Journal of Financial Economics

How do income-driven repayment plans benefit student debt borrowers?[☆]

Sylvain Catherine[a,*], Mehran Ebrahimian[b], Constantine Yannelis[c]

[a] Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, 19104, PA, USA
[b] Stockholm School of Economics, Sveavägen 65, Stockholm, 113 83, Sweden
[c] University of Cambridge, Austin Robinson Building, Sidgwick Ave, Cambridge, CB3 9DD, United Kingdom

ARTICLE INFO

JEL classification:
D14
D31
H52
H81
I24

Keywords:
Student debt
Income-driven repayment
Fiscal policy

ABSTRACT

Using credit bureau data, we show that nearly half the increase in student debt since 2010 is due to deferred payments and the expansion of income-driven repayment (IDR) plans. These plans help borrowers smooth consumption, insure income risk, and reduce the effective debt cost. Using a life-cycle model, we quantify the welfare gains from this payment deferment and the channels through which welfare increases. We show that an optimally calibrated plan can achieve similar welfare gains at a much lower cost to taxpayers, and without encouraging additional borrowing. Finally, we use our quantitative framework to evaluate recent proposals to reform IDR rules.

1. Introduction

Student debt has risen more than sixfold between 2002 and 2020. While some of this rise is due to increases in the number of borrowers and borrowing amounts, a significant portion 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, many of them have deferred payments from early to later in the life cycle. As of 2020, more than half of loan balances are in IDR plans. This paper studies the mechanisms through which borrowers can benefit from deferring payments, the distribution of these benefits, and the cost-efficiency of repayment rules.

Our first contribution is to document that the slowdown in repayment explains nearly one half of the increase in student debt between 2010 and 2020 – roughly 0.4 trillion dollars – a fact that has received little attention in policy debates. This finding highlights how changes in repayment timing can significantly affect both household balance sheets and public finances.

Our second contribution is to analyze the welfare and fiscal implications of IDR programs. We show that the retiming of payments generate large welfare gains of nearly $51,000 on average per borrower under existing repayment rules. Borrowers with large balances benefit the most. About half of these gains are due to subsidization by taxpayers, while the remainder comes from consumption smoothing and insurance against income risk. However, optimal income-linked payment plans could offer greater welfare gains to borrowers at a much lower cost to taxpayers by reducing payments early in life and low-income states in exchange for an extended repayment period. Relative to existing rules, the calibration of IDR maximizing the joint welfare of borrowers and taxpayers would increase per-borrower welfare by about $1000 while reducing the cost to taxpayers by about $20,000. We compare this result to recent policy proposals.

We begin our analysis by quantifying the effect of the rise in payment deferral on aggregated student debt levels. Using administrative data from TransUnion, a major US credit bureau, we document that deferred payments increased rapidly with the rollout of IDR plans.

[☆] We are grateful to numerous seminar participants at Wharton, the Federal Reserve Board, Cambridge, the Norwegian School of Economics, the Stockholm School of Economics, Bonn University and Mannheim University as well as numerous 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. M. Ebrahimian acknowledges funding from the Swedish House of Finance.

Received 16 July 2025; Received in revised form 26 January 2026; Accepted 2 February 2026

Available online 11 February 2026
0304-405X/

We estimate deferred payments by borrower cohorts and construct counterfactual loan balances under standard repayment plans. Our decomposition shows that deferred payments account for $391 billion, or 44%, of the growth in student loan debt since IDR plans were introduced.

Next, we set up a life-cycle model in which households enter the workforce and repay their student loans. Because markets are incomplete, borrowers are unable to borrow against future earnings and to insure against income risk, preventing them from realizing their ideal consumption plan. The model captures key features of the student loan program, and IDR plans can increase borrower welfare by smoothing payments and providing insurance. Our primary focus is on evaluating steady-state IDR rules, abstracting from the transitional period following the introduction of these programs, which has been marked by frequent policy changes, institutional frictions, evolving borrower awareness, and promises of broad debt forgiveness.

Using the model, we estimate aggregate welfare gains from IDR plans and decompose these gains into four channels. First, IDR reduces lifetime payments, improving borrower welfare at the expense of taxpayers. Second, IDR helps borrowers smooth consumption over the life cycle by deferring payments. Third, by lowering payments in low-income states, IDR provides insurance against idiosyncratic income risk. Finally, if it targets students with lower lifetime consumption, IDR can mitigate social welfare losses arising from inequality. We propose a simple mathematical decomposition of welfare gains that clarifies how the design and calibration of IDR plans serve different policy objectives. Our methodology can be applied to other contexts.

Our results show large welfare gains from existing IDR rules, with average per-borrower welfare increasing by $50,600. While, transfers from taxpayers to borrowers represent half of these gains (47.5%), the other half is due to life-cycle consumption smoothing (29.9%) and insurance against income shocks (21.3%). On the other hand, we find no welfare gains from reducing inequality between borrowers. Although IDR rules favor low-income borrowers, they also benefits those with large balances, who tend to be high-income graduate-school borrowers.

The subsidization of student debt borrowers by taxpayers under existing rules raises the question of the cost-efficiency of IDR. We evaluate this question in two steps.

First, we explore the space of IDR parameters to identify the calibration that maximizes borrower welfare for a given cost to taxpayers. The optimal policy features a significantly longer repayment period—extending up to retirement age.[1] A longer repayment period enables the government to recoup payments later in life. For a fixed fiscal cost, the present value of these additional payments allows for a higher income threshold above which payments are due. Raising this threshold protects borrowers' consumption earlier in the life-cycle and in low-income states, when the marginal utility of consumption is high. As a result, this tradeoff increases welfare by improving both insurance and intertemporal smoothing per taxpayer dollar.

Second, we raise the question of how much taxpayers should subsidize student debt through IDR, as offering more generous terms is only justified if the marginal welfare gains to borrowers exceed the additional budgetary costs. If the social planner puts equal weights on borrowers and taxpayers (considering also the fiscal externalities of raising taxes), we find that net welfare gains are maximized with a per-borrower subsidy of $4300—significantly lower than current rules ($24,100). Despite this reduction, the optimal policy increases borrower welfare by $1000 more than current rules. In addition to an extended repayment period, the optimal policy includes increasing the repayment rate from 10% to 24.4%, while substantially raising the threshold above which earnings are subject to repayment, from 150% to 237%. We assume that payments would remain capped by the ten-year, standard-plan formula. Under this policy, loans would be priced near fair value, mitigating moral hazard concerns. Our findings imply that if the government aims to subsidize students beyond the $4300 level, it should do so outside the scope of the student loan program—such as through grants—rather than through more generous IDR rules. Beyond this threshold, the marginal taxpayer dollar tends to increase consumption in below-average marginal utility states.

In light of our analysis, proposed policy changes such as the Saving on a Valuable Education (SAVE) plan appear sub-optimal. The SAVE plan, introduced in 2022 but later blocked by Federal courts, benefits undergraduate borrowers by lowering the fraction of income that they must pay and shortening time to forgiveness for those with smaller loans. We estimate that the SAVE plan increases borrower welfare by an additional $5300 but that these gains are fully explained by the incremental cost to taxpayers. The changes introduced by the SAVE plan would go in the opposite direction of the main takeaway of our analysis: it shortens the repayment period rather than extending it, and it raises the fiscal cost of IDR rather than minimizing it. By contrast, the Repayment Assistance Program (RAP) enacted in July 2025 extends the repayment period before forgiveness. Although the present value gains from this extension are partly captured by more generous and progressive repayment rules, RAP reduces borrower welfare by $1000 relative to existing IDR rules. However, this loss is substantially smaller than the fiscal savings of $7000 per borrower.

Our analysis assumes a constant distribution of student debt at graduation. However, if student debt is subsidized by taxpayers, students could be encouraged to borrow more, and schools to increase tuition (Eaton et al., 2020; Lucca et al., 2019). Under existing IDR rules, student debt is priced near fair value for most undergraduate borrowers, although some with high balances benefited from per-dollar 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. Therefore, the cost of the federal loan program under the SAVE rules could increase substantially if undergraduate students responded by borrowing more. On the other hand, because it strongly reduces implicit subsidization, the optimal calibration of IDR would also lower moral hazard. RAP would also slightly reduce the subsidization of student loans.

Related literature. This paper 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., 2019; Catherine and Yannelis, 2023), payment deferrals (Dinerstein et al., 2024; Hamdi et al., 2025) and limit increases (Black et al., 2023; Goodman et al., 2021). Though most of these studies are empirical, analyzing how repayment plans affect borrower outcomes, some papers study the student loans through the lens of a structural model (Lochner and Monge-Naranjo, 2011, 2016; Ebrahimian, 2020; de Silva, 2025). 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).

Our main contribution here is to study student debt repayment through the lens of a life-cycle model, as well as new positive and normative facts about trends in deferred payment, and welfare.

Our analysis of potential policy improvements is closely related to an important study by Boutros et al. (2024). This study finds that deferring the repayment period by ten years improves borrowers' ability to repay, thereby reducing the cost of IDR for the government. These budgetary savings can be returned to borrowers through lower interest rates, generating substantial welfare gains. Our policy analysis, however, differs in several ways. First, rather than limiting our scope to a few discrete policy experiments, we examine a continuous space

[1] This design resembles the Australian student loan system, which defaults all borrowers into an income-driven plan without offering forgiveness. By contrast, the UK offers forgiveness after 25 years, and the US after 20 to 30 years, depending on the plan.

of policy parameters to identify optimal calibrations. Intuitively, reallocating budgetary savings from an extended repayment period by increasing the repayment threshold may yield greater welfare benefits than reducing interest rates, as it specifically supports borrowers in low-income states, whereas the latter benefits those with larger debt, including those at the bottom of the marginal utility distribution. Second, while Boutros et al. (2024) focus on welfare gains from a 10-year deferral under the same taxpayer cost as existing rules, our analysis considers the full budgetary spectrum. This enables us to answer an additional important question: how much should the government subsidize IDR to maximize welfare for both borrowers and taxpayers? We find that this optimal budget represents only a fraction of the taxpayer cost of existing rules. Therefore, our analysis suggests that most of the fiscal gains from extending the repayment period should be returned to taxpayers. Despite returning only a fraction of fiscal gains to borrowers, the optimal policy calibration would improve their welfare because raising the repayment threshold, rather than reducing interest rates, allocate each dollar to high-marginal-utility states.

More broadly, our paper also contributes to the literature on the design of public policies in the presence of liquidity constraints and uninsurable income risk. 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., 2025; 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 play a decisive role in easing liquidity constraints and providing insurance against income risk for young workers.

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

2. Institutional background

In this section, we first present institutional background on student debt in the United States and then describe the evolution of repayment rules over the past twenty years.

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]

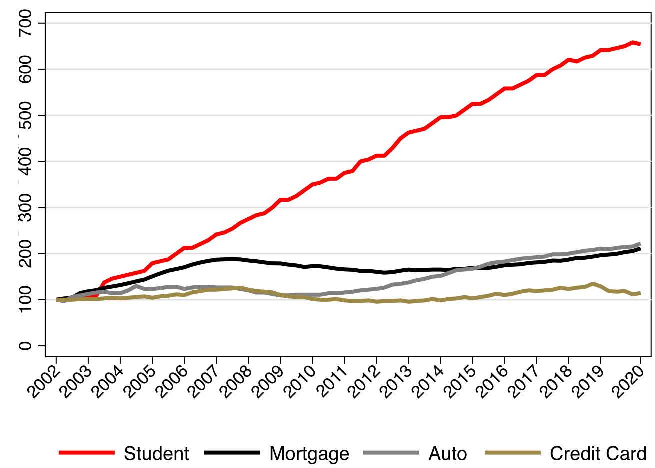

Fig. 1. Household debt in the United States 2004–2020.
This figure displays the growth in mortgage, auto loan, student loan, and credit card balances for consumers from 2002 to 2020. Balances are reported from the end of each fiscal quarter. Aggregate consumer balances are normalized to 100 in 2002.
Source: Data from the Federal Reserve Bank of New York.

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.

Fig. 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 2002, and plots the relative increase over time. For all categories of household debt, loan balances increase by less than 200%. 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.

2.2. Repayment plans

Historically, the vast majority of student loan borrowers were in a repayment option termed the Standard Plan, under which they 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. 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.

In 2008, the Department of Education began a dramatic expansion of IDR plans. Until then, one IDR plan existed, Income-Contingent Repayment but take-up was low because terms were fairly onerous for borrowers. 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 25 years in repayment. In subsequent years, more generous IDR options such as Pay as You Earn and Revised Pay as You Earn were introduced. These new plans made IDR more generous by reducing payment rates above the FPL from 15% to 10% and allowing forgiveness after 20 rather than 25 years in repayment.

[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.

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 an increased from $46,800 to $70,000 in 2024 for a family of four, below which borrowers would pay nothing. Second, the percentage of their income above that threshold 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 $1000, 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 when payments do not even cover interests.

The SAVE plan was blocked by courts in June 2024, and approximately one year later in July 2025 the Repayment Assistance Plan (RAP) was introduced. The RAP extends the period of repayment before forgiveness to 30 years. Under this new plan, borrowers pay a flat amount of $50 per month if they earn below $10,000, and then, for income above $10,000, they face a graduated payment scheme, in which they pay a percentage of 1%–10% of their income. Starting from 1% at the income of $10,000, the payment rate increases by increments of 1 percentage point for every $10,000 increase in income, until it reaches a maximum of 10% for an income of $100,000 and above. Borrowers also receive a payment subsidy of $600 per child.

