

NBER WORKING PAPER SERIES

STUDENT LOAN FORGIVENESS

Michael Dinerstein
Samuel Earnest
Dmitri K. Koustas
Constantine Yannelis

Working Paper 33462
<http://www.nber.org/papers/w33462>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2025

We are grateful to Michael Boutros, Sylvain Catherine, Tim De Silva, Caroline Hoxby, Erik Hurst, Adam Looney, Lesley Turner and Michael Weber for helpful discussions and comments as well as seminar participants at NBER Public Economics, NBER Education, Cambridge, the University of Chicago, MIT Sloan, the University of Illinois Gies School of Business, Vanderbilt Owen School of Management, the Tulane Freedman School of Business, Hong Kong University of Science and Technology, the Federal Reserve Bank of Philadelphia, the Federal Reserve Bank of Cleveland, Hong Kong University, the University of Bath School of Management, the Chinese University of Hong Kong, University of Houston Bauer College of Business, CUHK Shenzhen, ABN AMRO Financial Transaction Data conference and City University of Hong Kong. Koustas gratefully acknowledges support from the Becker Friedman Institute (BFI) and the Peter G. Peterson Institute Pandemic Response Policy Research Fund Award. 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. 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. We also rely on consumer credit data and employment data from a second large credit bureau. Due to legal limitations, including data sharing agreements, we cannot disclose the name of the specific credit bureau. The consumer credit information provided was anonymized and delivered in a format for analytical purposes only. This credit bureau has the right to review the research for legal compliance, but did not comment on the use of the data, analysis, and conclusions drawn from said data. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Michael Dinerstein, Samuel Earnest, Dmitri K. Koustas, and Constantine Yannelis. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Student Loan Forgiveness

Michael Dinerstein, Samuel Earnest, Dmitri K. Koustas, and Constantine Yannelis

NBER Working Paper No. 33462

February 2025

JEL No. G51, I22, I28, J29

ABSTRACT

Student loan forgiveness has been proposed as a means to alleviate soaring student loan burdens. This paper uses administrative credit bureau data to study the distributional, consumption, borrowing, and employment effects of the largest event of student loan forgiveness in history. Beginning in March 2021, the United States federal government ordered \$132 billion in student loans cancelled, or 7.8% of the total \$1.7 trillion in outstanding student debt. We estimate that forgiven borrowers' predicted monthly earnings were \$115 higher than borrowers who did not receive forgiveness and \$193 more than the general population. We find that student loan forgiveness led to increases in mortgage, auto, and credit card debt by 9 cents for every dollar forgiven. Borrowers' monthly earnings and employment fell, at increasing rates for each month post forgiveness. The implied Marginal Propensities for Consumption (MPC) and Earnings (MPE) are 0.27 and -0.49, respectively.

Michael Dinerstein
Department of Economics
Duke University
213 Social Sciences Bldg.
Box 90097
Durham, NC 27708
and NBER
michael.dinerstein@duke.edu

Samuel Earnest
Massachusetts Institute of Technology
searnest@mit.edu

Dmitri K. Koustas
Harris Public Policy
University of Chicago
1307 E. 60th St. Room 3055
Chicago, IL 60605
dkoustas@uchicago.edu

Constantine Yannelis
Faculty of Economics
University of Cambridge
Austin Robinson Building
Cambridge, CB3 9DD
United Kingdom
cny21@cam.ac.uk

A data appendix is available at <http://www.nber.org/data-appendix/w33462>

Student Loan Forgiveness*

Michael Dinerstein¹ Samuel Earnest²

Dmitri Koustas³ Constantine Yannelis⁴

February 2, 2025

Abstract

Student loan forgiveness has been proposed as a means to alleviate soaring student loan burdens. This paper uses administrative credit bureau data to study the distributional, consumption, borrowing, and employment effects of the largest event of student loan forgiveness in history. Beginning in March 2021, the United States federal government ordered \$132 billion in student loans cancelled, or 7.8% of the total \$1.7 trillion in outstanding student debt. We estimate that forgiven borrowers' predicted monthly earnings were \$115 higher than borrowers who did not receive forgiveness and \$193 more than the general population. We find that student loan forgiveness led to increases in mortgage, auto, and credit card debt by 9 cents for every dollar forgiven. Borrowers' monthly earnings and employment fell, at increasing rates for each month post forgiveness. The implied Marginal Propensities for Consumption (MPC) and Earnings (MPE) are 0.27 and -0.49, respectively.

*First draft: July 24, 2024. We are grateful to Michael Boutros, Sylvain Catherine, Tim De Silva, Caroline Hoxby, Erik Hurst, Adam Looney, Lesley Turner and Michael Weber for helpful discussions and comments as well as seminar participants at NBER Public Economics, NBER Education, Cambridge, the University of Chicago, MIT Sloan, the University of Illinois Gies School of Business, Vanderbilt Owen School of Management, the Tulane Freedman School of Business, Hong Kong University of Science and Technology, the Federal Reserve Bank of Philadelphia, the Federal Reserve Bank of Cleveland, Hong Kong University, the University of Bath School of Management, the Chinese University of Hong Kong, University of Houston Bauer College of Business, CUHK Shenzhen, ABN AMRO Financial Transaction Data conference and City University of Hong Kong. Koustas gratefully acknowledges support from the Becker Friedman Institute (BFI) and the Peter G. Peterson Institute Pandemic Response Policy Research Fund Award. 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. 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. We also rely on consumer credit data and employment data from a second large credit bureau. Due to legal limitations, including data sharing agreements, we cannot disclose the name of the specific credit bureau. The consumer credit information provided was anonymized and delivered in a format for analytical purposes only. This credit bureau has the right to review the research for legal compliance, but did not comment on the use of the data, analysis, and conclusions drawn from said data.

¹Duke University, National Bureau of Economic Research, CESifo

² MIT, Sloan School of Management

³University of Chicago, Harris School of Public Policy

⁴ University of Cambridge, National Bureau of Economic Research

1 Introduction

Since 2010, outstanding student debt and debt per borrower in the United States (US) have increased by 115% and 73%, respectively. To address the current debt burden, policymakers have called for broad-based student loan forgiveness, with goals ranging from redistributing toward low earners to providing economic stimulus. Despite this policy momentum, neither policy-makers nor academics have reached consensus on whether forgiveness will increase or decrease outcomes like spending¹ and earnings.²

In this paper, we analyze the largest event of student loan discharge in history. Beginning in March 2021, the US federal government ordered \$132 billion in student loans cancelled, or 7.8% of the total \$1.7 trillion in outstanding student debt. We assess how forgiveness is targeted and how it affects forgiven borrowers' consumption, debt, and earnings. We estimate that forgiveness targets borrowers with \$115 and \$193 higher predicted monthly earnings than non-forgiven borrowers and the general population of individuals with a credit history, respectively. We estimate that student loan forgiveness has large effects on household consumption and debt. For each dollar of student loan forgiveness, other debt increases by 8.8 cents. This increase is primarily driven by mortgage debt. We find little effect on loan delinquencies, other than for student loans. For labor market outcomes, we estimate that forgiveness leads to decreased earnings and increased switching across industries, with effects increasing over

¹For example, Senator Elizabeth Warren [called student debt relief](#) the “single most effective executive action available to provide a massive stimulus to our economy.” On the other hand, former Council of Economic Advisers Chair Jason Furman [stated that](#) “student loan debt forgiveness likely has a multiplier close to zero” and “even be net negative”.

²For example, eliminating debt overhang may increase earnings (Di Maggio et al., 2019; Donaldson et al., 2019), while lower repayment needs may lead credit-constrained borrowers to decrease earnings (Hampole, 2022). Removing debt for borrowers on income-driven plans may also increase labor supply (de Silva, 2023; Boutros et al., 2022). See Yannelis and Tracey (2022) for a discussion of the literature on student debt and household outcomes.

time. Adding up the effects, we estimate a marginal propensity for consumption of 0.27 and a marginal propensity for earnings of -0.49.

We start in Section 2 by describing the policy environment since March 2021. Forgiveness has largely been enacted through various administrative reforms and expansions of existing programs. In particular, administrative procedures for discharge under Income-Driven Repayment (IDR), Public Sector Loan Forgiveness (PSLF), and Borrower Defense were relaxed, which allowed about 3 million borrowers to benefit from cancellation. We characterize forgiveness as primarily an unexpected moderately-sized wealth shock. A subset of borrowers, however, also receive increased liquidity and potential changes in the tax rate of labor market earnings.³

We identify which borrowers had their student debt balances forgiven with administrative credit panel data, which we introduce in Section 3. We use comprehensive national administrative data from TransUnion, one of the largest credit bureaus, complemented by employment records obtained from a second large credit bureau. The TransUnion panel data comprise a ten percent sample of all individuals who have a credit history in the United States, and the employment data cover around 1/3 of the U.S. workforce. We estimate that 6% of borrowers in our sample have received forgiveness since March 2021, with an average of \$32,000 discharged. We validate our inference of forgiveness by comparing it to public statistics released by the Department of Education. The pattern of inferred forgiveness tracks reported forgiveness both over time and across US states.

We organize the empirical analysis around the policy arguments for forgiveness. The first set of arguments involves the policy's targeting. Many advocates argue that broad-based debt forgiveness might close wealth gaps if students from less wealthy backgrounds are more likely to take out debt or pay it back more slowly. Forgiveness might also compensate low-earning borrowers for low returns on their educational investments.⁴ In Section 4, we construct bor-

³While most debt forgiveness shocks would also imply a liquidity change, during much of our sample many loans were under a repayment pause (Dinerstein et al., 2024).

⁴Low returns could reflect students' choices like dropping out or schools' actions like false advertising. In other countries with broader public financing of higher education, progressive taxation might tie the debt burden more closely to returns.

rower predicted earnings based on lifecycle characteristics, gender, education, and employment industry. We then compare predicted monthly earnings between forgiven borrowers and several comparison sets. Relative to student loan borrowers who did not receive forgiveness, forgiven borrowers have \$703 higher average monthly predicted income. Much of this difference reflects that forgiven borrowers are older and at the higher-earning part of their lifecycle, but when we control for age and when the borrower took out loans, a difference of \$115 remains. When comparing forgiven borrowers to the general population of individuals with a credit history, we find more regressivity. Forgiven borrowers have \$504 higher monthly predicted earnings on average, which drops to \$193 when adding lifecycle controls. Unsurprisingly, the gap is even larger when comparing forgiven borrowers to individuals who did not attend college. Thus, the forgiven borrowers are positively selected on income.

The implemented forgiveness reflects discharges within the current administrative structure. We compare the targeting of this forgiveness to the targeting that President Joe Biden initially proposed but that was blocked by the courts. We estimate that Biden's proposal would have produced more progressive forgiveness, with the average forgiven dollar accruing to a borrower with \$4,868 in predicted monthly income compared to the \$5,562 of the actual policy.

A separate set of arguments for forgiveness emphasizes causal effects: borrowers might react in a way that provides macroeconomic stimulus or less distorted labor supply. If increased wealth leads to increased consumption, especially through durables with upfront payments, then broad-based forgiveness could have large stimulus effects and potentially close generational gaps in outcomes like homeownership. If borrowers had been choosing jobs based in part on how they translated into student debt burdens, then removing the burden might lead to a less constrained choice. We describe these possible channels in Section 5 and lay out our empirical strategy for estimating causal effects.⁵

We develop two complementary empirical strategies to identify the causal effects of loan

⁵A third argument is procedural. Some forgiven borrowers may have deserved forgiveness based on their loan contracts but for administrative reasons had not yet received it.

forgiveness on a borrower's outcomes. In our main analysis, we estimate a dynamic difference-in-difference model based on the timing of forgiveness. We compare how a borrower's outcomes change over time upon receiving forgiveness, relative to borrowers who have not yet received forgiveness. We thus absorb all time-invariant individual specific factors that could be correlated with who received forgiveness. There are two main concerns with this approach. First, we may incorrectly measure the treatment of loan forgiveness if some individuals see their loans discharged through processes other than loan forgiveness. Second, loan forgiveness events may be related to a borrower's potential outcomes, through policy targeting or through borrowers taking endogenous actions to receive forgiveness. To address these concerns, we offer a complementary instrumental variables approach. We use the fact that loan servicers are quasi-randomly assigned (Cornaggia and Xia, 2024). Different loan servicers were more or less likely to process loan forgiveness, due to various administrative procedures or information campaigns.⁶ While we do not observe individual loan servicers in the data, we observe anonymized servicer keys and hence can compute masked individual servicers' propensity to process forgiveness.

We find substantial effects on household consumption and debt and present the results in Section 6. Borrowers experiencing forgiveness increase mortgage borrowing by \$2,300, auto loan borrowing by \$230, and credit card borrowing by \$220 over the six months following forgiveness. This large mortgage effect is driven primarily by increases in the probability borrowers had any mortgage rather than balance changes for pre-existing mortgages. We also see an extensive margin effect on whether borrowers have any auto loan, which is counteracted by *reduced* balances on pre-existing debt. For credit card debt, we see increases on both the extensive and intensive margins. The increase in whether borrowers have a loan of any type suggests a fairly high change in durable ownership, with a large change in debt and somewhat smaller change in current consumption.

We find little to no effect on non-student loan delinquencies, although student loan delin-

⁶For example, [Mohela](#) was criticized for failing to process forgiveness applications.

quencies fall mechanically. Forgiven borrowers are also no more likely to move to higher earnings metropolitan areas and credit scores hardly move. Consumption effects are larger for young workers. Reassuringly, our two empirical strategies yield similar results.

We investigate the effect on labor market outcomes in Section 7. We estimate drops in monthly earnings of \$44 (or 2.3%) pooled over the first six months post-forgiveness, with the decrease exceeding \$75 in the sixth month. The decreased earnings suggests that borrowers may have been in higher-paying jobs or provided increased labor supply in part to pay back debt. We estimate a small drop in whether the borrowers are in the employment data, though when we condition on a balanced sample of individuals always in the data, our intensive margin estimates are very similar to the unconditional estimates. While we do not see an increase in the rate of job switching, we see that the nature of job switching changes: more switching across industries and out of public service. This latter effect is perhaps related to forgiveness under PSLE, as forgiven borrowers no longer need to work in public service to qualify for future forgiveness, and indeed we estimate larger earnings drops for individuals initially employed in public service jobs. Finally, among hourly workers, we estimate a drop in hours worked, where the total earnings drop comes half from an hours reduction and half from a wage reduction. Labor market effects are largest for younger workers with lower earnings, hourly workers, and public service workers. When we focus on previously defaulted borrowers, we find earnings increases, consistent with Di Maggio et al. (2019).

In Section 8, we translate our results into Marginal Propensities for Expenditure (MPX), Consumption (MPC), and Earnings (MPE), key structural parameters in macroeconomic and lifecycle models. Normalized by the value of forgiveness annuitized over the remaining lifetime, the MPC and MPE are 0.27, and -0.49, respectively. The MPX is 27 times the monthly increase in permanent income, which mostly reflects the up-front increase in durables balances. These estimates imply mixed stimulus effects. On the one hand, the consumption response is small and the reduction in labor earnings implies less economic activity. On the other hand, the expenditure associated with the consumption is concentrated right after forgiveness.

