Jain Family Institute

April 2024

A First Look at Student Debt Cancellation

By Laura Beamer, Jack Landry, Eduard Nilaj, Sérgio Pinto, and Marshall Steinbaum

Contents

View JFI's series of reports on Millennial Student Debt here.

568 Broadway, Suite 601
New York, NY, 10012

Introduction

There has been a great deal of debate among policymakers, scholars, and the general public about whether or not to cancel student debt, including a Supreme Court ruling that the Department of Education lacks the unilateral power to do so on a blanket basis. Part of that debate centers on speculation about who would benefit and what the economic effect of canceling debt would be. To date, however, there has been little concrete empirical data about who actually benefits from student debt cancellation, or what effect that cancellation has on the economic status and household wealth of beneficiaries. In this publication, we fill both of those gaps.

In the past two years, the Biden administration expanded debt relief to student borrowers seeking Public Service Loan Forgiveness (PSLF) (for borrowers with 10 years of enrollment in Income-Driven Repayment working at a qualified 'public service' employer, including government agencies and nonprofits), Borrower Defense to Repayment (which refers to reasons why borrowers should not have to repay their loans, including fraudulent originations), and closed-school discharge (when the institution a borrower attended is no longer eligible to issue new loans and, consequently, the old loans issued to attend them are no longer payable). Here we report on the first comprehensive study that looks at who benefited from these programs and what effect it had on their financial lives. On the first question, we find that student debtors who had their loans canceled are worse-off prior to cancellation than borrowers who repay their debts. Hence, the assumption that across-the-board loan forgiveness is regressive, or that cancellation benefits the best-off student borrowers because they bear the most student debt, is definitively false. We also show that borrowers who repaid their loans entered repayment in the mid-to-late 2000s (i.e. before the Great Recession), while those who ultimately had their loans canceled entered repayment a few years later. Our previous work shows that this life cycle/repayment timing difference is crucial—enrollment in higher education during and after the Great Recession is associated with significantly worsened odds of repaying student debt, as well as other signs of diminished economic well-being.

We then analyze the outcomes of borrowers who had their loans canceled, relative to those who did not. This is a policy-driven quasi-experiment in which individuals' student debt was removed, but

their education level and other individual characteristics were held constant. Larger discussions of the 'effect of student debt' tend not to be informed by such conceptually clean experimental evidence, because any analysis of the 'effect of student debt' is usually confounded by either the effect of the education it pays for, or other omitted variables that predict both debt and education simultaneously (for example, family socioeconomic background). Here we have the opportunity to study outcomes in an individual-level panel dataset before and after cancellation, when the amount of student debt an individual has is varied but other determinants of financial status, like educational attainment and family background, are held constant. As far as we are aware, this is the first quasi-experimental evidence about the effect of policy-driven student debt cancellations that have been enacted to date. (Di Maggio, Kalda, and Yao also study student debt cancellation using credit reporting data, stemming from individual judgments about chain-of-title errors on the part of a private sector holder of student debt. They find the affected borrowers took advantage of the financial windfall to pay down other debts, move, and take new jobs, leading to a permanent increase in income.)

We find that individuals who had their student debt canceled between 2020 and 2023 were able to take out secured loans signifying asset ownership, including home mortgages and auto loans. An implication of our findings is that unsecured student loans preclude asset accumulation and social mobility, consistent with the aforementioned study. This is significant for several reasons: first, it implies that operating a higher education system financed by individually-held unsecured loans incurs significant costs to borrowers as well as to the economy overall, crowding out asset purchases until later in life—after the unsecured debt has either been paid off or written off by the lender. This result is not consistent with the view that unsecured student loans finance higher education, which "pays off" in the labor market and, in turn, contributes to asset ownership and social mobility.

