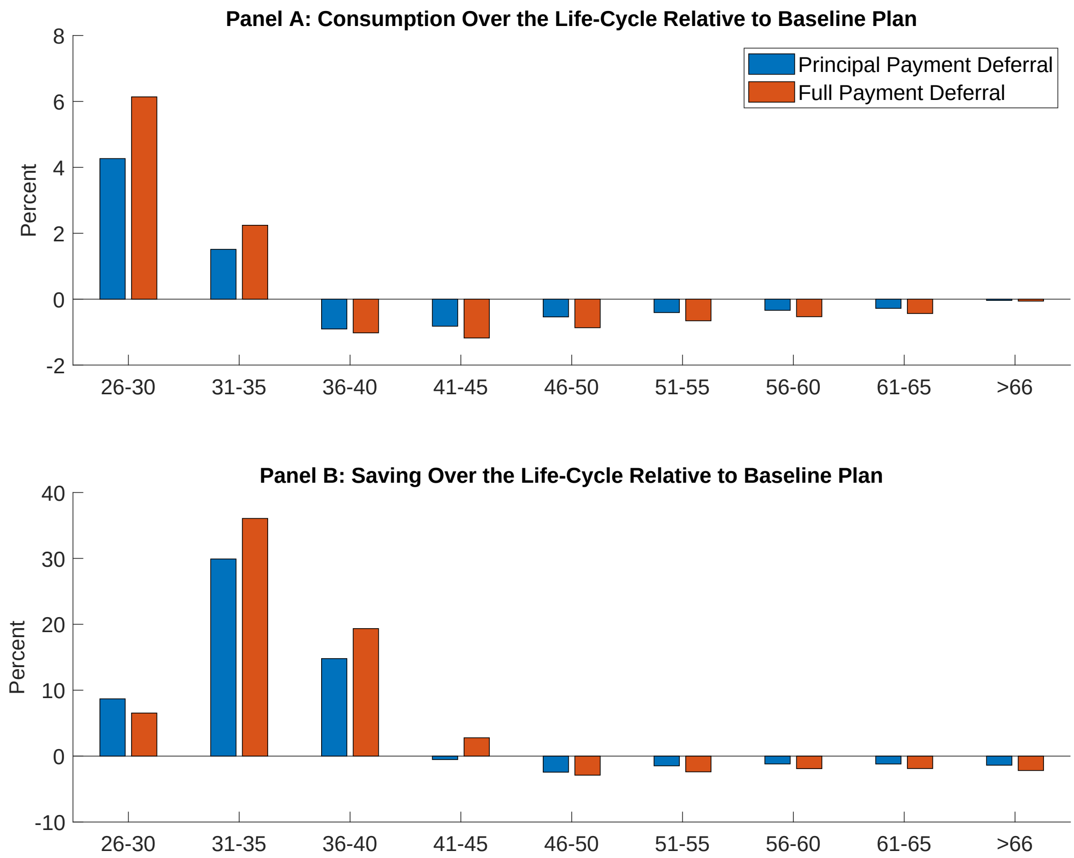

B Tables

Table 1: Income Process Parameters

This table shows parameters governing the income process detailed in Section 2.3.2. Panel (a) contains parameters for the deterministic components of income: the household fixed effect, the lifecycle age profile, and the retirement replacement rate. Panel (b) contains parameters for the unemployment shock, such as the replacement rate. Panels (c) and (d) contain parameters for the persistent and transitory shocks, respectively. The income process and parameters follow closely Guvenen et al. (2021) for the working life and Cocco et al. (2005) during retirement. Over the working life, the variance of the persistent income process is scaled down to match that in Cocco et al. (2005).