The estimates further reveal how consumption and earnings responses vary with the type of income or wealth shock. Compared to mortgage balance reductions (Agarwal et al., 2017) or stock appreciation (Chodorow-Reich et al., 2021; Di Maggio et al., 2020), student debt forgiveness has a much larger MPC. Compared to transitory transfers, our MPC of 0.27 is similar to the baseline estimate from Boehm et al. (2025) but much smaller than MPCs for expiring transfers. The most similar shock is perhaps the unearned income from winning the lottery (Imbens et al., 2001), which generates a very similar MPE as student debt forgiveness.

Finally, we combine our estimates to test whether borrowers act consistently with the permanent income hypothesis, which predicts that they would perfectly annuitize the value of the forgiven debt. We combine our MPC and MPE estimates, along with adjustments for down payments, labor income taxes, and novel data on the change in borrowers' expectations of forgiveness, and estimate that borrowers transfer 0.64 of the annuitized shock into the current period. When we focus, however, on the subset of borrowers who were subject to a payment pause and thus forgiveness is primarily a wealth shock, we estimate that borrowers transfer 0.82 of the annuitized shock into the current period. Behavior so consistent with the permanent income hypothesis may reflect our focus on somewhat older borrowers who face fewer constraints (e.g., credit or collateral constraints) than in other settings and perhaps that the repayment structure of student debt meant that borrowers might have already thought of their debt in annuitized terms.

The paper makes three main contributions. First, we provide a policy evaluation for the largest forgiveness of student debt in US history. Compared to other forgiveness policies, this policy provides a larger sample size for heterogeneity analysis and occurs during a period of macroeconomic stability rather than during a recession or major crisis. Importantly, our results suggest that forgiveness disproportionately went to higher-income borrowers. Second, we provide causal evidence on how people respond to student debt forgiveness. The nature of the consumption and earnings response points toward specific economic channels that are not necessarily present with other forms of debt. We especially highlight that student debt

has large income effects on labor supply and potentially locks borrowers into jobs based on repayment benefits. Third, we estimate an MPC and MPE that allow us to benchmark the size of borrowers' responses to estimates from the literature from other types of shocks and to test whether borrowers act consistently with the permanent income hypothesis.

This paper primarily joins a growing literature on student loans. Most work in this area studies the effects of repayment plan parameters. For example, several studies explore the effects of repayment plans (Mueller and Yannelis, 2019, 2022; Herbst and Hendren, 2021), maturity extension (Boutros, Clara and Gomes, 2022), loan limits (Black et al., 2020; Goodman et al., 2021), or payment pauses (Dinerstein et al., 2024; Hamdi et al., 2024; Chava et al., 2023). Hampole (2022) studies the effects of “no loan policies” that change whether students originate loans. In an important study, de Silva (2023) assesses the moral hazard effects of student loan forgiveness. Amromin and Eberly (2016), Lochner and Monge-Naranjo (2016), Yannelis and Tracey (2022), and Looney and Yannelis (2024) provide recent reviews of the literature.

A small number of studies focus directly on loan forgiveness. This paper is most closely related to Di Maggio et al. (2019), who study loan discharge that occurred after a private loan servicer lost title chains on defaulted loans. Our study differs conceptually from Di Maggio et al. (2019) for two reasons. First, their study focuses on borrowers already in default. Wealth and liquidity shocks may have very different effects (Ganong and Noel, 2020). Second, we focus on federal student loan borrowers. The vast majority of student loan borrowers in the United States take government loans, and hence private borrowers are likely a different sample. Jacob et al. (2024) study loan forgiveness for teachers. Our work is also closely related to Catherine and Yannelis (2022), who study the distributional effects of student loan forgiveness. Our analysis of distributional effects compares the targeting of actually implemented policy to a variety of policy benchmarks.

This paper also joins a literature in household finance, which studies the effect of debt relief programs. Much of this literature focuses on debt relief during the aftermath of the 2008

financial crisis, particularly analyzing mortgage relief policies. For example, Agarwal et al. (2017) study the Home Affordable Modification Program (HAMP), which provided relief to underwater mortgage borrowers. Many of these studies explore similar outcomes, and how debt and consumption respond to income and wealth shocks (Agarwal, Liu and Souleles, 2007; Mian and Sufi, 2009; Agarwal and Qian, 2014; Baker, 2018; Aydin, 2022). A smaller literature also focuses on mortgage relief (Cherry et al., 2021) and stimulus checks (Baker et al., 2020; Coibion et al., 2020).

More generally, our paper is related to a broader research in macroeconomics and labor economics estimating consumption and employment responses to income changes. As surveyed in Jappelli and Pistaferri (2010), there is a large literature focused on consumption responses to income changes. Other papers using natural experiments tend to focus on small or transitory shocks, such as the receipt of stimulus checks. Moreover, most of the surveyed literature on the consumption-side assumes that labor supply does not respond to the income shock. In contrast, student loan forgiveness can be viewed as a wealth shock, and may have effects on labor supply. While a number of studies estimate the consumption and employment responses to shocks in components of household wealth, these responses can be difficult to interpret without knowing whether consumers view the shocks as transitory or permanent. In the case of student loan forgiveness, we know the precise value of the shock to lifetime wealth, allowing us to follow the approach taken in the literature on lottery windfalls annuitizing the forgiveness over the remaining lifetime (Imbens et al., 2001; Cesarini et al., 2017; Golosov et al., 2023). Moreover, we are able to study individual consumption and labor supply responses to the shock in a consistent sample, which few papers have been able to do due to data availability. The policy we study is also unique for two reasons. First, unlike many government policies that occurred in the aftermath of the financial crisis or the pandemic, the student debt relief that we explore occurred during a boom. Second, many of the borrowers who received relief were relatively affluent, and most debt relief policies target lower-income borrowers.

2 Policy Environment

Student loans play an important role in the funding of higher education in the United States, with around half of undergraduate and graduate students in recent years receiving at least some loans as part of aid packages (National Center for Education Statistics, 2023). As of 2020, nearly 43 million individuals held approximately \$1.6 trillion in outstanding student loans, which represents the largest source of household debt after mortgages (U.S. Department of Education, 2024b; Board of Governors of the Federal Reserve System, 2024).

The federal government directly holds or guarantees over 92% of outstanding student loan debt.⁷ Federal student loans may come with several benefits, including fixed interest rates, income-driven repayment plans, and potential loan forgiveness options. The remaining 8% of student loan debt is held in private student loans, which also includes federal student loans that have since been refinanced with a private lender. Importantly, private student loans are not eligible for any federal student loan relief. Whether federal or private, student loan debt is difficult to discharge in bankruptcy. Delinquency of 270 days or more will typically result in wage garnishment up to 15 percent of disposable income. In the years prior to the pandemic, over 10 percent of student loan borrowers defaulted within 3 years (U.S. Department of Education, 2019).

As a result of rising student debt and the perceived burdens, calls for student loan forgiveness have increased in recent years. While President Biden campaigned in part on student loan forgiveness, the 2020 pandemic was a watershed moment for student loan payment relief. Most immediately, student loan required payments and interest were paused as part of bipartisan pandemic relief in early 2020, resulting in automatic rehabilitation of federal loans in default.⁸ In the end, the payment moratorium lasted 3.5 years until October 2023.

⁷Guaranteed loans were issued under the Federal Family Education Loan Program (FFEL). Under FFEL, private lenders provided the loans, while the federal government guaranteed the loans against default and provided subsidies to lenders. The FFEL program was discontinued in 2010 and replaced with the Direct Loan Program, while extant FFEL loans continued to be serviced by loan servicers. Yet another type of loan is the Perkins loan, issued by participating schools, with funds provided by the federal government and the participating schools. The Perkins loan program was discontinued in 2017.

⁸FFEL loans did not qualify for the repayment pause (Dinerstein et al., 2024).

In August 2022, Biden proposed broad-based cancellation of \$10,000-\$20,000 in student debt. Initially justified under legal authority granted to the Department of Education during national emergencies, the broad-based student loan forgiveness policy quickly encountered political opposition and legal challenges. Ultimately, a Supreme Court decision in June 2023 declared this plan unconstitutional, and no broad-based forgiveness was issued under this plan.

Over this same period, the Biden administration implemented several significant changes to existing student loan policies that can lead to forgiveness, by expanding eligibility for existing programs, retroactively counting payments, and simplifying processes. These changes have since led to forgiveness for around 3 million borrowers (U.S. Department of Education, 2024a). In most cases, borrowers receiving forgiveness under these programs receive forgiveness for the full amount of their outstanding student loan balances, which in some cases total in the hundreds of thousands of dollars. The timeline of key announcements is shown in Appendix Table B.1, along with counts of affected borrowers. We provide further details on each program below.

2.1 Income-Driven Repayment

An Income-Driven Repayment (IDR) plan for student loans is a repayment option that ties monthly student loan payments to income and family size. These plans are designed to make loan repayment more manageable for borrowers who may have difficulty affording standard 10-year repayment plans by extending the repayment term to 20 or 25 years depending on year of loan issuance and/or specific IDR plan, after which any remaining balance may be forgiven. Historically, forgiveness has required the borrower to be enrolled in an IDR plan over the whole time period and for qualifying payments to cover the full scheduled payment. In part due to the length of time necessary to be enrolled in an IDR plan and other verification requirements, just 50 borrowers had received forgiveness through IDR plans by 2020 (U.S. Department of Education, 2024a).

2.1.1 Payment Count Adjustment

In July 2023 the Department of Education announced payment count adjustments to count previously ineligible monthly payments towards IDR forgiveness. Any borrowers with loans making the required number of eligible repayments received automatic forgiveness—even if the payments were not made while on an IDR plan. Additionally, 12 or more consecutive months of forbearance would also qualify. As of 2024, the Department of Education reported that these IDR-payment-count adjustments resulted in approximately \$46 billion of discharged student debt across approximately 930,000 borrowers, or around \$49,000 per borrower (U.S. Department of Education, 2024a). Payment count adjustments did not require borrowers to take any actions and likely came as a surprise for many borrowers who thought they had not accrued any qualifying payments.

2.1.2 Public Service Loan Forgiveness Program

Introduced in 2007, the public service loan forgiveness program (PSLF) offers student loan forgiveness to government and non-profit employees who made 120 qualifying monthly payments on an IDR plan. Starting in September 2017, the first borrowers became eligible and began applying to have their loans forgiven through the PSLF program. By 2018, just 55 borrowers received loan forgiveness, in part due to program complexity (U.S. Government Accountability Office, 2018). Despite some improvements, by January 2020, only around 3,000 individuals had received forgiveness under PSLF (Federal Student Aid, 2024).

Following an announcement in October 2021, the Department of Education made significant temporary adjustments to the program, including loosening the requirements for what payments could count towards PSLF. Suddenly, payments made on previously ineligible loans, partial payments, and late payments counted toward forgiveness (The White House, 2024). In order to qualify for this program, known as the “PSLF waiver,” individuals needed to submit a PSLF waiver form on or before October 31, 2022; after this date borrowers could no longer qualify for these adjustments, though many subsequently qualified for the aforemen-

tioned payment count adjustment.⁹ As of 2024, \$62.5 billion for nearly 872,000 borrowers, with an average discharge amount of around \$72,000, was forgiven under the PSLF waiver (U.S. Department of Education, 2024a).¹⁰

2.1.3 Saving on a Valuable Education (SAVE) Plan

In August 2023, the Biden administration announced an updated IDR plan, known as the Saving on a Valuable Education (SAVE) Plan. Borrowers making at least 10 years of payments who have originally taken out \$12,000 or less for college could receive forgiveness amounts up to \$12,000. For every \$1,000 borrowed above \$12,000, a borrower would receive forgiveness after an additional year of payments. Full forgiveness would occur after 20 years for undergraduate loans, and 25 years for those with graduate student loans. As of 2024, 153,000 borrowers had already received \$1.2 billion in forgiveness through new SAVE plans (U.S. Department of Education, 2024a).

2.2 Borrower Defense Forgiveness

Borrower defense is a legal ground for discharging direct loans when schools engage in misconduct related to providing federal loans or educational services. For example, borrowers whose schools close before they can complete their education or borrowers whose schools misrepresent the value of their services may qualify for borrower defense forgiveness. The current borrower defense regulation has existed since 1995 but was rarely used until 2015 when Corinthian Colleges, a publicly traded company operating numerous post-secondary institutions, filed for bankruptcy. Since 2015 borrowers from more than 60 institutions have qualified for loan forgiveness under borrower defense. High profile cases include Westwood College (\$1.5 billion), ITT Technical Institute (\$3.9 billion), Ashford University (\$72 million),

⁹Borrowers with FFEL and Perkins loans, which had not historically qualified for PSLF, were required to first consolidate into a Direct Loan, then apply for the waiver.

¹⁰For all IDR plans, including PSLF, each month of the 3.5-year payment pause counted as a qualifying payment toward forgiveness, even though no payments were being made.

and University of Phoenix (\$37 million).

In June 2022, the Department of Education settled the *Sweet v. Cardona* case. Among other elements, the settlement provided forgiveness to borrowers from a long list of schools; for borrowers who applied for borrower defense from a school not on the list, the settlement guaranteed a decision within a certain period of time, or discharge was automatic. In July 2022 the Department of Education announced that they had approved \$5.8 billion of debt discharged for the remaining 560,000 borrowers who had attended Corinthian.

As of December 2023 the Department of Education has reported total borrower defense debt discharges under the Biden Administration to be \$22.5 billion for approximately 1.3 million borrowers, or around \$17,000 per borrower. Forgiven borrowers include a mix of borrowers who actively applied for borrower defense and borrowers who received forgiveness automatically when the Department of Education implemented a group discharge.

2.2.1 Total and Permanent Disability (TPD) Discharge

The total and permanent disability (TPD) discharge program was designed to relieve debts for individuals with permanent disabilities that prevented them from being able to pay off their loans. Under the programs, individuals who provided documentation of their disability would receive debt forgiveness but would be monitored for the following three years. During this period borrowers whose earnings exceed certain thresholds and borrowers who do not report their income will have their loans reinstated. In March 2021 the Department of Education announced significant changes to the TPD debt discharge program, which included removing the income monitoring period. To justify this action, they cited a 2016 report by the Government Accountability Office which found that 98 percent of reinstated disability discharges occurred not because earnings were too high, but because borrowers simply did not submit the requested documentation (U.S. Department of Education, 2023; U.S. Government Accountability Office, 2016). This action resulted in an immediate \$1.3 billion of forgiveness for 41,000 individuals. In August 2021, data matches with the Social Security Administration led to a series of auto-

matic discharges. As of 2024, total TPD discharges have grown to \$14.1 billion of forgiveness across more than 548,300 borrowers, or around \$25,000 per borrower (U.S. Department of Education, 2024a).