3. Payment deferrals and the rise in balances

In this section, we use administrative data to quantify the increase in student loan balances that can be attributed to slowing payment. 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 can account for almost half of the rise in balances between 2010 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] We assign borrowers to repayment cohorts by the first time that they are scheduled to enter repayment.

3.2. Trends in student loan balances

As Fig. 2 illustrates, the rise in student debt is driven by the rising number of borrowers and the increase in average balances. The left panel shows aggregate balances, the middle panel shows the number of borrowers and the right panel shows average loan balances. Between 2001 and 2020, the number of borrowers tripled, increasing from 15 million to 45 million. The number of enrolled students increased by approximately 30% between 2000 and 2020, which is significant but not comparable in magnitude to the very large increase in outstanding borrowers.[5] Over the same time period, the average balance doubled, rising from roughly $20,000 in 2001 to nearly $40,000 in 2020.

A major contributor to the growth in balances is slowing repayment rates, which is shown in Fig. 3. The left panel shows the fraction of the initial balance still outstanding, for cohorts of borrowers in odd years from 2001 to 2019. The evolution of balances reveals a marked slowdown in repayment speed: The 2001 cohort had only 45% of their initial balance still outstanding after five years of payments, as would be expected in a 10-year repayment plan, compared to 80% for the 2015 cohort.[6] This is the combined result of disbursement amounts increasing and repayment significantly slowing down.

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 fraction of individuals in a cohort negatively amortizing is shown in the right panel of Fig. 3.[7] 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.

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.[8] To 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

How would aggregate student debt have evolved had borrowers made standard plan payments? To answer this question, we construct counterfactual balances by assuming that borrowers made standard plan payments. In that case, the counterfactual balance of borrower i would have evolved as:

$$L^{CF}_{it+1} = L^{CF}_{it}(1 + r_{L,i}) - R_{\text{std},it}, \tag{1}$$

where $L^{CF}_{it_0}$ is the balance at the beginning of the repayment period, $R_{\text{std},it} = \frac{r_{L,i} L^{CF}_{it_0}}{1-(1+r_{L,i})^{-10}}$ the yearly payment under the ten-year standard plan, and $r_{L,i}$ the loan interest rate. We sort borrowers by cohort of repayment, and proxy each cohort's interest rate as the average over the four years preceding repayment. Therefore, the counterfactual balance L^{CF}_{it} computes the evolution of balances as if standard plan payments were made.

[3] Keys et al. (2023) provide further details regarding 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 C.1 compares our aggregates with these other data sources. Aggregates and trends line up closely.

[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).

[6] The fraction of debt owned by older borrowers in each year has increased over time, as shown in Appendix Figure D.1.

[7] These figures are shown for longer time periods and in different iterations in Appendix Figures D.2 through D.5. They largely paint a similar portrait.

[8] Some of the slowdown in repayment around 2008 may be driven by the Great Financial Crisis. However, if slowing repayment were primarily caused by the financial crisis, we would expect effects to dissipate over time. This is not what we see, rather cohorts graduating a decade later take even more time to repay. This is consistent with IDR rules becoming more widely known and more generous over time.

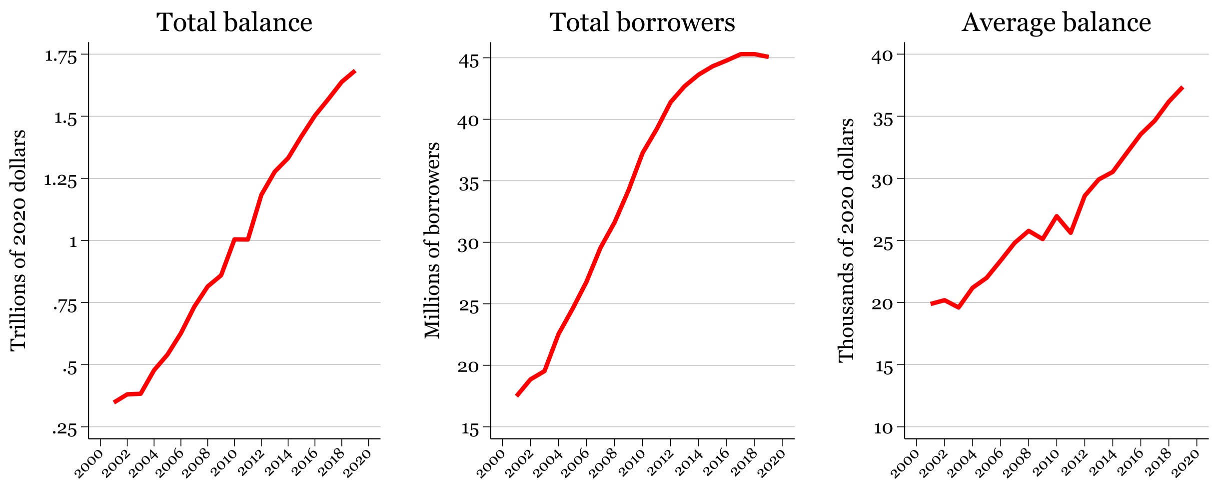

Fig. 2. Student loan aggregates.
This figure displays the aggregate student loan balance in the United States, the total number of borrowers by year, and the average balance by year, from 2001 to 2020. Numbers are as of December of each year.
Source: Data from TransUnion.

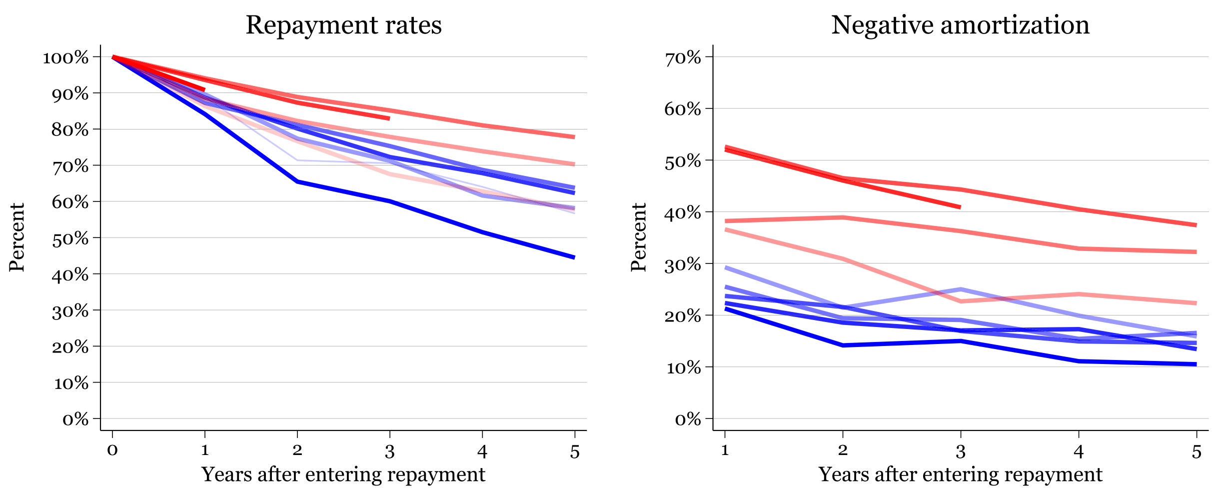

Fig. 3. Repayment rates and Negative amortization.
This figure shows the repayment behavior of every other repayment cohort from 2001 (blue) to 2019 (red). The left-hand panel indicates the repayment rate, or what percentage of the total balance owed by borrowers when entering repayment remains. The right-hand panel displays negative amortization, or the percentage of people who have an equal or higher balance than they did when entering repayment. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
Source: Data from TransUnion.

Fig. 4 presents the evolution of aggregate counterfactual balances, along with its real-world counterpart. The two series evolve similarly until 2008, suggesting that on average borrowers were up to that point making payments equivalent to standard plan payments. In reality, some borrowers default whereas other prepay but these two offsetting deviation from the standard plan seemed to nearly cancel each other. This result is consistent with the evidence from Fig. 3 that earlier cohorts paid down roughly 40% of their loan within five years.

The two series begin to diverge in 2008, the year IBR is introduced. The divergence remains modest up to 2011, but accelerates over time. By 2020, the actual balance is $1.7 trillion while the counterfactual one is $1.31 trillion.[9] Student loan balances would be $391 billion smaller if borrowers had continued to pay down their loans at the standard rate. Since, between 2008 and 2020, student loan balances rose by $885 billion, our analysis suggests that 44%, or roughly half of the increase in student loan balances is driven by slower repayment over time.

Fig. 5 presents a further decomposition of the trends shown by staking the cumulative value of deferred payments by cohort. The main red line represents the total, inclusive of cohorts repaying faster than under the standard plan. Before 2008, we see little evidence of deferred payments for any cohort. This begins to change modestly in 2009, with the introduction of the IBR plan. Take-up of the initial IBR plan was low, as it was not particularly generous and was poorly advertised by the federal government. We see sharp upticks in deferred balances with the introduction of PAYE in 2012 and then again with REPAYE in 2015. These plans were more generous, and were coupled by initiatives by the Obama administration to enroll and keep borrowers in these plans. This decomposition shows that the timing of deferred payments closely aligns with cohorts that had more generous IDR plans available.

We do not claim that the entirety of the rise in balances is explained by deferred payments, and many other factors likely played a role. Nonetheless, the fact that deferred payments can explain approximately half of the increase over time is consistent with the fact that other

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

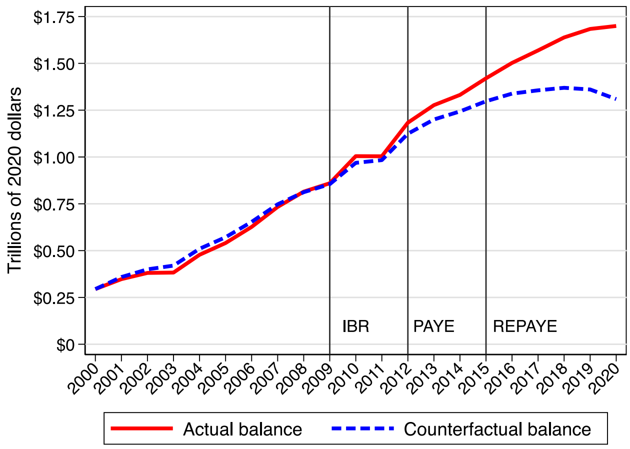

Fig. 4. Actual and counterfactual balance.
This figure displays the actual aggregate student loan balance and the counterfactual balance that would result if all students paid their loans down according to the 10-year standard plan. Vertical lines show the introduction of IDR plans.
Source: Data from TransUnion.

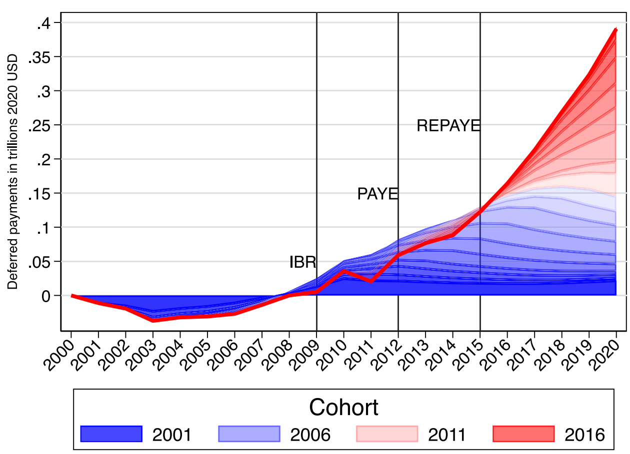

Fig. 5. Deferred payments by cohort.
This figure displays cumulative deferred payments by cohort, relative to the standard ten-year repayment plan. Vertical lines show the introduction of IDR plans. The red line represents aggregate deferred payments across cohorts. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
Source: Data from TransUnion.

trends cannot explain the entirety of the rise in student loan balances. Factors like increases in tuition, interest rates, enrollment or shifts in the composition of borrowers may affect the rise in balances, but they are accounted for in our counterfactual. The counterfactual balance, in which everyone repays under the standard plan, fixes for each cohort and time period the initial balance and interest rate and assumes that each borrower remains in the standard plan. Our comparison of the actual and counterfactual balances thus isolates the factors that affected the rise in balances beyond deferred repayment.[10]

4. Model

To study the welfare consequences of payment deferral, and the introduction of new IDR plans, 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 realizing 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 IDR 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 and maximizes expected utility:

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

where the period utility function is:

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

with γ denoting 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 W_{it} evolves 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.

4.2. Student loan

Borrowers graduate with a student debt L_{it_0}. We consider two counterfactual economies that differ in their student loan repayment rules. First, a benchmark economy without IDR. Then, an economy with IDR repayment rules similar to the US system, but in which enrollment in IDR remains optional. Our goal is to evaluate the welfare effects of the availability of IDR relative to the benchmark economy.

[10] Direct evidence also shows that factors beyond deferred repayments cannot explain the entirety of the rise in the aggregate student debt balance. While tuition increased, the amounts were not large enough to explain borrowing increases. The College Board reports that schools tuition increased by 33% at public four-year colleges and 25% at private four-year colleges between 2008 and 2020. Default patterns also cannot explain the rise in balances. The three-year cohort default rate was 12.2% in 2008. While this rate rose slightly to 14.7% in 2014, the rate fell to 9.7% by 2020. Trends in student enrollment patterns also seem unlikely to explain the entirely of the rise in balances. Total enrollment remained constant over that time period. Borrowing rates also remained relatively steady and even declined. The number of active borrowers taking loans actually declined over the same period, and average annual borrowing remained roughly constant or declined slightly. The College Board reports that undergraduate borrowers borrowed $6700 in the 06–07 academic year, and $6440 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. The number of active graduate student borrowers during the same time period increased from 1.3 million to 1.8 million. This 40% increase in the number of graduate borrowers seems unlikely to explain the sixfold increase in balances.