Second, the individuals in our panel dataset who benefited from loan cancellation also benefited from the student loan repayment pause that was in effect from 2020–2023. Moreover, the vast majority of them seem to have been enrolled in Income-Driven Repayment (IDR) programs for years prior to that (a precondition for cancellation under PSLF). Yet the increased asset purchases we document did not occur until after outright loan cancellation, as opposed to when the repayment pause took effect and the borrowers' monthly payments

dropped to zero or near-zero. This is significant since the primary alternative to alleviating the burden of student debt debated in policy circles and thus far enacted (aside from the batch cancellations we study) emphasizes the reduction in monthly payments through IDR, not principal write-downs. Our previous findings show that IDR is indeed effective at reducing delinquency and other forms of credit distress (when borrowers manage to navigate the enrollment process successfully), but it is not a substitute for cancellation when it comes to building wealth.

The programs that have been adopted to date to alleviate student debt are qualitatively different from the proposal for wide-scale cancellation that the Supreme Court struck down in 2023. As the Biden-Harris administration re-attempts a broader relief plan via negotiated rulemaking, these results provide the first empirical evidence both about who is likely to benefit and what the economic effect of such a program would be.

This report proceeds as follows: we first explain our data sources and our methodology for determining which borrowers, among those with zero balances, had their student loans canceled. We then report summary statistics about those borrowers and their financial status, compared to borrowers who achieve zero balances by other means. Finally, we report the results of an event study analysis of loan cancellation, using our panel of borrowers.

I. How We Classify Zero Balance Loans and Why

We used two proprietary credit bureau datasets for this study. The first is a repeated cross-section of one million student borrowers, aged 18–34, sampled annually at the end of June from 2009 through 2023. The second is a panel of one million student borrowers, aged 18–34, first sampled in 2009 and then annually through 2023, regardless of whether they continue to have outstanding student debt. In both the cross-section and panel datasets, for each archive from 2020–2023, we first classify all the student loans held by sampled borrowers as either zero-balance as of the archive date (June 30 of each calendar year), or positive-balance. We then use monthly data on loan balances over the previous two years leading up to the archive date, as well as loan-level

status indicators, to determine why each zero-balance loan was zeroed out. The categories of zero-balance loans that we assign are:

- Discharged in bankruptcy/charged off by the creditor.
- Repaid.
- Canceled under one of the federal cancellation programs (PSLF, Defense to Repayment, or closed-school discharge).

Before this assignment, we assess whether the zero-balance loans we observe were transferred to another servicer, refinanced, or consolidated into another student loan. We expect that these loans are "replaced" by another loan in the borrower's credit record. Hence, for our purposes, they are not "true zeroes." Many student loans in this category had to be identified using a series of data science techniques, since they may look similar to zeroed-out loans where the final payment was a waterfall payment (in which case, the loan should be properly categorized as a zero-balance loan that was repaid). Once we determine that a zero-balance loan was transferred, refinanced, or consolidated, we remove it from further analysis.

Second, we used the loan-level status codes to identify zero-balance loans fully or partially discharged in bankruptcy and loans subject to creditor charge-offs. The loan-level data includes, for example, flags that read "Bankruptcy Chapter 13 Petitioned, or Discharged and Bankruptcy Chapter 13 Reaffirmation of Debt Rescinded." We include charge-offs—loans written off as a financial loss by the creditor, as opposed to expunged by a court—in the same category as loans discharged in bankruptcy. While charged-off loans are legally different from loans discharged through bankruptcy, with different implications for the borrower's credit record, the financial characteristics of the loans and borrowers with bankrupt and charged-off loans are similar enough that we combine them into the same category for this study's purpose.

Third, we categorize a zero-balance loan as repaid under four circumstances:

1) if its balance history indicates a declining balance in the two years prior to reaching zero-balance status. We conduct this test at both the loan level and the borrower level, across all of a borrower's outstanding student loans;
2) the loan and borrower have a stagnant balance history over the years preceding the sampling date, and the pre-zero balance is

less than or equal to the loan's origination amount (indicating that repayment has taken place);

3) there are no flags on the loan indicating non-payment and the borrower has been in repayment for less than five years, regardless of balance history, or
4) the borrower made significant payments on the loan, amounting to half or more of the total outstanding balance at the start of the two-year lookback period. We permit significant divergence between the observed payments and the decline in balance due to sparseness in the actual-payments data.

