How Do Income-Driven Repayment Plans Benefit Student Debt Borrowers?*

Sylvain Catherine[†] Mehran Ebrahimian[‡] Constantine Yannelis[§]

October 7, 2024

Abstract

The rapid rise in student loan balances has raised concerns among economists and policymakers. Using administrative credit bureau data, we find that nearly half of the increase in balances from 2000 to 2020 is due to deferred payments, largely driven by the expansion of income-driven repayment (IDR) plans, which link payments to income. These plans help borrowers by smoothing consumption, insuring against labor income risk, and reducing the present value of future payments. We build a life-cycle model to quantify the welfare gains from this payment deferment and the channels through which borrower welfare increases. New, more generous IDR rules increase this transfers from taxpayers to borrowers without yielding net welfare gains. By lowering the average marginal cost of undergraduate debt to less than 50 cents per dollar, these rules may also incentivize excessive borrowing. We demonstrate that an optimally calibrated IDR plan can achieve similar welfare gains for borrowers at a much lower cost to taxpayers, and without encouraging additional borrowing, primarily through maturity extension.

*We are grateful to numerous seminar participants and colleagues for helpful comments. TransUnion (the data provider) has the right to review the research before dissemination to ensure it accurately describes TransUnion data, does not disclose confidential information, and does not contain material it deems to be misleading or false regarding TransUnion, TransUnion's partners, affiliates or customer base, or the consumer lending industry.

[†]University of Pennsylvania, Wharton School

[‡]Stockholm School of Economics

[§]University of Cambridge and NBER

1 Introduction

Student debt has risen more than sixfold between 2002 and 2024, a rapid rise unseen in other consumer debt markets. While some of this rise is due to increases in the number of borrowers and average borrowing amounts, a significant portion of this trend is explained by another factor—borrowers taking longer to pay down their debts. Following the introduction of new income-driven repayment (IDR) plans, which tie the amount that borrowers pay to their income, large numbers of borrowers have transferred payments from early to later in the life-cycle. As of 2020, more than half of loan balances are in income-driven repayment (IDR) plans (CBO, 2020).

This paper explores the mechanisms through which borrowers benefit from deferring payments, the distribution of these benefits, and evaluates the cost-efficiency of existing repayment rules. We document that the slowdown in loan repayment explains approximately one-half of the increase in student loan balances between 2000 and 2020, a fact that has received little attention in policy debates. We show that the retiming of payments leads to significant welfare gains, equivalent to $50,600 on average per borrower. Borrowers with large balances benefit the most. About half of these welfare gains are due to redistribution from taxpayers to borrowers, while most of the remainder comes from consumption smoothing and insurance against income risk. We also show that new and more generous payment plans do not increase welfare beyond taxpayer transfers. Furthermore, an optimal income-linked payment plan would have a long maturity, providing substantial smoothing and insurance benefits while significantly reducing taxpayer costs.

We begin our analysis by quantifying how the rise in payment deferral and increases in the uptake of IDR plans affected aggregated student debt levels. Using administrative data from TransUnion, a major US credit bureau, we show that deferred payments increased rapidly with the rollout of IDR plans. We estimate deferred payments by borrower cohorts and construct counterfactual loan balances under standard repayment plans. Our decomposition shows that deferred payments and IDR uptake account for $390.6 billion, or 44%, of the growth in student loan debt since IDR plans were introduced.

Next, we present a life-cycle model in which graduates enter the workforce and repay their

student loans. Because markets are incomplete, borrowers are unable to borrow against future earnings and are unable to insure against income risk, preventing them from realizing their ideal consumption plan. The model captures key features of the student loan market, and income-driven plans can increase borrower welfare by smoothing payments and providing insurance.

The model estimates aggregate welfare gains from IDR plans and decomposes these gains into four key channels. First, IDR reduces lifetime payments, improving borrower welfare at the expense of taxpayers. Second, IDR helps borrowers smooth consumption over the lifecycle by deferring payments. Third, by automatically lowering payments during adverse economic conditions, IDR provides insurance against idiosyncratic labor income risk. Finally, by assisting students with lower expected lifetime incomes, IDR can mitigate welfare losses arising from inequality. This decomposition of economic channels clarifies how the design and calibration of IDR plans serve different policy objectives.

We apply our framework to quantify how linking payments to income affects borrower welfare. Our results show significant welfare gains from IDR, with average per-borrower welfare increasing by $50,600. While a substantial part of this increase represents a transfer from taxpayers, a significant portion stems from payment deferral and more favorable repayment terms. Notably, large shares of the welfare gain are due to lifecycle consumption smoothing (29.9%) and insurance against income shocks (21.3%). On the other hand, because borrowers with large balances tend to have higher expected lifetime earnings, we find no net welfare gains or losses from reducing inequality between borrowers.

We also study the new and more generous IDR plan, the Saving on a Valuable Education (SAVE) plan, introduced in 2022 but currently blocked by Federal courts. This plan increases benefits for borrowers, for example by lowering the fraction of income that borrowers must pay and shortening time to forgiveness for those with smaller balances at graduation. The SAVE plan increases borrower welfare by an additional $5,300, though this entire gain results from a transfer from taxpayers, with gains and losses in life-cycle smoothing and insurance offsetting each other.

We further study changes in student borrowing incentives following the introduction of the SAVE plan. If borrowers expect to repay only a fraction of the amount borrowed, this could

encourage more borrowing, with schools potentially capturing the subsidies (Eaton, Howell and Yannelis, 2020; Lucca, Nadauld and Shen, 2019). Before SAVE, under previous IDR rules, student debt was priced near fair value for most borrowers, although some with very high balances benefited from debt costs below one. With SAVE, many undergraduate borrowers with even moderate balances are likely to repay two-thirds or less of their borrowed amounts. We show that the behavioral response from undergraduate students could substantially increase the cost of the federal loan program.

We conclude by using our framework to study optimal policy rules. We explore the space of income-driven plan parameters to identify the calibration that maximizes welfare. For a given cost to taxpayers, the optimal policy includes a much longer repayment period, that would extend up to retirement age.[1] A longer repayment period allows the government to recoup payments later in life, enabling greater generosity earlier in life and in low-income states. Consequently, it maximizes welfare gains by enhancing both insurance and intertemporal smoothing per taxpayer dollar.

An important aspect of optimal policy design is the overall cost to taxpayers, as offering more generous repayment terms is justified only if the marginal welfare gains to borrowers exceed the additional budgetary costs. We find that net welfare gains are maximized with a per-borrower taxpayer subsidy of $6,250—significantly lower than both current rules ($24,050) and the SAVE plan ($30,000). Despite this reduction, gross welfare gains to borrowers would exceed those under current rules. In addition to an extended repayment period, the optimal policy includes increasing the repayment rate from 10% to 33%, while doubling the threshold at which earnings are subject to repayment. We assume that payments would remain capped by the standard-plan formula. Under this policy, loans would be priced near fair value, mitigating moral hazard concerns. An important implication of our findings is that if the government aims to subsidize students beyond the $6,250 threshold, it should do so outside the scope of the student loan program—such as through grants—rather than through more generous repayment terms.

Our paper makes four main contributions. First, we present new facts about the drivers of

[1]This is similar to the Australian student loan repayment system, which defaults all borrowers into an income-driven plan, but does not include any forgiveness. In contrast, in the UK, all borrowers are defaulted into IDR plans, but loans are forgiven after 25 years, and in the US after 20 or 25 years depending on the plan.

student loan balance growth, a topic frequency discussed by academics, policymakers, and the general public. We show that approximately half of the increase in student loan debt between 2000 and 2020 is driven by payment deferrals, an often neglected fact in the public discourse. Second, we present a new life-cycle model of student loan debt, which can be used both for welfare analysis and counterfactuals. Third, we present new normative results that show why borrowers benefit from payment deferrals. This provides important context for the design of student loan repayment plans, as changing parameters such as repayment rates or maturity can have different effects on borrowers welfare through each channel studied. Finally, we provide new results on optimal student loan repayment policy and find that improvements in the calibration of their parameters could lead to large efficiency gains.

This paper primarily joins a growing literature on student loans in the United States. In particular, several papers study repayment plans with a focus on IDR plans (Mueller and Yannelis, 2019, 2022; Herbst, 2018; Monarrez and Turner, 2024). A related literature also explores loan forgiveness (Di Maggio et al., forthcoming; Catherine and Yannelis, 2022), payment deferrals (Dinerstein et al., 2024; Hamdi et al., 2024) and limit increases (Black et al., 2020; Goodman et al., 2021). Though most of these studies are empirical, analyzing how repayment plans affect borrower outcomes, some papers study the student loan market through the lens of a structural model (Lochner and Monge-Naranjo, 2011, 2016; de Silva, 2023). Recent theoretical work also studies how the design of loan programs and tax systems can foster human capital formation (Findeisen and Sachs, 2016; Stantcheva, 2017). Amromin and Eberly (2016), Avery and Turner (2012), Yannelis and Tracey (2022) and Looney and Yannelis (2024) provide recent reviews of the literature. Our main contribution here is to provide a new model to analyze student debt repayment, as well as new positive and normative facts about trends in deferred payment, and welfare.

Our analysis of potential policy improvements is most closely related to an important study by Boutros, Clara and Gomes (2022), which uses a life-cycle model to show that shifting student loan payments by ten years can result in large welfare gains without increasing costs to taxpayers. Our structural analysis, however, differs from theirs in several ways. First, we develop a framework that decomposes welfare gains into distinct economic channels and examines how the calibration of repayment rules can support various policy objectives. Second,

rather than evaluating a few discrete alternative policies, we explore the continuous space of policy parameters to identify the optimal calibration. Both studies share the insight that differing payment schedules can improve the efficiency of the student loan program. In Boutros et al. (2022), this is achieved by delaying payments by ten years, whereas our optimal policy accomplishes this by raising the repayment threshold and extending the repayment period.

More broadly, this paper also contributes to the literature on the design of public policies in the presence of liquidity constraints and uninsurable income risks faced by households. A substantial body of research has examined the optimal provision of unemployment insurance (Baily, 1978; Hopenhayn and Nicolini, 1997; Chetty, 2006; Shimer and Werning, 2007) or welfare gains from public health programs (Finkelstein et al., 2019). Another branch of the literature has focused on the role of liquidity in retirement systems (Laibson et al., 1998; Beshears et al., 2024; Catherine et al., 2020). We extend this line of inquiry to student loan programs, highlighting that, given the rapid rise in student loan balances, the design of repayment rules can play a decisive role in easing liquidity constraints and providing insurance against income risk for young workers.

Finally, we extend the rich literature studying household finances through the lens of life-cycle models. Gomes, Haliassos and Ramadorai (2021) offer a comprehensive review of this literature.

The remainder of this paper is organized as follows. Section 2 discusses the institutional background of student lending in the United States. Section 3 presents descriptive facts, as well as a decomposition demonstrating how deferred payments contributed to the rise in student loan balances. Section 4 presents our life-cycle model of student loan payments. Section 5 discusses our welfare accounting measures. Section 6 presents estimates of payment deferment on aggregate welfare, and decomposes welfare components into the four channels. Section 7 discusses optimal policy for student loan repayment. Section 8 concludes.

2 Institutional Background

In this Section, we present institutional background on student debt in the United States and recent trends in loan balances and repayment patterns. We show that student debt has increased steadily over the past 25 years, and that a significant portion of this increase is driven by payment deferral. Recently, the rise in aggregate student debt is driven by increases in balances and deferred repayment related to the growing popularity of income-driven repayment programs.

2.1 Student Loans in the United States

The vast majority of student debt in the United States is directly disbursed or guaranteed by the federal government. Modern federal student loan programs began in 1965, with the passage of the Higher Education Act. Historically, there were two large federal student loan programs in the United States. The first was the Federal Family Education Loan Program (FFEL), which began in 1965, and which was terminated in 2010. The FFEL program was a guarantee program, under which private lenders provided capital for highly regulated loans. These funds were in turn guaranteed by the government. The William D. Ford Federal Direct Loan Program (DL) was authorized in 1992. Under the DL program, the US Treasury directly provided funds for student loans. Borrowers take either Subsidized and Unsubsidized loans. All borrowers are eligible for Unsubsidized loans, while borrowers from lower-income families are eligible for Subsidized loans. While the loans are quite similar, for Subsidized borrowers, interest does not accrue while borrowers are in school.

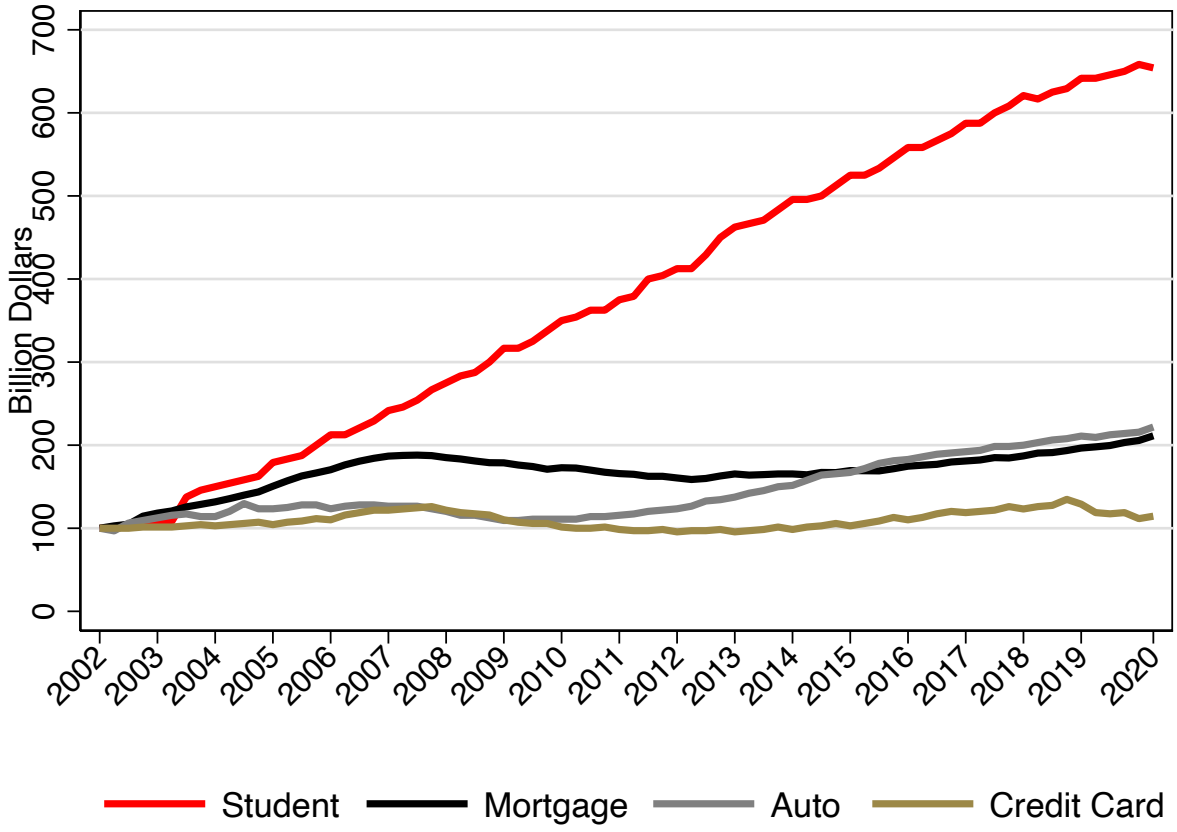
Federal student loans are highly regulated, with interest rates and borrowing limits set by Congress. Pricing does not vary based on risk, and all students of the same level face the same interest rate.[2] Borrowing limits vary by class level, and are higher for upper level and graduate students. Loans are serviced by private companies, with contracts from the Department of Education. There is a small private student loan market, the CFPB estimates that this accounts for approximately 8% of all student loans.

[2]There are slight differences in effective interest rates based on whether borrowers are Subsidized or Unsubsidized. Additionally, in some years subsidized borrowers had lower interest rates. Interest rates also differ for graduate and undergraduate borrowers.

Figure 1 shows a primary motivating fact for our analysis. Student debt has grown sharply over the past two decades– at a rate much faster than any other form of household debt. More precisely, the figure normalizes loan balances in 2003, and plots the relative increase over time. For all categories of household debt, loan balances increase by less than 200% over the relevant time period. This is even true of mortgages, which grew sharply in the run-up to the 2008 financial crisis. In contrast, student debt increases by more than six-fold during the same time period. What explains this tremendous and unprecedented rise in a category of household debt? As is explored in the remainder of this section, part of these trends are explained by payment deferment which occurred following the introduction of new repayment plans.

Figure 1: Household Debt in the United States 2004-2020

This figure displays the growth in mortgage, auto loan, student loan, and credit card balances for consumers from 2002 to 2020. Balances are reported from the end of each fiscal quarter. Aggregate consumer balances are normalized to 100 in 2002. Data from the Federal Reserve Bank of New York.



2.2 Repayment Plans

Historically, the vast majority of student loan borrowers were in a repayment option termed the Standard Plan, under which borrowers would repay their loans over 10 years by making 120 monthly payments of an equal fixed size. In some cases, students could defer or forbear repayment for a set period of time, due to events such as unemployment, economic hardship or enrollment in a graduate program. For cohorts borrowing after the 2006-07 academic year, interest rates are fixed. Some borrowers also chose the Extended Repayment Plan, which increased the loan maturity to twenty-five years. If borrowers did not make payments for more than 270 days, loans would become in default and the federal government could garnish 15% of wages above a threshold.

Beginning in 2008, the Department of Education began a dramatic expansion of IDR plans. Prior to 2008, one IDR plan existed, Income-Contingent Repayment, but the terms were fairly onerous for borrowers and as a consequence take-up was low. In 2008 Income-Based Repayment was introduced, which allowed borrowers to pay 15% of their income above 150% of the Federal Poverty Line (FPL). Remaining balances would be forgiven after twenty-five years in repayment. In subsequent years, more generous income-driven repayment options such as Pay as You Earn and Revised Pay as You Earn were introduced. These new plans made IDR more generous, for example making payment ten rather than fifteen percent above the FPL and allowing forgiveness after twenty rather than twenty-five years in repayment.

In 2022 a new and extraordinarily generous repayment plan was introduced, the Saving on a Valuable Education (SAVE) Plan. The SAVE plan increased generosity across four parameters. First, the threshold above which borrowers would make payments was raised from 150% to 225% of the FPL. This corresponds to in increased from $46,800 to $70,000 in 2024 for a family of four, below which borrowers would pay nothing. Second, the amount that undergraduate borrowers pay was cut in half to 5%. Graduate borrowers would continue to pay 10%, and borrowers with both types of loans would pay a weighted average. Third, the time to forgiveness was decreased to ten years for balances below $12,000. Beyond this, time to forgiveness would increase by one year for each additional $1,000, until a maximum of twenty years of undergraduate debt and twenty-five years for graduate debt. Finally, negative amortization was eliminated, meaning that balances would no longer grow by the interest rate if

borrowers failed to make payments.

3 Decomposing Payment Deferrals and the Rise in Balances

In this section, we quantify the increase in student loan balances that is attributed to slowing payment using administrative data. We construct a simple counterfactual, which shows how student loan balances would have evolved had borrowers made payments under the standard plan instead of slowing repayment under IDR and other forbearance programs. We find that payment deferral accounts for almost half of the increase in student loans balances between 2000 and 2020.

3.1 Data

Our main data source is the Booth TransUnion Consumer Credit Panel, an anonymized 10% sample of all TransUnion credit records from 2000 to 2020. Individuals who were in the initial sample in 2000 have their data continually updated, and each year an additional 10% of new first time individuals in the credit panel are added. A small fraction of individuals also leave the panel each year, for example due to death or emigration.[3] We can observe basic information about student loans, including the original balance, the current balance, scheduled payments, and maturity of the loan.[4]

3.2 Trends in Student Loan Balances

On a simple level, the rise in student debt show in Figure 1 is driven by both increases in the number of borrowers, and increases in average balances. The rise in student loan balances is shown in figure 2. The left panel shows aggregate balances, the middle panel of figure 2 shows the number of borrowers and the right panel shows average loan balances. During the

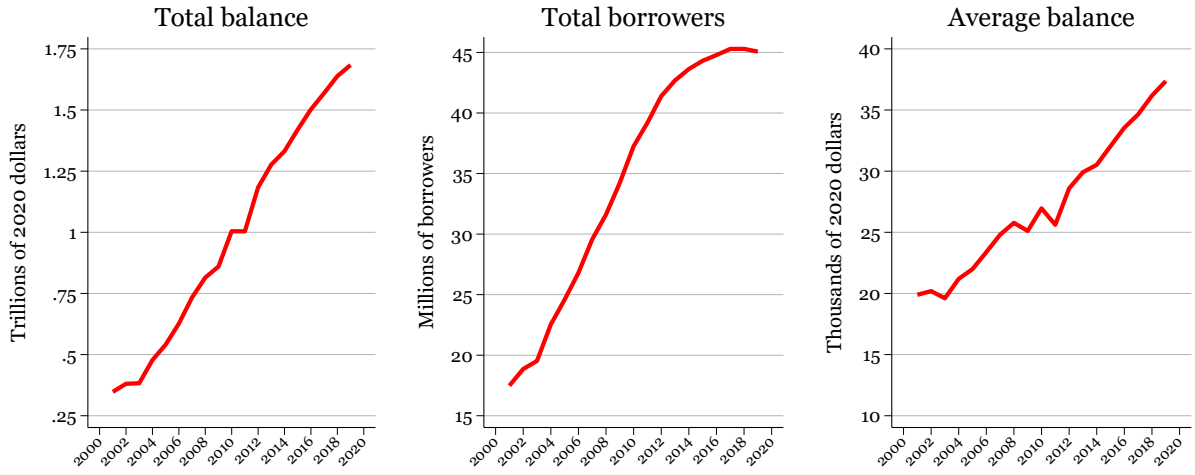
[3] Keys, Mahoney and Yang (2020) provide further details regardig the Booth TransUnion Consumer Credit Panel. All tables and figures that list TransUnion as a source have statistics calculated (or derived) based on credit data provided by TransUnion, a global information solutions company, through a relationship with the Kilts Center for Marketing at the University of Chicago Booth School of Business.