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:

$$L_{it+1} = L_{it}(1 + r_L) - R_{L,it}, \tag{5}$$

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:

$$R_{\text{std},it} = \frac{r_L L_{it_0}}{1 - (1 + r_L)^{-10}}, \tag{6}$$

until their 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:

$$R_{\text{garnish},it} = \min\left\{0.15 \times (Y_{it} - \Gamma_{it}), R_{\text{max garnish},it}\right\} \tag{7}$$

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. We model this second bound as a fraction λ_{garnish} of the average wage in the economy, called $Y_{1,t}$.

$$R_{\text{max garnish},it} = \min\left\{0.25 \times (Y_{it} - \Gamma_{it}), \lambda_{\text{garnish}} Y_{1,t}\right\} \tag{8}$$

We assume that borrowers fall in garnishment when their cash on hand $W_{it} + Y_{it}$ falls below a threshold λ_{default}.[11]

4.2.2. Economy with IDR

In reality, under the current system, borrowers can remain in the standard plan or enroll in the income-driven repayment program at any point in time. Borrowers in IDR can 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.[12] Hence, repayment in IDR is:

$$\begin{cases} R_{\text{idr},it} = \min\left\{\theta_{\text{idr}}\left(Y_{it} - \lambda_{\text{idr}} Y_{\text{pl},t}\right)^+, R_{\text{std},i}\right\} & \text{if } t \le t_{F_{\text{idr}}}, \\ R_{\text{idr},it} = 0 & \text{if } t > t_{F_{\text{idr}}}, \end{cases} \tag{9}$$

where discretionary income is defined as the share of their income above $\lambda_{\text{idr}} = 150\%$ of the federal poverty line, $Y_{\text{pl},t}$. Assuming that the poverty line represents a constant multiple of the average wage index, the income threshold $\lambda_{\text{idr}} Y_{\text{pl},t}$ is proportionate to $Y_{1,t}$. Until forgiven or fully repaid, balances evolve as in Eq. (5). Finally, we assume that garnishment no longer exists in the economy with IDR.[13]

[11] 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.

[12] Consistent with current law, we assume that forgiveness is untaxed. The American Rescue Plan Act of 2021 exempted student loan forgiveness from being treated as taxable income through 2025. If this treatment is not extended, treating forgiveness as taxable income would reduce the generosity of IDR plans and expected subsidy.

[13] 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.

4.2.3. Prepayment

In both the economies with and without IDR, 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 an idiosyncratic component $Y_{2,it}$:

$$Y_{it} = Y_{1,t} \cdot Y_{2,it}. \tag{10}$$

We model the idiosyncratic component $Y_{2,it}$ as in Guvenen et al. (2021). It depends on a deterministic function of age $g(t)$, a persistent z_i and a transitory component ε_i. We allow $g(t)$ to be correlated with student debt borrowing as in the data. 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 (11) summarizes the dynamic of $Y_{2,it}$ Eq. (11) is given in Box I.

We do not explicitly model the labor supply effects of income-driven repayment (IDR) programs. While the evidence on these effects is 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 (2025) report some effects in Australia. However, the labor supply elasticity estimated in de Silva (2025) is 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. In that case, it should not reduce labor supply (Manso et al., 2025). Unlike in Australia where student loans have zero interest rates, differing payments is costly in the U.S. 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 payment caps and zero interest rates.

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. Therefore, benefits from SNAP are:

$$B_{it}^{\text{SN}} = \left(0.06 Y_{1t} - 0.3 Y_{it}\right)^+ \qquad \text{if } W_{it} < 0.05 Y_{1t}. \tag{12}$$

These benefits allow households to maintain a consumption level above $0.06 Y_{1t}$ 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:

$$\Gamma_{S,it} = \tau \min\left\{Y_{it}, 2.5 Y_{1,t}\right\}. \tag{13}$$

Workers retire at age t_R, and their yearly retirement pension is:

$$B_{it \ge t_R}^{\text{SS}} = \begin{cases} 0.9 \times \text{AIYE}_{it_R} & \text{if } \text{AIYE}_{it_R} < 0.2 Y_{1,t_R} \\ 0.116 \times Y_{1,t_R} + 0.32 \times & \text{AIYE}_{it_R} \\ & \text{if } 0.2 Y_{1,t_R} \le \text{AIYE}_{it_R} < Y_{1,t_R} \\ 0.286 \times Y_{1,t_R} + 0.15 \times & \text{AIYE}_{it_R} \\ & \text{if } Y_{1,t_R} \le \text{AIYE}_{it_R}, \end{cases} \tag{14}$$

$$\begin{array}{ll} \text{Level of idiosyncratic earnings:} & Y_{2,it} = (1 - \nu_{it}) e^{\left(g(t) + \alpha^i + z_{it} + \epsilon_{it}\right)} \\ \text{Persistent component:} & z_{it} = \rho z_{it-1} + \eta_{it} \\ \text{Innovations to AR(1)}: & \eta_{it} \sim \begin{cases} \mathcal{N}(\mu_{\eta,1}, \sigma^2_{\eta,1}) & \text{with probability } p_z \\ \mathcal{N}(\mu_{\eta,2}, \sigma^2_{\eta,2}) & \text{with probability } 1 - p_z \end{cases} \\ \text{Initial condition of } z_{it}: & z_{it_0} \sim \mathcal{N}(0, \sigma^2_{z,t_0}) \\ \text{Transitory shock:} & \epsilon_{it} \sim \begin{cases} \mathcal{N}(\mu_{\epsilon,1}, \sigma^2_{\epsilon,1}) & \text{with probability } p_\epsilon \\ \mathcal{N}(\mu_{\epsilon,2}, \sigma^2_{\epsilon,2}) & \text{with probability } 1 - p_\epsilon \end{cases} \\ \text{Unemployment probability:} & p_{\nu,it}(z) = \dfrac{e^{\xi_{it}}}{1 + e^{\xi_{it}}}, where \xi_{it} = a + bt + c z_{it} + d z_{it} \end{array} \tag{11}$$

Box I.

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:

$$\text{AIYE}_{it} = \sum_{k=t_0}^{t} \frac{\min\left\{Y_{2,ik}, 2.5\right\}}{t - t_0 + 1} \times Y_{1,t}. \tag{15}$$

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

4.4. Optimization program

The optimal choice of a borrower i is obtained by solving the Bellman equation backward. The borrower sets the optimal consumption C_{its} and makes loan repayments $R_{L,its}$ under the optimal repayment plan $x_{its} \in \{\text{std}, \text{idr}\}$, at each age t and state of the world s, given the level of wealth (W_{its}) and the current loan balance (L_{its}) as well as permanent labor income profile term (z_{its}) as state variables. The optimal action of the agent also depends on the *initial* loan balance (L_{i0}) which determines the standard plan payment. Here is the recursive representation of the optimization program:

$$\begin{aligned} &V_{it}(W_{it}, L_{it}, z_{it}; L_{i0}) \\ &= \max_{C_{it}, R_{L,it}, x_{it}} \frac{1}{1-\gamma} \left(\frac{C_{it}}{\sqrt{N_{it}}} \right)^{1-\gamma} + \\ &\beta(1 - m_t)\mathbb{E}[V_{it+1}(W_{it+1}, L_{it+1}, z_{it+1}; L_{i0})], \end{aligned} \tag{16}$$

We summarized the notation here by dropping dependencies to the state of the world (index s), which is based on the realization of income shocks. The dynamics of state variables follow

$$W_{it+1} = \left(W_{it} - R_{L,it} + Y_{it} + B^{SS}_{it} + B^{SN}_{it} - \Gamma_{it} - C_{it}\right)(1 + r_f) \tag{17}$$

and

$$L_{it+1} = L_{it}(1 + r_L) - R_{L,it}. \tag{18}$$

The dynamic of z_{it} is exogenously set by the labor income process in Eq. (11).

The loan repayment $R_{L,it}$ needs to be *at least* equal to the required values of $R_{\text{std},it}$ (and $R_{\text{garnish},it}$ in case of defaults) under the standard plan: $x_{it} = \text{std}$, and $R_{\text{idr},it}$ under the IDR plan: $x_{it} = \text{idr}$. Larger repayments than the required ones are allowed, reflecting the option to prepay. Appendix Section A.2 provides details on how we numerically solve this optimization program.

5. Welfare

In this section, we present a framework to quantify how public policy can increase social welfare by improving consumption allocation and reducing inefficiencies. We proceed in two steps. First, we propose a quantification of economic inefficiencies by decomposing the distance between the socially-optimal and actual allocation of consumption. This decomposition allows us to identify and quantify the sources of welfare losses. In a second step, we use this decomposition to understand how policy can reduce each of these inefficiencies.

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

Distinguishing these mechanisms is important for two reasons. First, policymakers may value welfare gains differently based on their nature. Providing liquidity and insurance improves welfare without additional taxpayer cost, while redistribution may or may not align with a government's philosophy. Second, some policy goals might be better addressed with instruments beyond the scope of a specific study. 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:

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

subject to the intertemporal budget constraint:

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

where $P_{t_0,t} = \frac{1}{(1+r_f)^t}$ is the price in t_0 of one dollar in period t. We use the notation $\mathbb{E}_j[X] = \mathbb{E}[X|j]$ for the expected value of X conditional on j, and do the same for other moments.

In complete markets, at time t_0, households trade the present value of their lifetime wealth for consumption coupons in future years and states to fully hedge their consumption plan. The optimal portfolio of coupons is such that the utility of spending one more present-value dollar in time t and state s, which we denote v'_{its}, is equal across all periods and states. In other words, the first-best consumption plan implies $\text{Var}_i\left[v'\right] = 0$.

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:

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

Assuming no bequest, the present value cost of the consumption plan remains the same. If $\mathbb{E}_i\left[(C^* - C)P_{t_0}\right] = 0$, the welfare loss from market incompleteness can be approximated as:

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

This covariance can be further decomposed as:

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

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. The aggregate welfare of a cohort of borrowers is the sum of their expected utilities at graduation $\boldsymbol{V} = \sum_i V_{it_0}$.

A strictly utilitarian social planner wants borrowers to smooth consumption over time and across states, but also seeks to equalize marginal utilities across them. Following the same logic as in Eq. (22), the welfare loss between the optimal and actual distributions of the aggregate budget can be approximated as:

$$\boldsymbol{V}^\star - \boldsymbol{V} \approx \overbrace{\sum_i V^*_{it_0} - V_{it_0}}^{\text{Market incompleteness}} + \overbrace{I \cdot T \cdot \mathrm{cov}\left(\mathbb{E}_i[v'], \mathbb{E}_i[(C^\star - C)P_{t_0,t}]\right)}^{\text{Inequality}}. \tag{24}$$

The first term aggregates losses from incomplete markets over all individuals. The second term represents the loss from lifetime expected consumption inequality. Here $C^\star_{its}$ denotes the socially optimum consumption and $C^\star_{its} - C_{its}$ is the deviation of the actual consumption from this optimum one.

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:

$$\Delta^{\$}\boldsymbol{V} = \frac{\Delta \boldsymbol{V}}{\mathbb{E}[v']} \tag{25}$$

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 consumption, equally distributed across states, time and borrowers.

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 average welfare gain per borrower from the policy p is:

$$\begin{aligned}\Delta_{\mathrm{p}}\boldsymbol{V} &= \sum_i \sum_t \mathbb{E}_{it}\left[v(C + \Delta C) - v(C)\right]P_{t,t_0} \\ &\approx \sum_i \sum_t \mathbb{E}_{it}\left[v'\Delta C P_{t_0,t}\right].\end{aligned} \tag{26}$$

Finally, the net welfare gain, subtracting the cost to taxpayers, is:

$$\Delta_{p,\mathrm{net}}\boldsymbol{V} = \overbrace{\Delta_{\mathrm{p}}\boldsymbol{V}}^{\text{benefit to borrowers}} - \mathrm{MCPF} \cdot \overbrace{\sum_{t=t_0}^{T} -\mathbb{E}\left[v'_T\right]\Delta R_{L,t}P_{t_0 t}}^{\text{cost to taxpayers}}, \tag{27}$$

where $\mathbb{E}\left[v'_T\right]$ is to taxpayers what $\mathbb{E}\left[v'\right]$ is to borrowers, and MCPF denotes the Marginal Cost of Public Funds. Because of its distortionary effects, raising one dollar of public funds from taxpayers to finance the policy imposes more than a dollar costs for the aggregate economy. The associated fiscal externalities are captured by an MCPF of greater than one, which scales up the cost to taxpayers in the net welfare calculation (Browning, 1976; Dahlby, 2008).

In the rest of the paper, we assume borrowers and taxpayers to have the same average marginal utility over their lifetime, i.e. $\mathbb{E}\left[v'_T\right] = \mathbb{E}\left[v'\right]$. Readers can easily relax this simplifying assumption by multiplying the monetary transfer from taxpayers to borrowers by their estimate of the marginal utility wedge between the two groups to compute welfare gains or losses associated from transfers. For our baseline optimal policy design in Section 7.2, we set MCPF = 1.4 as the middle range value in the literature (Poterba, 1996; Olken, 2007; Heckman et al., 2010; Finkelstein and Hendren, 2020). Section 7.3 solves for the optimal policy under different calibrated MCPFs.