Finally, we assign the canceled category. The COVID repayment pause complicates assignment to the cancellation category, because payments were voluntary and interest was set to zero for most federal loans starting in 2020. For example, borrowers who had their loans canceled through PSLF in 2023 were likely eligible for the repayment pause in the preceding years. Hence, their loan balance may not have been increasing, since monthly interest was no longer due. We assigned zero-balance loans to the canceled category under four circumstances:

1) that loan's balance was increasing leading up to the sampling date, the borrower had an overall increasing balance across all of their active student loans, and the loan was flagged as distressed; or
2) same conditions as (1) except, in the absence of a flag indicating distress, the pre-zero-balance amount on the loan was higher than the loan's origination amount. By pre-zero-balance amount, we mean the last-observed positive balance during the two-year lookback, prior to observing a zero balance.
3) same as (1), except in the sampling years of 2022 and 2023, the loan's balance was stagnant leading up to the sampling date, as opposed to increasing, and the loan was flagged as distressed prior to the pause;
4) same conditions as (3) except, in the absence of a flag, the pre-zero-balance amount on the loan was higher than the loan's origination amount.

Further, for borrowers with multiple loans reaching zero balance in a single sampling year, every loan had to reach a zero-balance status in the same month for any of the loans to be classified as canceled, with their newest loan not less than five years into repayment. Lastly, borrowers observed as having a zero-balance loan but whose current

balance on all student loans was equal to or exceeded the total from the previous year were not eligible to have their zero-balance loans categorized as repaid or canceled. We did this to account for zero-balance loans that were likely to have been refinanced.

According to these definitions, the zero-balance categories are neither mutually exclusive nor exhaustive. In a minority of cases, ambiguity remains around how each loan reached its zero balance. Loans that do not meet any conditions for the aforementioned categories—about 25 and 30% of zero-balance loans in our cross-section and panel analysis, respectively—were placed in a fourth zero-balance category, "unknown." Loans that meet conditions for multiple categories—roughly 0.5 percent of the analyzed zero-balance loans—are subject to a categorization hierarchy: the discharged categorization trumps the repaid categorization which trumps the cancellation categorization. For the loans that are classified as 'unknown,' we believe they are probably transferred or refinanced, but we were unable to identify a 'replacement' loan in the borrower's credit record. The total import of all of these choices is a conservative categorization of loans as canceled, indicating a lower prevalence of cancellation in our panel relative to what aggregate statistics released by the Department of Education would imply.

II. Economic and Demographic Characteristics of Zero-Balance Borrowers

Over the last four years, the Department of Education has discharged $136.6 billion of federal student debt for over 3.7 million borrowers—a significant number, but still under 10% of student borrowers. What we know about who benefited has thus far been limited to aggregated data, news released through the Department, and anecdotal experiences shared in the press. This section seeks to understand much more about this group of borrowers by comparing them to peers who otherwise reached a zero dollar student loan balance through repayment or bankruptcy discharge.

To begin, Figure 2.1 visualizes the share of student loans and borrowers in our samples who have reached a zero balance in the

previous 12 months. While there was a decrease in loan balances reaching zero during the pandemic repayment pause (because principal payments were optional), it was slight.