[4] Total loan volumes in our data are also comparable to measures from other datasets such as Department of Education data from Looney and Yannelis (2015) and the Federal Reserve Bank of New York. Appendix figure B.1 in appendix B compares our aggregates with these other data sources. Aggregates and trends line up closely.

same time period, the number of borrowers more than tripled, increasing from 15 million to 45 million. The number of enrolled students increases by approximately 30% between 2000 and 2019, which is significant but not comparable in magnitude to the very large increase in outstanding borrowers. The rise in borrowers and enrollment is shown in figure 2.[5] The right most panel of figure 2 shows average balances doubled over the same time period, rising from roughly $20,000 in 2001 to $40,000 in 2020.

Figure 2: Student Loan Aggregates

This figure displays the aggregate student loan balance in the United States, the total number of borrowers by year, and the average balance by year, from 2000 to 2020. Numbers are as of December of each year. Source: TransUnion.



A major factor contributing to the growth in balances is slowing repayment rates, which is shown in Figure 3. The left panel of Figure 3 shows more detailed repayment patterns by cohort. The figure shows the fraction of the initial repayment balance outstanding, for cohorts in odd years from 2001 to 2019. Earlier to later cohorts transition in color from blue to red. The figure shows a marked decline in repayment in later cohorts, with borrowers who begin repayment in more recent years seeing larger balances outstanding relative to earlier cohorts. Indeed, some recent cohorts even see balances grow in initial years following repayment. This is in part due to the fact that, while disbursement amounts increased, repayment significantly

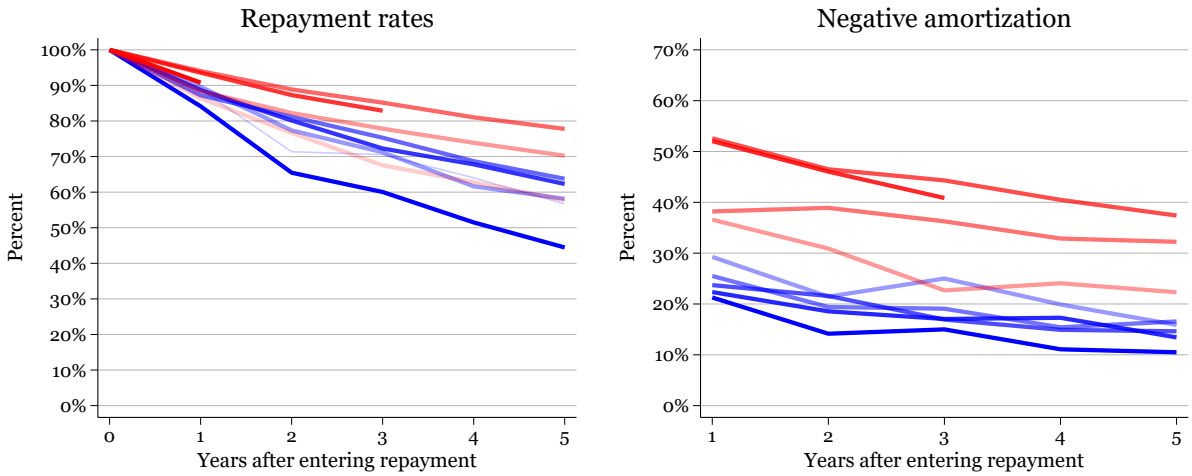
[5]Both undergraduate and graduate enrollment increases during the time period, and the proportion of graduate students as a share of all students remains relatively flat (NCES 2023).

slowed. The 2001 cohort had on average 45% of their repayment balance outstanding after five years, while the 2015 cohort had 80% of their repayment balance outstanding after five years.[6]

Much of this pattern is driven by negative amortization, or borrowers loan balances increasing due to payments not covering interest. Put simply, many borrowers are not making enough payments to reduce their balances. The faction of individuals in a cohort negatively amortizing is shown in the right panel of Figure 3. The panel again shows cohorts in odd years from 2001 to 2019 and with earlier to later cohorts transition in color from blue to red. More recent cohorts see a larger share of borrowers in negative amortization, with close to a third of borrowers negatively amortizing in the 2017 cohort.

Figure 3: Repayment Rates and Negative Amortization

This figure shows the repayment behavior of every other repayment cohort from 2001 to 2019. The left-hand panel indicates the repayment rate, or what percentage of the total balance owed by borrowers when entering repayment remains. The right-hand panel displays negative amortization, or the percentage of people who have an equal or higher balance than they did when entering repayment. Data from TransUnion.



While some of the increase in borrowers and borrowing can be explained by more enrollment, a significant portion of the rise in student loan balances, particularly since 2008, is driven by slowing repayment rates over time. A large part of these deferred payments are due to the rise of IDR plans, as well as increased use of deferment and forbearance plans. To

[6]The fraction of debt owned by older borrowers has increased over time, which is shown in appendix figure C.2.

what extent did these patterns contribute to the rise in student debt balances? To answer this question we decompose the amount of the rise due to these deferred payments by constructing counterfactual loan balances, assuming borrowers made payments under the standard plan.

3.3 Counterfactual Aggregate Student Debt

In this section, we show how aggregate student debt would have evolved had borrowers made standard plan payments. We construct counterfactual balances by first sorting borrowers by cohort of repayment. We then compute interest rates by taking the average over the four years preceding repayment. We compute standard plan payments $R_{1,it}$ using the amortization formula $R_{1,it} = \frac{r_L L_{i0}}{1-(1+r_L)^{-10}}$, where r_L is the interest rate and L_{i0} is the balance at repayment. We then compute the evolution of balances assuming standard plan payments are made and the counterfactual balance B_{it}^{CF} is given by

$$B_{it}^{CF} = L_i - \sum_{i=1}^{T} R_{1,it} \tag{1}$$

The counterfactual balance B_{it}^{CF} essentially computes the evolution of balances, assuming that standard plan payments are made. Figure 4 presents the average counterfactual balance $\bar{B}_{it}^{CF}$, along with actual balances $\bar{L}_t$ in each year. The two series evolve similarly until 2008, suggesting that on average borrowers were up to that point making payments equivalent to standard plan payments.[7] This is consistent with the evidence in Figure 3, that for earlier cohorts they are paying down a substantial portion of their debt as repayment progress.

In 2008, the year in which IBR is introduced, we see the series begin to diverge. The gap between the actual and counterfactual balance widens over time. By 2020, the last year in our sample, the actual balance is $1.6998 trillion while the counterfactual balance is $1.3091 trillion.[8] Thus this exercise suggests that student loan balances would be $390.6 billion smaller if borrowers had continued to pay down their loans as they did prior to the introduction of new IDR plans. Between 2008 and 2020 student loan balances rose by $884.9 billion, from $814.9 billion to $1.6998 trillion. Our exercise thus suggests that 44%, or roughly half of the increase

[7] Despite the fact that the series are similar, many borrowers default and other prepay. However on average payments are similar to those in the standard plan.

[8] We end the sample in 2020, as at that point repayment was paused and interest accrual was frozen.

in student loan balances is driven by slower repayment over time as opposed to other factors.

The fact that deferred payments can explain approximately half of the increase over time is consistent with the fact that other trends cannot explain the entirety of the rise in student loan balances. While tuition increased, the amounts were not large enough to explain borrowing increases. The College Board reports that at four-year public schools tuition increased by 33% at public four-year colleges and 25% at private four-year colleges between 2008 and 2020. Enrollment and borrowing patterns also cannot explain increases in borrowing, the same source reports that the total number of enrollments remained constant over the same time period. The total number of active borrowers actually declined over the same period, and average annual borrowing remained roughly constant or declined slightly.[9]

Figure 4: Actual and Counterfactual Balance

This figure displays the actual aggregate student loan balance plotted against the counterfactual student loan balance that would result if all students paid their loans down according to the 10-year standard plan. Data from TransUnion.

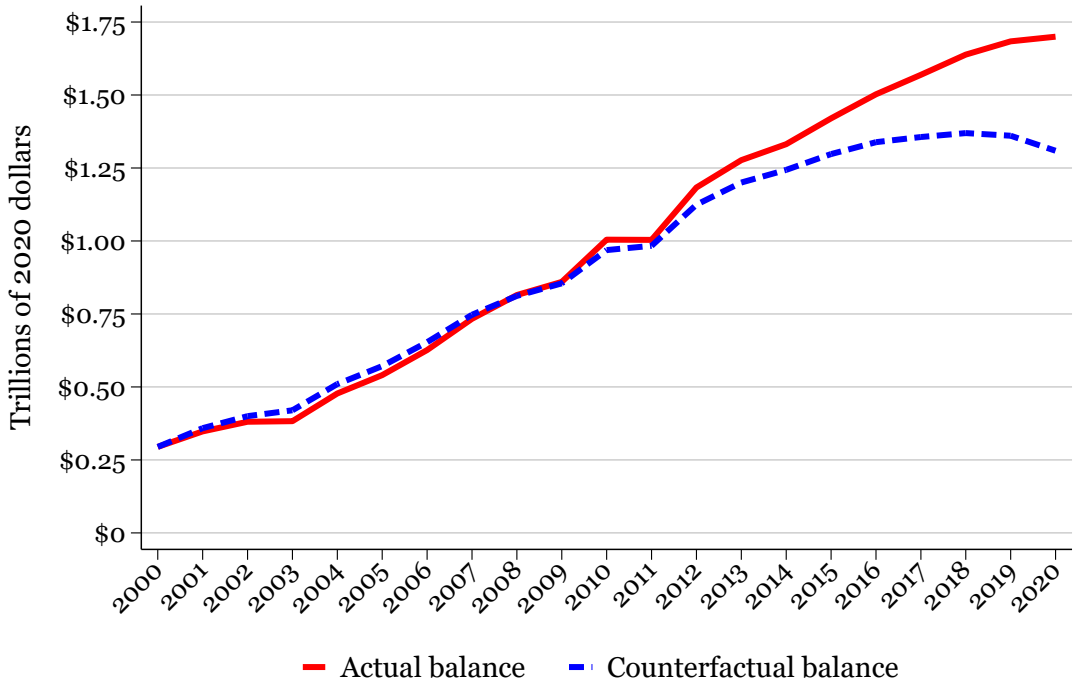


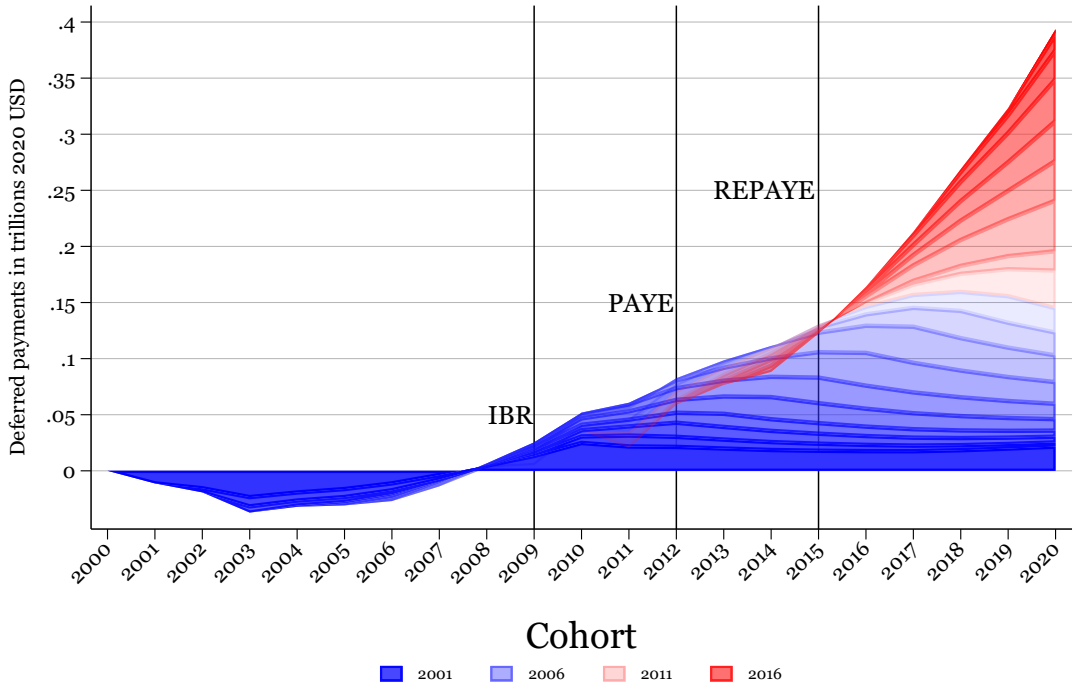
Figure 5 presents a further decomposition of the trends shown in Figure 4. The figure

[9]The College Board reports that undergraduate borrowers borrowed $6,700 in the 06-07 academic year, and $6,440 in the 21-22 academic year. Graduate borrowers borrowed $21,560 in the 06-07 academic year, and $18,970 in the 21-22 academic year.

shows the cumulative $300 billion increase in deferred balances by cohort, over time between 2000 and 2020. The figure plots $\sum_{i=1}^{T} R_{1,it} - P_{it}$, where P_{it} are payments made by individual i in year t. Before 2008, we see little evidence of deferred payments. This begins to change in 2009, with the introduction of the IBR plan. For these cohorts deferred payments continue to grow, and we see sharp upticks with the introduction of PAYE and REPAYE. This decomposition shows that the timing of deferred payments closely aligns with cohorts that had more generous IDR plans available.

Figure 5: Deferred Payments by Cohort

This figure displays cumulative deferred payments by cohort, based on a standard ten-year repayment plan. The figure shows $\sum_{i=1}^{T} R_{1,it} - P_{it}$, where P_{it} are payments made by individual i in year t and $R_{1,it}$ are payments under the standard plan. Vertical lines show the introduction of IDR plans. Source: TransUnion.



While this section established that payment deferral is a significant contributor to the increase in student loan balances during the first two decades of the 2000s, this does not tell us whether slowing repayment benefited or harmed borrowers, and to what degree. The consequences of payment deferral in economic terms are ambiguous, operate through multiple channels and may be small or large. On the one hand, borrowers pay less earlier in the life-cycle when marginal utility is higher, on the other hand they pay more over the life of a

loan. Borrowers may also benefit from insurance, transferring consumption to periods when marginal utility is higher and redistribution between among borrowers. Finally, there are potentially large transfers from taxpayers to borrowers, as payments later in the life of a loan lower its present value. To study the welfare consequences of payment deferral, and the introduction of new IDR plans, we build a life-cycle model and decompose welfare gains from policy counterfactuals in the remainder of the paper.

4 Model

In this section, we set up a life-cycle model in which borrowers enter the labor market at graduation and start repaying their student debt. We start from a standard consumption model in the spirit of Gourinchas and Parker (2002). In this framework, incomplete markets prevent households from borrowing against future earnings and from insuring against income risk, preventing them from realizeing their ideal consumption plan. We augment the model to capture key institutional features of federal student loans. By default, borrowers enroll in the "standard plan" and reimburse their loan with fixed payments over the next ten years. Alternatively, they enroll in an income-driven repayment program, where borrowers pay a fraction of their earnings above a threshold. The model serves three purposes. First, it allows us to decompose the welfare effects of policies into different channels. Second, the model can be used to study the welfare effects of policy counterfactuals. Finally, we can compute optimal policy parameters.

4.1 Agent

Borrowers are indexed by $i \in \{1,\dots,I\}$, year since graduation by $t \in \{t_0 = 1,\dots,T\}$, and states/earning trajectories by $s \in \{1,\dots,S\}$. The agent has constant relative risk aversion (CRRA) and maximizes expected utility, which is given by:

$$V_{it_0} = \mathbb{E}\sum_{t=t_0}^{t_R-1} \beta^{t-t_0}\left(\prod_{k=t_0}^{t-1}(1-m_k)\right)u(C_{it}) \quad + \mathbb{E}\left[V_{t_R}\right] \tag{2}$$

where $\mathbb{E}\left[V_{t_R}\right]$ is expected utility at retirement and the period utility function is:

$$u(C_{it}) = \frac{1}{1-\gamma}\left(\frac{C_{it}}{\sqrt{N_{it}}}\right)^{1-\gamma}, \tag{3}$$

where γ is the coefficient of relative risk aversion, m_k the mortality rate at age k, and β the discount factor. The agent's gross financial wealth evolves W_{it} as:

$$W_{it+1} = \left(W_{it} - R_{L,it} + Y_{it} + B_{it}^{SS} + B_{it}^{SN} - \Gamma_{it} - C_{it}\right)(1 + r_f) \tag{4}$$

Where $R_{L,it}$ is the student loan payment, Y_{it} is income, B_{it}^{SS} and B_{it}^{SN} are respectively benefits from social security and food stamps and Γ_{it} are taxes. We approximate the expected utility at retirement using the solution from Merton (1971):

$$\mathbb{E}\left[V_{t_R}\right] = b\frac{\overline{W}_{iR}^{1-\gamma}}{1-\gamma} \tag{5}$$

where $\upsilon = \left[(1-\beta) - (1-\gamma)r_f\right]/\gamma$ and $b = \left[\left(1 - e^{-\upsilon(T-t_R)}\right)/\upsilon\right]^{\gamma}$, and $T - t_R$ is life expectancy in retirement. $\overline{W}_{iR}$ is total wealth at retirement, defined as financial wealth W, the present value of Social Security benefits, discounted at the risk-free rate, net of remaining student debt, if any.

4.2 Student Loan

Borrowers graduate with a level of student debt L_{it_0}. We consider three counterfactual economies that differ in their student loan repayment rules. First, a benchmark economy without IDR. Then, two economies with IDR with sets of repayment rules similar to the US system before and after the introduction of SAVE. Our goal is to evaluate the welfare effects of these two IDR systems relative to the benchmark economy.

4.2.1 Benchmark Economy Without IDR

In the benchmark economy, student loans are reimbursed over a ten-year period through a fixed-payment schedule. We refer to this schedule as the "Standard Plan." If borrowers fail

to make the standard plan payment, the government garnishes part of their earnings to repay their loan. Denoting $R_{L,it}$ the payment, loan balances evolve as:

$$L_{it+1} = L_{it}(1 + r_L) - R_{L,it}, \tag{6}$$

and increases when repayments do not cover interest. $R_{L,it}$ can take different values, depending on whether the borrower is in default.

Standard Plan Under the standard plan, borrowers pay a fixed amount over the ten years following their graduation. This fixed yearly repayment is:

$$R_{\text{std},it} = \frac{r_L L_{it_0}}{1 - (1 + r_L)^{-10}} \tag{7}$$

until the debt is fully repaid.

Garnishment In the benchmark economy, borrowers in default fall into wage garnishment. In that case, the government garnishes 15% of disposable income, defined as earnings net of taxes Γ. Their repayment is therefore:

$$R_{\text{garnish},it} = \min\left\{0.15 \times (Y_{it} - \Gamma_{it}), R_{\text{max garnish},it}\right\} \tag{8}$$

where $R_{\text{max garnish},it}$ is a legal limit on the garnishment amount. Specifically, garnishment cannot exceed 25% of disposable income or 30 times the hourly federal minimum wage per week.

$$R_{\text{max garnish},it} = \min\left\{0.25 \times (Y_{it} - \Gamma_{it}), \lambda_{\text{garnish}} Y_{1,t}\right\} \tag{9}$$

We assume that borrowers fall in garnishment when their cash on hand $W_{it} + Y_{it}$ falls below a threshold λ_{default}.[10]

[10] We choose not to model the choice of going into garnishment. Doing so would require us to model all the costs associated with garnishment, such as lower credit score.