5.3. Decomposition of welfare gains

To understand the mechanisms through which IDR improves welfare, we can decompose the net welfare gains from policy p into five components. These components can be interpreted as a sequence of four differences in differences which decomposes welfare gains into distinct economic mechanisms:

$$\begin{aligned}\frac{\Delta_{p,\mathrm{net}}\boldsymbol{V}}{I} &= \overbrace{\frac{1}{I}\sum_i\sum_t \mathbb{E}_{it}\left[v'\ \Delta C P_{t_0,t}\right] - \frac{1}{I}\sum_i\sum_t \mathbb{E}_{it}\left[v'\right]\ \mathbb{E}_{it}\left[\Delta C P_{t_0,t}\right]}^{\text{insurance}} \\ &+ \overbrace{\frac{1}{I}\sum_i\sum_t \mathbb{E}_{it}\left[v'\right]\ \mathbb{E}_{it}\left[\Delta C P_{t_0,t}\right] - \frac{1}{I}\sum_i\sum_t \mathbb{E}_{i}\left[v'\right]\ \mathbb{E}_{i}\left[\Delta C P_{t_0,t}\right]}^{\text{intertemporal smoothing}} \\ &+ \overbrace{\frac{1}{I}\sum_i\sum_t \mathbb{E}_{i}\left[v'\right]\ \mathbb{E}_{i}\left[\Delta C P_{t_0,t}\right] - \sum_t \mathbb{E}\left[v'\right]\ \mathbb{E}\left[\Delta C P_{t_0,t}\right]}^{\text{transfer progressivity}} \\ &+ \overbrace{\sum_t \mathbb{E}\left[v'\right]\ \mathbb{E}\left[\Delta C P_{t_0,t}\right]}^{\text{mean transfer}} \\ &- \mathrm{MCPF} \cdot \overbrace{\sum_t -\mathbb{E}\left[v'_T\right]\Delta R_{L,t}P_{t_0 t}}^{\text{cost to taxpayers}}.\end{aligned} \tag{28}$$

Insurance. The first line is the difference between welfare gains to borrowers in the model, and in a counterfactual in which, for each borrower i and year t, v' is equalized across states of the world. In this counterfactual, there is no room for welfare gains from insurance against income risk. However, policies can still redistribute consumption over the life cycle, between borrowers, and from taxpayers to borrowers. Therefore, the difference between the two terms isolates welfare gains from insurance.

Intertemporal smoothing. The second line represents the differences in welfare gains between a world where v' is only equalized across states, and one in which it is also equalized 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. Under the simplifying assumption that $\mathbb{E}\left[v'_{\Gamma}\right] = \mathbb{E}\left[v'\right]$, and if the marginal cost of public funds, MCPF, is equal to one (no distortionary costs of taxation) the welfare gains of this transfer for borrowers is equal to what it effectively costs taxpayers, and this last line washes away. The implication is that, in a complete-market setting, a transfer from taxpayers to borrowers cannot mechanically generate welfare gains. Changes in government revenues and changes in borrower consumption have the same present values and are weighted by the same coefficient $\mathbb{E}\left[v'_{\Gamma}\right] = \mathbb{E}\left[v'\right]$.

Decomposition of policy gains. Assuming $\mathbb{E}\left[v'_{\Gamma}\right] = \mathbb{E}\left[v'\right]$, we can write Eq. (28) as a sum of covariances, net of the fiscal externalities of raising public funds for the policy:

$$\begin{aligned}\frac{\Delta_{p,\text{net}} V}{I} = & \overbrace{\frac{1}{I}\sum_i \sum_t \text{cov}_{it}\left(v', \Delta C P_{t_0,t}\right)}^{\text{insurance}} \\ & + \overbrace{\frac{T}{I}\sum_i \text{cov}_i\left(\mathop{\mathbb{E}}_{it}[v'], \mathop{\mathbb{E}}_{it}\left[\Delta C P_{t_0,t}\right]\right)}^{\text{intertemporal smoothing}} \\ & + \overbrace{T \cdot \text{cov}\left(\mathop{\mathbb{E}}_i[v'], \mathop{\mathbb{E}}_i[\Delta C P_{t_0,t}]\right)}^{\text{transfer progressivity}} \\ & - \overbrace{\sum_t -(\text{MCPF}-1)\cdot \mathbb{E}\left[v'\right] \Delta R_{L,t} P_{t_0 t}}^{\text{fiscal externality}} .\end{aligned} \tag{29}$$

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. The last term presents fiscal externalities of transfers from taxpayers to an average borrower.

Note that, if the distribution of consumption is already at the first-best, then no variation in v' exist and each covariance term in Eq. (28) is zero by construction. Therefore no welfare gain can be expected from policy changes. Also, with MCPF $= 1$ transfers from taxpayers to borrowers do not affect net welfare, while for the realistic case of MCPF > 1 such transfers would reduce the net welfare.

Implementation. Our decomposition is conceptually similar to the one proposed by Dávila and Schaab (2025). The main advantage of our approach lies in its ease of implementation. Our method uses a first-order approximation of how utility gains arise from redistributing consumption across states, time, and individuals. This allows us to implement the decomposition directly from the simulated data of two model solutions only—the baseline and the policy counterfactual. As detailed in Appendix Section A.3, the algorithm is simple: it requires only taking averages sequentially across states, time, and individuals to compute the components of Eq. (28).

6. Existing IDR rules

In this section, we calibrate and simulate our model to examine the welfare effect of existing IDR rules.

6.1. Model calibration

Earnings. The income process is calibrated following Guvenen et al. (2021) with two differences. First, to make model more tractable, we assume unemployment shocks to last an entire year, which is a very close approximation of what these authors estimate. Second, 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. Appendix Section A.1 reports the parameters of this function and the stochastic process.

Dependents and poverty line. We calibrate the number of persons in the households N_{it} as a deterministic function of age. 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 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 national wage index when a borrower graduated. Fig. 6 reports the initial debt for each group. Finally, we calibrate the initial level of wealth to match the relationship between student debt and wealth among fresh graduates in the SCF.

Terminal conditions. We solve the model by dynamic programming, starting from retirement year t_R. We approximate the expected utility at retirement using the solution from (Merton, 1971):

$$\mathbb{E}\left[V_{t_R}\right] = b\frac{\overline{W}_{iR}^{1-\gamma}}{1-\gamma} \tag{30}$$

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

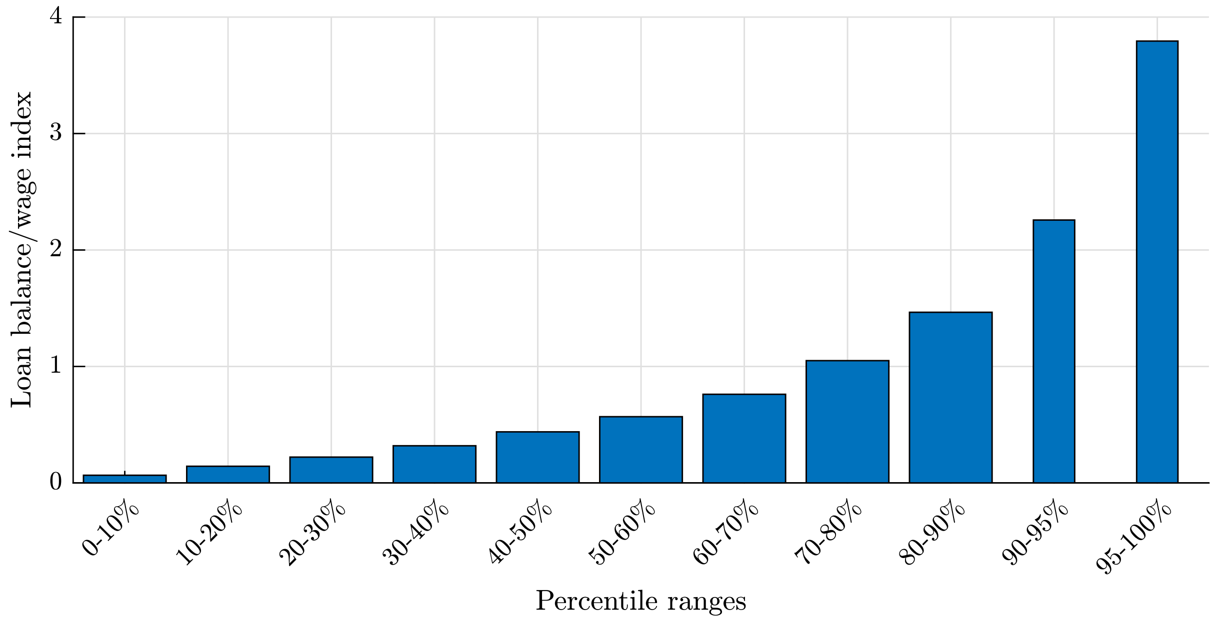

Fig. 6. Initial balance by centile of loan size.
This figure displays the median student loan balance at graduation (or exit) for the cohort that completed their program in 2019 and 2020, segmented by debt decile. The final decile is divided into two equal parts.
Source: TransUnion.

Table 1
Welfare gains to borrowers from IDR.

Total change	Financial frictions		Transfer	
in borrower welfare	Insurance	Intertemporal smoothing	Progressivity	Mean
0.793	0.169	0.237	0.002	0.377

This table reports the decomposition of aggregate welfare gains for borrowers under existing rules. The decomposition is defined in Eq. (28). Welfare gains are converted in dollar equivalent per borrower using Eq. (25), and normalized by the national wage index ($63,795 in 2022).

6.2. *Welfare gains*

Table 1 reports welfare gains from IDR plans for borrowers, in dollars (scaled by the national wage index), and broken down by economic mechanism. Under existing rules, improvement in consumption smoothing and better insurance against income risk explains slightly more than half of the welfare gains of borrowers. Subsidization from taxpayers explains the rest.

Relative to the standard ten-year plan, IDR increases borrower welfare by 0.793 times the average wage, or approximately $50,600. The largest single component of this welfare increase, accounting for 48% of the gain, comes from a transfer from taxpayers. IDR reduces the average present value of future student loan repayments at graduation by 0.377 of average wage ($\approx \$24,100$). Most of the rest of the welfare gains come from the fact that IDR also helps borrowers smooth consumption over the life cycle and provides insurance against income risk.[14] These two sources represent, in unit of national wage index, 0.237 ($\approx \$15,100$) and 0.169 ($\approx \$10,800$) of welfare gains, respectively. Overall, existing IDR rules do not constitute a zero-sum game that only improves the welfare of borrowers at the expense of taxpayers. In fact, assuming an MCPF = 1.4, IDR generates $0.793/(1.4 * 0.377) \approx 1.5$ times more welfare in dollar terms for borrowers than the cost to taxpayers, including fiscal externalities.

To interpret magnitudes, consider that the average balance at graduation is 0.806 ($\approx \$51,400$). First, present value gains from IDR represent 47% of this amount. Importantly, these gains are relative to the standard-plan payments, which, due to high interest rates, sum up to a present value higher than the loan itself. Second, welfare gains to borrowers are twice as large as present value gains. To generate such efficiency gains, present value gains need to increase consumption in periods×states for which v' is twice as large as $\mathbb{E}[v']$. Within a given period, and for $\gamma = 2$, this corresponds to a 30%-below-average consumption level.[15]

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 distribution of lifetime consumption.[16]

6.3. *Distribution of welfare gains*

Panel A of Fig. 7 reports how the average per-borrower welfare gains vary as a function of their lifetime earnings and their debt level at graduation. 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, welfare gains are decomposed from the perspective of the social planner, who converts utility gains to a monetary equivalent

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

[15] As $u'(C) = C^{-\gamma}$ and $0.7^{-2} \approx 2 * 1^{-2}$. Panel A of Appendix Figure A.2 shows that, over the life-cycle, consumption gains are concentrated below age 35 and relative consumption levels consistent with this back-of-the-envelope calculation.