(a) Deterministic Type & Lifecycle Components

Parameter	Value
α_i	0.99
a_0	-2.0317
a_1	0.3194
a_2	$-0.0577/10$
a_3	$-0.0033/100$
ω	0.94

(b) Unemployment Shock

Parameter	Value
λ	0.52
a_ν	-2.495
b_ν	-1.037
c_ν	-5.051
d_ν	-1.087

(c) Persistent Process

Parameter	Value
ρ	0.991
p_z	0.176
$\mu_{\eta,1}$	-0.524
$\sigma_{\eta,1}$	0.113
$\sigma_{\eta,2}$	0.046
κ_η	0.470

(d) Transitory Shock

Parameter	Value
p_ϵ	0.044
$\mu_{\epsilon,1}$	0.134
$\sigma_{\epsilon,1}$	0.762
$\sigma_{\epsilon,2}$	0.055

Table 2: Repayment Plan Parameters

This table shows parameters related to the repayment plans in the model.

Parameter		Value
Interest rate premium on student debt	φ	0.035
Standard Repayment Plan length	N_{SRP}	10
Income-Driven Repayment Plan maximum length	N_{IDRP}	25
Fraction of Federal Poverty Level (FPL) for constructing discretionary income	ω_{DI}	1.50
Fraction of Discretionary Income for IDRP payment	ω_{IDRP}	0.10

Table 3: Model and Data

Panel A shows the calibrated utility costs of delinquency and IDRP enrolment, which are chosen to match the model and data moments in Panel B. Panel C compares four untargeted moments: average loan balances, average income, debt to income, and net wealth. Debt-to-income is student debt to income, while net wealth is the overall net wealth of the household. All data moments from the SCF.

Panel A: Calibrated Parameters

Parameter		Value	Source/Target
IDRP utility penalty	ξ^{IDRP}	0.00001	Proportion of IDRP in SCF
Delinquency utility penalty	ξ^{D}	0.00590	Proportion of delinquency in SCF

Panel B: Targeted Moments

Moment	Data	Model
Proportion of SRP	0.507	0.555
Proportion of IDRP	0.313	0.285
Proportion of Delinquency	0.180	0.160

Panel C: Untargeted Moments

	Data			Model		
Moment	SRP	IDRP	Delinquency	SRP	IDRP	Delinquency
Median Loan Balance	14.000	22.000	24.000	14.626	26.855	32.604
Median Income	76.359	55.996	30.544	81.135	45.821	17.869
Median Debt-to-Income Ratio	0.179	0.426	0.655	0.172	0.580	1.859
Median Net Wealth	10.480	-3.790	-10.470	7.473	-17.015	-29.453

Table 4: Standard Repayment plan vs Income-Drive Repayment Plan During the Lifecyle

This table shows average income, income growth, consumption, consumption growth, probability of being enrolled in SRP/IDRP/Delinquency, average debt outstanding, average debt payments, average net wealth and average debt payments to income for households during their working age over the lifecycle. Panel A report the statistics when agents choose a standard repayment plan, panel B report the same when agents choose an income driven repayment plan and the last panel when agents become delinquent. All statistics are conditional on the agent having debt outstanding (with the exception of the second line of the table). On the second line we report the probability that agents choose SRP, enroll in IDRP, or become delinquent, conditional on having debt outstanding at age 25.

	Panel A: Enroll in SRP				Panel B: Enroll in IDRP				Panel C: Delinquency			
	26-30	31-35	36-40	41-65	26-30	31-35	36-40	41-65	26-30	31-35	36-40	41-65
Probability	0.431	0.654	0.693	0.316	0.462	0.236	0.136	0.240	0.107	0.109	0.171	0.444
Probability (as a fraction of age 25)	0.431	0.654	0.372	0.020	0.462	0.236	0.073	0.015	0.107	0.109	0.092	0.027
Income	75.476	96.288	92.928	80.958	44.441	46.990	42.068	41.270	17.395	19.522	20.966	20.890
Income growth	0.134	0.080	0.049	0.063	0.118	0.030	0.011	-0.008	-0.313	-0.063	0.005	0.005
Consumption	44.742	55.824	52.775	43.955	25.275	26.508	23.770	24.015	13.874	14.010	14.408	14.102
Consumption growth	0.123	0.077	0.049	0.068	0.056	0.016	0.004	-0.005	-0.172	-0.028	0.009	0.005
Leverage outstanding	25.285	14.298	6.868	17.027	27.805	21.951	19.996	32.147	29.311	30.583	34.365	44.288
Net wealth	-16.079	13.064	39.738	55.048	-19.385	-7.993	3.952	9.281	-28.061	-28.087	-29.738	-35.146
Total payments	3.697	3.697	3.135	3.285	2.900	3.205	3.202	2.925	0.000	0.000	0.000	0.000
Principal payments	2.560	3.054	2.826	2.518	1.649	2.217	2.302	1.478	-1.319	-1.376	-1.546	-1.993
Interest payments	1.138	0.643	0.309	0.766	1.251	0.988	0.900	1.447	1.319	1.376	1.546	1.993
Debt payments to debt outstanding	0.148	0.294	0.745	0.461	0.106	0.158	0.208	0.179	0.000	0.000	0.000	0.000
Debt payments to income	0.052	0.044	0.043	0.047	0.068	0.074	0.087	0.081	0.000	0.000	0.000	0.000