4.2.2 Economy With IDR, Before SAVE

In reality, under the current system borrowers can enroll in the standard plan or in the income-driven repayment program. IDR is governed by two key parameters, a payment rate $\theta_{\text{idr}} = 10\%$ and a threshold λ_{idr} above which payments are made. Borrowers in IDR pay a fraction $\theta_{\text{idr}} = 10\%$ of their discretionary income, but no more than they would repay in the standard plan, until their balance is fully repaid. Moreover, student debt is forgiven after $t_{F_{\text{idr}}} = 20$ years of payments. Hence, repayment in IDR is:

$$\begin{cases} R_{\text{idr},it} = \min\left\{\theta_{\text{idr}}\left(Y_{it} - \lambda_{\text{idr}} Y_{1,t}\right)^{+}, R_{\text{std},i}\right\} & \text{if } t \leq t_{F_{\text{idr}}}, \\ R_{\text{idr},it} = 0 & \text{if } t > t_{F_{\text{idr}}}. \end{cases} \tag{10}$$

Discretionary income is defined as the share of their income above 150% of the federal poverty line. Assuming that it represents a constant multiple of the average wage index, we denote this threshold $\lambda_{\text{idr}} Y_{1,t}$. Until forgiven or fully repaid, balances evolve as in Equation (6). Finally, we assume that garnishment no longer exists in the economy with IDR.[11]

4.2.3 Economy After SAVE

SAVE introduces important changes to the IDR program. First, payments are no longer capped by the standard plan, so the repayment formula becomes:

$$\begin{cases} R_{\text{save},it} = \theta_{\text{save},i}\left(Y_{it} - \lambda_{\text{save}} Y_{1,t}\right)^{+} & \text{if } t \leq t_{F_{\text{save}},i} \\ R_{\text{save},it} = 0 & \text{if } t > t_{F_{\text{save}},i}. \end{cases} \tag{11}$$

Second, parameters depends on how much students borrowed and whether their loan financed an undergraduate or a graduate degree. The payment rate $\theta_{\text{save},i}$ is 5% for undergraduate debt, 10% for graduate debt and a weighted average for borrowers with a mixture of undergraduate and graduate debts. Discretionary income starts at 225% of the federal poverty line, instead of 150% under previous rules. Debt is forgiven after $t_{F_{\text{save}},i} = 10$ years for undergraduate students

[11] In the past, IDR was associated with significant administrative costs, notably because borrowers had to get their earnings re-certified every year. Thanks to recent policy changes, borrowers in IDR will be automatically recertified using tax data. Consequently, we assume that, moving forward, borrowers will systematically choose to enroll in IDR over going into garnishment, following the dominating strategy.

who borrowed less than $L_{it_0} < \$12,000$. Each additional \$1,000 increases the forgiveness clock by one year up to 20 years. Borrowers with graduate debt must make $t_{F_{\text{save}},i} = 25$ years of payment before their debt is forgiven.[12]

Finally, SAVE introduces interest-rate subsidization: if $R_{\text{save},it}$ is below interests $r_L L_{it}$, the government finances the difference. In other words, balances can no longer increase over time:

$$L_{it+1} = \min\left\{L_{it}(1+r_L) - R_{\text{save},it}, L_{it}\right\}. \tag{12}$$

4.2.4 Prepayment

In all cases, and at any point in time, borrowers can choose to make a greater payment towards reimbursing their loan, partially or completely.

4.3 Income and Taxes

Labor earnings Earnings Y_{it} are the product of an aggregate $Y_{1,t}$ and a idiosyncratic component $Y_{2,it}$:

$$Y_{it} = Y_{1,t} \cdot Y_{2,it}. \tag{13}$$

We model the idiosyncratic component $Y_{2,it}$ as in Guvenen et al. (2021). Specifically, we assume that it depends on a deterministic function of age $g(t)$, a persistent z_i and a transitory component ε_i. The persistent component follows an AR(1) process. To reproduce the skewness and kurtosis observed in the data, innovations to the persistent and transitory components have normal-mixture distributions. Finally, the agent faces a small probability of unemployment, which depends on persistent income and age. The equation set (14) summarizes the dynamic of $Y_{2,it}$:

[12] Our understanding is that, conditional on having any graduate debt, this rule applies to undegraduate debt as well.

$$
\begin{aligned}
&\text{Level of idiosyncratic earnings:} && Y_{2,it} = (1-\nu_{it})e^{\left(g(t)+\alpha^i+z_{it}+\varepsilon_{it}\right)} \\
&\text{Persistent component:} && z_{it} = \rho z_{it-1} + \eta_{it} \\
&\text{Innovations to AR(1):} && \eta_{it} \sim \begin{cases} \mathcal{N}(\mu_{\eta,1}, \sigma^2_{\eta,1}) & \text{with probability } p_z \\ \mathcal{N}(\mu_{\eta,2}, \sigma^2_{\eta,2}) & \text{with probability } 1-p_z \end{cases} \\
&\text{Initial condition of } z_{it}\text{:} && z_{it_0} \sim \mathcal{N}(0, \sigma^2_{z,t_0}) \\
&\text{Transitory shock:} && \varepsilon_{it} \sim \begin{cases} \mathcal{N}(\mu_{\varepsilon,1}, \sigma^2_{\varepsilon,1}) & \text{with probability } p_\varepsilon \\ \mathcal{N}(\mu_{\varepsilon,2}, \sigma^2_{\varepsilon,2}) & \text{with probability } 1-p_\varepsilon \end{cases} \\
&\text{Unemployment probability:} && p_{\nu,it}(z) = \frac{e^{\xi_{it}}}{1+e^{\xi_{it}}}, \text{ where } \xi_{it} = a + bt + cz_{it} + dz_{it}
\end{aligned}
\tag{14}
$$

We do not explicitly model the labor supply effects of income-driven repayment (IDR) programs. While the evidence on these effects is somewhat mixed, it generally suggests that any labor supply responses are small or negligible. Britton and Gruber (2020) find no evidence of labor supply responses to income-contingent loans in the UK, while de Silva (2023) report some effects in Australia. However, the labor supply elasticity estimated in de Silva (2023) is quite small, at 0.11.

Moreover, in the U.S., IDR payment amounts are capped by the standard repayment plan, implying that the effective marginal payment rate is zero for income levels at which the IDR formula exceeds the standard payment. Furthermore, in many cases, IDR primarily shifts payments across different time periods rather than altering total lifetime payments. Consequently, labor supply responses under IDR in the U.S. are likely to be much smaller than those observed in the tax literature or in IDR systems in countries without such payment caps.

Social safety net Borrowers with very low wealth and earnings qualify for Supplemental Nutrition Assistance Program (SNAP, or "food stamps"). SNAP-eligible individuals receive income compensation equal to 6% of the wage index minus 30% of pre-transfer income. Qualifying for SNAP requires that wealth is less than 5% of the national wage. Specifically, benefits from

from SNAP are:

$$B_{it}^{\text{SN}} = (0.06Y_{1t} - 0.3Y_{it})^{+} \qquad \text{if } W_{it} < 0.05Y_{1t}. \tag{15}$$

where W_i is the household's wealth, defined below. These benefits allow households to maintain a consumption level above $0.06Y_{1t}$ in all circumstances. Therefore, their existence prevents welfare effects of student loan repayment rules from being driven by rare disaster states.

Retirement Benefits Borrowers contribute $\tau = 6.2\%$ of their earnings towards the retirement system. Contribution only applies below the maximum taxable earnings limit, which has remained roughly equal to 2.5 times the average wage $Y_{1,t}$ over the past four decades. Hence, Social Security taxes are:

$$\Gamma_{S,it} = \tau \min\left\{Y_{it}, 2.5Y_{1,t}\right\}. \tag{16}$$

Workers retire at age t_R , and their yearly retirement pension, as a fraction of the average wage Y_{1,t_R} is:

$$B_{it\geq t_R}^{\text{SS}} = \begin{cases} 0.9 \times \text{AIYE}_{it_R} & \text{if } \text{AIYE}_{it_R} < 0.2Y_{1,t_R} \\ 0.116 \times Y_{1,t_R} + 0.32 \times \text{AIYE}_{it_R} & \text{if } 0.2Y_{1,t_R} \leq \text{AIYE}_{it_R} < Y_{1,t_R} \\ 0.286 \times Y_{1,t_R} + 0.15 \times \text{AIYE}_{it_R} & \text{if } Y_{1,t_R} \geq \text{AIYE}_{it_R}, \end{cases} \tag{17}$$

where Y_{1,t_R} is the value of the national average wage when an individual retires and AIYE_{it_R} is their average indexed yearly earnings at retirement. The AIYE keeps tracked of past earnings, indexed on the growth rate of the aggregate wages and follows:

$$\text{AIYE}_{it} = \sum_{k=t_0}^{t} \frac{\min\left\{Y_{2,ik}, 2.5\right\}}{t - t_0 + 1} \times Y_{1,t}. \tag{18}$$

Income Tax Households pay income tax on earnings and retirement benefits following the progressive schedule detailed in appendix A.1.2.

5 Welfare

In this section, we present a conceptual framework to quantify how public policy can increase social welfare by improving consumption allocation. While we focus on student loan repayment programs, our framework can be applied to other policies.

The first inefficiency arises from incomplete markets, which can prevent households from smoothing consumption over their life cycle and insuring against idiosyncratic risk. Our framework measures how policy changes—specifically loan repayment rules—can generate welfare gains by providing liquidity and insurance. The second inefficiency stems from inequality. For risk-averse households, a utilitarian planner can enhance aggregate welfare by transferring consumption from high- to low-income households. We can also assess how policy changes affect welfare through this mechanism.

Distinguishing and separately measuring these welfare gains is important for at least two reasons. Policymakers may value welfare gains differently based on their nature. For example, providing liquidity and insurance can improve welfare without additional taxpayer cost, while redistribution may or may not align with a government's philosophy. Moreover, certain policy goals might be better addressed with instruments beyond the scope of a specific study. For instance, policymakers may prefer other tools for income redistribution and prioritize reducing financial frictions in the design of student loan programs.

5.1 Sources of Welfare Loss

First-best consumption plan In complete and frictionless markets, the first-order condition for consumption of borrower i is:

$$\forall \{t,s\}, \quad \underbrace{\beta^{t-t_0}\left(\prod_{k=t_0}^{t-1}(1-m_k)\right)u'\left(C^*_{its}\right)\frac{1}{P_{t_0}(t)}}_{v'_{its}} = \underbrace{u'\left(C^*_{it_0}\right)}_{v'_{it_0}}, \tag{19}$$

subject to the intertemporal budget constraint:

$$\sum_t^T \mathop{\mathbb{E}}_{it}[C^*]P_{t_0,t} = \sum_t^T \mathop{\mathbb{E}}_{it}[Y-R]P_{t_0,t}, \tag{20}$$

where $P_{t_0,t} = \frac{1}{(1+r_f)^t}$ is the price in t_0 of a zero-coupon bond paid in period t. Throughout the rest of the paper, we use the notation $\mathbb{E}_j[X] = \mathbb{E}[X|j]$ to denote the expected value of X conditional on j, and do the same for other moments.

In a complete market agents would trade the present value of their lifetime wealth for consumption coupons in future periods and states of the world, resolving all uncertainty at time t_0. The optimal portfolio of coupons would be such that the marginal expected utility of spending one more present-value dollar in time t and state s, which we denote ν'_{its}, is equal across all periods and states.

Market incompleteness The lifetime efficiency loss relative to the complete-market benchmark can be approximated with a first-order Taylor expansion of V_{it_0} around the incomplete-market consumption path:

$$V^*_{it_0} - V_{it_0} \approx \sum_t^T \mathbb{E}_{it}\left[\nu'(C^* - C)\, P_{t_0,t}\right]. \tag{21}$$

Assuming that the agent leaves no bequest, the average per-period expected consumption remains the same in present value terms and $\mathbb{E}_i\left[(C^* - C)\, P_{t_0}\right] = 0$. Therefore, the welfare loss from market incompleteness can be approximated as:

$$V^*_{it_0} - V_{it_0} \approx T \cdot \mathrm{cov}_i\left(\nu', (C^* - C)\, P_{t_0,t}\right). \tag{22}$$

This covariance can be further decomposed as:

$$\overbrace{V^*_{it_0} - V_{it_0}}^{\text{Market incompleteness}} \approx \overbrace{\sum_t \mathrm{cov}_{it}\left(\nu', (C^* - C)\, P_{t_0,t}\right)}^{\text{imperfect insurance}} + \overbrace{T \cdot \mathrm{cov}_i\left(\mathbb{E}_{it}[\nu'], \mathbb{E}_{it}\left[(C^* - C)\, P_{t_0,t}\right]\right)}^{\text{imperfect intertemporal smoothing}}. \tag{23}$$

The first term represents losses from imperfect insurance against idiosyncratic risk, which prevents the agent from perfectly smoothing consumption across states. The second term represents losses from imperfect smoothing over the life cycle due to liquidity constraints.

Inequality From the perspective of the social planner, the aggregate welfare of borrowers is the sum of their expected utilities at graduation:

$$\boldsymbol{V} = \sum_i V_{it_0}. \tag{24}$$

A strictly utilitarian planner wants households to equalize the marginal utility from money over time and across states, but also seek to equalize marginal utilities across individuals. Following the same logic as in Equation (22), the welfare loss between the optimal and actual distributions of the aggregate budget can be approximated as:

$$\begin{aligned} \boldsymbol{V}^* - \boldsymbol{V} \approx & \overbrace{\sum_i \sum_t \operatorname{cov}_{it}\left(v', (C^* - C) P_{t_0,t}\right) + \sum_i T \cdot \operatorname{cov}_i\left(\mathop{\mathbb{E}}_{it}[v'], \mathop{\mathbb{E}}_{it}\left[(C^* - C) P_{t_0,t}\right]\right)}^{\text{Market incompleteness}} \\ & + \overbrace{I \cdot T \cdot \operatorname{cov}\left(\mathop{\mathbb{E}}_{it}[v'], \mathop{\mathbb{E}}_{it}[\Delta C P_{t_0,t}]\right)}^{\text{Inequality}}. \end{aligned} \tag{25}$$

The first two terms aggregate losses from incomplete markets, summing $V^*_{it_0} - V_{it_0}$ over all individuals, while the last term represents the loss from lifetime expected consumption inequality.

Monetary Measure of Welfare Variations To facilitate the quantitative interpretation of our findings, we scale changes in aggregate welfare $\boldsymbol{V}$ to report welfare gains and losses in dollar terms:

$$\Delta^{\$}\boldsymbol{V} = \frac{\Delta \boldsymbol{V}}{\mathbb{E}[v']} \tag{26}$$

where $\mathbb{E}[v'] = \frac{1}{I \times T \times S} \sum_i \sum_t \sum_s v'_{its}$ is the average marginal utility of present-value dollars across individuals, years and states. This normalization reports welfare gains as marginal dollars of increased endowment in the initial period, equally distributed across state, time and individual.

5.2 Welfare Gains From Policy

The total welfare gains of a policy can be written as the difference between benefits to borrowers and costs to taxpayers. To see this, define pre-policy consumption levels for individual i at state s at time t by C_{its}, and post-policy consumption by $C_{its} + \Delta C_{its}$. The pre-policy average

expected utility of borrowers at graduation is:

$$\boldsymbol{V} = \sum_i \sum_t \mathbb{E}_{it}\left[v(CP_{t,t_0})\right], \tag{27}$$

and the average welfare gain per borrower from the policy p is:

$$\Delta_{\mathrm{p}}\boldsymbol{V} = \sum_i \sum_t \mathbb{E}_{it}\left[v(C+\Delta C) - v(C)\right] P_{t,t_0}. \tag{28}$$

To a first-order Taylor approximation, the change in welfare is:

$$\Delta_{\mathrm{p}}\boldsymbol{V} \approx \sum_i \sum_t \mathbb{E}_{it}\left[v'\Delta C\right], \tag{29}$$

whereas the net welfare gain, subtracting the cost to taxpayers, can be written as:

$$\Delta_{p,\mathrm{net}}\boldsymbol{V} = \overbrace{\Delta_{\mathrm{p}}\boldsymbol{V}}^{\text{benefit to borrowers}} - \overbrace{\sum_{t=t_0}^{T} -\Delta R_{L,t} P_{t_0 t}}^{\text{cost to taxpayers}} \tag{30}$$

where $\Delta R_{L,t}$ is change in loan repayment across all individuals/states at time t.

5.3 Decomposition of Welfare Gains

To understand the mechanisms through which IDR improves welfare, we decompose equation (30) to isolate the net welfare gains of a policy into five components:

$$\begin{aligned}\frac{\Delta^{\$}_{p,\text{net}} V}{I} = &\overbrace{\frac{\frac{1}{I}\sum_i \sum_t \mathbb{E}_{it}\left[\nu' \; \Delta CP_{t_0,t}\right]}{\mathbb{E}\left[\nu'\right]} - \frac{\frac{1}{I}\sum_i \sum_t \mathbb{E}_{it}\left[\nu'\right] \; \mathbb{E}_{it}\left[\Delta CP_{t_0,t}\right]}{\mathbb{E}\left[\nu'\right]}}^{\text{insurance}} \\ &+ \overbrace{\frac{\frac{1}{I}\sum_i \sum_t \mathbb{E}_{it}\left[\nu'\right] \; \mathbb{E}_{it}\left[\Delta CP_{t_0,t}\right]}{\mathbb{E}\left[\nu'\right]} - \frac{\frac{1}{I}\sum_i \sum_t \mathbb{E}_{i}\left[\nu'\right] \; \mathbb{E}_{i}\left[\Delta CP_{t_0,t}\right]}{\mathbb{E}\left[\nu'\right]}}^{\text{intertemporal smoothing}} \\ &+ \overbrace{\frac{\frac{1}{I}\sum_i \sum_t \mathbb{E}_{i}\left[\nu'\right] \; \mathbb{E}_{i}\left[\Delta CP_{t_0,t}\right]}{\mathbb{E}\left[\nu'\right]} - \frac{\sum_t \mathbb{E}\left[\nu'\right] \; \mathbb{E}\left[\Delta CP_{t_0,t}\right]}{\mathbb{E}\left[\nu'\right]}}^{\text{transfer progressivity}} \\ &+ \overbrace{\frac{\sum_t \mathbb{E}\left[\nu'\right] \; \mathbb{E}\left[\Delta CP_{t_0,t}\right]}{\mathbb{E}\left[\nu'\right]}}^{\text{mean transfer}} - \overbrace{\sum_t -\Delta R_{L,t} P_{t_0 t}}^{\text{cost to taxpayers}}.\end{aligned} \tag{31}$$

Equation (31) can be interpreted as a sequence of four differences in differences which decompose net welfare gains into economic channels.

Insurance The first line is the difference between welfare gains to borrowers in the model, and in a counterfactual in which the marginal utility is equalized across states of the world and so are policy-induced changes in consumption, for each borrower and year. In this counterfactual, there is no room for welfare gains from insurance against income risk. However, policies can still transfer wealth to borrowers and help them smooth consumption over the life cycle. Therefore, the difference between the two terms measures the welfare gains from insurance.

Intertemporal Smoothing The second line represents the differences in welfare gains between a world where marginal utility is equalized across states, and one in which borrowers also perfectly smooth consumption over the life cycle. In that second intermediate world, policy can only help borrowers through wealth transfers. Therefore, the difference in welfare gains isolate changes in borrowers' ability to smooth consumption over the life cycle.

Transfer Progressivity The third line is the difference between two worlds with different levels of progressivity. In both worlds, borrowers allocate the policy-induced change in their lifetime consumption uniformly across periods and states but, in the another world, all borrowers receive an equal share of the present value cost to taxpayers. The difference isolates the welfare gains from the progressivity of the transfer, that is its ability to distribute more to borrowers with higher marginal utility.

Mean Transfer The last line represents the welfare gain to borrowers from changes in wealth transfers between taxpayers and borrowers, as any subsidies to borrowers must be paid. By construction, the dollar welfare gains of this transfer is equal to what it costs taxpayers, that is $\sum_t \Delta CP_{t_0,t} - \Delta R_{L,t} P_{t_0 t}$. Implicitly, taxpayers and borrowers are assumed to have the same marginal utility, such that pure uniform transfers neither improve nor deteriorate welfare.

Noting that the average change in the present value of consumption is equal to that of repayments, we can write Equation (31) as a sum of covariances:

$$\frac{\Delta^{\$}_{p,\text{net}} V}{I} = \frac{\overbrace{\frac{1}{I}\sum_i \sum_t \text{cov}_{it}\left(v', \Delta CP_{t_0,t}\right)}^{\text{insurance}}}{\mathbb{E}[v']} + \frac{\overbrace{\frac{T}{I}\sum_i \text{cov}_i\left(\mathbb{E}_{it}[v'], \mathbb{E}_{it}\left[\Delta CP_{t_0,t}\right]\right)}^{\text{intertemporal smoothing}}}{\mathbb{E}[v']} + \frac{\overbrace{T \cdot \text{cov}\left(\mathbb{E}_i[v'], \mathbb{E}_i[\Delta CP_{t_0,t}]\right)}^{\text{transfer progressivity}}}{\mathbb{E}[v']}. \tag{32}$$

Gains from the insurance channel derive from the covariance between changes in consumption and levels of marginal utility across alternative realizations of the income process, for each borrower and year. Similarly, for each borrower, gains from intertemporal consumption smoothing come from the covariance between changes in consumption and marginal utility across periods. Finally, gains from progressivity come from the covariance between changes in consumption and marginal utility across borrowers.

6 Existing IDR Rules

In this section, we calibrate our model to examine the impact of existing IDR rules both before and after the implementation of the SAVE program. Prior to the introduction of SAVE, welfare gains for borrowers were evenly distributed between taxpayer subsidization and improvement in intertemporal consumption smoothing and insurance against income risk. The introduction of SAVE increases gains from subsidization, particularly for undergraduate students, without significant additional gains from easing financial frictions. Moreover, SAVE substantially lowers the marginal cost of debt for undergraduate students, which could encourage them to borrow more in the future.

6.1 Model Calibration

Earnings The income dynamics is calibrated following Guvenen et al. (2021) with two differences. We report the parameters of the stochastic process in Appendix Table A.1. The first difference is that we assume unemployment shocks to last an entire year, which is a very close approximation of what these authors estimate and makes the model more tractable. The second difference is that we estimate the deterministic life-cycle component of earnings as a cubic polynomial function of age, taking into account the correlation between lifetime earnings and how much workers borrowed as students. We report the parameters of this function in Appendix Table A.2.