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

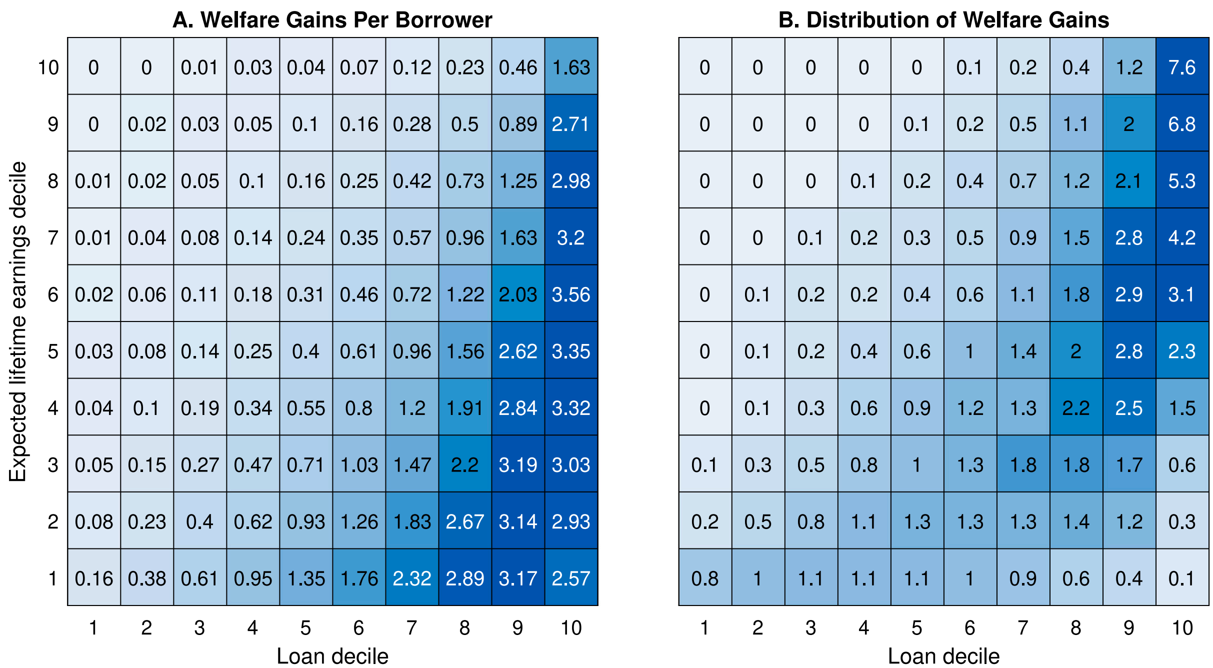

Fig. 7. Welfare Gains Per Borrower by decile of debt and income.
Panel A reports the average welfare gain from IDR per borrower as a function of their decile of expected lifetime earnings and decile of student loan at graduation. Gains are reported in multiples of the national average wage index ($63,795 in 2022). Panel B 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 total aggregate gains for borrowers.

using the same conversion rate $\mathbb{E}\left[v'\right]^{-1}$ for all borrowers. From the perspective of individual borrowers, the equivalent wealth 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 the social planner, the per-person welfare gains in the top decile of balances exceed 2 average wage (>$130,000). From these borrowers' perspective, the equivalent wealth variation would likely be higher since they tend to have higher expected lifetime consumption.

Paradoxically, borrowers in the top decile 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.

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.

To account for the distribution of borrowers, Panel B of Fig. 7 reports the percentage share of aggregate welfare gains by decile of expected earnings and debt at graduation. Approximately 24% of the total welfare gains go to borrowers in the top decile of the debt distribution and the top four deciles of the expected earnings distribution. This distribution explains why IDR does not deliver aggregate welfare by redistributing consumption between borrowers, as shown by the progressivity channel being nearly zero in Table 1. The fact that IDR benefits borrowers with high balances is offset by the their tendency to also be high earners.

7. Policy design

Can policymakers improve the calibration of IDR parameters to deliver higher welfare gains to borrowers or reduce its cost to taxpayers? In this section, we analyze the mechanisms through which the calibration of IDR plans impacts both borrower welfare and the 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 welfare net of cost to taxpayers and fiscal externalities. Our findings indicate that current IDR plans are not cost-efficient and that alternative calibrations could provide greater benefits to borrowers while substantially reducing costs for taxpayers, including fiscal externalities.

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. (2022). Appendix Section A.4 describes and validates this methodology.

7.1. The role of IDR parameters

IDR rules are governed by three key parameters: the number of payment years before forgiveness, the threshold above which earnings are deemed discretionary, and the percentage 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? Fig. 8 reports the change in the four components of welfare when we vary each of these four policy parameters independently, holding the remaining three to their baseline value.

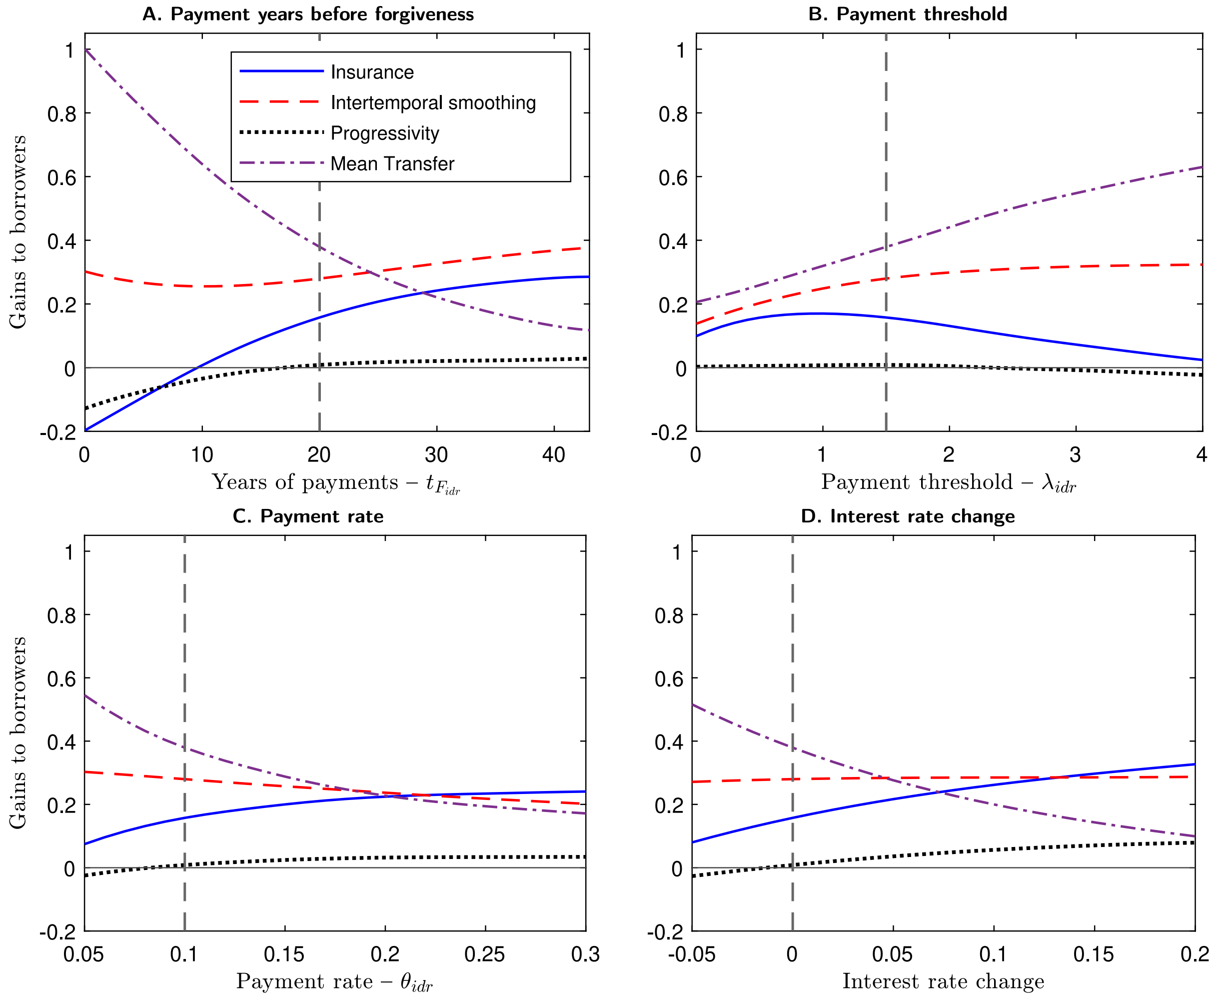

Fig. 8. Role of IDR parameters.
This figure illustrates how the four components of welfare gains evolve with four IDR parameters. Each parameter is varied independently, with the remaining three held constant at their baseline values under existing IDR rules, indicated by vertical dashed lines. Panel A reports how the four component of welfare depend on the number of payment years before forgiveness. That is the parameter $t_{F_{idr}}$ in Eq. (9). 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 Eq. (9). Panel C examines the effects of the discretionary income share allocated to payments. That is the parameter θ_{idr} in Eq. (9). Panel D explores the changes in welfare components in response to a uniform increase in the IDR loan interest rate for all loan amounts.

Years of payment before forgiveness $\left(t_{F_{idr}}\right)$. 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 at later periods, 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 examines the effect of the income threshold—expressed as a proportion of the federal poverty line—above which IDR payments begin. As expected, raising the threshold lowers payments and increases transfers from taxpayers to borrowers. More interestingly, the welfare insurance component exhibits a hump-shaped relationship with the threshold. At one extreme, if the threshold is zero, borrowers receive no downside protection since payments are owed from the first dollar of earnings. At the other extreme, if the threshold is infinite, borrowers never repay, and the link between IDR benefits and income risk disappears. In between, higher thresholds provide stronger protection in low-income states but also mechanically reduce payments from high-income borrowers by shrinking the share of their income subject to repayment. This dual effect explains the hump: as the threshold rises, additional protection no longer accrue to low-income borrowers but keep reducing payments of high-income borrowers. Importantly, the negative revenue effect of higher thresholds could in principle be offset by raising the payment rate.

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 closer to the payment threshold are less affected in proportion of their income. 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 transfer components of borrower welfare move inversely, and why raising the payment rate increases IDR's progressivity. This does not imply, however, that low earners gain more in absolute terms, but

Table 2
Jacobian matrix of welfare components with respect to policy parameters, normalized by fiscal cost.

	Financial frictions		Transfer		Total change
	Insurance	Intertemporal smoothing	Progressivity	Mean	in borrower welfare
Years of payment ($t_{F_{\mathrm{idr}}}$)	−0.405	−0.139	−0.085	0.714	0.086
Payment threshold (λ_{idr})	−0.246	0.299	−0.016	0.714	0.751
Payment rate (θ_{idr})	−0.345	0.145	−0.135	0.714	0.379
Loan interest rate	−0.408	−0.032	−0.187	0.714	0.087

This table reports the first-order derivative of the components of borrower welfare gains with respect to the IDR policy parameters, at the current IDR policy rule, and divided by the derivative of the fiscal cost. The components of borrower welfare gains come from the decomposition defined in Eq. (28). Welfare gains are estimated using the auxiliary model and reported per dollar cost to the government, which is equal to the change in the present value of loan payment multiplied by the MCPF of 1.4. Derivatives are based on 1 year increase in IDR years of repayment ($t_{F_{\mathrm{idr}}}$), 10 percentage point increase in the payment threshold (λ_{idr}), 1 percentage point increase in the payment rate (θ_{idr}), and 0.1 percentage point increase in the IDR loan interest rate, in rows 1–4, respectively.

rather that they capture a larger share of the program's benefits. Taken together, the results suggest that jointly raising the payment threshold and the repayment rate can deliver larger gains from intertemporal smoothing without eroding the insurance value of IDR.

Loan interest rate. Finally, Panel D examines the impact of increasing the interest rate on student loans. As anticipated, higher interest rates lower the transfers from 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.

Policy trade-offs. Setting policy parameters requires balancing borrower welfare gains against their fiscal cost. For a policy to be optimal at a given budget, increases in generosity across different parameters should deliver the same marginal welfare gains per additional taxpayer dollar, once fiscal externalities are taken into account. A second condition is that these marginal gains equal the marginal cost of public funds. Table 2 reports the Jacobian matrix of welfare with respect to policy parameters, scaled by the derivative of taxpayer costs and fiscal externalities. The matrix therefore measures the welfare benefits to borrowers of expanding generosity along a given policy dimension, normalized by the associated fiscal cost.

The last column of Table 2 shows along which policy margin generosity is the most cost-efficient. Reducing the number of years of payment or the loan interest rate are the least efficient ways to increase the cost of the student loan program, as it only transfers consumption from taxpayers to borrowers without actually generating efficiency gains. By contrast, the most efficient way to generate welfare is to increase the payment threshold, as it allows borrowers to differ payment in time.

Differences between rows show the consequence of being more generous on one parameter while keeping the budget constant by being less generous on another. For example, subtracting the third row from the second shows that financing an increase in the payment threshold by increasing the payment rate deliver net gains in terms of insurance (−0.246 + 0.345 = 0.099), intertemporal smoothing (0.299–0.145 = 0.154) and progressivity (−0.016 + 0.135 = 0.119), for a total net welfare gain of 0.372.

Overall, Table 2 indicates that, for a given program cost, the current calibration of IDR is suboptimal: welfare could be improved by lengthening the repayment period while raising the payment threshold.

Table 3
Optimal IDR parameters.

Policy	Years of payment $t_{F_{\mathrm{idr}}}$	Payment threshold λ_{idr}	Payment rate θ_{idr}
Budget-neutral	43	131%	23.6%
Highest net welfare	43	237%	24.4%
Status Quo	20	150%	10%

This table reports estimated IDR policy 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 and fiscal externalities. The Status Quo reports parameters under existing rules.

Furthermore, since all components in the last columns fall below one, our results suggest that cutting program costs would yield greater welfare gains for taxpayers than the corresponding losses borne by borrowers.

7.2. Optimal policy calibrations

So far, we have discussed how adjusting various policy parameters affects borrower welfare and program costs. 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, accounting for costs to taxpayers and fiscal externalities.

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 remaining key IDR parameters: the repayment period ($t_{F_{\mathrm{idr}}}$), the payment threshold (λ_{idr}), and the payment rate (θ_{idr}), while keeping the interest rate the same.

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 $t_{F_{\mathrm{idr}}} \leq 43$. Second, since federal and state income taxes can combine to reach a marginal tax rate of 40%, we impose $\theta_{\mathrm{idr}} \leq 0.6$, ensuring the overall marginal "tax" rate does not exceed 100%. Table 3 presents the optimal IDR parameters for the two budget-neutral and net-welfare-maximizing policies.

Table 4
Welfare gains from optimal IDR parameters.

	Total change in borrower welfare	Financial frictions		Transfer	
		Insurance	Intertemporal smoothing	Progressivity	Mean
Budget-neutral	0.604	0.307	0.274	0.023	0.000
Highest net welfare	0.808	0.329	0.374	0.037	0.068
Status Quo	0.793	0.169	0.237	0.002	0.377

This table reports the decomposition of borrower welfare gains under the optimal IDR policy parameters reported in Table 3. The decomposition is defined in Eq. (28). 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.