To summarize, around 3 million borrowers, approximately 7 percent of those who had outstanding student loans as of 2020, have received forgiveness under one of these adjustments to existing forgiveness programs. Over \$133 billion has been discharged, or an average of \$44,000 per borrower.

2.3 Characterizing the Policy Variation

Relative to other forgiveness policies, this policy variation has several distinguishing features. First, many forgiveness policies are responses to economic crises or recessions. This policy, instead, is occurring at a time of economic stability, especially toward the end of our sample. Second, many forgiveness policies target borrowers with very specific characteristics (e.g., teachers or borrowers in default). This policy, instead, is relatively broad-based.

For many borrowers, the debt forgiveness can be thought of as a moderately-sized wealth shock. While other shocks in components of household wealth, such as housing or the stock market, may have permanent and transitory components, the cancellation of the student debt balance is permanent and of a precise amount, analogous to lottery windfalls. For subsets of borrowers, the policy may entail additional features. First, while most borrowers during our period were subject to the payment pause, those ineligible for the pause would see a reduction in upcoming scheduled payments when their debt is forgiven. These borrowers thus might be subject to a combined wealth and liquidity shock. Second, some borrowers' labor market choices might have interacted with their student debt repayment contracts such that forgiveness removed constraints. If borrowers were in default and had their wages garnished, then forgiveness would lower the tax rate on wages. If borrowers were choosing lower-paying jobs so that IDR payments would be lower and eventual forgiven balances higher, then forgiveness would lower the shadow price on higher wages. Finally, if borrowers were in public service

jobs to qualify for PSLF, then forgiveness would eliminate the extra benefit of being in a public service job.

We expect that for many borrowers, the substantial policy uncertainty would have made it difficult to anticipate forgiveness. The Biden administration’s initial loan forgiveness attempts were immediately paused by courts, and announcements regarding payment count adjustments came a month after the Supreme Court permanently blocked initial broader forgiveness plans. Moreover, most servicers do not regularly report payment counts to borrowers, so there was little scope for borrowers to obtain information on how close they were to discharge. Similarly, changes to PSLF marked significant breaks from past program implementation and were likely not broadly anticipated. In our analysis, we will check for pretrends for signs of anticipation. In addition, Appendix A reexamines a survey of beliefs about forgiveness from Koustas et al. (2024) and finds forgiven borrowers have no baseline differences and no pretrends in their beliefs about forgiveness. Except around the announcement of Biden’s proposed policy, baseline forgiveness expectations are around 25%. In Section 8, we will thus incorporate baseline expectations when characterizing borrower behavior.

3 Data and Identifying Forgiveness

3.1 Data

Our first data source is the Booth TransUnion Consumer Credit Panel. The data are an anonymized 10% panel sample of all TransUnion credit records from 2000 to 2024. All individuals who were initially in the sample in 2000 have their data continuously updated, and each year 10% of new entrants in the TransUnion data are added. Our main sample consists of all borrowers who have an open student loan as of January 2021. We drop duplicate accounts as well as any accounts that are joint, cosigned loans, which are likely to be parent or private loans, or individuals missing birthdates. For computational purposes, we then take a random 25% sample of this list of borrowers, resulting in 992,289 student loan borrowers.

For each borrower we observe the open date, balances, payments due, payments made, and 30 day delinquency status on all trade lines reported to TransUnion on a monthly frequency. We classify student loan borrowing cohorts using the earliest recorded student loan open date. We consider all observations between January 2021 and March 2024, resulting in 38,699,271 unique borrower-month observations. Table 1 presents summary statistics for our main analysis variables, broken down by borrowers who receive forgiveness and those who do not. Borrowers who receive forgiveness tend to be slightly older than the median borrower, consistent with the details of forgiveness policies. Consequently, these borrowers tend to have lower levels of student loan debt, and slightly higher levels of other types of household debt.

We complement our analysis with credit and employment information obtained from a second large credit bureau (“employment records”).¹¹ These data are collected for employment verification and income verification purposes, such as when applying for credit or seeking new employment. We have access to employment records available from 2017 to present.¹² These data provide national coverage of employment history and monthly earnings from over 3 million employers, including government and non-profits, with around 50 million active records per month, which is just over 30 percent of total non-farm employment from the Bureau of Labor Statistics establishment survey. In total, around 370 million individuals have at least one employment record, and 76 percent of individuals with student loans have at least one employment record over the period 2017-2023. We have access to a 10% random sample of records.

Table 2 presents summary statistics for outcomes obtained from this second credit bureau, again broken out by borrowers who receive forgiveness and those who do not. Comparing student loan balances to TransUnion, the balances are reassuringly very similar. Balances for all student loan borrowers average \$38,004 over the period 2021-2023, compared to \$39,539 in TransUnion; among forgiven borrowers, balances average \$20,881, compared to \$23,454

¹¹ Due to legal limitations, including data sharing agreements, we cannot disclose the name of the specific credit bureau from which the consumer credit data was obtained for this analysis.

¹² Similar data from an earlier period have been previously used to study the effects of minimum wages (Gopalan et al., 2021).

in TransUnion.

Around 20% of student loan borrowers in our sample have an employment record in any month; 16 percent have earnings reported.¹³ Forgiven borrowers are also more likely to be employed in public service jobs, which is consistent with the PSLF program contributing to many of the new forgiveness events.¹⁴ Thirty-eight percent of forgiven borrowers are employed at a possible public sector employer, compared with 31 percent of overall borrowers and just 22 percent of individuals without student debt. Conditional on having reported earnings, monthly earnings of student loan borrowers are around \$4,700, or \$56,400 on an annualized basis, compared to around \$5,700 for forgiven student loan borrowers, or \$68,400 on an annualized basis. Median earnings on an annualized basis are around \$50,000 for both groups.

3.2 Identifying Forgiveness

We identify forgiven student loans via a multi-step process. First, we flag student loans which close or go from a non-zero balance to a zero balance in any two consecutive months during our sampling period. Many of these closures and drops are likely borrowers paying off their student loans either on time or via prepayments. To avoid misclassifying prepayments as forgiveness we next filter to borrowers for which the amount paid in the month of closure is less than the discharged balance amount. This means we consider forgiveness only among borrowers whose student loan balances decrease by amounts that cannot be explained by payments.¹⁵

In addition to prepayments, there are other credit bureau reporting practices that could result in misclassifying forgiveness. One of these practices is the way bureaus handle refinancing and loan consolidation. Many borrowers consolidate their loans to make payments simpler, or to take advantage of new payment options. When a borrower refinances, it is common in credit bureau data for a trade line to close and then reappear several months later due to reporting

¹³Not all employers report earnings.

¹⁴Public service is defined as being employed in public administration, health care, and social assistance, or educational services, based on the NAICS code of the employer. This is imperfect since a smaller share of for-profits also operate in these sectors.

¹⁵A vanishingly small number of student loans are discharged annually via bankruptcy.

lags between servicers and credit bureaus (Gibbs et al., 2023). To address this concern, we remove any flagged student loan closures for which 75% or more of the initial balance returns in the months following the debt disappearance. This removes approximately 46% of initially flagged borrowers.

Another reporting practice that could result in misclassifying forgiveness is the handling of deaths in the data. When a borrower dies it is also common in credit bureau data for some of the borrower’s trade lines to abruptly close without explanation even though many of the borrower’s other trade lines continue to be reported on (Gibbs et al., 2023). This occurs because not all servicers may learn of the borrower’s death at the same time. To avoid misidentifying deaths as loan forgiveness we remove any flagged student loan closures that coincide with unexplained drops in mortgage, auto loan, or credit card debt. We flag these coinciding drops as a 50% decline in non-student loan balances with total payments made on these lines in the three months approaching the drop aggregating to less than 25% of the missing balance. This removes approximately 6% of the initially flagged borrowers.

While the data are reported at a monthly frequency, how quickly credit activity is reported to TransUnion varies by lender or servicer (Gibbs et al., 2023). Some servicers may report immediately while others’ reporting lags may mean that TransUnion only records credit-related changes a month after they actually occurred. This potential reporting lag means that when we see a student loan balance decrease in month t , we have some uncertainty over whether the actual balance decrease occurred between months $t - 2$ and $t - 1$ or between months $t - 1$ and t . In either case, we know the decrease has occurred by month t (and not before $t - 1$). In descriptive analysis, we proceed with classifying events as occurring between $t - 1$ and t but discuss in Section 5.2 how we account in our empirical specification for potentially lagged reporting.

With these filters, we identify 58,538 borrowers (about 6% of borrowers in our sample) who receive student loan forgiveness in the TransUnion data.¹⁶ This lines up closely with the

¹⁶We employ the same process to identify forgiveness in the employment records data which yields nearly identical rates of student loan forgiveness.

Department of Education’s statement that 7% of borrowers have received forgiveness. These forgiveness rates are departures from historical rates. In Appendix Figure C.1, we show that prior to 2021, forgiveness rates were below 1.5%, and that this more recent period, especially 2023, has much higher rates. Because borrowers’ pre-forgiveness balances vary considerably, so does the amount forgiven. In Appendix Figure C.2, we show the histogram of amount forgiven. The modal forgiven borrower had a low balance, but many had balances above \$20,000 or \$30,000.

Figure 1 shows the number of borrowers forgiven in each month. The dashed vertical lines mark the major Department of Education (DOE) forgiveness announcements. The figure shows that our estimates of forgiveness closely align with policy announcements, which suggests that we are accurately measuring loan forgiveness following new policies. Appendix B.2 further validates our measure of forgiveness. We show that there is more forgiveness for earlier cohorts, which had greater eligibility for loan discharge under PSLF and IDR, and that our aggregate and per capita measured state-level forgiveness closely matches data from the DOE.

4 Targeting

We start by assessing the targeting of student loan forgiveness. Targeting is relative to some benchmark population that either could have received the wealth shock under a counterfactual policy or that is most likely to fund the wealth shock (e.g., future taxpayers). The population that receives the studied wealth shock is clear, but there are many possible benchmark populations. Here, we will examine four benchmark populations: (1) borrowers with student debt, (2) all individuals with a credit history, (3) individuals with a credit history who did not attend college, and (4) individuals who would have received forgiveness under Biden’s original proposal.

As we saw in Table 2, borrowers who receive forgiveness have higher earnings (both unconditional and conditional on positive earnings) on average than individuals in the other two

groups. These borrowers earn an average of \$900 (19%) more per month than the average student loan borrower. Earnings, though, could be endogenous with respect to qualifying for forgiveness. For example, if borrowers with a disability avoid working to increase the chance of qualifying for forgiveness, then we might find that the policy targets lower earners, but such an exercise would overlook the distortions in earnings caused by the policy.

We thus assess targeting based on an individual's *predicted* income. We split a borrower's observable characteristics into lifecycle characteristics (W_i) and non-lifecycle characteristics (X_i) and use them to predict December 2019 monthly earnings:

$$Earnings_i = \alpha'X_i + \beta'W_i + \epsilon_i. \quad (1)$$

X_i includes indicator variables for gender, industry, and education level. W_i includes indicator variables for age and the year when the borrower took out her first student loan, if any. We estimate this specification via OLS and construct predicted income ($\widehat{Earnings}_i$) based on the estimated coefficients.

We then regress predicted income on the lifecycle characteristics and an indicator variable for whether the individual received student loan forgiveness during our sample period:

$$\widehat{Earnings}_i = \gamma Forgiveness_i + \delta'W_i + \nu_i. \quad (2)$$

We control for W_i to show how much forgiveness is targeted to higher earners based on lifecycle versus within-lifecycle factors. We present the estimates in Table 3.¹⁷ We start with Panel A, which compares forgiven borrowers to borrowers with student debt who do not receive forgiveness. The mean difference in predicted monthly earnings (in December 2019) between forgiven borrowers and non-forgiven borrowers is positive and substantial: \$703 (column 1). Much of this difference, though, reflects that forgiven borrowers tend to be older and in the

¹⁷The sample for this analysis differs slightly from Table 2 since we restrict to the approximately 50% of individuals with complete demographics information.

point of their careers where they earn more. In column 2, we show the estimates where we control for W_i and find that the estimated gap reduces to \$115. This difference represents 2% of the mean earnings among student loan borrowers. In the remaining columns we show that the gap is slightly smaller among workers in public service oriented industries. Thus, even compared to other borrowers with student debt, those who receive forgiveness are positively selected on income and are not necessarily the ones in the most financial distress.

In Panel B, we repeat the analysis but instead of using non-forgiven student borrowers as the benchmark population, we switch to all individuals with a credit history. This comparison group might represent an unweighted set of individuals who might be most likely to pay for the forgiveness through taxation. We find that forgiven borrowers are even more positively selected on income relative to this comparison group: a mean monthly predicted earnings difference of \$193 and \$504, with and without lifecycle controls, respectively.

The increase in selection is perhaps unsurprising given student loan borrowers are more likely to have attended college and, specifically, higher tuition schools. Both of these characteristics are positively associated with earnings. We demonstrate this point further in Panel C, where we estimate that relative to a comparison group of individuals who did not go to college, forgiven borrowers are very positively selected on predicted monthly income. Even with lifecycle controls, forgiven borrowers make \$1,240 more in monthly predicted earnings, or 26% of the sample mean.

These results show that the implemented forgiveness is relatively regressive and suggests that there was room for more progressive policies. We end by examining the targeting of one such policy: Biden's original proposal. We approximate the Biden proposal by forgiving up to \$10,000 for every borrower with annual earnings less than \$125,000. This only approximates the proposal for two reasons. First, we do not observe household earnings so we use the individual cutoff for all borrowers. Second, we do not know if the borrower was originally eligible for Pell grants, which would have increased forgiveness up to \$20,000 under the Biden proposal. By ignoring this second component, we likely understate the progressivity of the

proposed forgiveness.

Because the Biden proposal offers forgiveness to nearly all borrowers, rather than comparing forgiven borrowers to another group, we estimate mean (monthly) predicted earnings weighted by the number of dollars forgiven. The mean dollar of forgiveness under the implemented forgiveness policy accrues to a borrower with \$5,562 in predicted monthly earnings. The mean dollar of forgiveness under the approximated Biden proposal accrues to a borrower with \$4,868 in predicted monthly earnings, for a difference of \$694. The Biden proposal was therefore considerably more progressive than the forgiveness that has been implemented.¹⁸ Most of this difference is driven by the earnings cutoff of \$125,000 rather than variation in balances below \$10,000.

In Appendix Table C.1 we repeat the exercises with actual earnings rather than predicted earnings. We see that positive earnings gaps persist between forgiven borrowers and various comparison groups. The actual earnings gaps tend to be larger than the predicted earnings gaps. The difference is especially large when comparing the actual forgiveness to the Biden proposal. The mean dollar of forgiveness under the implemented forgiveness policy accrues to a borrower with \$6,832 in actual monthly earnings. The mean dollar of forgiveness under the approximated Biden proposal accrues to a borrower with \$3,874 in predicted monthly earnings, for a difference of \$2,958.

Thus, relative to a variety of comparison groups, forgiven borrowers are selected on having higher earnings and higher predicted earnings. These differences partly reflect lifecycle effects but still remain large when comparing individuals of the same age.