Dependents and Poverty Line We calibrate the number of persons in the households N_{it} as a deterministic function of age. Specifically, we use the SCF to estimate the number of children per adult as a cubic polynomial of age and add one to the predicted value to obtain N_{it}. We assume $N_{it} = 1$ in retirement. We define age as 23 plus the number of years since graduation. This regression is reported in Appendix Table A.2. The federal poverty line is determined as a function of the predicted number of children, using federal guidelines and linear interpolation between integers.

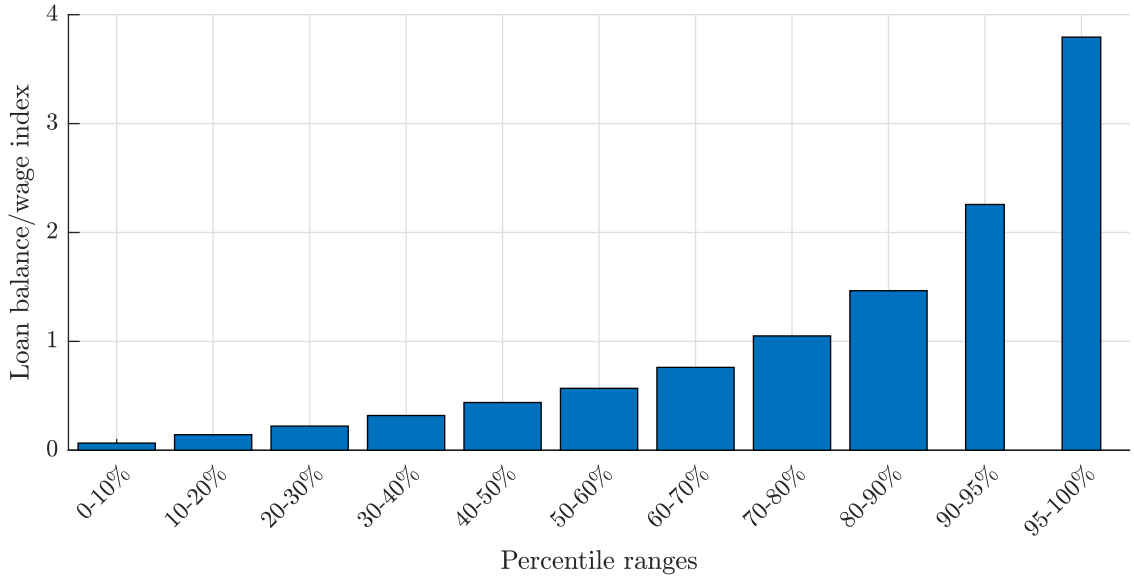
Preferences We calibrate preferences based on Gourinchas and Parker (2002)'s classic study on consumption over the life cycle in the presence of labor income uncertainty. Specifically, we set the discount factor to $\beta = 0.96$, their baseline estimate, and relative risk aversion to $\gamma = 2$, their estimate for college and graduate school graduates.

Interest Rates The real risk-free rate is $r_f = 0.02$. The interest rate on student loans should be a value-weighted average of rates on undegraduate, graduate, and PLUS debt. In the model, we assume that the interest rate only depends on total balance at graduation. For debt below the undegraduate limit of $0.705 \times Y_1$ (\$45,000), the real interest rate is $r_{L,i} = 0.04$. For the part of the debt between the undergraduate limit and the graduate school limit of $1.026 \times Y_1$ (\$65,500), the interest rate is set to 0.055. For any debt in excess of the graduate school limit, that is debt borrowed under the LOAN Plus program, the interest rate is 0.065.

Initial Conditions We simulate the model for eleven levels of initial loan size: we take the median of the first nine deciles, and the two halves of the highest deciles. To match the model, we scale loan amounts by the value of the wage index an individual graduated. Figure 6 reports the scaled level of initial debt for each group. Finally, we set the initial level of gross wealth as to match the relationship between student debt balances and wealth among fresh graduates in the SCF.

Figure 6: Initial Balance by Centile of Loan Size

This figure displays the median student loan balance at graduation (or exit) for the cohort that completed their program in 2019 and 2020, segmented by debt decile. The final decile is divided into two equal parts. Source: TransUnion.



6.2 Welfare Gains

Table 1 presents the distribution of welfare gains from Income-Driven Repayment (IDR) plans for borrowers, quantified in dollars (scaled by the wage index) and broken down by economic mechanism. Prior to the implementation of the SAVE program, improvement in consumption smoothing and better insurance against income risk explained slightly more than half of the welfare gains of borrowers. Subsidization from taxpayers explained the rest. With the introduction of SAVE, changes in welfare gains for borrowers are predominantly explained by increased subsidization from taxpayers, without additional gains arising from the easing of financial constraints.

Table 1: Welfare Gains to Borrowers from IDR

This table reports the decomposition of aggregate welfare gains for borrowers before and after SAVE. The decomposition is defined in Equation (31). Welfare gains are reported in dollar equivalent per borrower and as a multiple of the national wage index ($63,795 in 2022).

	Total change in borrower welfare	Financial frictions		Transfer	
		Insurance	Intertemporal Smoothing	Progressivity	Mean
Before SAVE	0.793	0.169	0.237	0.002	0.377
After SAVE	0.876	0.114	0.294	-0.001	0.470
Change	0.083	-0.055	0.057	-0.003	0.093

IDR Before SAVE Relative to the standard ten-year plan, IDR increases borrower welfare by .793 times the average wage, or $50,589. The largest single component of this welfare increase, accounting for 48% of the gain, comes from a transfer from taxpayers. IDR reduces the average present value of student loans at graduation by 0.377 average wage ($\approx \$24,000$). Most of the rest of the welfare gain come from the fact that IDR also helps borrowers smooth consumption over the life cycle and provides insurance against income risk.[13] These two sources of welfare gains represent 0.237 ($\approx \$15,100$) and 0.169 ($\approx \$10,800$) times the national wage index respectively. Overall, existing IDR rules do not constitute a zero-sum game that only improve the welfare of borrowers at the expense of taxpayers. In fact, in dollar terms, IDR generates $0.793/0.377 = 2.1$ times more welfare for borrowers than it costs taxpayers.

Present value gains from IDR are not distributed equally between borrowers but the progressivity of program is theoretically ambiguous. On the one hand, IDR favors borrowers with high debt-to-income ratios and borrowers with income close or below the payment threshold. On the other hand, students who borrowed more tend to earn more. Quantitatively, we find IDR to be neither progressive nor particularly regressive, with gains relatively equally distributed across the earnings distribution.[14]

[13] Appendix figure A.2 provides a way of visualizing these consumption smoothing gains over the life-cycle.

[14] Over the past decade, many borrowers failed to enroll in IDR even though they would have benefited from the program. Therefore, it is possible that IDR was more regressive than our model suggests. We assume no

IDR After SAVE The introduction of SAVE increases borrower welfare by 0.083 ($\approx \$5,300$). This gain reflects the transition to more generous repayment rules. However, the combined benefits from insurance and intertemporal smoothing does not contribute to these welfare gains. In fact, the welfare of borrowers improves mostly thanks to the redistribution from taxpayers to borrowers. The gains from intertemporal smoothing are larger under SAVE because the threshold under which borrowers start repaying their debt is raised from 150% of the poverty line to 225% and the share of discretionary earnings paid is significantly reduced. These changes lengthen the repayment timeline of the loans. However, they also reduce the program ability to provide insurance against income risk. This is because the reduction in payment are less concentrated in low-income states of the world. Instead, a greater part of these payments reduction are now expected under normal circumstances. As such, they are better described as a transfer from taxpayers to borrowers rather than an insurance program against adverse career shocks. Under SAVE, the average transfer from taxpayers to borrowers reaches 0.47 times the wage index ($\approx \$30,000$), a 0.093 ($\approx \$5,900$) increase relative to existing rules. Overall, the gains to borrowers are below the cost to taxpayers, suggesting a negative net welfare effect from the program.[15]

6.3 Distributional Consequences

Figure 7 reports how the average per-borrower welfare gains vary as a function of their lifetime earnings and their debt level at graduation. Panel A shows that the benefits from IDR are decreasing in earnings and increasing in debt. The relationship between welfare gains and loan size is particularly pronounced, with the highest gains observed among borrowers in the top decile of loan balances.

Importantly, Figure 7 decomposes welfare gains using the social planner's perspective, who converts utility gains to a monetary equivalent using a unique conversion rate $\mathbb{E}[v']^{-1}$ for all borrowers. From the perspective of individual borrowers, the equivalent monetary variation would be higher (lower) for borrowers with above (below) average lifetime income since they have lower (higher) marginal utility from consumption. For instance, from the perspective of

friction moving forward, as a consequence of the simplification of the enrollment and recertification processes.

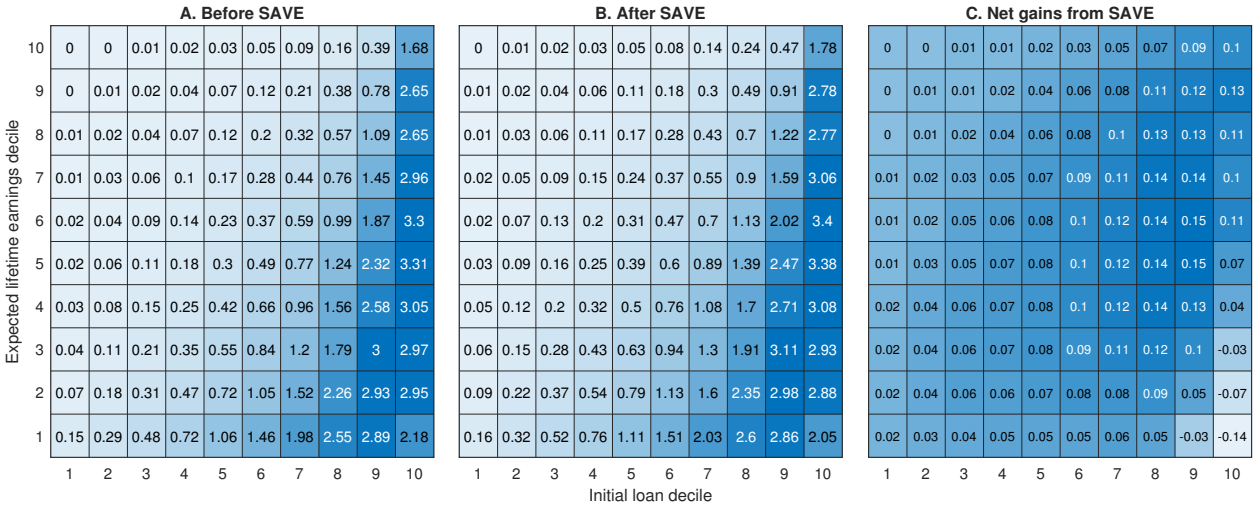
[15] Importantly, in Table 1, we assume that the debt distribution at graduation will not change in response to new IDR rules. In reality, it is likely that students will borrow more money, as we discuss in Section 6.4.

the social planner, the per-person welfare gains from individuals in the top decile of balances exceed 2 average wage ($> \$125,000$). Nonetheless, the equivalent monetary variation from these borrowers' perspective would likely be higher since they tend to have higher expected lifetime earnings.

Paradoxically, borrowers in the top two deciles of debt but at the bottom at the earnings distribution benefit slightly less from IDR than their counterparts with slightly better income expectation. In the benchmark economy, these borrowers fall in garnishment and cannot repay their debt, which limits the gains from IDR in consumption terms. A limitation of our model is that it does not account for the indirect costs of default, potentially underestimating the welfare gains from IDR for these borrowers.

Figure 7: Welfare Gains Per Borrower by Decile of Debt and Income

This figure reports the average welfare gain from IDR per borrower as a function of their decile of expected lifetime earnings and decile of student loan at graduation. Gains are reported in multiples of the national average wage index ($63,795 in 2022). Panels A and B reports the welfare gains before and after SAVE respectively. Panel C reports the difference. Panel C uses a different color scale from Panels A and B. Deciles of lifetime earnings are defined within the population of borrowers.



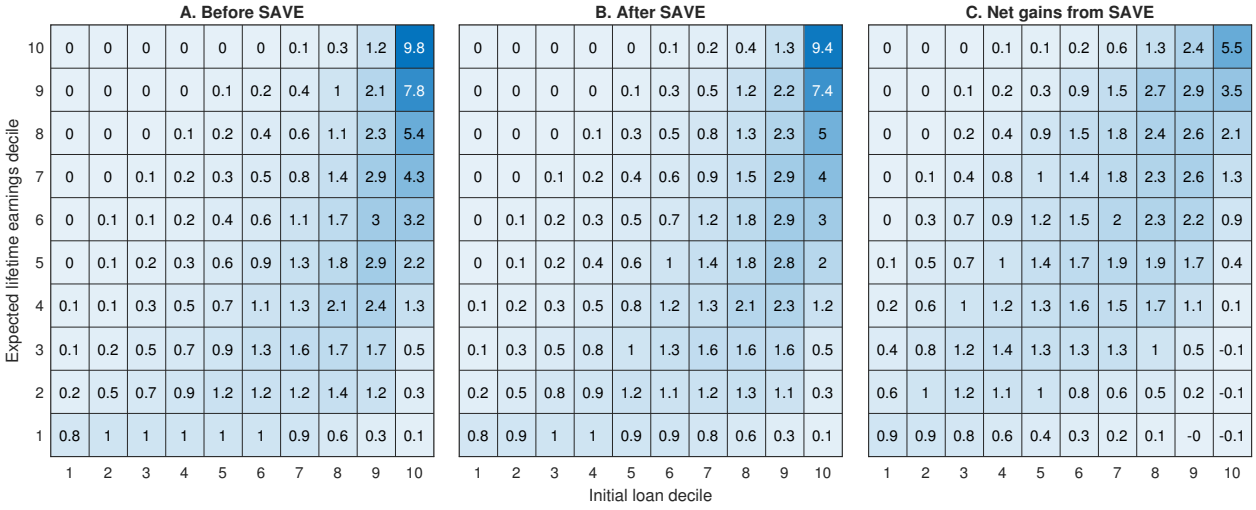
Panel B shows that gains from IDR are larger after SAVE and are similarly distributed. Panel C presents the difference between panels A and B, which are the net gains from SAVE for different borrowers. The panel shows that incremental gains from SAVE are slightly less concentrated. For instance, before SAVE, a borrower in the median decile of the earnings distribution saw welfare gains that were 10 times larger if they were in the top decile of student

debt than if they were in the median decile. In contrast, Panel C shows that the incremental gains from SAVE are more evenly distributed. This can be attributed to SAVE's increased generosity towards borrowers with undergraduate debt, whereas the primary beneficiaries of previous IDR rules were those with large balances accumulated through graduate school loans.

Interestingly, our findings suggest that SAVE actually reduces welfare for borrowers with very large balances and very low expected earnings. From the government's perspective, these student loans function in a fashion similar to a deeply out-of-the-money call option on the borrower's future earnings. For graduate school debt, the effect of SAVE on this call option is ambiguous. On one hand, SAVE raises the earnings threshold above which borrowers begin repaying their loans. On the other hand, IDR payments under SAVE are no longer capped by the standard repayment plan, and forgiveness occurs after 25 years rather than 20. In practice, however, very few graduate borrowers fall within the lowest deciles of expected lifetime earnings. The distribution of welfare gains on a per-borrower basis does not fully capture the distribution of aggregate welfare gains across the population. This is because borrowers with larger balances generally have higher expected earnings, meaning the population is not uniformly distributed across combinations of income and debt deciles.

Figure 8: Distribution of Welfare Gains by Decile of Debt and Income (in %)

This figure reports the share of total welfare gains from IDR by decile of expected lifetime earnings and decile of student loan at graduation, taking into account the population distribution, and in percentage of the aggregate gains. Panels A and B reports the distribution of welfare gains before and after SAVE respectively. Panel C reports the distribution of welfare gains from the incremental changes in IDR induced by SAVE. Deciles of lifetime earnings are defined within the population of borrowers.



To account for the distribution of borrowers, Figure 8 reports the percentage share of aggregate welfare gains by decile of expected earnings and debt at graduation. Panel A reveals that approximately 27% of the total welfare gains come from borrowers in the top decile of the debt distribution and the top four deciles of the expected earnings distribution. As shown in Panel B, this pattern remains similar, though somewhat less pronounced after the introduction of SAVE. Panel C demonstrates this shift by reporting the distribution of gains from the incremental changes in IDR induced by SAVE. This highlighs that the incremental gains from SAVE are distributed less disproportionately across borrowers.

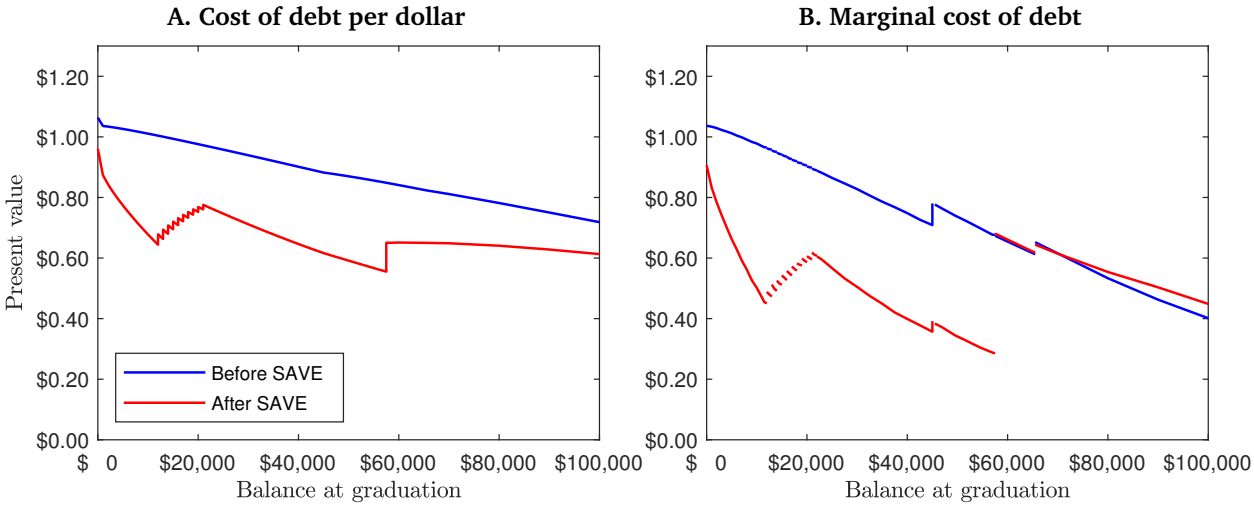
6.4 Cost of Debt and Borrowing Incentives

Table 1 shows that IDR reduces the cost of student debt by transferring taxpayer dollars to borrowers. This reduction introduces potential moral hazard, affecting the incentives of both students and schools. Students' incentives may change by inducing more students to borrow and borrowers to accumulate larger balances, as they will repay a smaller fraction of their debts. At the school level, this change in effective prices borne by students may lead schools

to increase program costs, potentially driving up tuition fees (Eaton et al., 2020).

Figure 9: Average and Marginal Cost of Student Debt

Panel A reports the average simulated present value of student debt at graduation, per dollar of debt, before and after SAVE. Panel B reports the mean marginal cost of student debt, defined as the present value of borrowing one more dollar just before graduation. The marginal cost of debt is infinite at points where the number of repayment years increases.



To explore this issue, Panel A of Figure 9 reports the expected cost of debt across borrowers as a function of their balance at graduation. The top solid line shows the cost of debt per dollar before SAVE, when less generous IDR plans existed. The bottom line shows the cost of debt following the introduction of the new plan. Prior to SAVE, student debt was a relatively fairly-priced source of financing. For instance, the average present value of debt at graduation slightly exceeded its face value for balances under $15,000. Even at balances around $40,000, borrowers who optimized their repayment strategies could expect to repay 90% of their debt. For unsubsidized loans, these repayment percentages would be slightly higher— and closer to 100%—when expressed as a percentage of the disbursed amount, since the balance at graduation includes interests accrued over the period of study.

The introduction of SAVE significantly altered these incentives for undergraduate debt by substantially lowering the cost of debt. For example, students graduating with the undergraduate debt limit of $57,500 can now expect to repay only 55% of this amount in present value, imposing a cost of $25,875 on taxpayers. From a present value perspective, this means students may forgo potential financial benefits by not borrowing as much as they could.

SAVE has also complicated the relationship between the present value of debt and its face value at graduation. First, the average cost of debt experiences discrete jumps every $1,000 between $12,000 and $20,000 because the repayment period extends by one year with each increment, up to 20 years. Second, students with debt from graduate school must make payments during 25 years before forgiveness. In Figure 9, we assume borrowers fully utilize their undergraduate loan limits before incurring graduate debt. Consequently, the repayment period jumps from 20 to 25 years at the undergraduate debt limit of $57,500, causing a jump in the present value of debt beyond this threshold. The slope also flattens significantly above the limit because the rate at which discretionary earnings are "taxed" is a weighted average of the undergraduate rate of 5% and the graduate school rate of 10%. Nevertheless, the cost of debt remains below its face values.

Panel B reports the average marginal cost of debt, highlighting the moral hazard at the intensive margin. Even before SAVE, the average marginal cost of debt turned negative for balances exceeding $7,000, potentially encouraging some students to borrow more. However, the implied subsidy was relatively small and could have been offset by factors such as accruing interest, inattention, debt aversion, or other frictions. Post-SAVE, the average marginal cost of debt remains below $1 and drops below 50 cents on the dollar for undergraduate students borrowing over $30,000. On the other hand, SAVE increases the marginal cost of graduate school debt by extending the repayment period from 20 to 25 years. Before SAVE, the marginal cost of debt only increased discontinuously at points where the marginal interest rate shifted. After SAVE, the marginal cost also exhibit discontinuities at points where changes occur in the payment formula or the length of the repayment period.