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 be slightly lower, at 131% of the federal poverty line, and the program would collect 23.6% of discretionary income, up from 10%. As a budget-neutral policy, this approach would reduce the expected cost of IDR for taxpayers from about $24,100 per borrower to zero.

Table 4 details the corresponding welfare gains and their components. The budget-neutral policy generates welfare gains of 0.581 ($37,100) from intertemporal smoothing and insurance, which exceeds the gains from these channels under the current IDR rules, which is 0.406 ($25,900). Additionally, it performs slightly better when welfare gains from progressivity are included, yielding gains of 0.604 ($38,500), 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 any 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 ($50,600), compared to 0.604 ($38,500) under the budget-neutral policy. Nonetheless, the incremental improvement of 0.189 ($12,100) represents only half of the 0.377 ($24,100) cost to taxpayers, making net welfare substantially higher under the budget-neutral policy, when taxpayers' interests and other fiscal externalities are taken into account. Our results then show that financial frictions can be effectively mitigated without subsidizing borrowers.

The decomposition of welfare gains provides two key insights. First, the budget-neutral policy exhibits slight progressivity, whereas current IDR rules were neither progressive nor regressive. Second, the budget-neutral policy delivers greater welfare benefits via both intertemporal consumption smoothing and insurance, with the latter being the dominant channel.

Policy maximizing net welfare. Tables 3 and 4 also report the IDR parameters and the decomposition of welfare gains for the policy that maximizes borrower welfare gains net of taxpayer costs and fiscal externalities (Eq. (29)). This policy would be more progressive, compared with the current IDR policy, featuring a higher repayment threshold; but it has a long repayment period and a higher repayment rate. The repayment period is at its maximum of 43 years (instead of 20), and the payment rate is 24.4% (instead of 10%). The additional budget is used to raise the payment threshold to 237% (up from 150%). This change enhances gains from intertemporal smoothing by reducing payments early in the life cycle, as it shields a larger portion of earnings. The taxpayer cost 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 existing IDR rules to the net-welfare-maximizing policy reveals that borrowers would be relatively indifferent between the two, despite the latter costing only 0.068 ($4300), i.e., 0.309 ($19,700) less to taxpayers. 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 and fiscal externalities. This is why the net-welfare maximizing policy in the end features some transfers.

Welfare gains frontier. Large subsidies in existing IDR plans raise the question of how much welfare can be increased for a given budget. Fig. 9 provides some insight into this question by showing welfare gains to borrowers as well as net welfare gains under optimally calibrated parameters, as a function of taxpayer cost. The blue solid line represents welfare gains from IDR for borrowers, while the dashed red line represents the net welfare gains after accounting for the taxpayer cost and fiscal externalities. Additionally, the figure highlights the welfare gains and costs associated with existing IDR, illustrating the gap between the Status-Quo and cost-efficient calibrations. Net welfare gains peak for a per-borrower budget of 0.068 ($4300), 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 and fiscal externalities.

Fig. 9 underscores several key findings. First, a significant portion of the welfare gains delivered by current IDR rules 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 Current IDR policy on the "welfare gains to borrowers" curve shows that the same level of welfare gains achieved under existing rules could be realized at a cost of approximately $4100 per borrower, representing a savings of $20,000 to taxpayers.

Why are optimal rules 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-marginal-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 utility loss to borrowers.

7.3. Marginal cost of public funds and optimal IDR

This section checks the robustness of our welfare results to the calibrated MCPF, which takes into account the fiscal externalities of

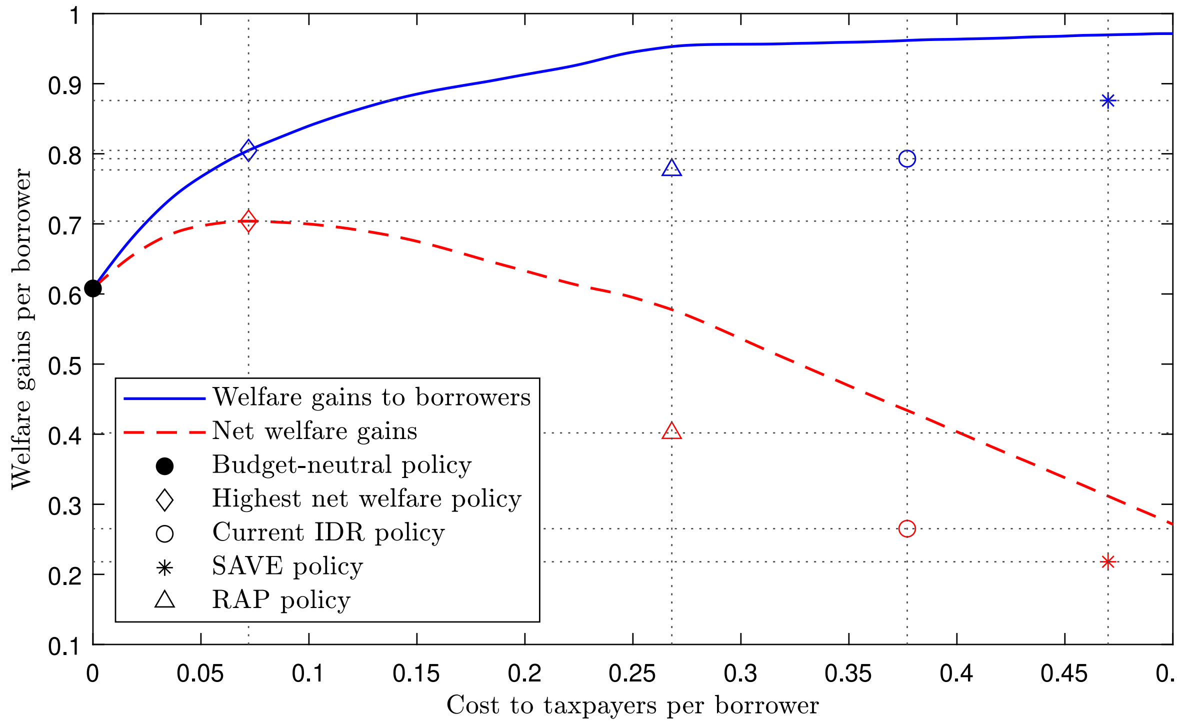

Fig. 9. Welfare Gains by policy budget under optimal calibration.
This figure illustrates the relation between the per-borrower budget allocated to IDR and welfare gains to borrowers (blue line) or net welfare gains (red dashed line) under the corresponding optimal calibration of IDR parameters. The black dot (•) marks the coordinates of the budget-neutral policy. Diamonds (◊) mark the coordinates of welfare gains under the net-welfare maximizing policy. Circles (∘), stars (∗) and triangle (△) represent current IDR rules, SAVE, and RAP 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). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the policy cost, including distortionary effects of taxation. Varying MCPF changes the weights of the cost to taxpayers and welfare gains to borrowers in the net-welfare Eq. (29). Table 5 reports the optimal IDR policy parameters as well as the resulting optimal transfers and borrower welfare gains when re-optimizing net welfare for different values of MCPF. Following estimates from the public finance literature and cost–benefit analyses of public policies (Ballard et al., 1985; Poterba, 1996; Kleven and Kreiner, 2006; Heckman et al., 2010), we consider values of MCPF between 1 and 2. MCPF = 1 represents a special case that taxation is not distortionary (no fiscal externalities), while MCPF = 2 represents an extreme upper case. MCPF = 1.4 refers to our baseline optimal policy results in Section 7.2.

As expected, with higher MCPFs the optimal size of transfers is reduced. Consequently, borrower welfare gains decreases. The optimal policy becomes less generous through a lower payment threshold λ_{idr}, falling from 297% under MCPF = 1, to 207% under MCPF = 2. The repayment horizon remains at its maximum of 43 years, since in either scenario shortening the repayment period does not generate enough welfare gains per dollar of extra fiscal costs. In the end, though, the resulting optimal welfare numbers do not materially change over the entire range of MCPF $\in [1, 2]$.

Even with the larger MCPFs, the optimal policy continues to deliver substantial welfare gains to borrowers. This is the case because the source of welfare gains in an optimal policy is not transfers from taxpayers, but instead improving insurance and intertemporal smoothing. This means that, through a public-policy perspectives, the optimal IDRs are highly efficient. We quantify this efficiency by the Marginal Value of Public funds (MVPF, Finkelstein and Hendren, 2020), measured as the welfare gains in dollar terms for borrowers per dollar of transfers from taxpayers when multiplied by MCPF to include fiscal externalities. Table 5, last Column reports large MVPFs of around 8 for the optimal IDR policies. For comparison, Hendren and Sprung-Keyser (2020) estimate MVPFs for a vast number of historical policy changes and find

Table 5
Marginal cost of public funds and optimal IDR policy.

MCPF	Optimal policy parameter			Optimal welfare results		
	$t_{F_{idr}}$	λ_{idr}	θ_{idr}	Total change, borrowers	Transfer, Mean	MVPF
1	43	297%	35.4%	0.850	0.104	8.2
1.2	43	246%	25.0%	0.817	0.075	9.1
1.4	43	237%	24.4%	0.808	0.068	8.4
1.6	43	226%	24.4%	0.795	0.060	8.3
1.8	43	226%	24.5%	0.795	0.059	7.5
2	43	207%	24.3%	0.767	0.045	8.5

This table reports estimated policy parameters that maximize net welfare, and the resulting optimal transfers and borrower welfare gains with varying values of the Marginal Cost of Public Funds (MCPF). The optimal policy in each row maximizes borrower welfare net of the cost to taxpayers time the MCPF, to include fiscal externalities (Eq. (29)). The value of MCPF = 1.4 resembles the baseline “Highest net welfare" results in Section 7.2 and Table 4. The last column shows the marginal value of public funds, MVPF, defined as the welfare gains to borrowers divided by direct cost to taxpayers (transfers) and the MCPF (to account for fiscal externalities).

values generally between 0.5 and 2 for policies that target adults. In fact, in our setting, even a zero-budget optimal policy generates welfare gains for borrowers (Table 4, first row), which implies that MVPF may be arbitrarily large.

This is not the case, though, with the current IDR rule, under which the policy cost averages 0.377, nearly half of the welfare gains of 0.793 for borrowers (Table 4, third row). With the benchmark calibration of MCPF = 1.4, the resulting MVPF for the current IDR rule is $0.793/(1.4 * 0.377) \approx 1.5$. The consideration of distortions in raising public funds, in particular, further supports the need to re-optimize IDR plans, to generate efficiency gains by addressing financing frictions, insurance and intertemporal smoothing for borrowers, while avoiding transfers from taxpayers.

8. Comparison to recent policy proposals

In this section, we apply our framework to evaluate the potential gains from recent policy proposals and benchmark their efficiency against our optimal IDR calibration. We focus on the two most prominent reforms: the Saving on a Valuable Education (SAVE) plan proposed by the late Biden administration, and the Repayment Assistance Program (RAP), enacted into law in July 2025.

8.1. Saving on a valuable education plan

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

$$\begin{cases} R_{\text{save},it} = \theta_{\text{save},i}\left(Y_{it} - \lambda_{\text{save}} Y_{\text{pl},t}\right)^{+} & \text{if } t \le t_{F_{\text{save}},i} \\ R_{\text{save},it} = 0 & \text{if } t > t_{F_{\text{save}},i}. \end{cases} \tag{31}$$

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 $\lambda_{\text{save}} = 225\%$ 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 who borrowed less than $L_{it_0} \le \$12,000$. Each additional \$1000 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.[17]

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:

$$L_{it+1} = \min\left\{L_{it}(1+r_L) - R_{\text{save},it}, L_{it}\right\}. \tag{32}$$

Welfare gains. Table 6 compares the welfare gains from SAVE relative to existing IDR rules. The SAVE proposal increases borrower welfare by 0.083 (≈\$5300). 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 only 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. However, they also reduce the program's 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 0.470 times the wage index (≈\$30,000), a 0.093 (≈\$5900) 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.[18] This is confirmed by the location of SAVE in the cost-efficiency space of Fig. 9. First, the overall subsidy of 0.470 times the wage index (≈\$30,000) embedded in the SAVE program is well beyond the optimal budget of 0.068 (≈\$4300) per borrower. Second, at this level of subsidy, SAVE is dominated by the optimal calibration of IDR parameters. This finding is not self-evident as SAVE has significantly more complicated rules and so gives policymakers more degrees of freedom to maximize welfare.

Table 6
Welfare gains to borrowers from recent policy proposals.

	Total change in borrower welfare	Financial frictions		Transfer	
		Insurance	Intertemporal smoothing	Progressivity	Mean
Status Quo	0.793	0.169	0.237	0.002	0.377
SAVE	0.876	0.114	0.294	−0.007	0.470
Change	0.083	−0.055	0.057	−0.009	0.093
RAP	0.777	0.216	0.276	0.016	0.268
Change	−0.016	0.047	0.039	0.014	−0.109

This table reports the decomposition of aggregate welfare gains for borrowers under existing IDR rules, SAVE and RAP. The decomposition is defined in Eq. (28). Welfare gains are converted in dollar equivalent per borrower using Eq. (25), and normalized by the national wage index (\$63,795 in 2022).