5 Potential Mechanisms and Empirical Strategy

The other set of arguments for forgiveness predict shifts in forgiven borrowers' behavior that may stimulate the economy or lead to more efficient allocations. We start by laying out some of

¹⁸However, it is still worth noting that the proposed forgiveness would have targeted individuals with higher average earnings, especially when considering the non-student borrower population.

the potential mechanisms that could lead to behavioral changes and then describe our empirical strategy for estimating causal effects of forgiveness on borrower outcomes.

5.1 Potential Mechanisms

5.1.1 Wealth Effects

The debt forgiveness we study can be thought of as a moderately-sized wealth shock. In the large literature estimating the Marginal Propensity to Consume (MPC) from small, transitory shocks, most such shocks are around \$500 and are typically assumed not to have wealth effects (see e.g., Jappelli and Pistaferri, 2010; Boehm et al., 2024). While a number of studies estimate the consumption response to shocks in components of household wealth such as housing or the stock market, these responses can be difficult to interpret without knowing whether consumers view the shocks as transitory or permanent. In the case of student loan forgiveness, we know the precise value of the shock to lifetime wealth, allowing us to follow the approach taken in the literature on lottery windfalls annuitizing the forgiveness over the remaining lifetime (Imbens et al., 2001; Cesarini et al., 2017; Golosov et al., 2023).¹⁹

Unlike winning the lottery, loan forgiveness does not provide an up-front windfall; however, it similarly frees up cash flow in the budget constraint that can be used for other purposes. Consider the household's intratemporal budget constraint:

$$c_t = -p_t + n_t + y_t + (1 + r)a_{t-1} - a_t$$

where p_t is a committed payment, n_t is unearned income, y_t is earned income, a_t are assets, and r is the interest rate. As can be clearly seen from the budget constraint, the change in consumption from a reduction in committed payments of one dollar is equivalent to an increase in unearned income of one dollar in the period, i.e. $\frac{\partial c_t}{\partial(-p_t)} = \frac{\partial c_t}{\partial n_t}$.²⁰

¹⁹Unlike lottery winnings, student loan forgiveness is not subject to taxation.

²⁰A commonly cited benchmark in the consumption literature is the permanent-income hypothesis, which predicts the MPC from the *annuity* to be large, approximately 1. We will return to this prediction in Section 8.

Another important distinction to make is the difference between consumption and expenditure, which can differ dramatically when households respond by making durable purchases. Non-student debt could either increase or decrease following student loan forgiveness. Debt could decrease if households use increased liquidity or wealth to pay down other debt. On the other hand, debt forgiveness could lead to more borrowing, and the aforementioned consumption effects might lead to debt increases. If households have more cash on hand, or anticipate having more cash on hand, they may be better able to service debt payments. This may encourage large durable purchases like homes or cars, financed by debt.

5.1.2 Additional Channels

Beyond wealth effects, additional channels could affect earnings in divergent directions. Empirical studies, usually focusing on small groups of borrowers, have found positive, negative or no effects of student debt on earnings (Yannelis and Tracey, 2022). While many of these channels could be relevant both for changes in whether debt is originated and changes in whether debt is forgiven, surprise forgiveness may dampen potential responses. Our empirical results will therefore speak to whether borrowers still have margins of adjustment at the point of forgiveness.

Constraints after college: One channel through which student debt may affect earnings is through liquidity constraints. Rothstein and Rouse (2011) argue that student borrowers need to earn more to meet payments, and hence will work longer hours or select into higher paying jobs. This mechanism would predict loan forgiveness lowers earnings, as students would no longer need higher earnings to make loan payments. We would expect this mechanism to be most relevant for borrowers not enrolled in income-driven repayment plans and not subject to the payment pause.

Debt overhang/wage garnishments: A channel leading to an increase in labor supply would be debt overhang. This mechanism would lead borrowers to reduce earnings in response to student debt, because additional earnings would go to creditors. Borrowers may believe that

their debts will be discharged eventually, perhaps through IDR or courts, and hence reduce earnings. Di Maggio et al. (2019) find evidence for debt overhang reducing labor supply for defaulted private student loan borrowers. Importantly, we would expect debt overhang effects to be dominant for borrowers who are more likely to have their loans discharged. These, however, are only a small share of our borrowers, many of whom likely expected to be making payments for many years in the future. Relatedly, borrowers in default on student loans are subject to wage garnishment, which is like a tax on working. Forgiveness would remove this tax leading to an increase in labor supply (substitution effect).

Job lock: Student loan debt may distort labor supply choices and match quality, similar to the “job/employment lock” studied in the health insurance context (e.g. Garthwaite et al., 2014). In particular, PSLF lock some borrowers into public service jobs to obtain forgiveness, and the forgiveness event may release these constraints. This mechanism could increase or decrease earnings depending on whether public service jobs are lower paying than alternatives or require living in higher cost of living metro areas (e.g., Washington, DC or New York).

5.2 Empirical Strategy

In estimating causal effects of forgiveness, we might naturally compare how a borrower’s outcomes change after receiving forgiveness relative to before, with an appropriate comparison group that does not receive forgiveness. Assuming, though, that forgiveness is an exogenous event may be inappropriate in some contexts. When borrowers have some control over the process by which they receive forgiveness, their actions may coincide with other household or balance sheet shocks that could confound estimation of causal effects. But as we described in Section 2, most forgiveness during our sample period is likely driven by a set of opaque policy changes rather than individuals’ actions. Though the opacity has some empirical downsides – e.g., we cannot translate policy announcements into designs that isolate finer sources of variation – it also limits the ability of individuals to sort endogenously into treatment.

In case forgiveness events are confounded by other shocks or we have attributed too many

balance drops to loan forgiveness policies, we will also conduct an instrumental variables analysis that uses another source of exogenous variation to predict forgiveness. This analysis entails additional strong assumptions and thus we treat it as extra evidence consistent with our main specification.

5.2.1 Difference-in-Differences Estimator

Let $i \in \{1, 2, \dots, I\}$ denote a borrower and $t \in \{0, 1, \dots, T\}$ denote a time period (month). All borrowers are present for all time periods such that we have a balanced panel. Following notation from De Chaisemartin and d’Haultfoeuille (2024), each i belongs to a group $g(i)$, where we will separate and label groups based on the timing of their forgiveness: e.g., $g(i) = g(i')$ for $i \neq i'$ if i and i' receive forgiveness in the same month, g . Borrowers who do not receive forgiveness during our panel are in their own group, with $g = T + 1$.

Our parameter of interest is the average treatment effect on the treated (ATT) for some Y_{it} for borrowers who received forgiveness l periods ago. We use the estimator from De Chaisemartin and d’Haultfoeuille (2024), though because we have a binary absorbing treatment and all borrowers begin the sample without treatment, the estimator coincides numerically with the estimator from Callaway and Sant’Anna (2021). Our comparison group consists of borrower groups who have not yet received forgiveness, and because forgiveness is a relatively rare event our comparison group is dominated by borrowers who have never received forgiveness. To estimate the ATT, we weight individual group ATTs by their relative number of borrowers.²¹ We estimate ATTs out to $l = 6$ and construct placebo ATT estimates back to $l = -5$. We cluster our standard errors at the forgiveness group level.²²

Because we have some uncertainty whether a balance decrease (forgiveness event) that shows up in the data in g occurred between $g - 2$ and $g - 1$ or between $g - 1$ or g , we specify

²¹As Callaway and Sant’Anna (2021) point out, comparing estimates across different horizons l and l' is complicated by differences in weights in addition to differences in ATTs. Because we have a balanced panel and nearly all groups receive treatment well before the end of the panel, the group weights hardly change across different values of l .

²²Our balanced panel means that if we were to define groups at the individual level, akin to using individual fixed effects, the estimates are numerically identical.

$g - 2$ as the last period when a group was untreated. We pursue a “doughnut hole” approach by excluding $g - 1$ from the analysis because its level of treatment is unclear. We also report pooled estimates of the average ATT from $l = 0$ to $l = 6$ (relative to a pooled average from $l = -5$ to $l = -2$).

5.2.2 Instrumental Variables

We pursue a second empirical strategy to assess how robust our results are to different assumptions. Rather than assume forgiveness events are exogenous, we create an instrumental variable that shifts whether an individual borrower receives forgiveness.

We use the institutional feature that when borrowers take out their students loans, they are randomly assigned to a servicer (Cornaggia and Xia, 2024). Some servicers may interact more directly with Federal relief programs or be more careful about recording borrowers’ eligibility for forgiveness programs such that borrowers randomly assigned to these servicers are differentially likely to receive forgiveness. We thus construct a servicer’s (leave-out) forgiveness rate as the fraction of the servicer’s borrowers who receive forgiveness during our sample period, excluding the borrower in question. Let s index servicers, with $s(i)$ capturing the assignment of i to servicer, as of the beginning of our sample.²³ Then our instrument is:

$$Leniency_i = \frac{1}{I_{s(i)} - 1} \sum_{i' \neq i: s(i) = s(i')} Forgiveness_{i'}, \quad (3)$$

where $Forgiveness_{i'}$ is a dummy variable for whether i' receives student loan forgiveness during our sample and I_s is the number of borrowers with servicer s .

We then estimate a long-difference specification that compares how an individual’s outcome changes from half a year after the event relative to half a year before the event. The long-difference specification mimics the type of cross-sectional analysis common in examiner IV

²³Though initial assignment of borrowers to servicers may be random, non-random sorting could occur over time as borrowers face certain events, such as enrolling in PSLF. We will also assess the randomness with balance checks.

papers. We estimate the following specification via two-stage least squares:

$$\begin{aligned}\Delta Y_i &= \beta \text{Forgiveness}_i + \lambda X_i + \epsilon_i \\ \text{Forgiveness}_i &= \gamma \text{Leniency}_i + \pi X_i + \mu_i.\end{aligned}\tag{4}$$

ΔY_i is the difference in outcome Y_{it} between six months after forgiveness and six months before forgiveness ($Y_{i,g+6} - Y_{i,g-6}$). For borrowers who never receive forgiveness, we randomly assign a placebo forgiveness date. These borrowers are the comparison group that controls for changes over time unrelated to treatment. In our preferred IV specification, we will include age and forgiveness year (or placebo forgiveness year) fixed effects in X_i .

Our estimates will identify the local average treatment effect for compliers – i.e., borrowers whose forgiveness outcome depends on their servicer. If some of the forgiveness in our sample is borrower-driven, then we might expect that compliers might be quite different from borrowers whose forgiveness does not depend on their servicer and thus our difference-in-differences and IV estimates might diverge. But if, instead, the estimates coincide, this suggests that supply-side actions might be more important in driving forgiveness events.

Interpreting these estimates as a causal treatment effect of receiving forgiveness relies on a potentially strong exclusion restriction that servicers only affect forgiveness probabilities and not other factors related to borrower outcomes. If the same servicers that have organized paperwork that changes borrowers' forgiveness rates also help borrowers stay on track with payments, then we might attribute changes in outcomes to forgiveness rather than these other factors. Because the exclusion restriction is strong, we leave the instrumental variables analysis as a robustness check for our preferred difference-in-differences specification.

6 Effects on Credit Outcomes

We now turn to the treatment effect estimates. Table 4 presents estimates on student loan outcomes following our difference-in-differences approach. The first row presents dynamic

difference-in-difference estimates following De Chaisemartin and d’Haultfoeuille (2024), while the second presents two-way fixed effects estimates for comparison.²⁴ We find that forgiveness leads to substantial effects on student loan balances, with the average balance dropping by \$31,800 following forgiveness. Payments due drop by approximately \$64, a relatively small number due to many borrowers being subject to the payment pause in much of our estimation window. We also see a sharp drop in delinquencies, by 2.4 percentage points. This drop is expected given the forgiven loans can no longer be delinquent. All the estimates are statistically significant at the 1% level.

Figure 2 presents the corresponding event study coefficient estimates, which assist us in evaluating the main identifying assumptions. For each of the three outcomes, we see fairly flat trends prior to forgiveness, with sharp and immediate drops in student loan outcomes upon forgiveness. The estimated decreases are persistent.

Table 5 presents the estimates on credit bureau outcomes. For the six months following forgiveness, we estimate increases in mortgage, auto, and credit card balances totaling approximately \$2,800. The majority of this increase, \$2,320, is driven by increases in mortgage balances, while the remainder is equally split between auto and credit card balances. We see corresponding increases in payments due for all three credit categories. All the balance and payment estimates are statistically significant at the 1% level. There are few effects on delinquencies, except a statistically significant but very small increase in credit card delinquencies.

Figure C.5 presents event study coefficient estimates analogous to the estimates in Table 5. For each of the three loan categories, we see flat pre-trends and then a gradual increase in balances and payments due following loan forgiveness. The increases in balances and payments due level off within four months. Consistent with Table 5, we see little evidence of changes in delinquencies, other than perhaps a very small increase in credit card delinquencies.

The increased balances on non-student debt could reflect borrowers opening new lines of

²⁴The two-way fixed effects estimates may use negative weights on treatment effects (Callaway and Sant’Anna, 2021). We include the two-way fixed effects estimates for comparison, and find that differences are small, likely because the comparison group includes mostly borrowers who do not receive forgiveness during our sample.

credit (e.g., a new mortgage or financing a new car) or increased use of credit on existing lines. We examine the causal effect of forgiveness on the number of credit lines in columns 1-2 of Table C.2 (with event study graphs in Figure C.14). As expected, we estimate a decrease in the number of student loans lines. For the other forms of credit, we see a statistically significant increase in the number of lines in each credit type. We also show estimates of the treatment effect on whether the borrower has any open line of credit of each type. While the increase in credit card lines is largely driven by additional credit cards, the mortgage and auto results come from opening a first mortgage or auto loan. This extensive margin response is consistent with borrowers using credit to finance consumption of new durables.

We further assess the effects on the intensive margin of credit usage. We examine balances, payments due, and delinquencies on credit lines that were already open at the time of the forgiveness event in columns 3-4 of Table C.2 (with event study graphs in Figure C.15). Here, we see a (noisy) decrease in borrowing on pre-existing mortgage lines and a large and statistically significant decrease in auto loan balances. For credit cards, we estimate an increase in credit usage on lines that were already open. This effect for credit cards explains about half of the overall treatment effect, while the effects on mortgages and auto loans move in the opposite direction of the total effect.

While the additional mortgage line likely reflects a new property purchase, most of these properties are in the borrower's original zip code. We estimate the effect of forgiveness on relocation and present the estimates in Table C.3 (event study graphs in Figure C.7). We see a precise null effect.²⁵

In Table C.3 (and Figure C.8) we estimate the effect of forgiveness on credit scores. We estimate a small (1-2 point) decrease in credit scores. Thus, the forgiveness event does not affect credit supply, at least the extent of supply that varies with credit score. This result is reminiscent of the results in Dinerstein et al. (2024), where increased borrowing from a payment pause came in a sample with minimal change in credit scores. Because credit score

²⁵We define relocation as a change in the associated with a borrower's primary address.

might go down in response to the increased borrowing, the small effect may hide an initial increase in credit supply that is counteracted by a subsequent decrease.