Annual Share of Loans and Borrowers with Balances Reduced to $0, 2020 to 2023

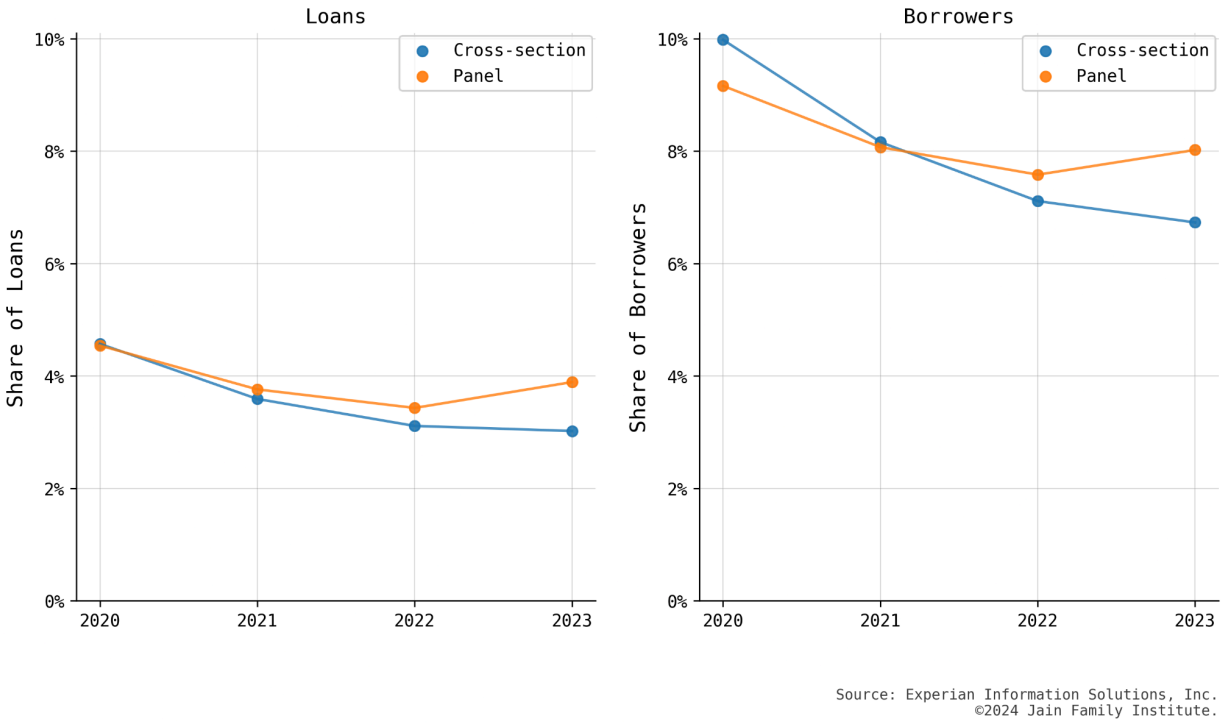

Source: Experian Information Solutions, Inc.
©2024 Jain Family Institute.

Figure 2.1 *- Share of each annual credit archive to have reached a $0 balance in the 12 months prior to the archive date (June 30 of each calendar year).*

Out of the loans/borrowers that reached a zero dollar balance during this period, over 90% of those in the cross-section—and over 85% of those in the panel—did so via repayment. In both the cross-section and the panel, the share of zero balance loans categorized as canceled generally increases year-over-year, with a clustering of canceled debt in 2022 and 2023, consistent with the aggregate statistics released by the Department of Education. The reason is a combination of declining repayments during the pause, when they weren't mandatory, and the ramping up of mostly PSLF-derived cancellations over the course of the Biden administration (both by expediting the processing of applications for PSLF and crediting borrowers for time in repayment, rectifying past servicing errors). We observe a larger share of canceled loans in the panel compared to the cross-section, because the panel consists of older

borrowers who have been in repayment for longer and hence are more likely to be eligible for the cancellations that have so far been enacted.

One inconsistency between our categorization and the aggregate statistics reported by the Department of Education is that we classify a significantly smaller share of observed zero-balance loans as canceled than the aggregate statistics would imply. This may be because of the age cap in place in both of our datasets: the cross-sections are capped at 35, and the panel was capped at 34 when first sampled in 2009, meaning the effective age cap is now 48. If older borrowers—which we cannot observe in our data—are more likely to have had their loans canceled, that would contribute to finding a lower share of canceled loans. But the more important reason is probably that our methodology for determining that a zero-balance loan was canceled is conservative. Specifically, we require that the balance on the loan was either increasing prior to cancellation or that its pre-cancellation balance was higher than its origination amount. In reality, it's perfectly possible for a loan on which significant repayment had occurred to have nonetheless been canceled. The upshot is that our analysis of cancellation should be taken as pertaining to the subset of borrowers for whom we can be most certain that their zero-balance loan was canceled as opposed to repaid, and there are probably additional borrowers whose loans we classify as repaid but which were actually canceled.