Comparative statics reported in Fig. 8 help explain why SAVE fails to improve upon existing IDR rules. First, SAVE reduces the time to forgiveness for undergraduate borrowers. Panel A of Fig. 8 shows that this policy change lowers the gains from insurance, intertemporal smoothing, and progressivity. Of course, borrowers benefit from this change, but the increased subsidization it entails more than offsets the associated efficiency losses. From the combined perspective of borrowers and taxpayers (including fiscal externalities), this change represents a negative-sum game, especially when considering the marginal cost of public funds, MCPF of greater than 1. SAVE also subsidizes undergraduate borrowers by raising the payment threshold. Panel B of Fig. 8 shows that, all else equal, this change has an ambiguous effect in terms of efficiency: it enhances intertemporal smoothing by deferring payments, but reduces the value of insurance. The net effect of these two channels appears small—particularly when compared to the direct increase in subsidization. Overall, Fig. 9 suggests that SAVE would re-allocate resources from taxpayers to borrowers without net welfare gains.

8.2. Repayment assistance program

Rules. RAP rules contrast with SAVE in many respects. First, rather than shortening it, RAP extends the repayment period to $t_{F_{\text{rap}}} = 30$ years. Second, while repayment remains convex in income, the percentage of earnings going to reimbursement $\theta_{\text{rap},i}$ no longer depends on the federal poverty line. For borrowers below \$10,000, yearly payment is set to \$600. Between \$10,000 and \$20,000, borrowers pay 1% of their earnings. This percentage jumps every \$10,000 up to a maximum of 10%. Finally, borrowers with kids receive a payment subsidy δ_{kids} of \$600 a year per child. We implement these rules by scaling these parameters by the national wage index and assuming that they will follow the growth of per capita earnings.[19]

$$\begin{cases} R_{\text{rap},it} = \left(\theta_{\text{rap},i}\left(Y_{it}\right) \cdot Y_{it} - \delta_{\text{kids}} \cdot n_{\text{kids}}\right)^{+} & \text{if } t \le t_{F_{\text{rap}}} \\ R_{\text{rap},it} = 0 & \text{if } t > t_{F_{\text{rap}}}. \end{cases} \tag{33}$$

Like SAVE, RAP prevents balances from increasing and further insures that they decrease by a dollar amount $\delta = \$600$ every year. Hence, balances evolve as:

$$L_{it+1} = \min\left\{L_{it}(1+r_L) - R_{\text{rap},it}, L_{it} - \delta\right\}. \tag{34}$$

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

[18] Importantly, in Table 6, 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 9.1.

[19] At this stage, there is no clear indication regarding how the parameters will evolve in the future.

Welfare gains. Table 6 reports welfare gains under the RAP rules. The changes relative to existing IDR rules differ substantially from those under the SAVE proposal. Rather than increasing by 0.093, the cost to taxpayers decreases by 0.109 (≈$7000), indicating that RAP is fiscally less generous than the current IDR system. However, the average welfare gain per borrower declines by only 0.016 (≈$1000). This relatively smaller drop reflects the fact that part of the fiscal savings from extending the repayment period are recycled into more generous repayment terms for low-income borrowers, allowing for greater payment deferral and improved insurance against income shocks. These benefits are captured by the efficiency gains of 0.047 and 0.039 through the insurance and intertemporal smoothing channels, respectively.

Overall, Fig. 9 shows that RAP remains cost-inefficient. Nonetheless, the policy is directionally consistent with the central implication of our model: extending the repayment horizon can enhance welfare if paired with more progressive repayment rules. This is evident in the net welfare gains, as borrower welfare declines by significantly less than the reduction in the program's budget and associated fiscal externalities.

9. Discussion

In this section, we discuss the implications of two simplifying assumptions in our model. First, we treat the distribution of debt at graduation as exogenous. We argue that concerns about moral hazard are likely to be smaller under the optimal calibration of IDR than under existing rules or SAVE, as the former reduce the subsidization of student loans. Second, we abstract from potential effects of IDR rules on labor supply. Drawing on both theoretical models and empirical evidence, we argue that labor supply responses are likely to be small and, if anything, would reinforce our findings.

9.1. Cost of debt and borrowing incentives

Table 1 shows that IDR significantly reduces the effective cost of student debt for borrowers. Therefore, it introduces potential moral hazard for both students and schools. A lower effective cost of debt can encourage more students to borrow and borrowers to accumulate larger balances. 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). Changes in repayment rules that increase the subsidization of student loans—such as SAVE, but not RAP—may induce students to borrow more. We start by comparing current IDR rules and SAVE rules to illustrate the risk of moral hazard in borrowing behaviors. We then show that this risk would be substantially reduced in our optimal calibration.

Current IDR rules. To explore the over-borrowing issue, Panel A of Fig. 10 reports the expected cost of debt across borrowers as a function of their balance at graduation under existing rules and SAVE. The top (blue) line shows the cost of debt per dollar under existing IDR rules. Under existing rules, student debt is a relatively fairly-priced source of financing for debt below the undergraduate limit. 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 graduation balances include interests accrued over the period of study. Very large debts, typically accrued during graduate school are highly subsidized.

Panel B reports the average marginal cost of debt, highlighting the moral hazard at the intensive margin. Even under existing rules, the average marginal cost of debt is below $1 for balances exceeding $7000, potentially encouraging some students to borrow more. However, the implied subsidy was relatively small and could have been offset by factors such as accruing interests before graduation, inattention, debt aversion, or other frictions.

SAVE. The SAVE rules complicate the relationship between the present value of debt and its face value at graduation. First, the average cost of debt experiences discrete jumps every $1000 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 Fig. 10, 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.

With 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. Under existing IDR rules, the marginal cost of debt only increases discontinuously at points where the marginal interest rate shifts. With SAVE, the marginal cost also exhibit discontinuities at points where changes occur in the payment formula or the length of the repayment period.

The total cost of SAVE for taxpayers would strongly depend on the behavioral response of undergraduate students. In Appendix B, we estimate its cost to be $15bn per cohort of new borrowers if the distribution of loans at graduation does not change. This cost would rise to $25bn if the number of borrowers does not change but undergraduate borrowers borrow as much as they could, and to over $70bn if all undergraduate students become borrowers.

Optimal policy calibrations. Moral hazard concerns on borrowed amounts would be very limited under our optimal policy calibrations. Fig. 11 shows that under both the optimal budget-neutral and net-welfare maximizing calibration of IDR, the average and marginal cost of debt remains 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.

9.2. Labor supply

Another potential concern with income-driven repayment rules that increase repayment rates is the possibility that higher effective marginal tax rates discourage labor supply. Empirically, Britton and Gruber (2020) finds no effect of IDR on labor supply in the United Kingdom, but de Silva (2025) documents more response from Australian workers. Moreover, Fu et al. (2025) show that IDR can impact job choices through moral hazard.

In this section, we try to estimate the potential fiscal cost of changes to IDR rules caused by labor supply responses. We proceed in two steps.

First, for each borrower i we compute a lifetime marginal tax rate (LMTR), defined as the ratio of the change in the present value of loan repayments and taxes paid to the change in the present value of lifetime earnings following a marginal increase in earnings:

$$\text{LMTR}_i \equiv \frac{\sum_t (1+r)^{-t}\, \mathbb{E}_t\left[\Delta \Gamma_{i,t}\right]}{\sum_t (1+r)^{-t}\, \mathbb{E}_t\left[\Delta Y_{i,t}\right]}, \tag{35}$$

where $\Gamma_{i,t}$ denotes loan repayments and taxes at time t, $Y_{i,t}$ denotes labor earnings, and r is the risk-free interest rate used to discount future flows. This formulation captures the intertemporal nature of IDR repayment schedules: a lower payment today generally implies higher payments later, unless part of the loan is ultimately forgiven. So what matters for labor supply is not so much variation in the effective marginal tax rate in any given year but overall. This approach is also better suited to capture potential effects on career choices at graduation, which have long-run effects on earnings.

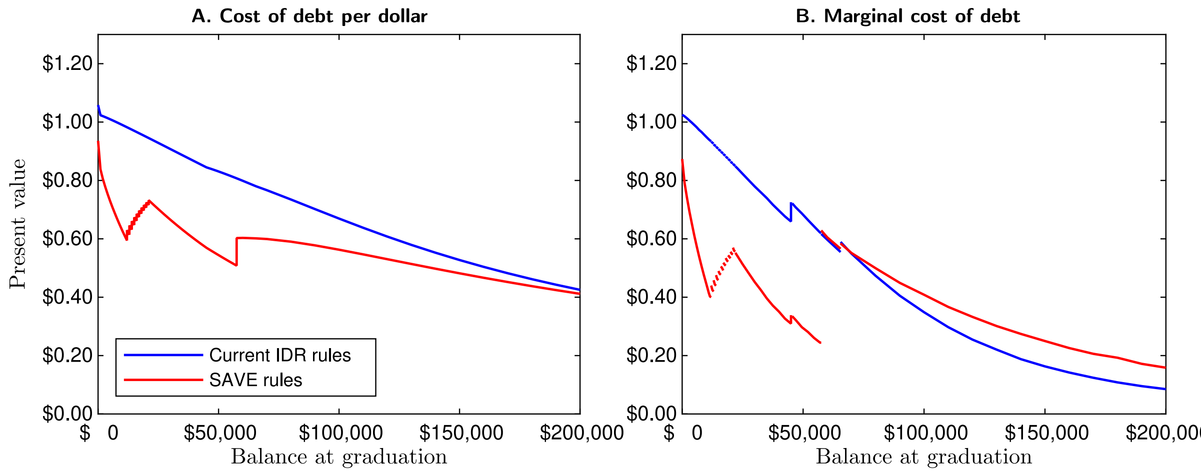

Fig. 10. Cost of student debt under existing rules and SAVE.
Panel A reports the average simulated present value of student debt at graduation, per dollar of debt, under existing IDR rules and with SAVE. Panel B reports the mean marginal cost of student debt, defined as the present value of borrowing one more dollar just before graduation. The marginal cost of debt is infinite at points where the number of repayment years increases. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

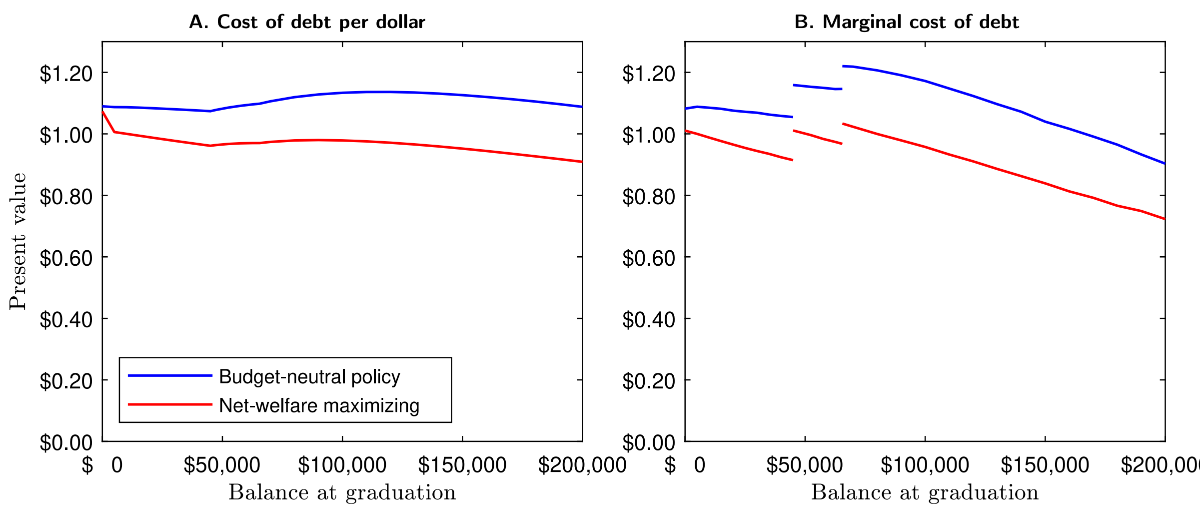

Fig. 11. Cost of student debt with optimal policy parameters.
Panel A reports the average simulated present value of student debt at graduation, per dollar of debt, under the budget-neutral and net welfare maximizing policies presented in Section 7.2. Panel B reports the mean marginal cost of student debt, defined as the present value of borrowing one more dollar just before graduation.

Second, given an elasticity of labor supply with respect to the net-of-tax earnings rate, ε, the percentage change in labor supply (hours worked) for individual i under policy p relative to the status quo policy is obtained by:

$$\frac{\Delta h_{i,p}}{h_{i,0}} = \left(\frac{1-\text{LMTR}_{i,p}}{1-\text{LMTR}_{i,0}}\right)^{\varepsilon} - 1. \tag{36}$$

Assuming earnings are proportional to hours worked, this maps directly into a change in lifetime earnings,

$$\Delta Y_{i,p} = Y_{i,0} \cdot \frac{\Delta h_{i,p}}{h_{i,0}}. \tag{37}$$

Finally, the behavioral tax and loan revenue effect of policy p is obtained by applying the policy's LMTR to the incremental earnings:

$$\Delta R_p^{\text{beh}} = \sum_i \text{LMTR}_{i,p} \cdot \Delta Y_{i,p}. \tag{38}$$

We report $-\Delta R_p^{\text{beh}}$ as the tax and loan revenue loss due to discouraged labor supply, relative to current IDR rules.