The estimates we have presented show pooled effects out to 6 months post-forgiveness. In Appendix Table C.4, we show that the effects appear quite stable out to 12 months post-forgiveness, as the coefficient estimates are quite similar to those in Table 5. We caution, though, that our panel becomes somewhat imbalanced as we extend the post-forgiveness period such that earlier forgiveness cohorts contribute larger weight to the estimated effects.

6.1 Results Using Servicer Variation

We next consider our complementary IV strategy outlined in section 5.2.2. Before jumping to estimates, we assess how balanced our instrument (servicer's forgiveness rate) is across a variety of observable borrower characteristics. Table C.5 shows that a borrower's age and credit score are not strongly related to her servicer's (leave-out) forgiveness rate. Similarly, we show a minimal relationship between this forgiveness rate and characteristics of the state where the borrower lives: share of residents who are Black, share of residents who are Asian, share of residents who are female, mean household income, share of residents that are registered Democrats, and share of residents with a college degree.

Table 6 presents the IV estimates. Specifically, the table shows estimates of the coefficients γ and β from equation (4). The top panel shows the first stage, the effect of servicer leniency on the probability of receiving forgiveness. Coefficients are quite close to 1, and the first stage F-statistics all exceed 8,000. The bottom panel presents estimates of β , the effect of loan forgiveness on outcomes. Generally, the magnitudes are quite similar to those in Tables 4 and 5 suggesting that both strategies capture similar treatment effects of loan forgiveness on credit outcomes.

6.2 Heterogeneous Effects and Robustness

Our estimates pool over different policies that generate forgiveness. These policies affect different populations of borrowers and may have different treatment effects. Some of the key variables that would clarify which program generated the forgiveness (e.g., disability status; institution attended) are unfortunately not recorded in the data. We thus use the timing of the forgiveness event as a proxy for which program generated it. Specifically, we use the government’s announced forgiveness events from Appendix Table B.1 and for any month in which a single policy accounts for at least 90% of forgiven borrowers, we classify borrowers that received forgiveness in that month as under the majority policy. For example, in March 2021 the only announcement is for Permanent Disability, so any borrower who we estimate receives forgiveness in March 2021, we will classify as receiving it under Permanent Disability. For December 2023, both the IDR Adjustment and PSLF have announcements accounting for large numbers of forgiven borrowers, so we do not classify borrowers in that month.

We present the estimates by imputed program type in Appendix Table C.6. While we estimate some differences – larger mortgage debt responses for forgiveness under Permanent Disability and Borrower Defense – the estimates are qualitatively similar across all four programs. We estimate positive effects on mortgage, auto, and credit card balances for all programs, and statistically significant for all but auto loans for Disability and Borrower Defense. The similarity in estimates might imply that the effects are driven by economic channels, such as wealth effects, that are common to forgiveness under all of the programs.

We further examine mechanisms by focusing on borrowers who receive a wealth, but not liquidity, shock. In Appendix Table C.7, we present estimates for the subsample of borrowers that were subject to the payment pause. For forgiven borrowers, these are borrowers who received forgiveness at least six months before the end of the payment pause and had no payment due in the month prior to forgiveness. For borrowers who did not receive forgiveness (the comparison group), these are borrowers who had no payments due in the month before the end of the payment pause. When we isolate this subsample, we see in Panel A that there

remains a large decrease in student loan balance post-forgiveness but that the payment due (mechanically) does not meaningfully change. Further, we see no causal effect on delinquency, which makes sense given loans could not go delinquent during the payment pause. Panels B, C, and D show the estimates on mortgage, auto, and credit card debt, respectively. The estimates are quite similar to our baseline estimates, which suggests that most of the policy effects in the full sample are driven by wealth effects rather than liquidity effects. This finding is an interesting contrast to mortgage forgiveness where borrowers are more responsive to liquidity changes than balance reductions (Ganong and Noel, 2020).

Finally, we show robustness to inference when clustering our standard errors at different levels. In our main analysis, we cluster at the forgiveness cohort level, under the assumption that the timing of the policies (Appendix Table B.1) separates borrowers who are subject to common shocks. We expect that this form of clustering may be conservative, in part because the never treated group is pooled into a single cluster and requires estimate a large number of within-cluster parameters for inference. In Appendix Table C.8, we show how standard errors change if we cluster at the borrowing cohort (when the borrower first took out student loans), a factor related to when borrowers might receive forgiveness under IDR and PSLF. We indeed find that this form of clustering generates smaller standard errors and that our baseline clustering may be conservative.

7 Effects on Labor Market Outcomes

We now examine the causal effect of forgiveness on labor market outcomes. Table 7 and Figure C.6 show treatment effects for earnings-related outcomes. We start by focusing on the intensive margin by using the sample of workers who have reported earnings for all of our sample period. We estimate a drop in monthly earnings of \$44, or 2.3% when we use log earnings. As with the credit outcomes, the dynamic difference-in-difference and two-way fixed effect estimators deliver very similar results. We then look at the extensive margin by using the full sample of

workers, not conditional on having reported earnings for the entirety of the sample period. For these workers, we estimate a 0.4 percentage point decrease in the probability of having any earnings reported or being employed.

Figure C.6 shows minimal evidence of pre-trends. The decreases in earnings and the probability of being employed may reflect standard wealth effects, as well a lower need for immediate earnings to pay back debt. But labor supply adjustments may take some time to implement. Consistent with a staggered response to the initial shock, we see that the effects on the labor market outcomes are increasing in magnitude over time. This time path differs from the credit outcome effects, which were fairly steady by 6 months post-forgiveness.²⁶

In Table 7 we further investigate the labor supply response to by looking at job switching. We do not see any change in the rate at which borrowers have new employment post-forgiveness, but we instead see that the nature of job switching changes. We estimate an increase in switching industries, and especially switching out of industries associated with public service. The increase in industry switching is more than 10% of the baseline switching rate, while the outflow from public service jobs is statistically significant but small. This latter effect may reflect job lock related to PSLF, as forgiven borrowers no longer need to work in public service to qualify for future forgiveness.

7.1 Heterogeneous Effects

In Table 8, we explore heterogeneity by different subgroups of borrowers. Columns (1) and (2) split the sample by hourly and salaried workers. We find larger effect for hourly workers, who are likely more elastic and are more able to request more hours of overtime, or cut hours in response to changed incentives. For hourly workers, we also see a drop in the hourly wage in column (7) of Table 7. When we decompose the total earnings effect for hourly workers into hours and wages, we attribute half of the decreased earnings to reduced hours and half to

²⁶In Appendix Table C.9 we show estimates for the period out to 12 months post-forgiveness and find them continuing to increase in magnitude, but we again suggest caution in interpreting the results as the panel becomes relatively unbalanced.

reduced wages. The labor supply response may be a combination of working less and switching to lower-wage jobs.

7.1.1 Plan Type

Like with the credit outcomes, our data do not allow us to attach each forgiveness event to a specific program. For earnings, we focus on public sector employees who are more likely to receive forgiveness under PSLF. Columns (3) and (4) split the sample by workers in public service. We find that earnings effects are larger for workers in public service. This is consistent with the earlier drop in public service workers seen in Table 7, and likely reflect workers leaving public service.

7.1.2 Debt Overhang

Columns (5) and (6) of Table 8 split the sample by whether borrowers had ever defaulted. Our main results are driven entirely by the set of borrowers who never defaulted. Interestingly, we find opposite signed results for borrowers who had previously defaulted. This finding is consistent with Di Maggio et al. (2019), who find positive earnings effects for a sample of defaulted borrowers, and also with much of the literature on consumer bankruptcy discharge, such as Dobbie and Song (2015). Why might effects differ for borrowers in default? Di Maggio et al. (2019) argue that their negative employment effects are consistent with a debt overhang channel, which is consistent with the theory literature (Donaldson et al., 2019). Debt overhang arises from a shift in incentives when borrowers are close to discharge, either through bankruptcy or other forgiveness, and a reduction in earnings due to those being paid to creditors. This effect is much more likely to be much more relevant for borrowers in default who anticipate some form of future discharge.

8 MPX, MPC and MPE out of Student Loan Forgiveness

We have so far presented reduced form borrowing and earnings responses to student loan forgiveness. We next map our estimates to Marginal Propensities for Expenditure (MPX), Consumption (MPC), and Earnings (MPE).

When considering consumption responses, an important distinction is the difference between consumption and spending, which can differ dramatically when households respond by making expensive durable purchases. Durables are consumed and typically financed over multiple periods. Spending responses including durable purchases are estimates of the MPX, which will exceed estimates of the MPC (Laibson et al., 2022). In addition, the MPX is the relevant estimate for understanding the macroeconomic effects, since they align with the timing of expenditure for GDP calculations. In our context, we can see balances on two of the largest household durables: homes and autos; we also see credit card balances, which could include a mix of non-durable, semi-durable, and durable goods. Because the stock of balances most closely maps to the purchase of a durable, we will use changes in balances for our calculation of the MPX.²⁷ In addition, we see the flow of payments resulting from these balances. We will use these payment flows in our MPC calculations.

We consider three different normalizations of our reduced form responses, mapping to different ways the corresponding MPCs and MPEs are typically discussed in the literature. First, we report the response normalized by the total size of the wealth shock, i.e. the reduction in student loan balances, which is approximately \$32,000 on average. Second, we annuitize (at the monthly level) the reduction in student loan balances:

$$b = \frac{r}{1+r} \left(1 - \left(\frac{1}{1+r} \right)^{T-k+1} \right)^{-1} B, \quad (5)$$

where b is the annuitized monthly value and B is the size of the balance reduction. r is the

²⁷Balances on durables exclude down payments, which will cause us to underestimate spending using the increase in balances. If the average downpayment is 10%, we would therefore underestimate durable spending by 10%. More generally, our spending responses miss cash and debit card spending.

post-tax monthly interest rate, assumed to be $\frac{2.5\%}{12} = 0.21\%$ as in Golosov et al. (2023), and T is life expectancy in months, which we assume to be $80 * 12 = 960$. k is the borrower's age (in months) at the time of forgiveness. Given the average forgiveness amount is around \$32,000, and the age of the typical borrower in our sample is around 39, the typical borrower has about 41 years of remaining life expectancy. Spreading the wealth shock equally every month over 41 years is a savings of around \$104 per month.

Finally, we directly use our estimates of the reduction in payments due. This is likely the most relevant denominator for hand-to-mouth or rule-of-thumb consumers. From the perspective of the standard lifetime consumption-savings model, the short-run reduction in payments due would be too large given that the modal households would not expect to be paying their student loans off for the rest of their remaining lifetimes. In our context, the average reductions in payment due in Table 4 happens to be smaller, rather than larger, than the evenly-spread wealth shock, likely due to factors like the payment pause, forbearance, and especially IDR.

The calculations so far assume that the changes are unexpected. But as discussed in Appendix A, borrowers' beliefs that they would receive forgiveness, prior to actually receiving forgiveness, averaged 25%. If borrowers had incorporated these beliefs into pre-forgiveness decisions about consumption or labor supply, then the responses to forgiveness would be based on the difference in the wealth shock and the expected wealth shock. We therefore multiply each normalization – the size of the wealth shock, the annuitized amount, and the reduction in payments due – by 0.75.

We report the implied total MPX, monthly MPC, and monthly MPE estimates in Panel A of Table 9. The point estimate shown in Columns (2)-(4) in the Table is equal to the reduced-form effect on the outcome, divided by the first-stage effect on the treatment. We bootstrap this ratio to obtain standard errors.

Our estimated MPX, based on the combined increase in mortgage, auto, and credit card balances, and normalized by the reduction in average student loan balances, is around 0.12. Our estimate of the MPX using the annuity method is 27.4 and using the savings from student

loan payments is 57.5. In other words, households spend 27.4 to 57.5 times the monthly wealth shock, as households front-load spending on durables like housing. We now turn to our MPC estimates. Our estimate of the monthly MPC using the annuity method is 0.27 and using the flow values of payments is 0.57. Note that the MPX is around 100 times our MPC estimates. This ratio is useful for converting MPC to MPX, and vice versa.

The next rows in the Table focus on our earnings sample. The first set of MPE estimates restrict to positive earnings (intensive margin), whereas the second MPE estimates fills in zeros when someone has no earnings being reported in the month.²⁸ In either case, we find an MPE of around -0.5 when annuitizing the reduction in student loan balances; responses are higher, -0.88 to -0.97, when considering only the reduction in payment due. Our estimates of the MPE, using the annuity method, are comparable with Golosov et al. (2023), who also find an MPE around -0.5 among U.S. lottery winners.

These MPX, MPC, and MPE estimates provide summary measures of the consumption and labor supply responses. Our estimates appear broadly in line with other recent estimates in the literature. While our estimates are only partial-equilibrium, they provide some insight on the macroeconomic implications of broad-scale forgiveness. Our MPX estimate suggests that forgiveness so far has led to an increase in upfront spending of \$11.5 billion.²⁹ For the labor market consequences, our MPE estimates imply this forgiveness reduced earnings by around \$1.4 billion per year.³⁰ Since this represents around 7% of outstanding debt, forgiving all student loan debt would increase these aggregate effects by over 10 times. Large-scale forgiveness, however, is more likely to have general equilibrium effects, and the overall economic benefits for the economy will likely to depend on how the forgiveness is financed by the government.

²⁸ See data limitations discussed above.

²⁹ \$132 billion in forgiveness * 0.116 * 0.75 change in beliefs.

³⁰ 3 million borrowers receiving forgiveness * -\$1,248 in savings when annuitized over remaining lifetime * 0.75 change in beliefs * MPE of -0.49

8.1 Permanent Income Hypothesis

We end by testing whether borrowers' responses are consistent with the permanent income hypothesis. We return to an intratemporal budget constraint with labor taxes:

$$c_t - y_t + T(y_t) = n_t,$$

where c_t is consumption, y_t are pre-tax labor earnings, $T(y_t)$ are labor taxes, and n_t is all unearned income, which includes negative payments ($-p_t$) for student loans. Consider the student debt forgiveness event. If borrowers act consistently with the permanent income hypothesis, then they would allocate the full shock B equivalently to each period t . Specifically, they allocate b (from Equation 5) to each period. Let αb represent the amount borrowers actually allocate, where $\alpha = 1$ is behavior consistent with the permanent income hypothesis.³¹ If Δ represents the change in response to forgiveness,

$$\Delta c_t - \Delta y_t + \Delta T(y_t) = \Delta n_t = \alpha b.$$

Dividing by b , we estimate:

$$\hat{\alpha} = \widehat{MPC} - \widehat{MPE} + \widehat{MPT}, \quad (6)$$

where the MPT is the marginal change in labor earnings taxes and the MPC, MPE, and MPT all use the annuity normalization.³²

We present the estimates of α in the last column of Panel B of Table 9. The top row shows estimates for our full sample, the same sample as in Panel A. We estimate $\hat{\alpha} = 0.64$, a bit below 1. We show how $\hat{\alpha}$ changes in the next rows if we assume our data are missing down payments of 10% (row 2) or include the extensive margin of employment and assume individuals exiting

³¹Given we are using a positive wealth shock and most borrowers have many years left in their working life, we ignore precautionary savings considerations (Blundell et al., 2008).