6.5 Potential Cost of SAVE to Taxpayers

How will changes in borrowing incentives impact the fiscal cost of IDR programs? This crucially depends on the behavioral response of borrowers and schools in the face of new incentives. Understanding the effect of IDR rules on borrowing behavior lies outside the scope of our model. However, the significant reduction in the cost of borrowing under the SAVE plan is likely to have an impact on its overall cost to taxpayers, not only through direct budgetary costs, but also by changing who borrowers and how much they borrow.

To estimate this potential fiscal impact, we consider three distinct scenarios, with a detailed discussion of our methodology provided in Appendix Section A.4. Appendix Table A.3 reports the decomposition of our computations. In the first baseline scenario, we assume no incentives effects of changes in IDR rules. In other words, we assume that neither the number of borrowers nor the distribution of student loans at graduation changes. Under this assumption, the cost of implementing the SAVE plan is equal to the change in net present value (approximately $6,200) per borrower, multiplied by the total number of borrowers (approximately 2.5 million). This yields an aggregate cost of nearly $15.3bn per cohort per cohort.

In the second scenario, we explore the potential for moral hazard among current borrowers. In this case, borrowers maximize their undergraduate debt but the number of borrowers does not change. As illustrated in Figure 9, the introduction of the SAVE plan creates a strong incentive for college students to increase their borrowing, though it does not make graduate debt more affordable. In this analysis, we assume that the lower four deciles of student loan balances correspond to individuals graduating from community colleges or those who did not complete college, which represent approximately 35% of borrowers in the data. We further assume that these students would increase their borrowing to $20,000, the limit for two years of higher education. Similarly, we assume that the next three deciles represent college graduates, who would borrow at least the four-year limit of $45,000. Graduate students, occupying the upper deciles, are assumed to have borrowed at least this amount during their undergraduate years. Under these assumptions, the additional cost of the SAVE plan would rise to $22.8bn.

Finally, we consider a scenario in which there is both moral hazard among current borrowers, and a change in the selection of borrowers. In this setting, all undergraduate students become borrowers and adopt the borrowing behavior described in the second scenario. The cost per new borrower is calculated as the difference between the expected present value of their payments and their outstanding balance at graduation. This scenario introduces approximately 3.2 million new borrowers, with an average cost to taxpayers of nearly $17,000 per borrower, resulting in a total cost of up to $77.3bn per cohort. This scenario represents an upper bound on the potential fiscal impact of the SAVE plan.

In this section, we have discussed how existing plans affect welfare, and through which channels. We showed that new plans did little to increase welfare, and mainly served as a

transfer from taxpayers to (generally affluent) borrowers. This analysis raises the natural question of how student loan repayment plans should be designed. We explore this question in the next section.

7 Policy Design

In this section, we analyze the mechanisms through which the calibration of IDR plans impacts both borrower welfare and the fiscal cost to taxpayers. We begin by examining how variations in key program parameters—such as the repayment duration, payment threshold, and payment rate—affect borrower welfare through different mechanisms. Next, we identify the parameter vector that maximizes borrower welfare without increasing the taxpayer burden, as well as the vector that maximizes borrower surplus net of cost to taxpayers. Our findings indicate that current IDR plans are not cost-efficient and that alternative calibrations could provide greater benefits to borrowers while reducing costs for taxpayers.

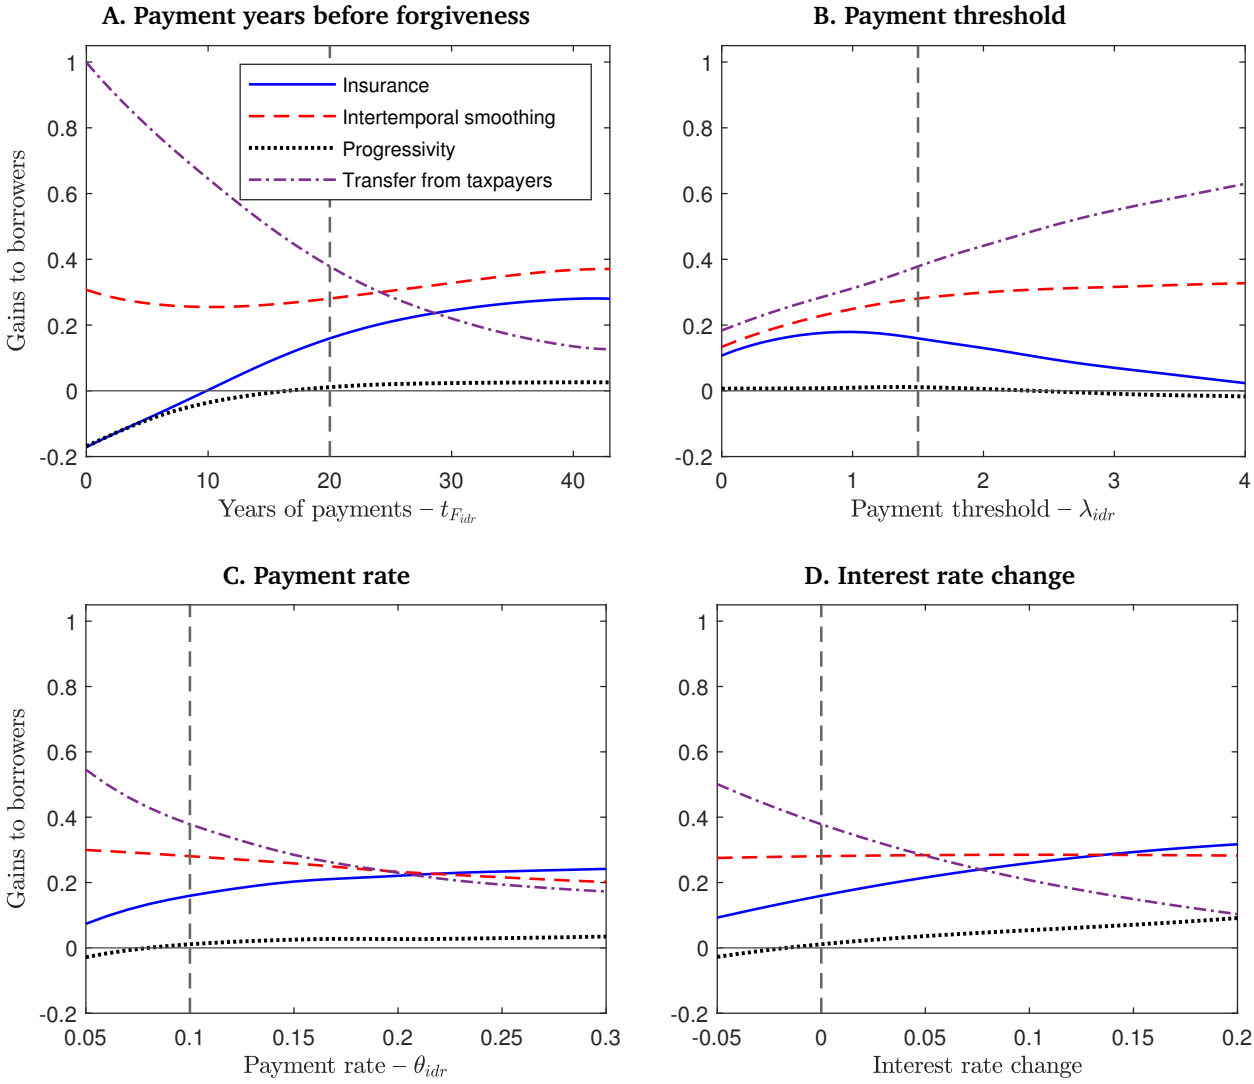
To circumvent the high computational cost of this analysis, we rely on an auxiliary model of the relationship between model parameters and welfare moments, adapting the method developed in Catherine et al. (2023). Appendix A.3 describes and validates this methodology.

7.1 The Role of IDR Parameters

The effects of IDR are governed by three key parameters: the number of payments years before forgiveness, the threshold above which earnings are deemed discretionary, the share of discretionary earnings allocated to payments. The interest rate on loans can also be used to balance the cost of the student debt program. How do these policy parameters influence the way IDR improve borrower welfare and its cost to taxpayers? Figure 10 reports the evolution of the four components of welfare when we vary each of these four policy parameters independently, holding the remaining four to their baseline value before the introduction of the SAVE program.

Figure 10: Role of IDR Parameters

This figure illustrates how the four components of welfare gains evolve with four IDR parameters. Each parameter is varied independently, with the remaining three held constant at their baseline values prior to the implementation of the SAVE program, indicated by vertical dashed lines. Panel A reports how the four component of welfare depend on the number of payment years before forgiveness. Panel B illustrates the relationship with the payment threshold, expressed as a percentage of the federal poverty line. That is the parameter λ_{idr} in equation (10). Panel C examines the effects of the discretionary income share allocated to payments. That is the parameter θ_{idr} in equation (10). Panel D explores the changes in welfare components in response to a uniform increase in interest rate for all loan amounts.



Years of Payment Before Forgiveness $\left(t_{F_{\text{idr}}}\right)$ Panel A shows that as the number of years before forgiveness increases, the transfer from taxpayers to borrowers decreases. This inverse relationship levels off as the repayment period extends, since many borrowers have repaid

their loans in full by this point, making additional years in the repayment period largely irrelevant. Importantly, while the generosity of IDR declines with a longer forgiveness timeline, it does not negatively impact welfare through the intertemporal smoothing channel. In fact, the welfare of borrowers improves through this channel as they are less likely to face borrowing constraints in later years, meaning they are more capable of managing extended repayment periods without significant distress. This means that IDR primarily benefits borrowers with lower-than-expected lifetime earnings, while for others, the main advantage lies in deferring repayments further into the future, a period when financial constraints are less binding.

Payment Threshold (λ_{idr}) Panel B presents the effect of the income threshold, expressed as a proportion of the federal poverty line, above which IDR payments begin. As expected, raising the payment threshold results in lower payments and a greater transfer from taxpayers to borrowers. Perhaps more interestingly, the relationship between the welfare insurance component and the payment threshold is hump-shaped. To understand this, consider the extremes: if the threshold is set to zero, there is no downside protection for borrowers, as they are taxed from the first cent of earnings. Conversely, if the threshold is infinite, borrowers will never make a payment, disconnecting the benefits of IDR from their earnings trajectory. Thus, as the payment threshold increases to a certain point, the present value gains it generates tend to accrue even under typical circumstances, rendering them less dependent on borrowers facing negative income shocks. Consequently, these gains function more as a pure subsidy from taxpayers rather than as insurance. This explains why the introduction of the SAVE plan diminishes the IDR program's capacity to provide insurance against income risk, as shown in Table 1. Panel B also reveals that increasing the threshold enhances the benefits of intertemporal smoothing. As earnings generally rise with age, a higher payment threshold effectively defers payments to later in the life cycle.

Payment Rate (θ_{idr}) Panel C shows that increasing the share of discretionary earnings allocated to payments reduces the generosity of IDR, lowering taxpayer transfers but enhancing the insurance component, as borrowers below the payment threshold remain unaffected. A low payment rate leads to a deficit, shifting costs to taxpayers, whereas a higher rate recoups more

costs from high-income borrowers. This dynamic explains why the insurance and taxpayer transfer components of borrower welfare evolve inversely and why increasing the payment rate enhances IDR's progressivity. However, it does not imply that low earners benefit more in absolute terms, but rather that they receive a larger share of the program's overall benefits.

Loan Interest Rate Finally, Panel D examines the impact of increasing the interest rate on student loans. As anticipated, higher interest rates lower the cost of student debt for taxpayers. Additionally, we find that higher interest rates enhance the insurance value of IDR. When IDR payments are lower than those of the standard plan, payments are deferred and accrue interest at the loan rate. For borrowers with modest earnings, accruing interest may have minimal impact if they reach the forgiveness period before fully repaying their loan. However, for borrowers who experience a significant rise in earnings over time, the cost of deferring payments becomes substantial. As a result, higher interest rates disproportionately increase the debt burden for those whose earnings exceed expectations. For similar reasons, higher interest rates also increase the progressivity of IDR.

7.2 Optimal Policy Calibrations

So far, we have discussed how adjusting various policy parameters affects welfare and transfers. The next natural step is to determine the optimal policy parameters. To explore potential improvements over existing IDR rules, we analyze the space of IDR parameters to maximize welfare gains. We consider two scenarios: first, a budget-neutral policy where expected aggregate repayments match those under the standard ten-year plan, and second, a policy that maximizes net welfare while accounting for costs to taxpayers. Additionally, we show that borrower welfare can be increased without imposing additional costs on taxpayers by adjusting program parameters.

We make several assumptions in determining optimal IDR policies. Consistent with previous IDR reforms, we keep the student loan interest rate unchanged, ensuring the terms of the standard repayment plan remain intact. This reflects institutional realities, where student loan interest rates are set by Congress, while new IDR plans are created by the Department of Education. By maintaining these terms, we ensure that no borrower is disadvantaged by the IDR

program. Therefore, we optimize welfare with respect to three key parameters: the repayment period ($t_{F_{\text{idr}}}$), the payment threshold (λ_{idr}), and the payment rate (θ_{idr}), while keeping other aspects of the program the same as before the introduction of SAVE.

We impose two constraints on the values of these parameters. First, consistent with current law, households cannot borrow against Social Security benefits, so we assume earnings beyond retirement age cannot be used to repay student debt, imposing $t_{F_{\text{idr}}} \leq 43$. Second, since federal and state income taxes can combine to reach a marginal tax rate of 40%, we impose $\theta_{\text{idr}} \leq 0.6$, ensuring the overall marginal "tax" rate does not exceed 100%. Table 2 presents the optimal IDR parameters for the two policies.

Table 2: Optimal IDR Parameters

This table reports estimated parameters maximizing welfare under two policy constraints. The "budget-neutral" policy maximizes borrower welfare at no cost to taxpayers relative to the benchmark economy without IDR. The "highest net welfare" policy maximizes borrower welfare net of the cost to taxpayers.

	Years of payment	Payment threshold	Payment rate
Policy	$t_{F_{\text{idr}}}$	λ_{idr}	θ_{idr}
Budget-neutral	43	150%	33%
Highest net welfare	43	284%	32%

Budget-neutral Policy We first consider a budget-neutral version of IDR, where the aggregate expected repayment matches that of the standard plan. Under this constraint, the optimized parameters differ significantly from current IDR rules. Specifically, borrowers would face a 43-year repayment period before loan forgiveness, compared to the current 20 years. The payment threshold would remain at 150% of the federal poverty line, and the program would collect 33% of discretionary income, up from 10%. As a budget-neutral policy, this approach would reduce the expected cost of IDR for taxpayers from $24,000 per borrower to zero.

Table 3 details the corresponding welfare gains and their components. The budget-neutral policy generates welfare gains of 0.579 ($36,937) from intertemporal smoothing and insurance, which exceeds the gains from these channels under the current IDR rules, 0.406 ($25,900). Additionally, it performs slightly better when welfare gains from progressivity are included,

yielding gains of 0.598 ($38,149), which is again greater than under the existing rules. These results demonstrate that improvements in consumption smoothing over time and across states delivered by IDR can be achieved without imposing additional costs on taxpayers.

Total borrower welfare would decline, but this is only due to the transfer from taxpayers being eliminated. When taxpayer subsidization is considered, borrowers benefit more under the existing IDR rules, with welfare increasing by 0.793 ($50,589), compared to 0.598 ($38,149) under the budget-neutral policy. Nonetheless, the incremental improvement of 0.195 ($12,440) represents only half of the 0.377 ($24,051) cost to taxpayers, making aggregate welfare substantially higher under the budget-neutral policy when taxpayers' interests are taken into account. Furthermore, our results show that financial frictions can be effectively mitigated without subsidizing borrowing.

Table 3: Welfare Gains from Optimal IDR Parameters

This table reports the decomposition of aggregate welfare gains under the IDR policy parameters reported in Table 2. The decomposition is defined in Equation (31). Welfare gains are estimated using the auxiliary model and reported in dollar equivalent per borrower and as a multiple of the national wage index ($63,795 in 2022). The Status Quo reports welfare gains under current rules (before SAVE).

	Total change in borrower welfare	Financial frictions		Transfer	
		Insurance	Intertemporal Smoothing	Progressivity	Mean
Budget-neutral	0.598	0.318	0.261	0.019	0.000
Highest net welfare	0.840	0.315	0.387	0.040	0.098
Status Quo	0.793	0.169	0.237	0.002	0.377

The decomposition of welfare effects provides two key insights. First, the budget-neutral policy exhibits slight progressivity, whereas current IDR rules are mildly regressive. Second, the budget-neutral policy delivers greater welfare benefits through both intertemporal consumption smoothing and insurance, with the latter being the dominant channel. The high payment threshold ensures that a smaller proportion of early-career earnings is allocated to loan repayment, offering liquidity during the post-graduation period. It also limits payments for low-income borrowers, thereby providing insurance against idiosyncratic income risks.

Policy Maximizing Net Welfare Tables 2 and 3 also report the IDR parameters and the decomposition of welfare gains for the policy that maximizes net welfare gains. This policy would be more progressive, featuring a long repayment period, a high repayment threshold, and a high repayment rate. The repayment period remains at its maximum of 43 years, and the payment rate is nearly unchanged. The additional budget is allocated to significantly raising the payment threshold, which increases from 150% to 284%. This higher threshold enhances gains from intertemporal smoothing by effectively reducing payments early in the life cycle, as it shields a larger portion of earnings. The taxpayer cost of this change remains moderate because the extended 43-year repayment period provides ample time for borrowers to repay. Under a shorter repayment timeline, the same change in the payment threshold would increase taxpayer losses more substantially, as a larger share of deferred payments would remain unpaid by the time of forgiveness.

Comparing the Pre-SAVE IDR rules to the net-welfare-maximizing policy reveals that borrowers would be relatively indifferent between the two, despite the latter costing approximately $18,000 less to taxpayers. Borrowers would prefer the SAVE plan, but mainly due to the large taxpayer subsidy. Our finding that the budget-neutral policy delivers higher net welfare gains than current IDR rules does not necessarily imply that the optimal IDR calibration excludes an implicit subsidy. While the government can transfer additional funds directly through grants or other mechanisms, policymakers may still choose to subsidize student borrowers if the gains for borrowers outweigh the costs to taxpayers.

Large subsidies in existing IDR plans raises the question of much much welfare can be increased given current subsidy rates. Figure 11 provides some insight into this question. The figure shows the magnitude of welfare gains that can be delivered to borrowers as a function of the cost of policy. More precisely, Figure 11 represents welfare to borrowers and net welfare gains under optimally calibrated parameters, as a function of taxpayer cost. The blue solid line represents welfare gains from IDR, while the dashed red line represents the net welfare gains after accounting for the taxpayer cost. Additionally, the figure highlights the welfare gains and costs associated with IDR both before and after the introduction of the SAVE program, illustrating the gap between these policies and cost-efficient calibrations. Net welfare gains peak for a per-borrower budget of .098 ($6,250), which indicates that policies reducing the

present value of debt by more than this amount increase marginal borrower welfare less than the cost to taxpayers.

Figure 11: Welfare Gains and Policy Budget Under Optimal Calibration

This figure illustrates the relation between the per-borrower budget allocated to IDR and welfare gains to borrowers (blue line) or net welfare gains (red dashed line) under the corresponding optimal calibration of IDR parameters. The black dot (•) marks the coordinates of the budget-neutral policy. Diamonds (◊) mark the coordinates of welfare gains under the net-welfare maximizing policy. Circles (∘) and stars (∗) represent Pre-SAVE and Post-SAVE IDR rules respectively. Welfare gains and cost to taxpayers are reported per borrower and as a multiple of the national wage index ($63,795 in 2022).