Table 7 reports our findings for an elasticity of $\varepsilon = 0.25$. Overall, labor supply varies by less than one percent across IDR policies. Counterintuitively, lifetime effective marginal tax rates are lower under the optimal IDR policies than under the status quo, even though total payments to the government are higher when holding lifetime income fixed. The optimal calibrations of IDR therefore increase the labor supply. On the other hand, SAVE reduces labor supply. This arises because the repayment period is extended from 20 to 43 years under optimal rules, while it is reduced under SAVE. By reducing the probability that any balance is ultimately forgiven, the optimal policies makes repayment more unavoidable and therefore mitigates moral hazard. Manso et al. (2025) provide a formal theoretical discussion of this argument. The result is also consistent with the findings of Boutros et al. (2024) regarding a 10-year repayment deferral.

Variations in government revenues induced by labor supply responses across different IDR policies are on the order of 1% of the national wage index, or roughly $600, which is comparable to the differences reported between IDR policies in de Silva (2025). Relative

Table 7
Fiscal cost from labor supply response.

	Status Quo	SAVE	RAP	Budget-neutral	Highest-welfare
Earnings	42.23	42.23	42.23	42.23	42.23
Taxes and loan repayments	10.64	10.53	10.72	10.97	10.91
Effective marginal tax rate (%)	30.44	30.61	30.38	29.63	30.03
Change in labor supply (%)	–	−0.06	0.02	0.28	0.13
Change tax and loan revenue	–	−0.01	−0.00	0.01	−0.00

This table displays the effect of labor supply responses to changes in IDR rules. Row 1 reports the present value of lifetime earnings. Row 2 reports the present value of taxes and loan payments under alternative repayment rules, holding labor supply fixed. Row 3 shows the income-weighted lifetime effective marginal tax rate, computed numerically by increasing lifetime earnings. Row 4 presents the change in lifetime earnings induced by changes in marginal tax rates, weighted by income. Row 5 reports the resulting change in government tax and loan revenues. Rows 4 and 5 can differ in sign because of covariance terms in the joint distribution of labor supply responses and marginal tax rates. All variables, except in row 3, are in dollar units, and are normalized by the national wage index ($63,795 in 2022).

to the overall fiscal cost of IDR, these variations are small. By contrast, de Silva (2025) finds that labor supply responses represent a large fraction of the overall cost of switching between different IDR calibrations in Australia. This difference can likely be explained by the fact that the average student loan balance in Australia is around $7250, whereas in the United States the average borrower owes about $51,400 at graduation. As a result, changes in IDR rules have much larger direct budgetary consequences for the government in the U.S. than in Australia.

9.3. Other costs of indebtedness

Another limitation of our model is that it overlooks the potential indirect cost of having large student debt balances, such as the inability 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 or RAP, preventing balances from ballooning.

10. Concluding remarks

This paper studies a hitherto under-explored channel accounting for much of the dramatic rise in student debt: the slowdown in repayments. Using administrative data, we find that payment deferral accounts for almost half of the increase in balances between 2010 and 2020. We show that this payment slowdown increases borrowers' welfare by providing insurance and liquidity but also because they transfer resources from taxpayers to borrowers. We find that welfare gains from insurance and consumption smoothing can be increased while substantially reducing costs to taxpayers. These gains can be achieved by increasing plan maturities, while at the same time making plans more progressive by raising repayment thresholds.

CRediT authorship contribution statement

Sylvain Catherine: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Mehran Ebrahimian:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Constantine Yannelis:** Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

Authors have nothing to disclose.

Acknowledgments

1. The author has not received no direct financial support for this project.
2. The author has not received any interested party.

References

Baily, M.N., 1978. Some aspects of optimal unemployment insurance. J. Public Econ. 10 (3), 379–402. http://dx.doi.org/10.1016/0047-2727(78)90053-1.

Ballard, C.L., Shoven, J.B., Whalley, J., 1985. General equilibrium computations of the marginal welfare costs of taxes in the United States. Am. Econ. Rev. 75 (1), 128–138, URL: https://www.jstor.org/stable/1812708.

Beshears, J., Choi, J.J., Clayton, C., Harris, C., Laibson, D., Madrian, B.C., 2025. Optimal illiquidity. J. Finance Econ. 165, 103996. http://dx.doi.org/10.1016/j.jfineco.2025.103996.

Black, S.E., Denning, J.T., Dettling, L.J., Goodman, S., Turner, L.J., 2023. Taking it to the limit: Effects of increased student loan availability on attainment, earnings, and financial well-being. Am. Econ. Rev. 113 (12), 3357–3400. http://dx.doi.org/10.1257/aer.20210926.

Boutros, M., Clara, N., Gomes, F., 2024. Borrow now, pay even later: A quantitative analysis of student debt payment plans. J. Finance Econ. 159, 103898. http://dx.doi.org/10.1016/j.jfineco.2024.103898.

Britton, J., Gruber, J., 2020. Do income contingent student loans reduce labor supply? Econ. Educ. Rev. 79, 102061. http://dx.doi.org/10.1016/j.econedurev.2020.102061.

Browning, E.K., 1976. The marginal cost of public funds. J. Political Econ. 84 (2), 283–298. http://dx.doi.org/10.1086/260432.

Catherine, S., Ebrahimian, M., Fereydounian, M., Sraer, D., Thesmar, D., 2022. Robustness Checks in Structural Analysis. Working Paper, National Bureau of Economic Research, http://dx.doi.org/10.3386/w30443.

Catherine, S., Miller, M., Sarin, N., 2020. Relaxing household liquidity constraints through social security. J. Public Econ. 189, 104243. http://dx.doi.org/10.1016/j.jpubeco.2020.104243.

Catherine, S., Yannelis, C., 2023. The distributional effects of student loan forgiveness. J. Finance Econ. 147 (2), 297–316. http://dx.doi.org/10.1016/j.jfineco.2022.10.003.

Chetty, R., 2006. A general formula for the optimal level of social insurance. J. Public Econ. 90 (10), 1879–1901. http://dx.doi.org/10.1016/j.jpubeco.2006.01.004.

Dahlby, B., 2008. The Marginal Cost of Public Funds: Theory and Applications. MIT Press, http://dx.doi.org/10.7551/mitpress/9780262042505.001.0001.

Dávila, E., Schaab, A., 2025. Welfare assessments with heterogeneous individuals. J. Political Econ. 133 (9), 2918–2961. http://dx.doi.org/10.1086/736211.

de Silva, T., 2025. Insurance versus moral hazard in income-contingent student loan repayment. Q. J. Econ. 140 (4), 2851–2905. http://dx.doi.org/10.1093/qje/qjaf036.

Di Maggio, M., Kalda, A., Yao, V., 2019. Second Chance: Life Without Student Debt. Working Paper, National Bureau of Economic Research, http://dx.doi.org/10.3386/w25810.

Dinerstein, M., Yannelis, C., Chen, C.-T., 2024. Debt moratoria: Evidence from student loan forbearance. Am. Econ. Rev.: Insights 6 (2), 196–213. http://dx.doi.org/10.1257/aeri.20230032.

Eaton, C., Howell, S.T., Yannelis, C., 2020. When investor incentives and consumer interests diverge: Private equity in higher education. Rev. Finance Stud. 33 (9), 4024–4060. http://dx.doi.org/10.1093/rfs/hhz129.

Ebrahimian, M., 2020. Student Loans and Social Mobility. Working Paper, Stockholm School of Economics, http://dx.doi.org/10.2139/ssrn.3680159.

Findeisen, S., Sachs, D., 2016. Education and optimal dynamic taxation: The role of income-contingent student loans. J. Public Econ. 138, 1–21. http://dx.doi.org/10.1016/j.jpubeco.2016.03.009.

Finkelstein, A., Hendren, N., 2020. Welfare analysis meets causal inference. J. Econ. Perspect. 34 (4), 146–167. http://dx.doi.org/10.1257/jep.34.4.146.

Finkelstein, A., Hendren, N., Luttmer, E.F.P., 2019. The value of medicaid: Interpreting results from the Oregon health insurance experiment. J. Political Econ. 127 (6), 2836–2874. http://dx.doi.org/10.1086/702238.

Fu, C., Li, X., Zafar, B., 2025. Moral Hazard and the Sustainability of Income-Driven Repayment Plans. Working Paper, National Bureau of Economic Research, http://dx.doi.org/10.3386/w33411.

Gomes, F., Haliassos, M., Ramadorai, T., 2021. Household finance. J. Econ. Lit 59 (3), 919–1000. http://dx.doi.org/10.1257/jel.20201461.

Goodman, S., Isen, A., Yannelis, C., 2021. A day late and a dollar short: Liquidity and household formation among student borrowers. J. Finance Econ. 142 (3), 1301–1323. http://dx.doi.org/10.1016/j.jfineco.2021.04.033.

Gourinchas, P.-O., Parker, J.A., 2002. Consumption risk over the life-cycle. Econometrica 70 (1), 47–89. http://dx.doi.org/10.1111/1468-0262.00269.

Guvenen, F., Karahan, F., Ozkan, S., Song, J., 2021. What do data on millions of U.S. workers reveal about lifecycle earnings dynamics? Econometrica 89 (5), 2303–2339. http://dx.doi.org/10.3982/ECTA14603.

Hamdi, N., Kalda, A., Sovich, D., 2025. The labor market consequences of student loan forbearance. Manag. Sci. 0 (0), doi:doi/10.1287/mnsc.2023.04047.

Heckman, J.J., Moon, S.H., Pinto, R., Savelyev, P.A., Yavitz, A., 2010. The rate of return to the HighScope perry preschool program. J. Public Econ. 94 (1–2), 114–128. http://dx.doi.org/10.1016/j.jpubeco.2009.11.001.

Hendren, N., Sprung-Keyser, B., 2020. A unified welfare analysis of government policies. Q. J. Econ. 135 (3), 1209–1318. http://dx.doi.org/10.1093/qje/qjaa006.

Herbst, D., 2018. Liquidity and Insurance in Student-Loan Contracts: The Costs and Benefits of Income-Driven Repayment. Unpublished Mimeo, University of Arizona.

Hopenhayn, H.A., Nicolini, J.P., 1997. Optimal unemployment insurance. J. Political Econ. 105 (2), 412–438. http://dx.doi.org/10.1086/262078.

Keys, B.J., Mahoney, N., Yang, H., 2023. What determines consumer financial distress? Place-and person-based factors. Rev. Finance Stud. 36 (1), 42–69. http://dx.doi.org/10.3386/w26808.

Kleven, H.J., Kreiner, C.T., 2006. The marginal cost of public funds: Hours of work versus labor force participation. J. Public Econ. 90 (10–11), 1955–1973. http://dx.doi.org/10.1016/j.jpubeco.2006.03.006.

Laibson, D.I., Repetto, A., Tobacman, J., 1998. Self-control and saving for retirement. In: Brook. Pap. Econ. Act.. Vol. 29, (1), pp. 91–196. http://dx.doi.org/10.2307/2534671.

Lochner, L., Monge-Naranjo, A., 2011. The nature of credit constraints and human capital. Am. Econ. Rev. 101 (6), 2487–2529. http://dx.doi.org/10.1257/aer.101.6.2487.

Lochner, L., Monge-Naranjo, A., 2016. Student loans and repayment: Theory, evidence and policy. In: Handb. Econ. Educ.. Vol. 5, pp. 397–478. http://dx.doi.org/10.3386/w20849.

Looney, A., Yannelis, C., 2015. A crisis in student loans?: How changes in the characteristics of borrowers and in the institutions they attended contributed to rising loan defaults. In: Brook. Pap. Econ. Act.. Vol. 2, Johns Hopkins University Press, pp. 1–89. http://dx.doi.org/10.1353/eca.2015.0003.

Lucca, D.O., Nadauld, T., Shen, K., 2019. Credit supply and the rise in college tuition: Evidence from the expansion in federal student aid programs. Rev. Finance Stud. 32 (2), 423–466. http://dx.doi.org/10.2139/ssrn.2646076.

Manso, G., Rivera, A., Wang, H.G., Xia, H., 2025. Student Loans and Labor Supply Incentives. Working Paper, University of California, Berkeley.

Merton, R.C., 1971. Optimum consumption and portfolio rules in a continuous-time model. J. Econ. Theory 3 (4), 373–413. http://dx.doi.org/10.1016/0022-0531(71)90038-x.

Monarrez, T., Turner, L.J., 2024. The Effect of Student Loan Payment Burdens on Borrower Outcomes. Working Paper, Federal Reserve Bank of Philadelphia, http://dx.doi.org/10.21799/frbp.wp.2024.08.

Mueller, H., Yannelis, C., 2019. The rise in student loan defaults. J. Finance Econ. 131 (1), 1–19. http://dx.doi.org/10.1016/j.jfineco.2018.07.013.

Mueller, H., Yannelis, C., 2022. Increasing enrollment in income-driven student loan repayment plans: Evidence from the navient field experiment. J. Finance 77 (1), 367–402. http://dx.doi.org/10.1111/jofi.13088.

Olken, B.A., 2007. Monitoring corruption: Evidence from a field experiment in Indonesia. J. Political Econ. 115 (2), 200–249. http://dx.doi.org/10.1086/517935.

Poterba, J.M., 1996. Government intervention in the markets for education and health care: how and why? In: Individual and Social Responsibility: Child Care, Education, Medical Care, and Long-Term Care in America. University of Chicago Press, pp. 277–308. http://dx.doi.org/10.3386/w4916.

Shimer, R., Werning, I., 2007. Reservation wages and unemployment insurance. Q. J. Econ. 122 (3), 1145–1185. http://dx.doi.org/10.1162/qjec.122.3.1145.

Stantcheva, S., 2017. Optimal taxation and human capital policies over the life cycle. J. Political Econ. 125 (6), 1931–1990. http://dx.doi.org/10.1086/694291.