³²As we do not observe taxes directly in the data, we estimate marginal tax rates based on the single-filer Federal tax schedule. The mean estimated marginal tax rate among forgiven borrowers is 20%. We add the mean state income tax for mean annual earnings in our sample (\$68,000), for a total marginal tax rate of $20\% + 4\% = 24\%$. As we show below, our results are not sensitive to moderate deviations in these assumptions.

the data have 0 earnings (row 3). Neither of these changes meaningfully affects our estimate. The fourth row shows that not incorporating baseline forgiveness beliefs lowers $\hat{\alpha}$ to 0.48.

Our full sample, though, does not provide the cleanest test of the permanent income hypothesis for two reasons. First, if there are reaching new levels of consumption or labor supply require adjustments that take time to make, then our estimates that pool over months 0 to 6 post-forgiveness may understate the long-run response. In the fifth row, we estimate α when using the estimated consumption and earnings responses from the sixth month post-forgiveness. We estimate $\hat{\alpha} = 1.11$ and fail to reject $\alpha = 1$. Second, some of the borrowers in our sample are subject to more than just a wealth shock. For example, some borrowers had required payments and thus forgiveness provides liquidity in addition to wealth. Liquidity effects might predict deviations from perfect smoothing of the wealth shock. For a cleaner test, we therefore use the subsample of borrowers who had no required student loan payments.³³ We present our estimate of α on this subsample in the last row and find $\hat{\alpha} = 0.82$. We again fail to reject $\alpha = 1$. Had we not considered the earnings responses, we would have decisively rejected the permanent income hypothesis.

Why do borrowers appear to act so consistently with the canonical lifecycle model? We speculate that several factors related to student debt forgiveness make it particularly amenable to perfect smoothing of the shock. First, we focus on somewhat older and higher-income borrowers who face fewer constraints that can lead to violations of consumption smoothing. For example, Ganong and Noel (2020) find that collateral constraints prevent underwater borrowers from consuming much out of decreases to their mortgage balances. These types of constraints are less likely to bind for the borrowers subject to student loan forgiveness. Second, the repayment structure of student debt – monthly payments of similar amounts – meant that borrowers might have already thought of their debt in smoothed or annuitized terms. Thus, borrowers may have found it particularly natural to adjust their monthly consumption

³³Specifically, we keep borrowers who receive forgiveness and had no required payments in the six months prior to forgiveness and borrowers who do not receive forgiveness and had no required payments in the period of the payment pause.

or earnings based on a sense of their monthly debt obligation.

9 Concluding Remarks

We study the effects of student loan forgiveness, a common proposal to address soaring student loan debt burdens, using the largest period of student loan forgiveness in history. The study uses two complementary strategies to generate large-scale variation in student loan forgiveness. The results indicate that loan forgiveness increases consumption in the short term, with sharp increases in mortgage, auto, and credit card debt following loan forgiveness. We find a negative effect on earnings and the probability of being employed.

While our study contributes to researchers' and policymakers' understanding of the effects of loan forgiveness, there remain substantial avenues for further research. This is especially the case given ongoing policy and legal disputes regarding whether loans should be forgiven, and who should receive varying amounts of loan cancellation. Future work may study optimal relief for borrowers, and how insurance acts with distributional and macroeconomic consequences of loan forgiveness and other policies to assist student debtors.

References

- Agarwal, Sumit and Wenlan Qian**, “Consumption and debt response to unanticipated income shocks: Evidence from a natural experiment in Singapore,” *American Economic Review*, 2014, 104 (12), 4205–30.
- , **Chunlin Liu, and Nicholas S Souleles**, “The reaction of consumer spending and debt to tax rebates—evidence from consumer credit data,” *Journal of Political Economy*, 2007, 115 (6), 986–1019.
- , **Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru**, “Policy intervention in debt renegotiation: Evidence from the home affordable modification program,” *Journal of Political Economy*, 2017, 125 (3), 654–712.
- Amromin, Gene and Janice Eberly**, “Education financing and student lending,” *Annual Review of Financial Economics*, 2016, 8, 289–315.
- Aydin, Deniz**, “Consumption response to credit expansions: Evidence from experimental assignment of 45,307 credit lines,” *American Economic Review*, 2022, 112 (1), 1–40.
- Baker, Scott R**, “Debt and the response to household income shocks: Validation and application of linked financial account data,” *Journal of Political Economy*, 2018, 126 (4), 1504–1557.
- , **Robert A Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis**, “Income, liquidity, and the consumption response to the 2020 economic stimulus payments,” Technical Report, National Bureau of Economic Research 2020.
- Black, Sandra E, Jeffrey T Denning, Lisa J Dettling, Sarena Goodman, and Lesley J Turner**, “Taking it to the limit: Effects of increased student loan availability on attainment, earnings, and financial well-being,” Technical Report, National Bureau of Economic Research 2020.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston**, “Consumption inequality and partial insurance,” *American Economic Review*, 2008, 98 (5), 1887–1921.
- Board of Governors of the Federal Reserve System**, “Consumer Credit - G.19,” https://www.federalreserve.gov/releases/g19/HIST/cc_hist_memo_levels.html 2024. Retrieved June 3, 2024.
- Boehm, Johannes, Etienne Fize, and Xavier Jaravel**, “Five facts about MPCs: Evidence from a randomized experiment,” *Working Paper*, 2024.
- , – , and – , “Five facts about MPCs: Evidence from a randomized experiment,” *American Economic Review*, 2025, 115 (1), 1–42.
- Boutros, Michael, Nuno Clara, and Francisco Gomes**, “Borrow now, pay even later: A quantitative analysis of student debt payment plans,” *Pay Even Later: A Quantitative Analysis of Student Debt Payment Plans (October 12, 2022)*, 2022.

- Callaway, Brantly and Pedro HC Sant'Anna**, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Catherine, Sylvain and Constantine Yannelis**, "The distributional effects of student loan forgiveness," *Journal of Financial Economics*, 2022.
- Cesarini, David, Erik Lindqvist, Matthew J. Notowidigdo, and Robert Östling**, "The effect of wealth on individual and household labor supply: Evidence from Swedish lotteries," *American Economic Review*, December 2017, 107 (12), 3917–46.
- Chaisemartin, Clément De and Xavier d'Haultfoeuille**, "Difference-in-differences estimators of intertemporal treatment effects," *Review of Economics and Statistics*, 2024, pp. 1–45.
- Chava, Sudheer, Heather Tookes, and Yafei Zhang**, "Leaving them hanging: Student loan forbearance, distressed borrowers, and their lenders," 2023.
- Cherry, Susan F, Erica Jiang Jiang, Gregor Matvos, Tomasz Piskorski, and Amit Seru**, "Government and private household debt relief during Covid-19," Technical Report Fall, National Bureau of Economic Research 2021.
- Chodorow-Reich, Gabriel, Plamen T Nenov, and Alp Simsek**, "Stock market wealth and the real economy: A local labor market approach," *American Economic Review*, 2021, 111 (5), 1613–57.
- Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber**, "How did US consumers use their stimulus payments?," Technical Report, National Bureau of Economic Research 2020.
- Cornaggia, Kimberly and Han Xia**, "Who mismanages student loans, and why?," *The Review of Financial Studies*, 2024, 37 (1), 161–200.
- de Silva, Tim**, "Insurance versus moral hazard in income-contingent student loan repayment," Available at SSRN 4614108, 2023.
- Dinerstein, Michael, Constantine Yannelis, and Ching-Tse Chen**, "Debt moratoria: Evidence from student loan forbearance," *American Economic Review: Insights*, 2024.
- Dobbie, Will and Jae Song**, "Debt relief and debtor outcomes: Measuring the effects Of consumer bankruptcy protection," *American Economic Review*, 2015, 105 (3), 1272–1311.
- Donaldson, Jason Roderick, Giorgia Piacentino, and Anjan Thakor**, "Household debt overhang and unemployment," *The Journal of Finance*, 2019, 74 (3), 1473–1502.
- Federal Student Aid**, "Loan forgiveness reports," <https://studentaid.gov/data-center/student/loan-forgiveness> 2024. Retrieved June 3, 2024.
- Ganong, Peter and Pascal Noel**, "Liquidity versus wealth in household debt obligations: Evidence from housing policy in the great recession," *American Economic Review*, 2020, 110 (10), 3100–3138.

- Garthwaite, Craig, Tal Gross, and Matthew J. Notowidigdo**, “Public health insurance, labor supply, and employment lock*,” *The Quarterly Journal of Economics*, 03 2014, 129 (2), 653–696.
- Gibbs, Christa, Benedict Guttman-Kenney, Donghoon Lee, Scott Nelson, Wilbert van der Klaauw, and Jialan Wang**, “Consumer credit reporting data,” 2023.
- Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky**, “How Americans respond to idiosyncratic and exogenous changes in household wealth and unearned income*,” *The Quarterly Journal of Economics*, 10 2023, 139 (2), 1321–1395.
- Goodman, Sarena, Adam Isen, and Constantine Yannelis**, “A day late and a dollar short: Liquidity and household formation among student borrowers,” *Journal of Financial Economics*, 2021, 142 (3), 1301–1323.
- Gopalan, Radhakrishnan, Barton H. Hamilton, Ankit Kalda, and David Sovich**, “State minimum wages, employment, and wage spillovers: Evidence from administrative payroll data,” *Journal of Labor Economics*, 2021, 39 (3), 673–707.
- Grigsby, John, Erik Hurst, and Ahu Yildirmaz**, “Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data,” *American Economic Review*, February 2021, 111 (2), 428–71.
- Hamdi, Naser, Ankit Kalda, and David Sovich**, “The labor market consequences of student loan forbearance,” *Available at SSRN 4787183*, 2024.
- Hampole, Menaka V**, “Financial frictions and human capital investments,” Technical Report, Working paper 2022.
- Herbst, Daniel and Nathaniel Hendren**, “Opportunity unraveled: Private information and the missing markets for financing human capital,” Technical Report, National Bureau of Economic Research 2021.
- Imbens, Guido W, Donald B. Rubin, and Bruce I. Sacerdote**, “Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players,” *American Economic Review*, September 2001, 91 (4), 778–794.
- Jacob, Brian, Damon Jones, and Benjamin J Keys**, “The value of student debt relief and the role of administrative barriers: Evidence from the Teacher Loan Forgiveness Program,” *Journal of Labor Economics*, 2024, 42 (1).
- Jappelli, Tullio and Luigi Pistaferri**, “The consumption response to income changes,” *Annual Review of Economics*, 2010, 2 (Volume 2, 2010), 479–506.
- Koustas, Dmitri, Michael Weber, and Constantine Yannelis**, “Beliefs about student loan repayment,” Technical Report, Unpublished Working Paper 2024.
- Laibson, David, Peter Maxted, and Benjamin Moll**, “A simple mapping from MPCs to MPXs,”

- NBER Working Paper 29664*, 2022.
- Lochner, Lance and Alexander Monge-Naranjo**, “Student loans and repayment: Theory, evidence, and policy,” in “Handbook of the Economics of Education,” Vol. 5, Elsevier, 2016, pp. 397–478.
- Looney, Adam and Constantine Yannelis**, “What went wrong with federal student loans?,” *Journal of Economic Perspectives*, 2024, 38 (3), 209–236.
- Maggio, Marco Di, Amir Kermani, and Kaveh Majlesi**, “Stock market returns and consumption,” *The Journal of Finance*, 2020, 75 (6), 3175–3219.
- , **Ankit Kalda, and Vincent Yao**, “Second chance: Life without student debt,” Technical Report, National Bureau of Economic Research 2019.
- Mian, Atif and Amir Sufi**, “The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis,” *The Quarterly Journal of Economics*, 2009, 124 (4), 1449–1496.
- , **Kamalesh Rao, and Amir Sufi**, “Household balance sheets, consumption and the economic slump,” *The Quarterly Journal of Economics*, 2013, 128 (4), 1687–1726.
- Mueller, Holger and Constantine Yannelis**, “The rise in student loan defaults,” *Journal of Financial Economics*, 2019, 131 (1), 1–19.
- and —, “Increasing enrollment in income-driven student loan repayment plans: Evidence from the Navient field experiment,” *The Journal of Finance*, 2022, 77 (1), 367–402.
- National Center for Education Statistics**, *Loans for undergraduate students and debt for bachelor’s degree recipients*, U.S. Department of Education, Institute of Education Sciences, 2023. Condition of Education. Retrieved June 3, 2024 from <https://nces.ed.gov/programs/coe/indicator/cub>.
- Rothstein, Jesse and Cecilia Elena Rouse**, “Constrained after college: Student loans and early-career occupational choices,” *Journal of Public Economics*, 2011, 95 (1-2), 149–163.
- The White House**, “Public student loan forgiveness,” <https://www.whitehouse.gov/publicserviceloanforgiveness/> 2024. Retrieved June 3, 2024.
- U.S. Department of Education**, “National default rate briefings - FY 2016 official cohort default rates,” <https://fsapartners.ed.gov/knowledge-center/library/electronic-announcements/2019-09-25/national-default-rate-briefings-fy-2016-official-cohort-default-rates> 2019.
- , “Over 323,000 federal student loan borrowers to receive \$5.8 billion in automatic total and permanent disability discharges,” <https://www.ed.gov/news/press-releases/over-323000-federal-student-loan-borrowers-receive-58-billion-automatic-total-a> 2023. Retrieved June 21, 2024.

– , “Biden-Harris administration announces new plans to deliver debt relief to tens of millions of Americans,” <https://www.ed.gov/news/press-releases/biden-harris-administration-announces-new-plans-deliver-debt-relief-tens-million> 2024.

– , “Federal student aid data center: portfolio summary,” <https://studentaid.gov/data-center/student/portfolio> 2024.

U.S. Government Accountability Office, “Federal student loans: Education could improve direct loan program customer service and oversight,” Report GAO-17-45, U.S. Government Accountability Office 2016. Retrieved June 21, 2024.

– , “Federal student loans: Education could do more to help ensure borrowers are aware of repayment and forgiveness options,” Technical Report GAO-18-547 2018.