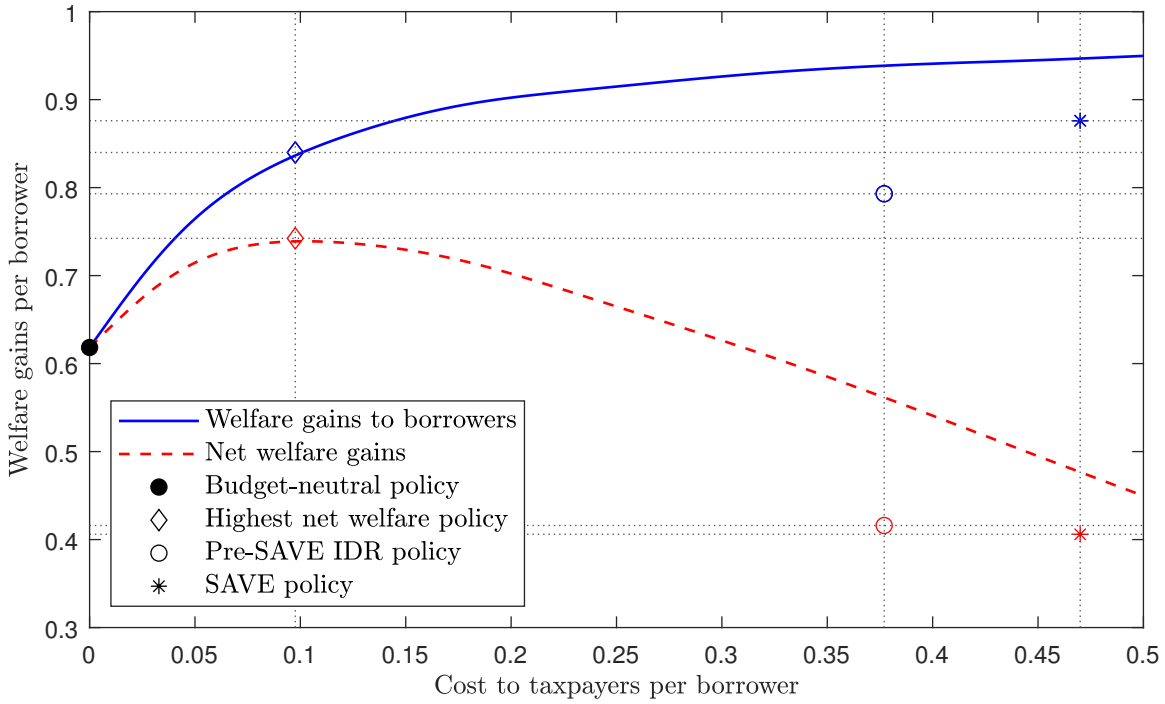


Figure 11 underscores several key findings. First, a significant portion of the welfare gains delivered by current IDR programs can be achieved at no cost to taxpayers. Second, while increasing taxpayer costs enables policymakers to deliver higher net welfare gains, the diminishing marginal returns cause net gains to peak at a budget significantly lower than the current policy's cost. For instance, comparing the position of Pre-SAVE IDR on the "welfare gains to borrowers" curve shows that the same level of welfare gains achieved before SAVE could be realized at a cost of approximately $4,150 per borrower, representing a savings of nearly $20,000 to taxpayers. Notably, despite introducing more complex repayment rules and additional de-

grees of freedom, the SAVE program still exhibits a significant degree of cost inefficiency.[16]

Why are optimal rules significantly more cost-efficient than existing ones? The answer has to do with the income trajectory of individuals over the life-cycle. Raising the payment threshold directs more financial relief to borrowers, while extending the repayment period and increasing the payment rate benefits taxpayers. From the taxpayers' perspective, every dollar holds the same value. However, borrowers place a much higher value on dollars saved through an increased payment threshold, as these savings occur in states of high marginal utility either early on in the life-cycle, or in other periods when income and consumption are low. Consequently, borrowers are willing to trade a dollar saved during low-income, high-utility periods or states for several dollars later in life or in more favorable earning scenarios. This trade enables the social planner to reduce the overall cost of IDR with minimal or no utility loss to borrowers.

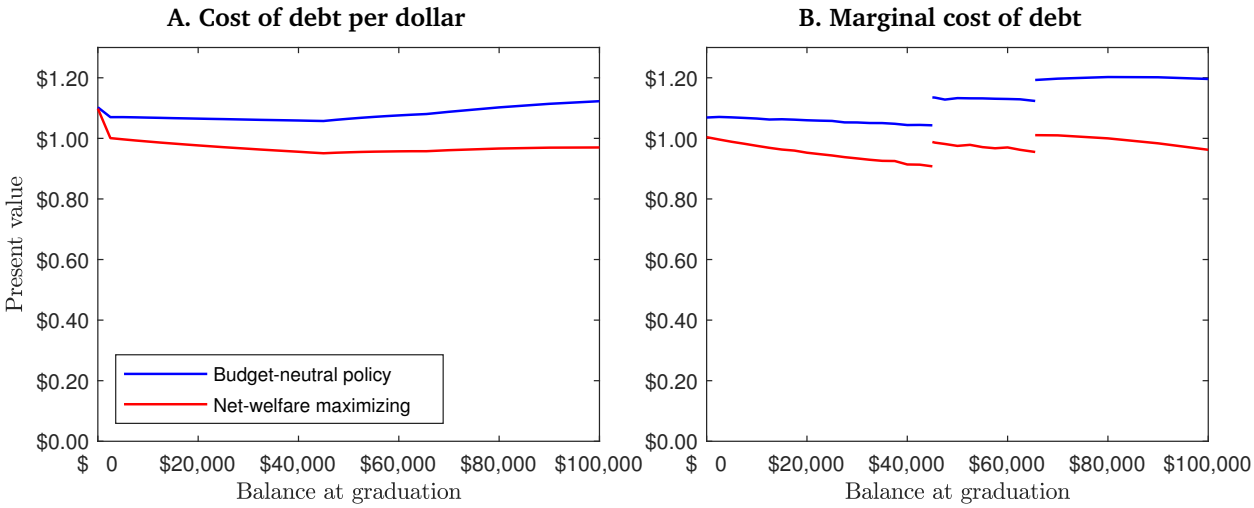
7.3 Discussion

Our policy analysis is subject to several caveats. First, we take the distribution of balances at graduation as given. However, as discussed previously, IDR can generate moral hazard by encouraging students to borrow more. Figure 12 shows that under both the optimal budget-neutral and net-welfare maximizing calibration of IDR, the average and marginal cost of debt remains above or close to one, thus muting the concern that borrowers would borrow more under these policies. This would also reduce the risk that part of the taxpayer's subsidization would be captured by colleges in the form of higher tuitions.

[16] Note that aggregate welfare can be decreasing, as we are limiting the policy tools here to three parameters. If we allowed lump sum transfers, such as one time targeted forgiveness, welfare would be weakly increasing (red dashed line could be kept flat.) Varying three policy parameters is likely closer to reality, as policy makers are unlikely to be able to make lump sum transfers due to legal constraints or administrative barriers leading to inability to target borrowers with high marginal utility. For example, the Supreme Court of the United States blocked student loan forgiveness in June 2023.

Figure 12: Average and Marginal Cost of Student Debt with Optimal Policy Parameters

Panel A reports the average simulated present value of student debt at graduation, per dollar of debt, under the budget-neutral and net welfare maximizing policies presented in Section 7.2. Panel B reports the mean marginal cost of student debt, defined as the present value of borrowing one more dollar just before graduation. The marginal cost of debt is infinite at points where the number of repayment years increases.



A second potential concern with our optimal policies, which feature significantly higher repayment rates than existing IDR rules, is the potential discouragement of labor supply. While we do not explicitly model labor supply, we believe this concern is largely unwarranted for two key reasons. First, under any repayment rate, borrower payments remain capped by the standard-plan formula. Second, empirical studies of the labor supply effects of IDR such as Britton and Gruber (2020) and de Silva (2023) tend to find no or very small effects. Finally, and more importantly, the two optimal policy calibrations we present are significantly less costly to taxpayers. High taxpayer costs are symptomatic of a policy calibration that leads to substantial loan forgiveness at the end of the repayment period. Under such calibration, many borrowers would expect a reduction in current payments to ultimately result in greater loan forgiveness, rather than increased future payments. This anticipation should reduce their optimal labor supply. In contrast, lower taxpayer costs suggest that most borrowers are expected to fully repay their loans. For these borrowers, a reduction in labor supply would not lead to forgiveness but rather to higher future payments. As a result, incorporating endogenous labor supply into the model might, in fact, widen the efficiency gap between our optimal calibration and existing IDR rules, further supporting our argument.

Another limitation of our model is that it does not take into account for the potential indirect cost of having large student debt balances for borrowers, such as the ability to obtain mortgage or the psychological cost of having a debt with large face value. One potential solution is to prevent negative amortization, something actually implemented under SAVE and preventing balances from ballooning.

8 Concluding Remarks

This paper studies a hitherto underexplored channel accounting for much of the dramatic rise in student debt over the past decades: the slowdown in repayments. Using administrative data, we find that payment deferral accounts for almost half of the six-fold increase in student loan balances between 2000 and 2020. We further explore how this payment slowdown affected borrowers' welfare. Through the lens of our model, we find that the majority of the welfare gains come from insurance, consumption smoothing over the life-cycle and redistributive transfers to student borrowers from taxpayers. Turning the model to optimal plan design, we find that the government can increase welfare gains from insurance and consumption smoothing at no additional cost to taxpayers. These gains could largely be achieved by increasing plan maturities, while at the same time making plans more progressive by raising repayment threshold and payment rates.

A natural avenue for future research is to further enrich and extend our framework to study optimal policy and the design of insurance programs for borrowers. Our results indicate that new and more generous IDR plans launched in 2023 increased borrower welfare, but almost entirely through transfers from taxpayers. Moreover, significant welfare gains can be achieved without fiscal costs, largely by extending payment maturities. Future work should determine optimal parameters for other parameters and contract schemes not studied in our setting. For example, a multitude of loan forgiveness schemes have been proposed by policy makers, current loan programs have various forbearance schemes and contracts could also include an equity rather than a debt component. Moreover, the framework may also be useful in measuring the welfare effects of government interventions in more general household loan programs, which frequently occur during recessions and crises.

References

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

Avery, Christopher and Sarah Turner, "Student loans: Do college students borrow too much–or not enough?," *Journal of Economic Perspectives*, 2012, *26* (1), 165–92.

Baily, Martin Neil, "Some aspects of optimal unemployment insurance," *Journal of Public Economics*, 1978, *10* (3), 379–402.

Beshears, John, James J. Choi, Christopher Clayton, Christopher Harris, David Laibson, and Brigitte C. Madrian, "Optimal Illiquidity," *Working Paper*, 2024.

Black, Sandra E, Jeffrey T Denning, Lisa J Dettling, Sarena Goodman, and Lesley J Turner, "Taking it to the limit: Effects of increased student loan availability on attainment, earnings, and financial well-being," Technical Report, National Bureau of Economic Research 2020.

Boutros, Michael, Nuno Clara, and Francisco Gomes, "Borrow now, pay even later: A quantitative analysis of student debt payment plans," *Pay Even Later: A Quantitative Analysis of Student Debt Payment Plans (October 12, 2022)*, 2022.

Britton, Jack and Jonathan Gruber, "Do income contingent student loans reduce labor supply?," *Economics of Education Review*, 2020, *79*, 102061.

Catherine, Sylvain and Constantine Yannelis, "The distributional effects of student loan forgivenesss," *Journal of Financial Economics*, 2022.

_ , Max Miller, and Natasha Sarin, "Relaxing household liquidity constraints through social security," *Journal of Public Economics*, 2020, *189*, 104243.

_ , Mehran Ebrahimian, David Sraer, and David Thesmar, "Robustness checks in structural analysis," 2023. Working paper.

CBO, "Income-driven repayment plans for student koans: Budgetary costs and policy options," *Congressional Budget Office Discussion Paper*, 2020, (14).

Chetty, Raj, "A general formula for the optimal level of social insurance," *Journal of Public Economics*, 2006, *90* (10), 1879–1901.

de Silva, Tim, "Insurance versus moral hazard in income-contingent student loan repayment," *Available at SSRN 4614108*, 2023.

Dinerstein, Michael, Constantine Yannelis, and Ching-Tse Chen, "Debt moratoria: Evidence from student loan forbearance," *American Economic Review: Insights*, 2024.

Eaton, Charlie, Sabrina T Howell, and Constantine Yannelis, "When investor incentives and consumer interests diverge: Private equity in higher education," *The Review of Financial Studies*, 2020, *33* (9), 4024–4060.

Findeisen, Sebastian and Dominik Sachs, "Education and optimal dynamic taxation: The

role of income-contingent student loans," *Journal of Public Economics*, 2016, *138*, 1–21.

Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer, "The value of medicaid: Interpreting results from the Oregon health insurance experiment," *Journal of Political Economy*, 2019, *127* (6), 2836–2874.

Gomes, Francisco, Michael Haliassos, and Tarun Ramadorai, "Household finance," *Journal of Economic Literature*, 2021, *59* (3), 919–1000.

Goodman, Sarena, Adam Isen, and Constantine Yannelis, "A day late and a dollar short: Liquidity and household formation among student borrowers," *Journal of Financial Economics*, 2021, *142* (3), 1301–1323.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption risk over the life-cycle," *Econometrica*, 2002, *70* (1), 47–89.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What do data on millions of U.S. workers reveal about lifecycle earnings dynamics?," *Econometrica*, 2021, *89* (5), 2303–2339.

Hamdi, Naser, Ankit Kalda, and David Sovich, "The labor market consequences of student loan forbearance," *Available at SSRN 4787183*, 2024.

Herbst, Daniel, "Liquidity and insurance in student-loan contracts: The costs and benefits of income-driven repayment," *Unpublished Mimeo*, 2018.

Hopenhayn, Hugo A. and Juan Pablo Nicolini, "Optimal unemployment insurance," *Journal of Political Economy*, 1997, *105* (2), 412–438.

Keys, Benjamin J, Neale Mahoney, and Hanbin Yang, "What determines consumer financial distress? Place-and person-based factors," Technical Report, National Bureau of Economic Research 2020.

Laibson, David I., Andrea Repetto, and Jeremy Tobacman, "Self-control and saving for retirement," *Brookings Papers on Economic Activity*, 1998, *29* (1), 91–196.

Lochner, Lance and Alexander Monge-Naranjo, "The nature of credit constraints and human capital," *American Economic Review*, 2011, *101* (6), 2487–2529.

_ and _, "Student loans and repayment: Theory, evidence, and policy," in "Handbook of the Economics of Education," Vol. 5, Elsevier, 2016, pp. 397–478.

Looney, Adam and Constantine Yannelis, "A crisis in student loans? How changes in the characteristics of borrowers and in the institutions they attended contributed to rising loan defaults," *Brookings Papers on Economic Activity*, 2015, (Fall), 1–68.

_ and _, "What Went Wrong With Federal Student Loans?," *Journal of Economic Perspectives*, 2024, *26*.

Lucca, David O, Taylor Nadauld, and Karen Shen, "Credit supply and the rise in college tuition: Evidence from the expansion in federal student aid programs," *The Review of Financial*

Studies, 2019, *32* (2), 423–466.

Maggio, Marco Di, Ankit Kalda, and Vincent Yao, "Second chance: Life without student debt," *Journal of Finance*, forthcoming.

Merton, Robert C., "Optimum consumption and portfolio rules in a continuous-time model," *Journal of Economic Theory*, December 1971, *3* (4), 373–413.

Monarrez, Tomás and Lesley J Turner, "The effect of student loan payment burdens on borrower outcomes," 2024.

Mueller, Holger and Constantine Yannelis, "The rise in student loan defaults," *Journal of Financial Economics*, 2019, *131* (1), 1–19.

_ and _ , "Increasing enrollment in income-driven student loan repayment plans: Evidence from the Navient field experiment," *The Journal of Finance*, 2022, *77* (1), 367–402.

Shimer, Robert and Iván Werning, "Reservation wages and unemployment insurance*," *The Quarterly Journal of Economics*, 08 2007, *122* (3), 1145–1185.

Stantcheva, Stefanie, "Optimal Taxation and Human Capital Policies over the Life Cycle," *Journal of Political Economy*, 2017, *125* (6), 1931–1990.

Yannelis, Constantine and Greg Tracey, "Student loans and borrower outcomes," *Annual Review of Financial Economics*, 2022, *14*, 167–186.

A Model appendix

A.1 Model calibration

A.1.1 Income process

Table A.1: Calibration of Income Process

Parameter	Value	Parameter	Value
ρ	0.983	$\sigma_{\varepsilon,1}$	0.294
p_z	0.267	$\sigma_{\varepsilon,2}$	0.065
$\mu_{\eta,1}$	-0.194	σ_α	0.467
$\sigma_{\eta,1}$	0.444	a	-2.8571
$\sigma_{\eta,2}$	0.076	b	-0.7788
$\sigma_{z_1,0}$	0.495	c	-4.5407
p_ε	.092	d	-1.3702
$\mu_{\varepsilon,1}$	0.352		

A.1.2 Income taxes

In working life, the marginal income tax rates paid on labor income Y_{it} are as follows:

$$\text{Marginal Tax Rate}_{it} = \begin{cases} 0.10 & \text{if } Y_{it} < 0.18Y_{1t}, \\ 0.12 & \text{if } 0.18Y_{1t} < Y_{it} < 0.72Y_{1t}, \\ 0.22 & \text{if } 0.72Y_{1t} < Y_{it} < 1.54Y_{1t}, \\ 0.24 & \text{if } 1.54Y_{1t} < Y_{it} < 2.94Y_{1t}, \\ 0.32 & \text{if } 2.94Y_{1t} < Y_{it} < 3.73Y_{1t}, \\ 0.35 & \text{if } 3.73Y_{1t} < Y_{it} < 9.32Y_{1t}, \\ 0.37 & \text{if } Y_{it} > 9.32Y_{1t}. \end{cases} \tag{1}$$

In retirement, the exact same marginal tax rates and tax brackets are applied to Social Security benefits B_{it}^{SS} instead of labor income.

A.1.3 Life cycle profiles

Table A.2: Regressions at the Household Level on Years Since Head Graduated (Age')

	Kids per adults	Log(Labor Inc)
Age'	0.3214***	0.0762**
	(0.0058)	(0.0237)
$\frac{\text{Age'}^2}{10}$	-0.0611***	-0.0011
	(0.0011)	(0.0056)
$\frac{\text{Age'}^3}{100}$	0.0035***	-0.0010*
	(0.0001)	(0.0004)
Log(Initial Balance)		0.2068***
		(0.0091)
Debt Dummy		0.0228
		(0.0147)
Constant	-4.4926***	-2.1048***
	(0.0911)	(0.3225)
R-sqr	0.174	0.073
N	40841	25419

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

A.2 Numerical resolution of the model

The persistent component of log earnings z can take 61 equally-spaced between $z_{\min} = -2.5$ and $z_{\max} = 2.5$. We discretize the motion of z as a Markov chain. Specifically, we use Tauchen's method to estimate the transition matrix conditional on either of the two normal distributions from which the persistent shock η can be drawn. Then, we use a probability-weighted average of the two conditional transition matrix to discretize the normal mixture. The transitory shock is modeled similarly but, because of its small quantitative importance, we assume that it can take only five equally-spaced values between -.7 and .7.

Wealth W is scaled by national wage index Y_1 and takes values between the logs of $w_{\min} = 0.01$ and $w_{\max} - 30$. Current log wealth can take 75 equally-spaced values between the limits. Optimal consumption is determined by the choice of next period log wealth, which can take 301 equally-spaced values between the same boundaries, plus points outside of the grids down to $w_{\min} - 0.75$ and up to $w_{\max} + 0.75$. When solving the Bellman equation backward, we compute the value of V_{t+1} on the thinner grid assuming that $\log(V_{t+1}(1-\gamma))$ is locally linear in $\log W$

between points of the coarse grid. We extrapolate outside of the coarse grid boundaries using a second-order Taylor expansion.

Current loan balance takes values between 0.01 and 3 times the initial borrowed amount. We discretize the log of L using between 21 and 53 equally-spaced points, depending on the initial amount borrowed. This corresponds approximately to at most increments of .10 on the log scale for the largest loans.

The model is then solved by dynamic programming starting from retirement age.

A.3 Auxiliary model

We conduct our policy analysis using an approximation of the relationship between IDR policy parameters and moments of the welfare function. Catherine et al. (2023) show how to build an auxiliary model that approximates the relationship between moments and model parameters and then use the auxiliary model to estimate structural models quickly by the method of simulated moments. Estimating a structural model or finding optimal policy parameters are somewhat similar exercises which consists in exploring a parameter space to minimize an objective function. Therefore, we adapt their methodology to run policy experiments.

Throughout the paper, we partition the population of borrowers into 11 groups by initial loan size. Then, we simulate data for each group before computing welfare for the entire population of borrowers. From equation (31), we can compute and decompose welfare for each of these groups g as:

$$
\begin{aligned}
\frac{\Delta^{\$}_{p,\text{net}} V_g}{I_g} = {} & \frac{\overbrace{\frac{1}{I_g}\sum_i \sum_t \mathbb{E}_{it}\left[\nu' \ \Delta CP_{t_0,t}\right]}^{\text{m1}_g}}{\mathbb{E}[\nu']} - \frac{\overbrace{\frac{1}{I_g}\sum_i \sum_t \mathbb{E}_{it}[\nu'] \ \mathbb{E}_{it}\left[\Delta CP_{t_0,t}\right]}^{\text{m2}_g}}{\mathbb{E}[\nu']} \\
& + \frac{\overbrace{\frac{1}{I_g}\sum_i \sum_t \mathbb{E}_{it}[\nu'] \ \mathbb{E}_{it}\left[\Delta CP_{t_0,t}\right]}^{\text{m2}_g}}{\mathbb{E}[\nu']} - \frac{\overbrace{\frac{1}{I_g}\sum_i \sum_t \mathbb{E}_{i}[\nu'] \ \mathbb{E}_{i}\left[\Delta CP_{t_0,t}\right]}^{\text{m3}_g}}{\mathbb{E}[\nu']} \\
& + \frac{\overbrace{\frac{1}{I_g}\sum_i \sum_t \mathbb{E}_{i}[\nu'] \ \mathbb{E}_{i}\left[\Delta CP_{t_0,t}\right]}^{\text{m3}_g}}{\mathbb{E}[\nu']} - \frac{\overbrace{\sum_t \mathbb{E}[\nu'] \ \mathbb{E}\left[\Delta CP_{t_0,t}\right]}^{\text{m4}_g}}{\mathbb{E}[\nu']} \\
& + \frac{\overbrace{\sum_t \mathbb{E}[\nu'] \ \mathbb{E}\left[\Delta CP_{t_0,t}\right]}^{\text{m4}_g}}{\mathbb{E}[\nu']} - \frac{\overbrace{\sum_t -\Delta R_{L,t} P_{t_0,t}}^{\text{m5}_g}}{\mathbb{E}[\nu']}.
\end{aligned}
$$

where I_g is the number of borrowers in group g. Our welfare decomposition is a function of five moments m1_g, m2_g..., m5_g. Each of the moments are averages across borrowers. We

denote a sixth moment $\text{m6}_g = \mathbb{E}[\nu']^{-1}$: the inverse of the average marginal utility of money. We denote $\mathbf{m} = \{m1...m6\}$ the vector of these six moments.