Table 5: Determinants of Debt Repayments

This table shows the determinants of debt repayments. The two left-hand columns show agents' choices that lead to positive amortization of debt (either payments under SRP or IDRP) whereas the two right-hand columns show agents' choices that lead to negative amortization (either IDRP or delinquency). The first row shows the probability. The second and third rows show debt-to-income and debt payments to income. The fourth, fifth and sixth row shows age, income and net wealth respectively. The last four rows show principal and interest payments (realized and counterfactual payments under a SRP).

	Positive Amortization ($\Delta S_{t+1} > 0$)		Negative Amortization ($\Delta S_{t+1} < 0$)	
	SRP	IDRP	IDRP	Delinquency
Probability	0.555	0.276	0.009	0.160
Debt-to-income	0.204	0.651	2.093	2.067
Debt payments to income	0.046	0.074	0.051	0.000
Principal payment	2.845	1.952	-0.558	-1.601
Interest payment	0.683	1.144	1.468	1.601
Comparable Principal payment under SRP	2.843	3.762	2.230	2.096
Comparable Interest payment under SRP	0.683	1.144	1.468	1.601

Table 6: Persistence in Debt Repayment Decisions

This table shows the persistence in agents's debt repayment choices. The columns of the table show the probability of agents choosing SRP, IDRP or delinquency at time $t+1$, conditional on what they choose at time t (rows). Therefore, each row of the table sums up to one.

	Transition probabilities		
	SRP_{t+1}	$IDRP_{t+1}$	$Delinquent_{t+1}$
SRP_t	0.940	0.054	0.007
$IDRP_t$	0.278	0.618	0.105
$Delinquent_t$	0.018	0.220	0.762

Table 7: Loan Premia

This table reports the endogenous loan premia for each student debt repayment plan under consideration. We calibrate the loan premia on the SRP/IDRP to the average student loan interest rate in the data. The loan premia on the other two contracts (PPD and FDP) are set to equalize average student loan NPVs across all plans. We give details of the fixed-point iteration price in section 5.2. The bottom row reports the difference in the interest rate obtained for each loan modification (interest only or deferred) relative to the same base case (SRP/IDRP)

.

	Baseline (SRP/IDRP)	PPD	FPD
Loan Premia	3.50%	1.73%	1.50%
Difference		-1.77%	-2.00%

Table 8: Debt Repayment, Wealth, Income, and Consumption in the Three Different Economies

This table shows average probability of agents being enrolled in the SRP/IDRP/Delinquency, conditional on having debt outstanding (first three rows), on taking a loan at age 25 (fourth to sixth rows), and on income, consumption, average debt outstanding, average net wealth, average debt payments, average debt payments to debt outstanding and average debt payments to income for households during their working age over the lifecycle. Panel A shows the statistics for the baseline case, and Panels B and C show the statistics for the two student debt restructuring policies under analysis: PPD and FPD, respectively. For the first 10 years of the lifecycle, agents only make interest payments under the PPD policy (Panel B), or defer payments for 10 years under the FPD policy (Panel C).