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

Table 1: Summary Statistics

	Mean	SD	Min	Max	Median
Panel A: All Borrowers					
Student Loan Balance	39539.63	51224.30	0	288108	21821
Student Loan Payment Due	76.80	181.63	0	1043	0
Student Loan Delinquency	0.03	0.16	0	1	0
Mortgage Balance	73648.35	141569.71	0	683252	0
Mortgage Payment Due	542.51	993.40	0	4590	0
Mortgage Delinquency	0.01	0.08	0	1	0
Any Mortgage	0.30	0.46	0	1	0
Auto Loan Balance	10595.06	16044.04	0	77047	0
Auto Loan Payment Due	276.41	363.80	0	1652	0
Auto Loan Delinquency	0.04	0.20	0	1	0
Any Auto Loan	0.49	0.50	0	1	0
Credit Card Balance	5628.06	9102.06	0	49281	1830
Credit Card Payment Due	158.25	234.97	0	1273	67
Credit Card Delinquency	0.10	0.29	0	1	0
Any Credit Card	0.82	0.38	0	1	1
Ever Forgiven	0.06	0.24	0	1	0
Borrowing Cohort	2011.73	5.33	1967	2019	2012
Observations	38,699,271				
Number of Individuals	992,289				
Panel B: Forgiven Borrowers					
Student Loan Balance Before Forgiveness	37451.62	52066.93	0	288108	16998
Student Loan Payment Due Before Forgiveness	59.83	143.27	0	1043	0
Student Loan Delinquency Before Forgiveness	0.05	0.22	0	1	0
Student Loan Balance	23454.78	44618.23	0	288108	4317
Student Loan Payment Due	40.61	121.19	0	1043	0
Student Loan Delinquency	0.03	0.18	0	1	0
Mortgage Balance	103278.60	161175.32	0	683252	0
Mortgage Payment Due	761.62	1119.93	0	4590	0
Mortgage Delinquency	0.01	0.08	0	1	0
Any Mortgage	0.42	0.49	0	1	0
Auto Loan Balance	11050.49	16721.07	0	77047	0
Auto Loan Payment Due	291.22	379.60	0	1652	0
Auto Loan Delinquency	0.03	0.17	0	1	0
Any Auto Loan	0.49	0.50	0	1	0
Credit Card Balance	6472.24	9797.53	0	49281	2445
Credit Card Payment Due	175.12	248.55	0	1273	78
Credit Card Delinquency	0.08	0.27	0	1	0
Any Credit Card	0.87	0.33	0	1	1
Borrowing Cohort	2009.32	5.94	1969	2019	2009
Observations	2,282,982				
Number of Individuals	58,538				

Notes: This table presents summary statistics of the main outcome variables used in the analysis. Panel A presents statistics for all borrowers and Panel B restricts to only borrowers who receive forgiveness. Balances and payments have been winsorized at the 99% level. Borrowing cohorts are defined as the earliest year student loan borrowing is observed in the data. Source: TransUnion

Table 2: Earnings Summary Statistics

	Mean	SD	Median
Panel A: All Student Loan Borrowers			
Student Loan Balance	37248	55795	20000
Borrowing Cohort	2014	5.25	2015
Ever Forgiven	0.063	0.244	0
Monthly Earnings	1838	3148	0
Monthly Earnings (Conditional on Any Earnings)	4729	3438	4053
Active Employment Reported	0.523	0.499	1
Public Service Conditional on Employment	0.329	0.470	0
Observations	57,226,857		
Individuals	1,395,777		
Panel B: Forgiven Student Loan Borrowers			
Student Loan Balance	21768	45718	3739
Borrowing Cohort	2011	5.85	2011
Ever Forgiven	1	0	1
Monthly Earnings	2505	3709	0
Monthly Earnings (Conditional on Any Earnings)	5658	3638	5149
Active Employment Reported	0.578	0.493	1
Public Service Conditional on Employment	0.389	0.487	0
Observations	3,643,137		
Individuals	88,857		

Notes: This table presents on student loan balances and earnings between January 2021 and March 2024 among borrowers with available earnings data. Panel A presents statistics on a random sample of 3% of all borrowers with open student loans in January 2021. Panel B limits the sample to student loan borrowers who have been identified as receiving student loan forgiveness at some point in the panel. Public service employment is defined as being employed in an industry that falls under the NAICS sectors of public administration, health care, and social assistance, or educational services. Source: Employment records.

Table 3: Forgiveness Targeting

	(1)	(2)	(3)	(4)	(5)	(6)
	All Earners		Excl. Public Service		Public Service	
Panel A: Control = Non-Forgiven Student Loan Borrowers (10% Sample)						
	702.91***	115.18***	737.21***	119.21***	589.63***	91.54***
	(12.11)	(10.29)	(19.44)	(16.72)	(11.23)	(8.61)
Observations	301,142	301,142	180,874	180,874	120,268	120,268
Panel B: Control = General Population (1% Sample)						
	503.82***	193.01***	540.10***	223.40***	370.70***	121.70***
	(9.65)	(8.55)	(15.01)	(13.41)	(8.99)	(7.49)
Observations	468,644	468,644	304,717	304,717	163,972	163,972
Panel C: Control = No College Education (1% Sample)						
	1562.68***	1240.19***	1652.43***	1334.92***	1265.66***	1006.75***
	(9.59)	(8.44)	(14.44)	(12.86)	(8.65)	(6.64)
Observations	151,305	151,305	101,775	101,775	49,530	49,530
Borrower Cohort (Panel A) and Age FE		X		X		X

Notes: This table presents estimates of β using the following OLS equation:

$$\hat{Earnings}_i = \beta Forgiveness_i + \gamma Age_i + \epsilon_i$$

where $\hat{Earnings}_i$ is estimated using the predicted values from the regression:

$$Earnings_i = \beta X_i + \mu_i$$

Age_i is a vector of age bin dummies where bins are split into 10 year intervals. $Earnings_i$ are the monthly earnings for individual i in December 2019 and $Forgiveness_i$ is an indicator for if individual i received student loan forgiveness as of time t . X_i are the fixed effects for age, borrowing cohort, gender, education level, and industry. Standard errors are robust and are presented in parentheses below each estimate. Panel A restricts to borrowers that had open student loans in January of 2021. Panel B adds a 1% sample of all borrowers with earnings data that do not have open student loans in January 2021 but have a credit record. We use sample weights to appropriately account for the 1% sample of non-student borrowers. The sample is limited to borrowers with available demographics and earnings information. Source: Employment records.

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 4: Effects on Student Loans

	Outcomes		
	(1) Balances	(2) Payments Due	(3) Delinquency
Dynamic DID	-31843.308*** (4004.998)	-64.293*** (12.088)	-0.024*** (0.004)
Two-Way FE	-32286.662*** (3267.574)	-72.704*** (7.601)	-0.029*** (0.006)
Calendar Month FE	✓	✓	✓
Individual FE	✓	✓	✓
Observations	38,699,271	38,699,271	38,699,271
Number of Individuals	992,289	992,289	992,289

Notes: This table presents estimates of the average treatment effect on the treated. The rows labeled “Dynamic DID” use De Chaisemartin and d’Haultfoeuille (2024) and rows labeled “Two-Way FE” use two way fixed-effects. Comparison borrowers are those who have yet to have forgiveness or never will have forgiveness in sample. In the dynamic DID rows group-specific average treatment effects on the treated are averaged with weights based on group size. Grouping is done at the forgiveness month level which is econometrically identical to the individual level. Standard errors are clustered by forgiveness cohort and are presented in parentheses below each estimate. Unique observations is the number of unique borrower-month observations in the data. Source: TransUnion

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 5: Effects on Credit Outcomes

	Outcomes		
	(1) Balances	(2) Payments Due	(3) Delinquency
Panel A: Mortgages			
Dynamic DID	2318.755*** (453.110)	17.468*** (2.653)	-0.000 (0.000)
Two-Way FE	2455.390*** (352.716)	18.541*** (2.067)	-0.000 (0.000)
Panel B: Auto Loans			
Dynamic DID	232.043*** (42.585)	4.804*** (1.087)	0.000 (0.001)
Two-Way FE	291.136*** (32.575)	5.808*** (0.768)	0.000 (0.000)
Panel C: Credit Cards			
Dynamic DID	222.048*** (41.486)	5.337*** (1.186)	0.004** (0.002)
Two-Way FE	189.086*** (36.551)	5.141*** (1.251)	0.003*** (0.001)
Calendar Month FE	✓	✓	✓
Individual FE	✓	✓	✓
Observations	38,699,271	38,699,271	38,699,271
Number of Individuals	992,289	992,289	992,289

Notes: This table presents estimates of the average treatment effect on the treated. The rows labeled “Dynamic DID” use De Chaisemartin and d’Haultfoeuille (2024) and rows labeled “Two-Way FE” use two way fixed-effects. Comparison borrowers are those who have yet to have forgiveness or never will have forgiveness in sample. In the dynamic DID rows group-specific average treatment effects on the treated are averaged with weights based on group size. Grouping is done at the forgiveness month level which is econometrically identical to the individual level. Standard errors are clustered by forgiveness cohort and are presented in parentheses below each estimate. Unique observations is the number of unique borrower-month observations in the data. Source: TransUnion

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 6: Long Difference IV

	(1)	(2)	(3)
Panel A: First Stage			
Servicer Leniency	0.993*** (0.007)	0.974*** (0.009)	0.964*** (0.010)
Age FE		✓	✓
Forgiveness Year FE			✓
Observations	723,983	723,983	723,983
Panel B: IV Estimates			
	Balances	Payments Due	Delinquency
Student Loans			
Ever Forgiven	-35612.402*** (4247.419)	-108.030*** (24.614)	-0.035*** (0.011)
Mortgages			
Ever Forgiven	3258.406*** (1197.088)	27.793*** (8.153)	-0.001 (0.001)
Auto Loans			
Ever Forgiven	474.646*** (132.187)	9.458*** (3.113)	0.002 (0.002)
Credit Cards			
Ever Forgiven	266.943** (119.998)	6.179* (3.385)	-0.001 (0.003)
Age FE	✓	✓	✓
Forgiveness Year FE	✓	✓	✓
Observations	723,983	723,983	723,983

Notes: This table presents IV estimates of β using a two stage least squares estimation of the following equation

$$\Delta Y_i = \beta \text{Forgiveness}_i + \lambda X_i + \epsilon_i$$

where the first stage regression is

$$\text{Forgiveness}_i = \gamma \text{Leniency}_i + \pi X_i + \mu_i$$

where ΔY_i is the change in outcome for individual i over the six months before and after forgiveness is received. Forgiveness_i is an indicator for if individual i ever receives forgiveness. Leniency_i is the share of borrowers that received forgiveness under individual i 's servicer. X_i are the set of controls indicated below each column. Only borrowers which are observed over the six months before and after forgiveness are included. Borrowers who never received forgiveness are randomly assigned a placebo forgiveness date which is used to calculate ΔY_i . Standard errors are clustered at the servicer level. Source: TransUnion

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 7: Effects on Earnings and Employment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Earnings	Log Earnings	Any Employment	New Employment	New Industry	Public Service	Hours Worked
Dynamic DID	-44.28*** (11.88)	-0.023*** (0.006)	-0.004*** (0.001)	0.0005 (0.0003)	0.0005** (0.0002)	-0.002** (0.001)	-1.02*** (0.466)
Two-Way FE	-43.31*** (4.75)	-0.023*** (0.002)	-0.006*** (0.001)	0.0002 (0.0002)	0.0001 (0.0002)	-0.004*** (0.0009)	-1.39*** (0.22)
Calendar Month FE	✓	✓	✓	✓	✓	✓	✓
Forgiveness Cohort FE	✓	✓	✓	✓	✓	✓	✓
Observations	22,245,614	22,245,614	57,226,857	57,226,857	57,226,857	57,226,857	12,666,655
Period -1 Average	5809.94	8.30	0.61	0.007	0.004	0.233	138.72

Notes: This table presents estimates of the average treatment effect on the treated. The rows labeled “Dynamic DID” use De Chaisemartin and d’Haultfoeuille (2024) and rows labeled “Two-Way FE” use two way fixed-effects. Comparison borrowers are those who have yet to have forgiveness or never will have forgiveness in sample. In the dynamic DID rows group-specific average treatment effects on the treated are averaged with weights based on group size. Grouping is done at the forgiveness month level which is econometrically identical to the individual level. Standard errors are clustered by forgiveness cohort and are presented in parentheses below each estimate. Effects are estimated over a 6-month horizon post-forgiveness. Public service employment is defined as being employed in an industry that falls under the NAICS sectors of public administration, health care, and social assistance, or educational services. Source: Employment records.

* $p < .1$, ** $p < .05$, *** $p < .01$

Table 8: Earnings Effects Across Worker Heterogeneity

	Hourly Worker (1)	Salary Worker (2)	Public Service (3)	Non Public Service (5)	Never Defaulters (4)	Ever Defaulters (6)
Dynamic DID	-73.77** (27.04)	-20.61** (11.05)	-51.14*** (16.85)	-31.90** (14.56)	-55.75*** (13.88)	63.65** (27.63)
Two-Way FE	-47.22*** (13.30)	-40.18*** (5.18)	-43.80*** (7.51)	-37.55*** (5.03)	-50.28*** (5.35)	25.46** (11.57)
Calendar Month FE	✓	✓	✓	✓	✓	✓
Forgiveness Cohort FE	✓	✓	✓	✓	✓	✓
Observations	12,666,655	9,578,959	5,336,284	16,909,330	20,249,017	1,996,597
Period -1 Average	4469.81	7015.89	5435.51	6599.34	6109.19	3909.25

Notes: Notes: This table presents estimates of the average treatment effect on the treated. The rows labeled “Dynamic DID” use De Chaisemartin and d’Haultfoeuille (2024) and rows labeled “Two-Way FE” use two way fixed-effects. Comparison borrowers are those who have yet to have forgiveness or never will have forgiveness in sample. In the dynamic DID rows group-specific average treatment effects on the treated are averaged with weights based on group size. Grouping is done at the forgiveness month level which is econometrically identical to the individual level. Standard errors are clustered by forgiveness cohort and are presented in parentheses below each estimate. Effects are estimated over a 6-month horizon post-forgiveness. Public service employment is defined as being employed in an industry that falls under the NAICS sectors of public administration, health care, and social assistance, or educational services. Source: Employment records.