Breakdown of Zero-Balance Loan Categories in Cross-Section Data, 2020 to 2023

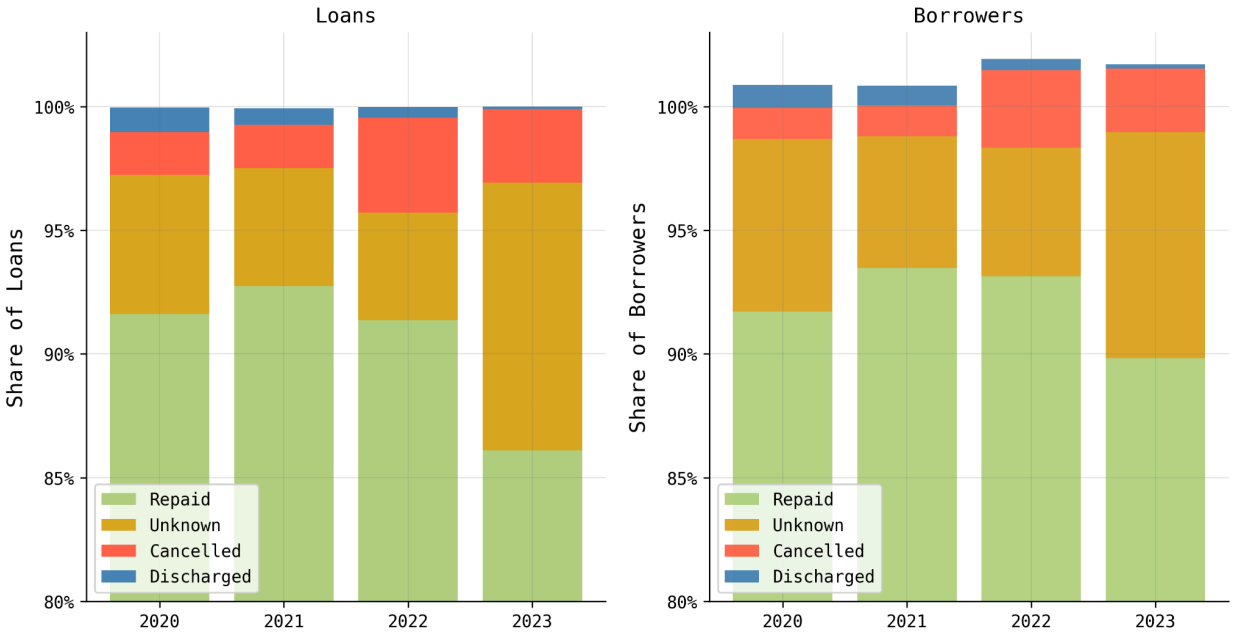

Source: Experian Information Solutions, Inc.

Figure 2.2 *- Apportioning the zero-balance loans to repaid, canceled, and discharged categories in the cross-section data. The combined total of zero balance borrowers exceeds 100% because borrower groupings are not exclusive—a borrower is counted in multiple groups if they have multiple zero-balance loans that were categorized differently. In order to make that subtlety visually apparent, the scale of the vertical axis is lengthened above 100%.*