	Panel A: Baseline				Panel B: PPD				Panel C: FPD			
	Overall	26-35	36-40	41-65	Overall	26-35	36-40	41-65	Overall	26-35	36-40	41-65
Pct of agents on SRP	0.555	0.555	0.693	0.316	0.798	0.962	0.752	0.621	0.807	1.000	0.752	0.660
Pct of agents on IDRP	0.285	0.337	0.136	0.240	0.100	0.000	0.156	0.188	0.109	0.000	0.150	0.187
Pct of agents Delinquent	0.160	0.109	0.171	0.444	0.102	0.038	0.092	0.191	0.084	0.000	0.098	0.152
Pct of agents on SRP (fraction of age 25)	0.188	0.555	0.372	0.020	0.430	0.962	0.752	0.174	0.495	1.000	0.752	0.262
Pct of agents on IDRP (fraction of age 25)	0.097	0.337	0.073	0.015	0.054	0.000	0.156	0.053	0.067	0.000	0.150	0.074
Pct of agents Delinquent (fraction of age 25)	0.054	0.109	0.092	0.027	0.055	0.038	0.092	0.054	0.051	0.000	0.098	0.060
Income	65.597	66.759	73.706	44.769	81.130	66.759	96.431	88.677	85.560	66.759	96.431	97.141
Consumption	38.160	39.040	42.269	25.918	47.407	40.080	54.738	51.591	50.416	40.533	54.673	57.239
Leverage outstanding	21.425	21.885	13.352	32.756	24.233	29.379	24.987	17.078	27.155	32.356	32.290	19.846
Net wealth	0.786	-6.403	22.997	4.031	33.100	-9.877	44.280	80.363	46.455	-12.159	39.725	103.038
Payments	2.819	3.068	2.609	1.741	1.898	0.771	3.026	2.539	1.723	0.000	2.976	2.655
Principal payments	1.855	2.083	2.008	0.267	1.235	-0.032	2.343	2.073	1.045	-0.808	2.170	2.160
Interest payments	0.964	0.985	0.601	1.474	0.662	0.803	0.683	0.467	0.678	0.808	0.806	0.495
Debt payments to debt outstanding	0.254	0.178	0.545	0.189	0.156	0.026	0.128	0.344	0.137	0.000	0.095	0.282
Debt payments to income	0.046	0.050	0.041	0.034	0.028	0.015	0.039	0.036	0.023	0.000	0.038	0.036

Table 9: Consumption, Net wealth and Debt Outstanding over the Lifecycle

This table shows average consumption, net wealth, and debt outstanding over the lifecycle. The first three columns show the statistics for the baseline case, the middle three columns for the PPD policy, and the last three columns for the FPD policy.

	Baseline			Principal Payment Deferral (PPD)			Full Payment Deferral (FPD)		
Age	Cons.	Net wealth	Debt outs.	Cons.	Net wealth	Debt outs.	Cons.	Net wealth	Debt outs.
26-30	31.343	-19.562	27.305	32.680	-20.825	29.241	33.267	-21.848	30.097
31-35	44.327	3.590	17.887	44.997	-1.539	29.441	45.321	-4.720	33.943
36-40	55.238	53.164	7.174	54.738	44.280	24.987	54.673	39.725	32.290
41-45	63.954	119.532	4.748	63.429	111.090	12.542	63.200	106.886	20.853
46-50	71.046	180.190	4.476	70.663	176.310	3.856	70.430	170.832	8.490
51-55	77.379	214.845	0.920	77.064	209.111	3.475	76.871	205.894	4.728
56-60	84.736	200.817	0.000	84.449	195.073	3.369	84.285	192.576	4.425
61-65	94.004	123.664	0.000	93.742	121.520	0.672	93.594	120.466	0.869
>66	75.793	3.986	0.000	75.764	3.931	0.000	75.748	3.899	0.000
Entire life-cycle average	70.172	60.276	4.167	70.141	57.502	7.172	70.125	55.807	9.046

Table 10: Welfare Gains Decomposition

This table reports the welfare gains of the two policies under analysis (PPD and FPD), relative to the baseline case (with the baseline loan premium, $\varphi^{Baseline} = 3.5\%$), as well as the welfare gains of a contract with the same features as the baseline contract but with longer maturity (LM). The different columns report results for different values of the loan premium in the PPD, FPD, and LM economies. In column 2, the loan premium is set at the value obtained in the baseline economy ($\varphi^{Baseline}$), while columns 3, 4 and 5 report the results with the equilibrium loan premium of the PPD, FPD, and LM economies, respectively.