* $p < .1$, ** $p < .05$, *** $p < .01$

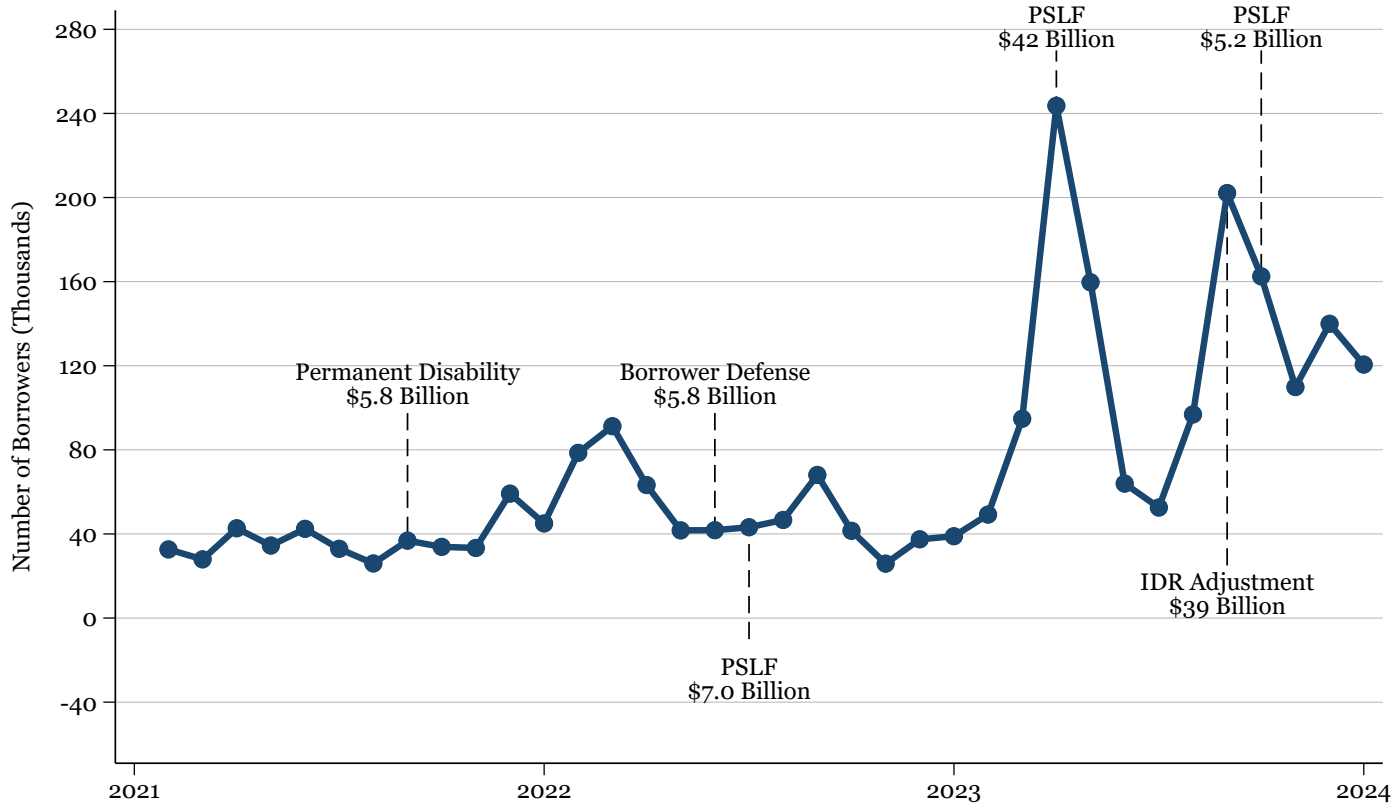
Table 9: MPX, MPC and MPE Estimates out of Student Loan Forgiveness

Panel A: MPX, MPC, and MPE				
	Reduced Form	Normalization		
		Raw Balance	Annuitized Balance	Payment Due
MPX	2772.847***	0.116***	27.352***	57.505***
(mortgage+auto+CC balances)	(474.467)	(0.029)	(8.467)	(10.016)
MPC	27.609***	0.001***	0.272***	0.573***
(mortgage+auto+CC balances)	(474.467)	(0.000)	(0.061)	(0.086)
MPE	-44.276***	-0.002***	-0.489**	-0.875***
(positive earnings)	(13.030)	(0.001)	(0.194)	(0.316)
MPE	-45.044***	-0.002***	-0.510***	-0.966***
(earnings including 0s)	(12.923)	(0.001)	(0.139)	(0.257)

Panel B: Test of PIH					
	MPC	MPE	MPT	MPC-MPE	MPC-MPE+MPT
Full sample	0.272***	-0.489**	-0.117**	0.762***	0.644***
	(0.061)	(0.194)	(0.047)	(0.204)	(0.160)
Including down payment	0.300***	-0.489**	-0.117**	0.789***	0.672***
	(0.067)	(0.194)	(0.047)	(0.206)	(0.162)
Including earnings extensive margin	0.272***	-0.510***	-0.122***	0.782***	0.660***
	(0.061)	(0.139)	(0.033)	(0.152)	(0.122)
Without expectations adjustment	0.204***	-0.367**	-0.088**	0.571***	0.483***
	(0.046)	(0.146)	(0.035)	(0.153)	(0.120)
6-months post-forgiveness	0.349	-0.994***	-0.239***	1.344***	1.105***
	(0.125)	(0.294)	(0.070)	(0.319)	(0.256)
Sample with no payments	0.315	-0.660***	-0.158***	0.974***	0.816***
	(0.069)	(0.245)	(0.059)	(0.255)	(0.199)


Notes: Panel A, Column (1) presents estimates of the average treatment effect on the treated using De Chaisemartin and d'Haultfoeuille (2024). Standard errors are clustered by forgiveness cohort and are presented in parentheses below each estimate. Panel A, Columns (2)-(4) show the ratio of the reduced-form effect, normalized by the (negative of the) first stage indicated in the column header. We calculate this ratio based on estimates from separate regressions of the average treatment effect on the treated using De Chaisemartin and d'Haultfoeuille (2024). Standard errors are calculated using a block bootstrap, where the blocks are forgiveness cohorts. The columns differ by the normalization. In each case, we divide the amount by 0.75 to reflect the change in beliefs of forgiveness, from receiving forgiveness. Panel B presents the MPC, MPE, and MPT. The “Full sample” row corresponds to the same sample as in Panel A. “Including down payment” increases consumption by 10%. “Including earnings extensive margin” uses the sample of all workers and imputes zero earnings when they leave the sample. “Without expectations adjustment” treats the forgiveness event as a complete surprise. “6-months post-forgiveness” estimates the MPC and MPE based on the response six months after the forgiveness. “Sample with no payments” uses the sample of borrowers who did not have payments due. and is the sample we use for our test of the permanent income hypothesis, where the null hypothesis is that MPC-MPE+MPT=1. Source: Transunion and Employment records. * $p < .1$, ** $p < .05$, *** $p < .01$

Figure 1: Borrowers Forgiven and Forgiveness Announcements



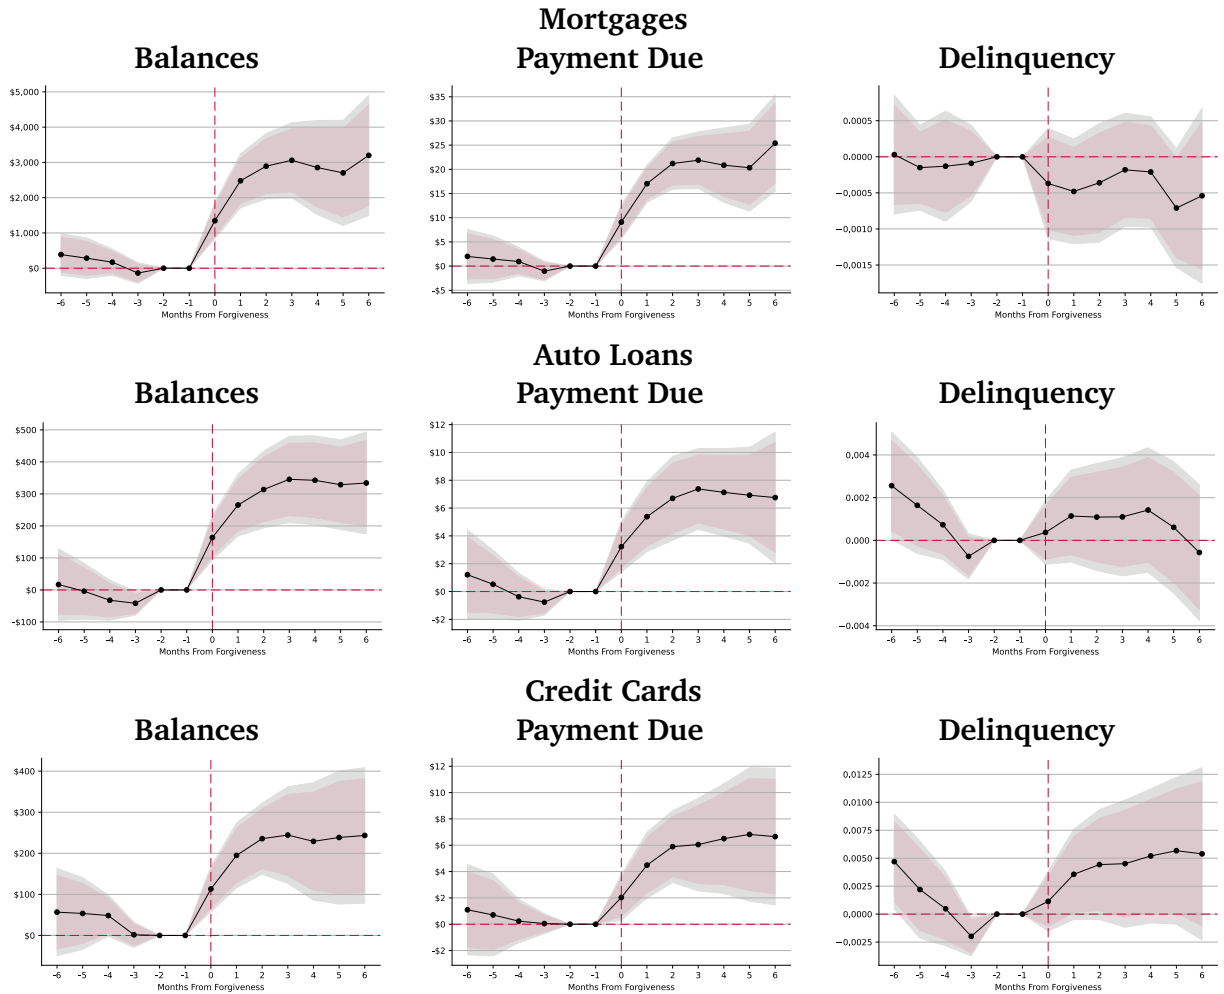
Notes: This figure plots the number of borrowers in the sample that have been identified as receiving student loan forgiveness in each month. The corresponding dashed lines mark the largest Department of Education (DOE) debt relief announcements along with the reported amount of debt to be discharged. These debt amounts represent amount of debt the DOE anticipates will be forgiven under each policy adjustment. TransUnion borrower numbers have been scaled by 10 to report a national estimate. For a more complete list of DOE debt relief announcements see Table B.1. Source: TransUnion & www.ed.gov

Figure 2: Effects on Student Loans



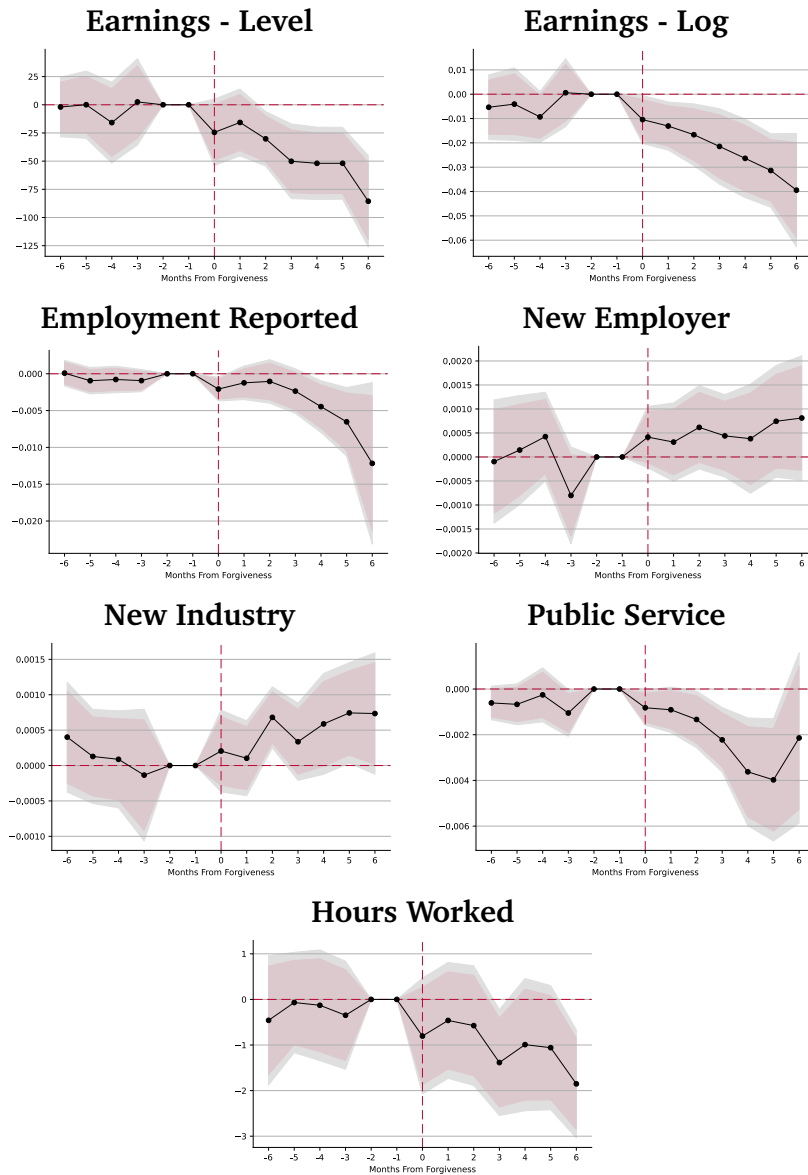
Notes: This figure plots estimates, 95% confidence intervals (grey) and 90% confidence intervals (red) of the average treatment effect on the treated in each of the six months leading up to and after forgiveness using De Chaisemartin and d'Haultfoeuille (2024). Comparison borrowers are those who have yet to have forgiveness or never will have forgiveness in sample. Forgiveness cohort-specific average treatment effects on the treated are averaged with weights based on cohort size. Standard errors are clustered by forgiveness cohort. We normalize period $T = -2$ to be 0 and omit period $T = -1$ as credit bureau reporting lags can cause the treatment event to occur in time 0 or -1. Source: TransUnion

Figure 3: Effects on Credit Outcomes



Notes: This figure plots estimates, 95% confidence intervals (grey) and 90% confidence intervals (red) of the average treatment effect on the treated in each of the six months leading up to and after forgiveness using De Chaisemartin and d’Haultfoeulle (2024). Comparison borrowers are those who have yet to have forgiveness or never will have forgiveness in sample. Forgiveness cohort-specific average treatment effects on the treated are averaged with weights based on cohort size. Standard errors are clustered by forgiveness cohort. We normalize period $T = -2$ to be 0 and omit period $T = -1$ as credit bureau reporting lags can cause the treatment event to occur in time 0 or -1. Source: TransUnion

Figure 4: Effects on Earnings and Employment



Notes: This figure plots estimates, 95% confidence intervals (grey) and 90% confidence intervals (red) of the average treatment effect on the treated in each of the six months leading up to and after forgiveness using De Chaisemartin and d'Haultfoeuille (2024). Comparison borrowers are those who have yet to have forgiveness or never will have forgiveness in sample. Forgiveness cohort-specific average treatment effects on the treated are averaged with weights based on cohort size. Standard errors are clustered by forgiveness cohort. We normalize period $T = -2$ to be 0 and omit period $T = -1$ as credit bureau reporting lags can cause the treatment event to occur in time 0 or -1. Source: Employment data