Breakdown of Zero-Balance Loan Categories in Panel Data, 2020 to 2023

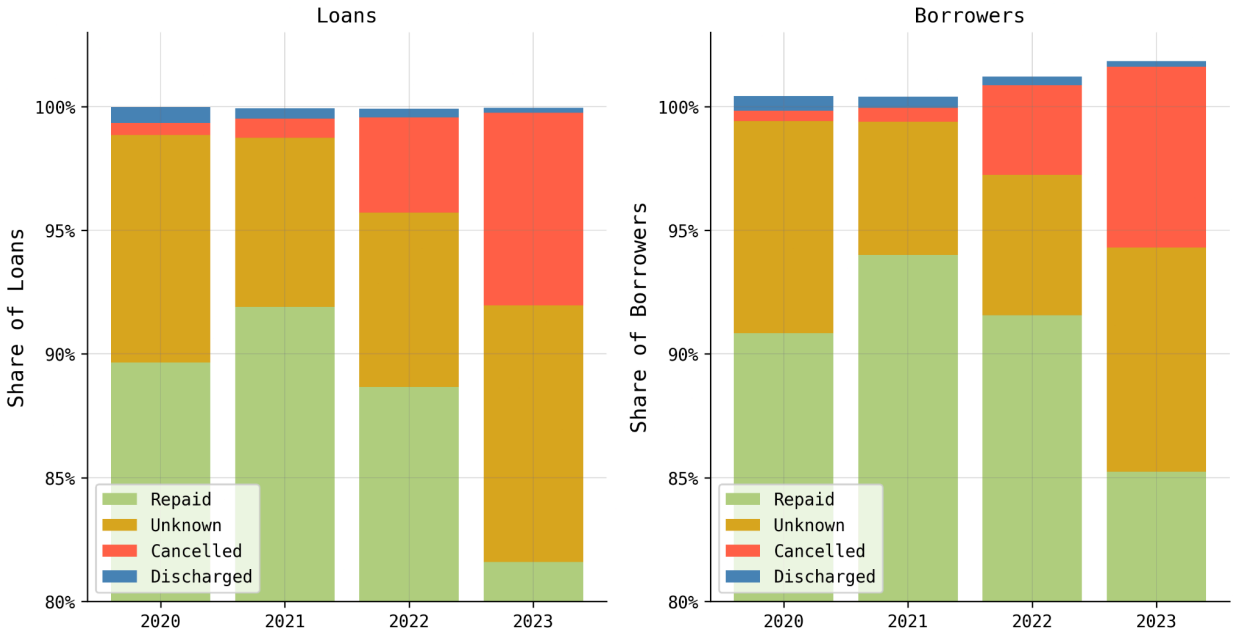

Source: Experian Information Solutions, Inc.

Figure 2.3 *- Apportioning the zero-balance loans to repaid, canceled, and discharged categories in the panel data. The combined total of zero balance borrowers exceeds 100% because borrower groupings are not exclusive—a borrower is counted in multiple groups if they have multiple zero-balance loans that were categorized differently. In order to make that subtlety visually apparent, the scale of the vertical axis is lengthened above 100%.*

Next, we group together the 2020–2023 cross-section and the 2020–2023 panel data to report the characteristics of borrowers who attained a zero balance on a student loan within that period. We report borrower characteristics by age, repayment cohort, income, gender, and race. All of this information about borrowers is available from their credit reports except for race. We use the borrower's residential Census tract racial composition (from the American Community Survey) to form an inference about the borrower's race.