Loan Premium (φ)	$\varphi^{Baseline}$	φ^{PPD}	φ^{FPD}	φ^{LM}
Principal Payment Deferral (PPD)	0.63%	**1.35%**		
Full Payment Deferral (FPD)	1.96%		**2.36%**	
Longer Maturity (LM)	0.26%			**0.63%**

Table 11: Loan Premia and Welfare Gains

This table reports loan premia and welfare gains for the two policies under consideration (PPD and FPD) when we require that average loan NPVs are 10% higher than in the baseline. Panel A reports the endogenous loan premia for each student debt repayment plan under consideration. Panel B reports the welfare gains.

Panel A: Loan premia			
	Baseline	PPD	FPD
Loan Premia	3.50%	1.87%	1.59%
Difference		-1.63%	-1.91%

Panel B: Welfare gains			
Loan Premium (φ)	$\varphi^{Baseline}$	φ^{PPD}	φ^{FPD}
PPD	0.63%	**1.30%**	
FPD	1.96%		**2.33%**

Table 12: Welfare Gains: Different Preference Parameters

This table reports welfare gains for the two policies under consideration Principal Payment Deferral (PPD) and Full Payment Deferral (FPD) for different preference parameters. The first column shows the welfare gains for our baseline calibration. The second and third column show the welfare gains for a lower and higher subjective discount factor, respectively. The last two columns show the welfare gains for a lower and higher elasticities of intertemporal substitution.

	Baseline	Lower Beta $\beta = 0.93$	Higher Beta $\beta = 0.97$	Lower EIS $\psi = 0.45$	Higher EIS $\psi = 0.55$
PPD	1.35%	1.85%	0.88%	1.48%	1.21%
FPD	2.36%	3.14%	1.59%	2.59%	2.06%

Table 13: Welfare Gains - Different Initial Levels of Debt

This table reports welfare gains of the PPD and FPD policies when households take lower and higher levels of debt. In particular, under the lower debt scenario (second row), households start their working life with 25% less debt, whereas under the higher debt scenario (last row), households start their working life with 25% more debt.

	PPD	FPD
Baseline	1.35%	2.33%
Lower debt	1.13%	1.83%
Higher debt	1.44%	2.64%

Table 14: Welfare Gains with Stock Market Participation Decision

This table reports moments of the model when agents can endogenously decide to participate in the stock market. The first row of this table reports the welfare gain. The bottom four rows report age, debt outstanding, income and net wealth, on the period agents decide to enter the stock market.

	Baseline	PPD	FPD
Welfare gain	n/a	1.65%	2.71%
Decision to participate:			
Age	31.622	30.734	30.566
Debt outstanding	23.484	29.069	31.232
Income	78.651	72.659	71.689
Net wealth	-3.080	-9.559	-12.073

Table 15: Comparison to Biden Administration's Policy Proposals

This table reports comparisons between the Biden Admininstration's proposed student debt repayment policies and the baseline, PPD, and FPD models. Panel A reports the amount of student debt forgiveness required to obtain the same welfare gain as the PPD and FPD plans. Panel B reports welfare gains of the two proposed IDRP plan modifications relative to the baseline.

Panel A: Loan Forgiveness		
	PPD	FPD
Welfare-Equivalent Debt Reduction	\$11,300	\$17,500

Panel B: IDRP Modifications		
	Welfare Gain	Loan Premium
(i) Shorter Repayment Period	0.31%	9.71%
(ii) Smaller Payments	1.94%	2.88%