If $\mathbf{m}$ is known for all groups of borrowers, its full-population counterpart $\mathbf{M}$ can be computed as:

$$\text{Mk} = \text{M6} \sum_g \omega_g \frac{\text{mk}_g}{\text{m6}_g} \qquad \forall k \leq 5, \tag{2}$$

where ω_g is the population weight of group g and M6 is the inverse of the average marginal value of money over the entire population of borrowers:

$$\text{M6} = \left(\sum_g \frac{\omega_g}{\text{m6}_g} \right)^{-1}. \tag{3}$$

We denote $f(\xi)$ the function that determines of the vector of moments $\mathbf{m}$ for borrowers with balance at graduation L_{t_0} and policy parameters $\xi = \left\{L_{t_0}, \theta_{\text{idr}}, \lambda_{\text{idr}}, t_{F_{\text{idr}}}, \Delta r_L\right\}$. An evaluation of f takes on average ten minutes and, for any IDR policy, the function must be computed for all 11 bins of L_{t_0}. Consequently, the computation cost of running comparative statics or finding optimal policy parameters is high. Therefore, we approximate f with an auxiliary function f^{approx} that can be computed without solving the dynamic model.

Following one of the strategies explored in Catherine et al. (2023), we assume f^{approx} to be a set of six cubic polynomials. Each polynomial predicts the value of a moment of $\mathbf{m}$ as a function of the vector of parameters ξ. The function $f^{\text{approx}}(\xi)$ is computed as follows:

1. We start by building a training sample that evaluate the function f from the full economic model for $N = 1{,}000$ combinations of ξ and L_{t_0}. Specifically, we use the Halton quasi-random sequence over an hypercube with boundaries $L_{t_0} \in [0.01, 4.5]$, $\theta_{\text{idr}} \in [0.05, 0.30]$, $\lambda_{\text{idr}} \in [0, 4]$, $t_{F_{\text{idr}}} \in [5, 40]$, and $\Delta r_L \in [-.04, 0.20]$. This training stage is computationally expensive but done only once.

2. For each query vector ξ_q to be evaluated, we estimate $f^{\text{approx}}(\xi_q)$ in two stages:

   (a) First, we compute the coefficients of the cubic polynomial functions by weighted-least-square using the training sample. To improve accuracy, we put more weight on training data points close to the area of interest. We define the distance between the vector of parameters being evaluated ξ_q and each training data point j as:

   $$D_j = \left(\frac{\xi_j - \xi_q}{\xi_0} \right)^T \left(\frac{\xi_j - \xi_q}{\xi_0} \right) \tag{4}$$

   where $\xi_0 = \{.5, .1, 1.5, 20, 1\}$ is a scaling vector. We use the inverse of D_j as regres-

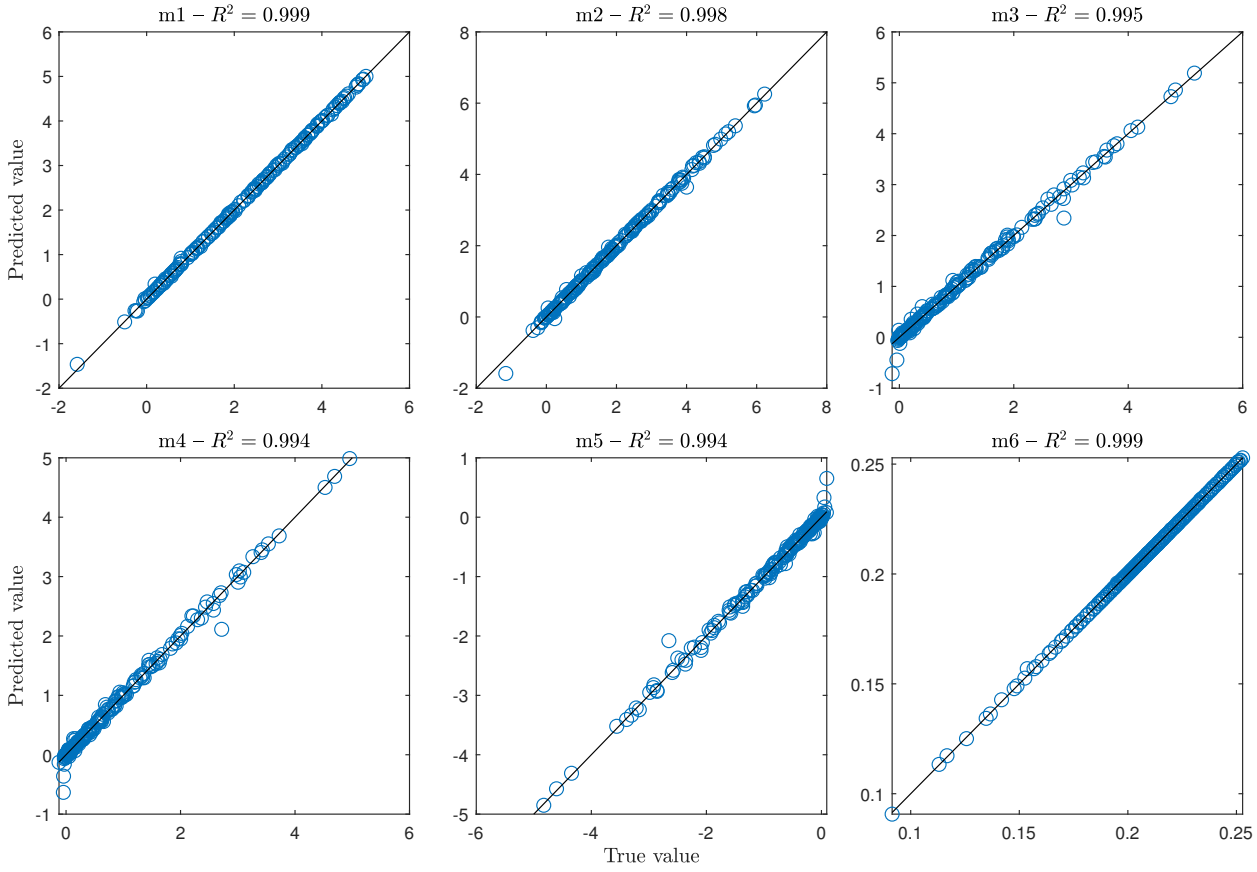
sion weights[1] to the estimate polynomial coefficients defining f^{approx}.

(b) Finally, we predict the vector of welfare moments as $\hat{\mathbf{m}}(\xi_q) = f^{\text{approx}}(\xi_q)$.

To validate the auxiliary model on 100 data points from the training data, using the 900 other points as training data to estimate f^{approx} and predict the six welfare moments. Figure A.1 illustrates the out-of-sample accuracy of our methodology by comparing predicted value from the auxiliary model $\hat{\mathbf{m}}$ to their economic model counterpart $\mathbf{m}$ for the 100 sets of parameters ξ. Overall, our methodology works very well with all R^2 approaching or exceeding 0.99.

Figure A.1: Out-of-Sample Performance of Auxiliary Model

This figure illustrates the out-of-sample performance of the auxiliary model. Each panel refers to one of the six welfare moments defined in equation (A.3). For each moment, we report the value predicted by the auxiliary model $\hat{m}$ as a function of its true value in the economic model, for 100 values of the vector of parameters ξ. The auxiliary model is calibrated using 900 other vectors ξ from the training data.



[1]This is a key difference with Catherine et al. (2023) who use the value of the SMM objective function for each training point as a measure of distance to construct weights. Because our goal is not to estimate a structural model, we follow a different strategy.

A.4 Potential cost of SAVE

This section details our methodology to evaluate the potential aggregate cost of the SAVE program per cohort of borrower. Table A.3 presents our estimation.

We start by splitting the population of students leaving a higher-education program by year into three categories, which broadly correspond to different borrowing limits:

1. Students graduating from two-year programs and students dropping out.
2. Students graduating from a four-year program
3. Students graduating from graduate schools.

We estimate the number in each group per year using data from the National Center for Education Statistics (NCES). The NCES also reports the share of students in each type of degree receiving federal financial aid, which allows us to compute the number of borrowers in each category given a fraction of students who borrow.

We compute the potential cost of SAVE for each of these three categories and under three scenarios:

1. In the first scenario, the number of borrowers and the distribution of balances is unchanged by the introduction of the SAVE plan.
2. In the second scenario, the number of borrowers remains the same, but all groups borrow the maximum allowed during their undergraduate study. We assume that students from two-year programs and drop outs borrow the maximum allowed over two years, which is currently $20,000. For college graduates and graduate schools students, we assume that their debt reaches at least the four-year limit of $45,000.
3. In the third scenario, all students become borrowers and adopt the behavior described in the second scenario.

We estimate the present value cost under each scenario using our dynamic model. To do so, we need to allocate simulated borrowers to the three groups. We assume that the first group corresponds to the lowest four deciles of graduation balances reported in Figure 6. The next three deciles represent graduates from four-year colleges and the top three deciles correspond to graduate school borrowers.

Under the first scenario, the cost of SAVE is the difference between the average net present values of payment under old and new IDR rules, multiplied by the number of borrowers in each group. For example, the cost for college graduates is equal to $20,460 - $10,096 = $10,364 per borrower multiplied by 0.84 million borrowers, representing $8.7bn.

Under the second scenario, we assume that each group of borrower increase their debt at graduation to the corresponding limits, but we leaves balances above that limit unchanged (for example, college students who took more than four year to complete their program). We simulate the model again and compare the cost of this debt under SAVE to the cost of unchanged balances under pre-SAVE IDR rules. Therefore, the incremental cost of SAVE per borrower is a combination of the change in repayment rules and the change in borrowing behavior. For instance, the incremental per college graduate borrower is $26,871 - \$10,096 = \$16,775$ which, multiplied by 0.84 million borrowers, represents an a total incremental cost of $14.70bn.

In the third scenario, the incremental cost of SAVE attributed to these borrowers is the same, but is complemented by all other students that were not borrowing before the introduction of the program. The cost of these new borrowers is equal to the difference between the net present value of their debt at graduation and the amount they borrowed, multiplied by the number of new borrowers.

Table A.3: Potential Incremental Cost of SAVE

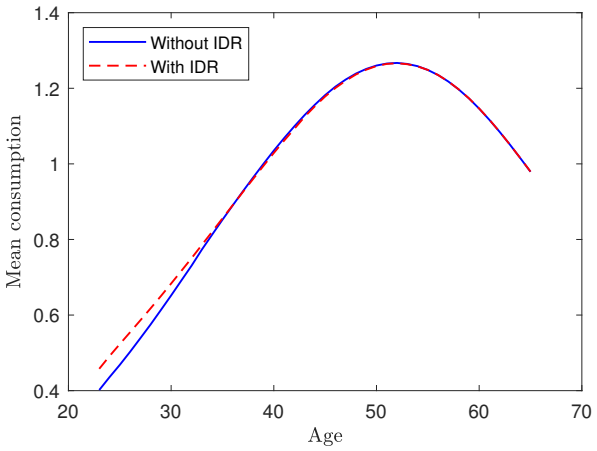
	Two-year college and drop outs	College	Graduate school	Total
Students (millions)	2.86	2.07	1.06	
Borrowers (millions)	0.87	0.84	0.77	
Cost of IDR per borrower: ($)				
- before Save	2,584	10,096	66,688	
- after Save, same borrowers and debt	6,839	20,460	70,439	
- after Save, undergraduate max out debt	11,373	26,871	68,146	
Total cost additional cost from SAVE by scenario ($ bn):				
- same borrowers, same debt	3.70	8.70	2.89	15.29
- same borrowers, undergraduates max out debt	7.65	14.07	1.12	22.84
- all undergraduates max out debt	26.48	40.94	1.12	68.54

A.5 Additional figures

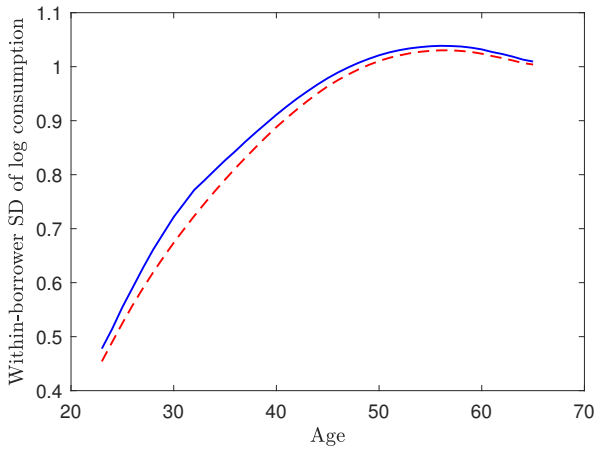
Figure A.2: Consumption Levels and Risk

Panel A reports the average consumption over the life cycle until retirement in the benchmark economy ("Without IDR") and in the economy with IDR before SAVE, reported as a multiple of the national wage index. Panel B represents the standard deviation of log consumption across realizations of the earnings process. The standard deviation is computed for each borrower, and the graph report the average standard deviation across borrowers.

A. Consumption level



B. Consumption risk

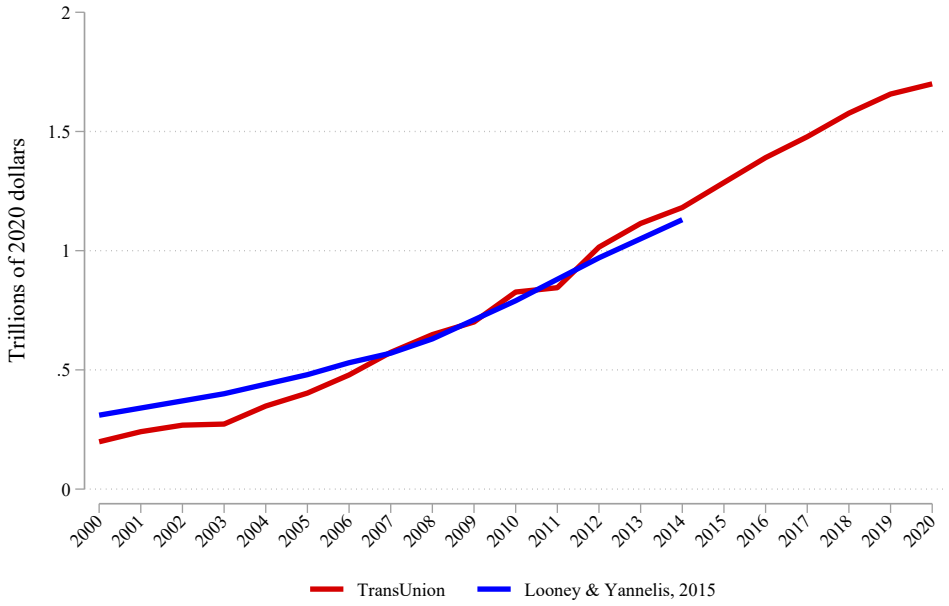


B Data Validation

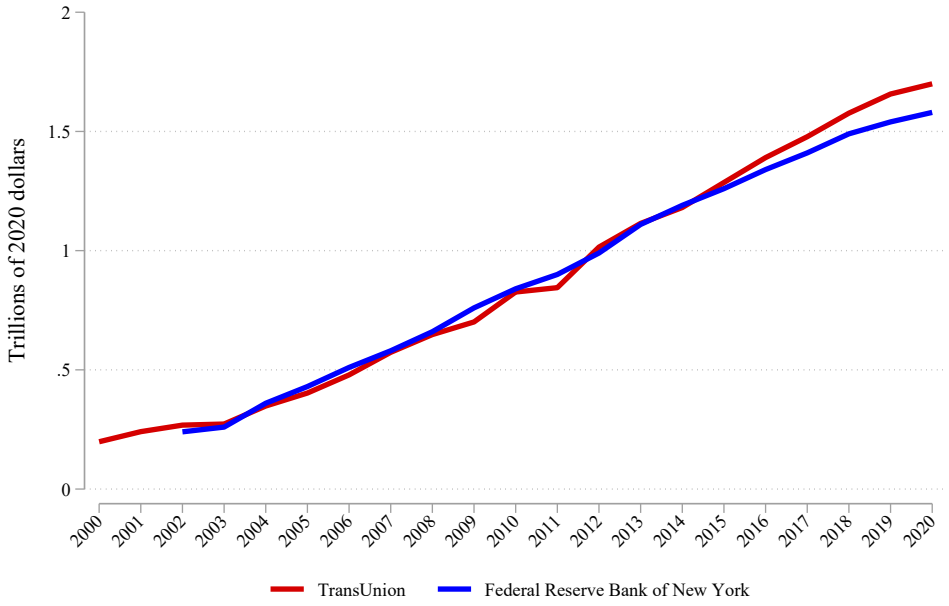
Figure B.1: Comparisons with LY and the Federal Reserve

This figure displays a collection of plots that compare the aggregates found for TransUnion to three external sources: the Federal Reserve Bank of New York, Federal Reserve Economic Data (FRED), and Looney & Yannelis, 2015.

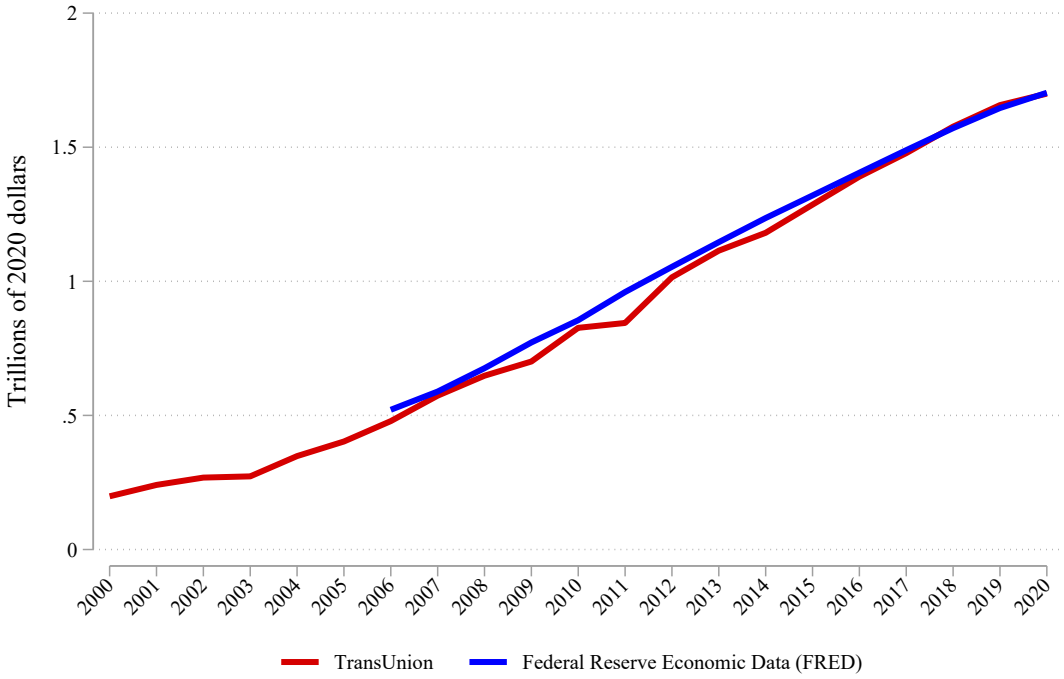
(a) Aggregate Balance: Looney & Yannelis, 2015



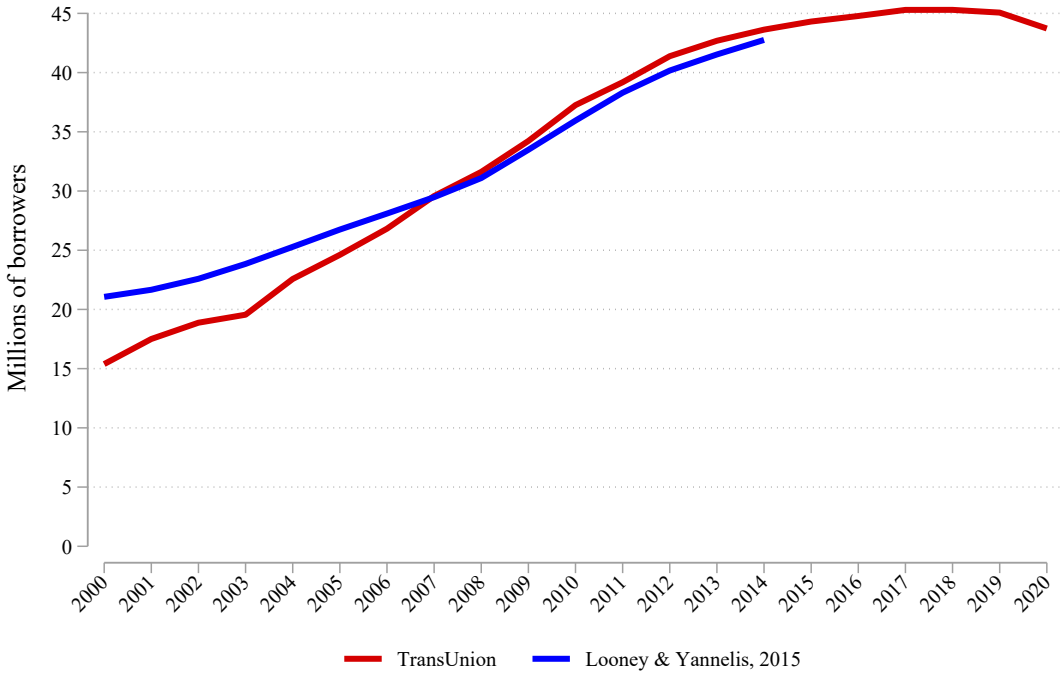
(b) Aggregate Balance: Federal Reserve Bank of New York