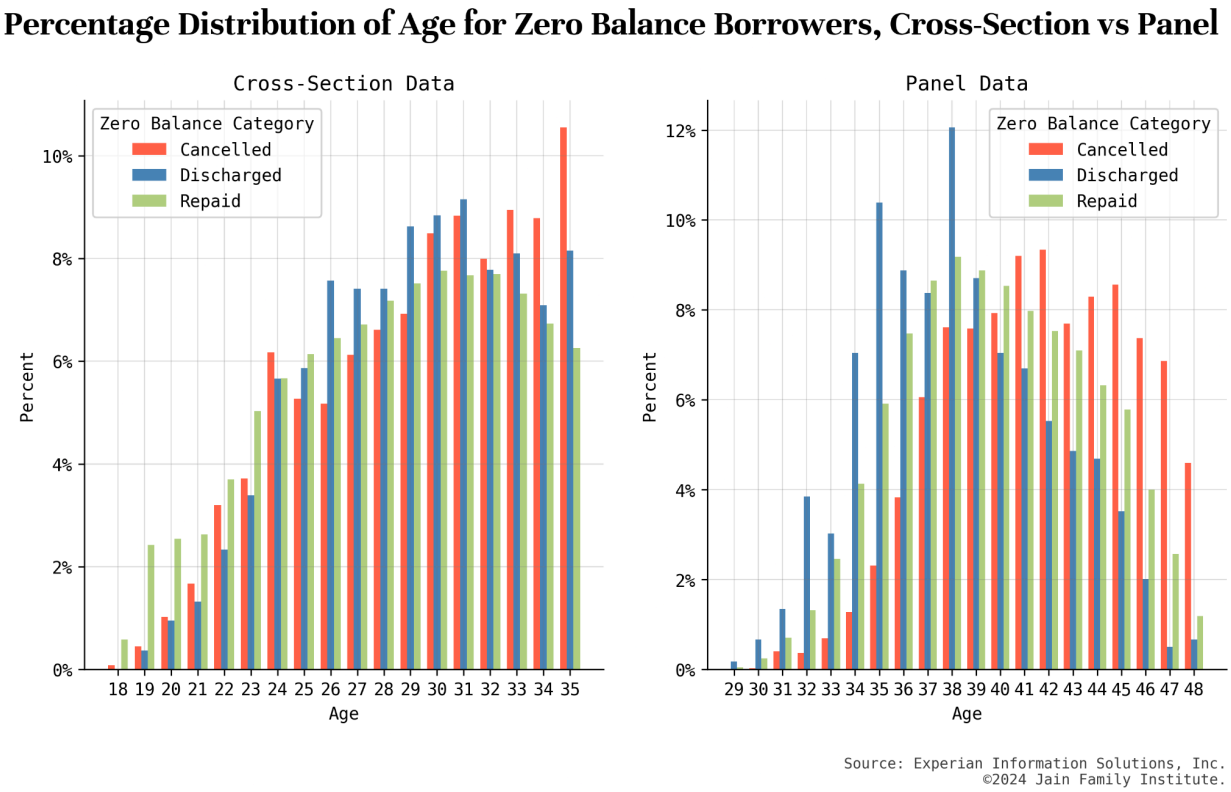

Figure 2.4 *- Age histograms for borrowers in each of the three zero-balance categories. Note that each category sums to 100% to visualize the within-category age distribution most clearly.*

Figure 2.4 plots the age distribution within each zero-balance category. The borrowers who received cancellation are generally older than those who repaid their loans in full or received a bankruptcy discharge. In the cross-section (which is age-capped at 35), the majority of these borrowers were over the age of 30. In the panel, canceled borrowers are older than those who repay.

We next report the distribution of zero-balance loans by borrower repayment cohorts, which is the year after the borrower's most recent original-issuance loan was originated. The vast majority of borrowers who received cancellation were those that entered repayment following the Great Recession, through 2016. This is a function of the changes to the student lending system around that time—2010 marked the point in which all federal student loans were issued directly by the Department of Education, which means they are more likely to eventually be eligible for cancellation. In the panel data, the borrowers who repaid, by contrast, are part of the mid-to-late 2000s repayment

cohorts, who are more likely to have gotten their start in the labor market before the onset of the Great Recession. This pattern is strikingly consistent with our study "The Long-Run Impact of the Great Recession on Student Debt," which showed that exposure to the Great Recession and its aftermath increased the probability of re-enrolling in higher education. This was due to labor market credentialization, reduced state funding for public universities, and institutional segregation, with dire consequences for eventual non-repayment. In this study, we show that those who took out loans after the Great Recession and subsequently could not repay them are disproportionately represented among the borrowers who had their loans canceled in the last several years.

Distribution of Repayment Cohorts for Zero Balance Borrowers, Cross-Section vs Panel