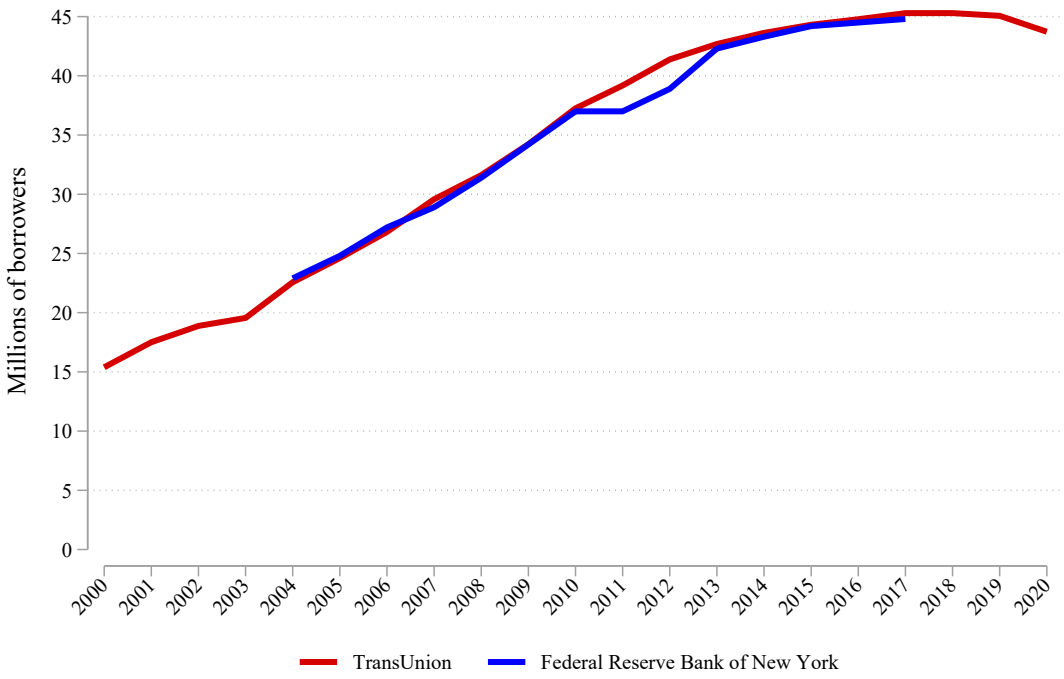
(c) Aggregate Balance: Federal Reserve Economic Data



(d) Aggregate Borrowers: Looney & Yannelis, 2015



(e) Aggregate Borrowers: Federal Reserve Bank of New York



C Additional figures

Figure C.1: Aggregate Balance, Split by Age Group

.

This figure breaks the aggregate balance down into agre groups. Data from TransUnion.

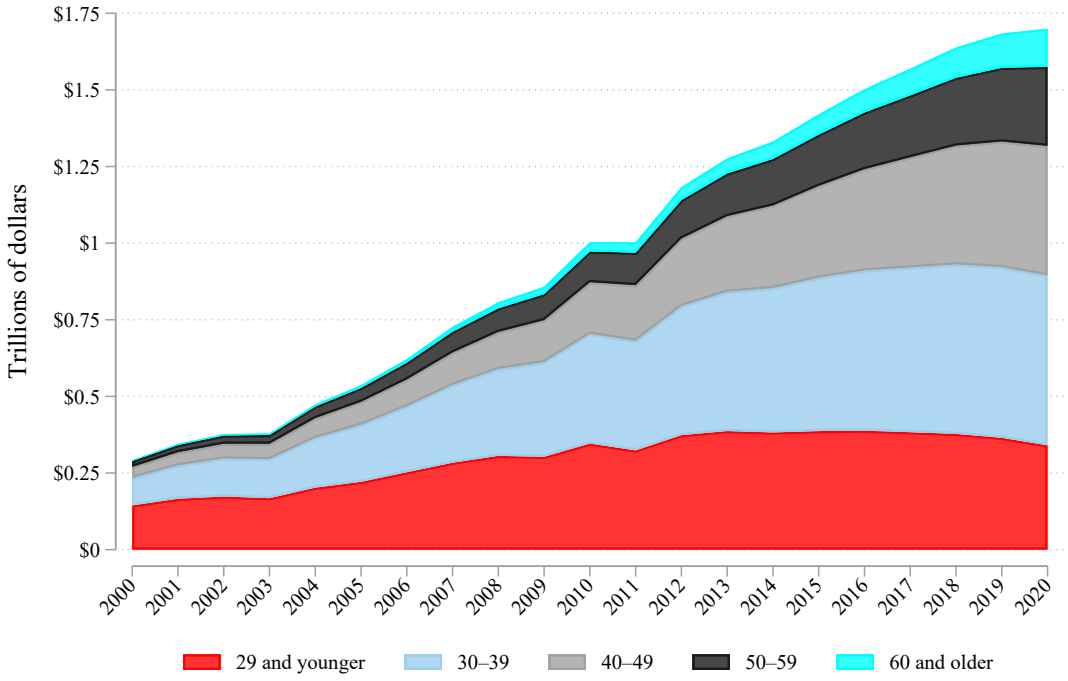
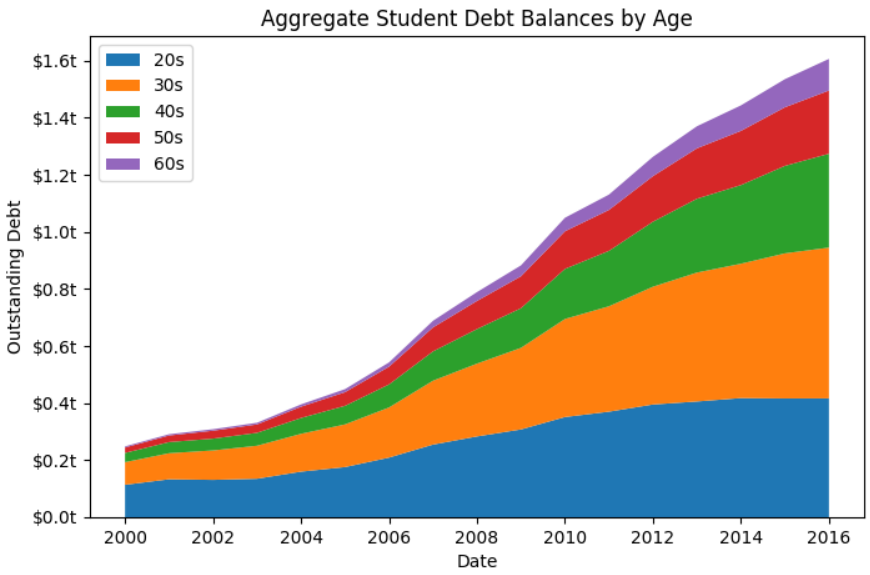
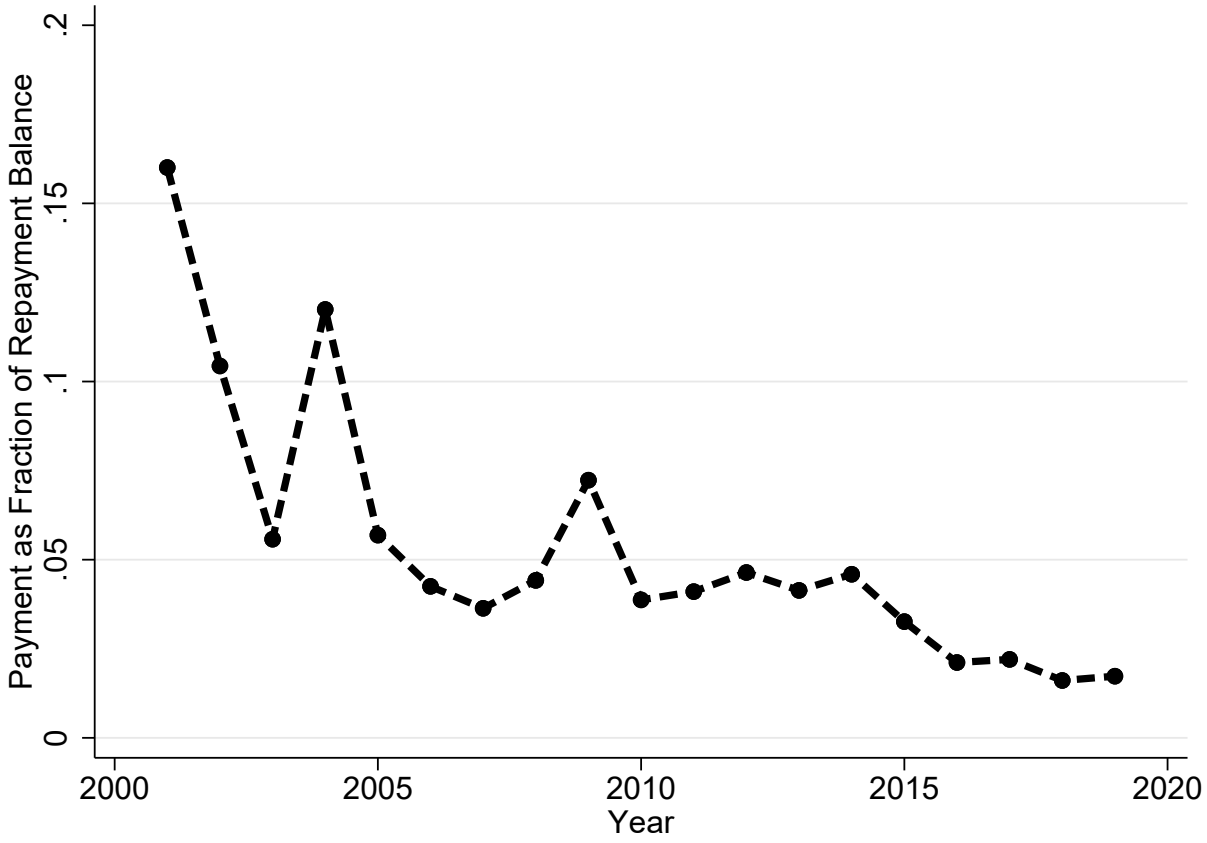


Figure C.2: Outstanding Balances by Age



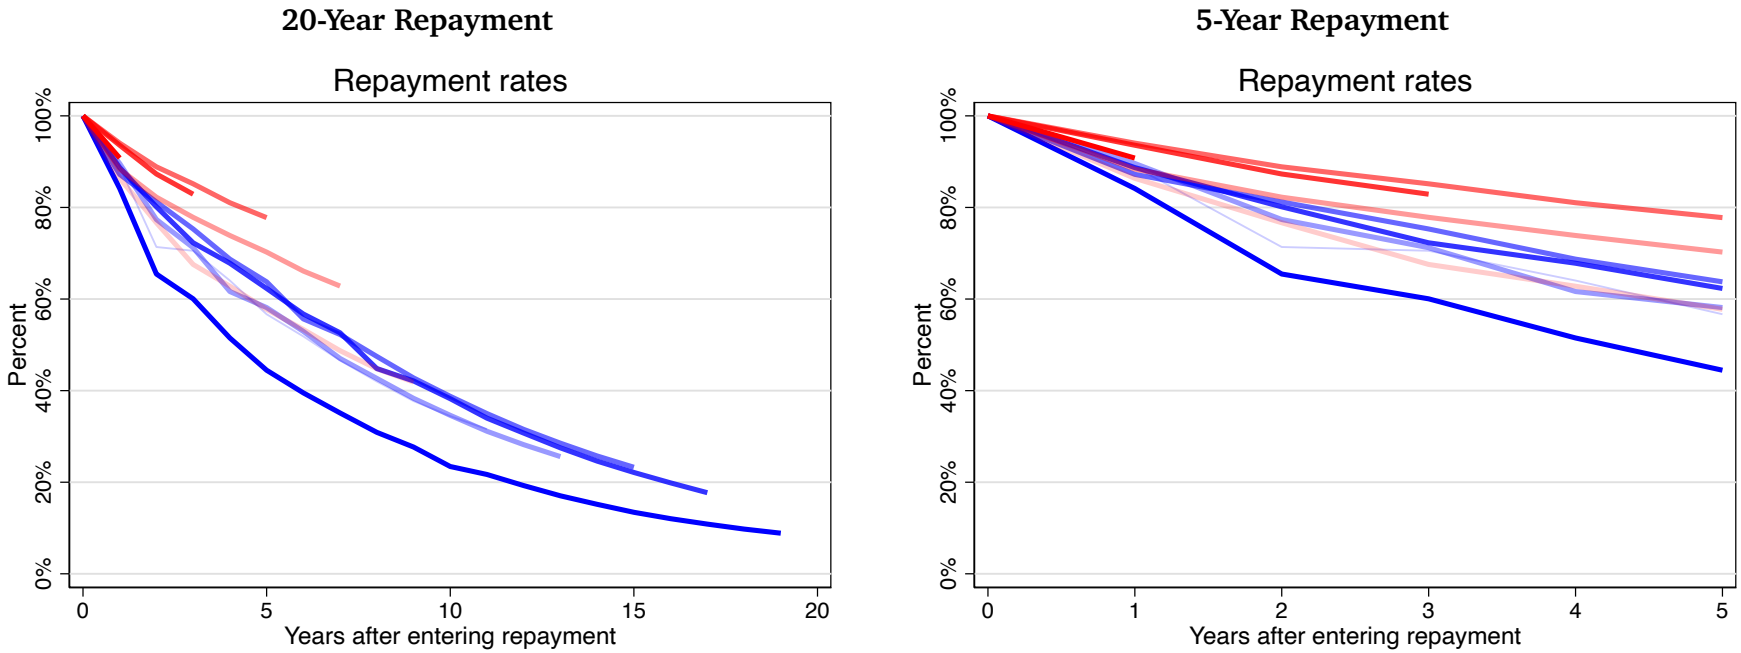
This figure depicts the aggregate outstanding balance in each year decomposed by the age of indebted borrower. The rise in student debt balances during the 2000s is driven mostly by rise in outstanding balances among borrowers in their 20s. In the 2010s, the rise in balances is driven by growing balances among older borrowers. Source: TransUnion.

Figure C.3: Annual Payment as a Fraction of Disbursement at Repayment



This figure depicts average annual payments as a fraction of repayment balance. Source: TransUnion.

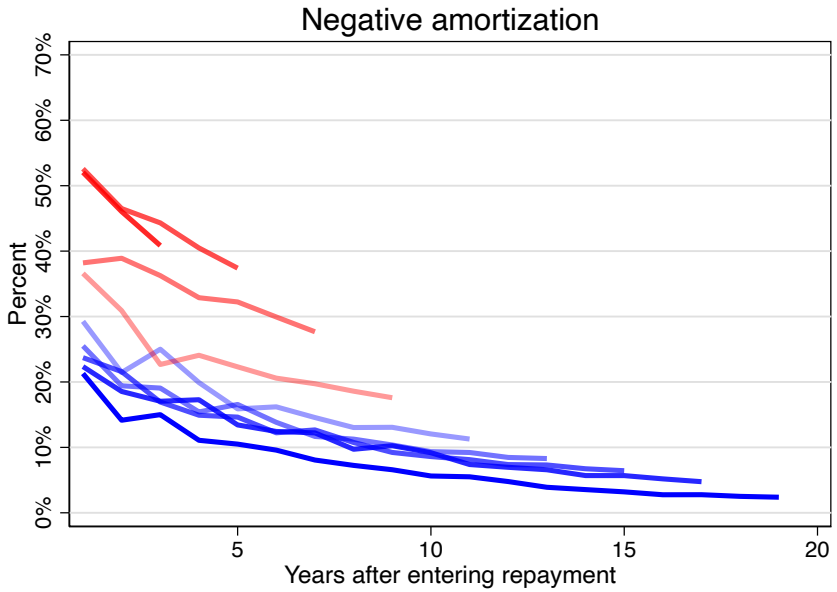
Figure C.4: Repayment Over Time



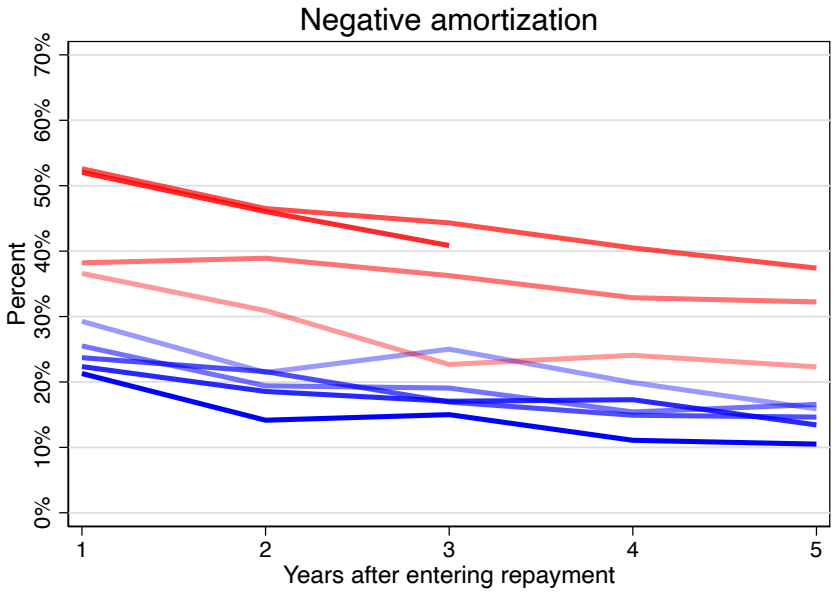
This figure depicts the fraction of borrowers repayment balance outstanding in a given repayment year, for cohorts between 2001 ad 2017, in odd years. Earlier to later cohorts progress from black to red. The fraction of repayment balance outstanding is defined as the fraction of the repayment amount which is outstanding. Source: TransUnion.

Figure C.5: Negative Amortization Over Time

20-Year Negative Amortization

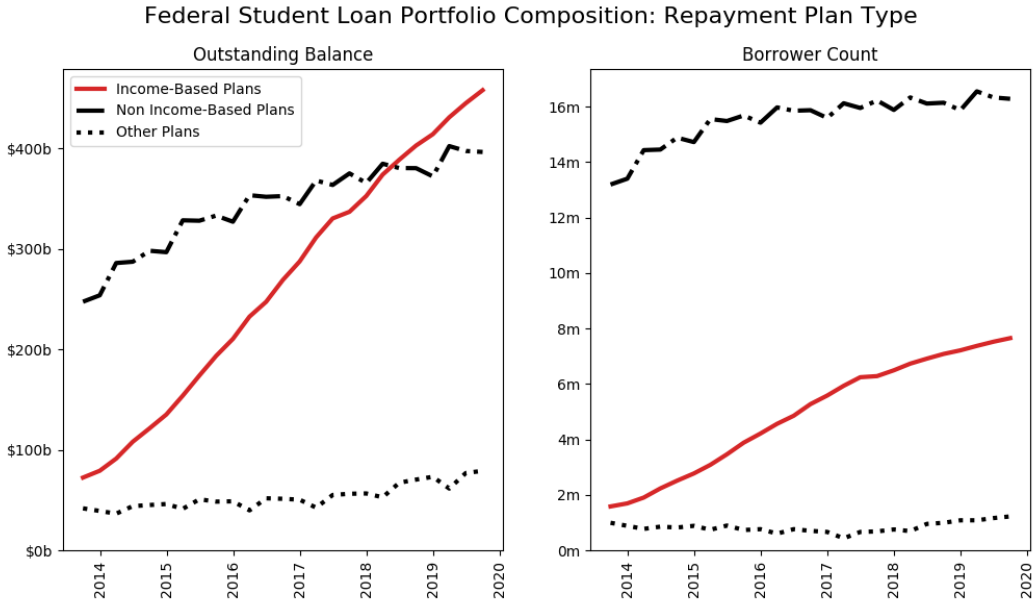


5-Year Negative Amortization



This figure depicts the fraction of borrowers negatively amortization in a given repayment year, for cohorts between 2001 ad 2017, in odd years. Earlier to later cohorts progress from black to red. Negative amortization is defined as owing more than balance at repayment. Source: TransUnion.

Figure C.6: Balance by Plan Type



This figure depicts the composition of the Direct Loan portfolio of the federal government according to the repayment plan of the borrower from mid-2013 to 2019. Income-based plans include income based repayment, income contingent repayment, PAYE and REPAYE plans. Non income-based plans include level and graduated plans. Other plans contain both income-based and non income-based plans. The rise of these plans is due almost entirely to the rise of "alternative plans", which are case-by-case plans that commonly give defaulted Parent PLUS Loan borrowers access to an otherwise unavailable income-based repayment plan (see: Edvisors). All balances include loans in deferment, repayment, and forbearance. Data are publicly reported and from the NSLDS.