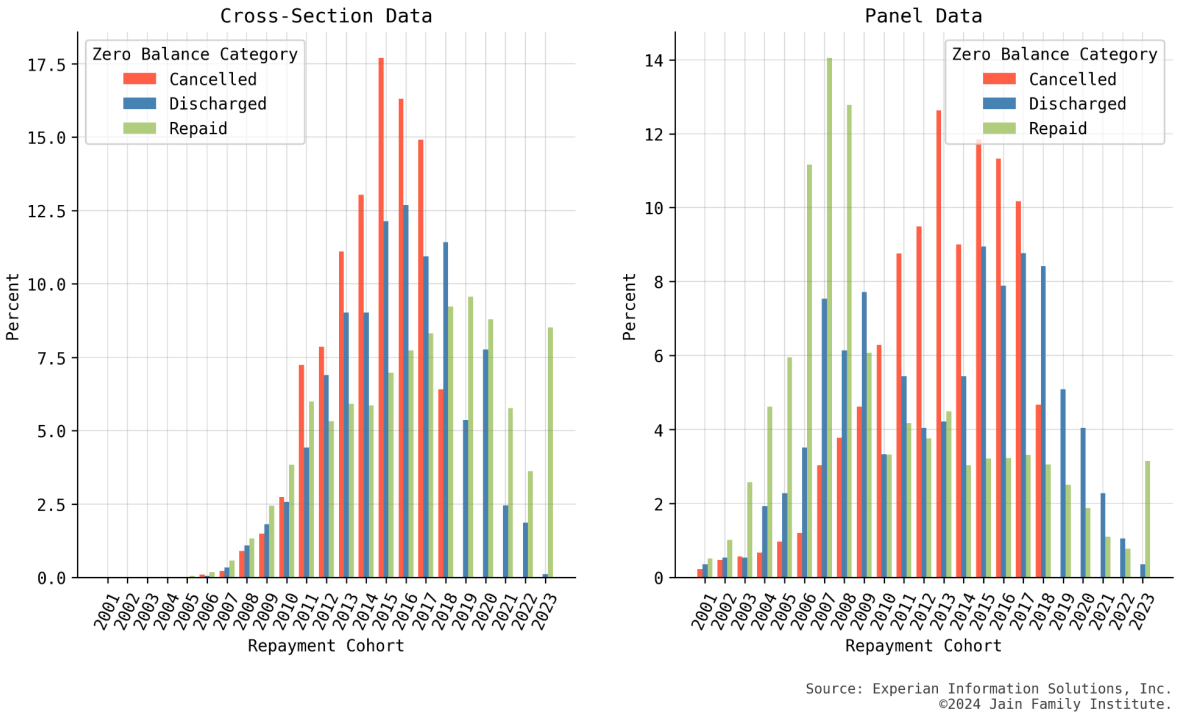

Figure 2.5 *- Repayment cohort histograms for borrowers in each of the three zero-balance categories. A borrower's repayment cohort is one year after their youngest original-issuance student loan was originated. Note that each category sums to 100% to visualize the within-category repayment cohort distribution most clearly.*

Comparing the income distributions of borrowers in each category reveals several interesting patterns. The overall takeaway is that borrowers who repay their loans have higher incomes than both the

cancellation and the discharge groups. On one level, this is unsurprising: higher-income borrowers can more easily repay their loans, so they do. The other groups can't, so they don't. Yet the public debate over student debt cancellation has broadly taken the view that cancellation benefits the best-off borrowers, because the best-off borrowers have the highest loan amounts. That idea has risen nearly to the level of conventional wisdom on the op-ed circuit, based on the prior belief that if education causes your income to go up, so must student debt. In fact, the higher your income, the more debt you're able to repay, so the less you benefit from the cancellation of un-repaid debt. Our previous work about the distribution of student debt in the population, including the income distribution of the people who don't have student loans, reveals that the better-off you are, the less you need to borrow, and the more you're able to repay. Given that, our finding that the people who don't repay and have their loans canceled are worse off in the pre-cancellation period than the people who do repay is unsurprising. Both groups have lower incomes than the people who didn't need student debt to finance their education in the first place.

The second pattern revealed in these income distributions for the zero-balance categories is that the character of the canceled group changed between 2020–2021 and 2022–2023. Namely, in the earlier period the borrowers who had their loans canceled were more similar to the borrowers who had their loans discharged, which is to say, low income. By contrast, in the later period, the cancellation group has incomes between the discharged group and the repaid group. The reason for the shift is probably the relative importance of the different cancellation programs over time. In the earlier period, cancellation arose primarily from closed-school discharge and Defense to Repayment, both of which would have been focused on students who had attended for-profit universities. The student populations of those institutions are disproportionately low-income compared to the overall higher education population. In the later period, the predominant source of cancellation would have been Public Service Loan Forgiveness, which is more focused on borrowers who attended traditional not-for-profit universities, who may have graduate or professional degrees, and who have worked for longer. That group has higher incomes, but still not as high as the borrowers who repaid their loans.

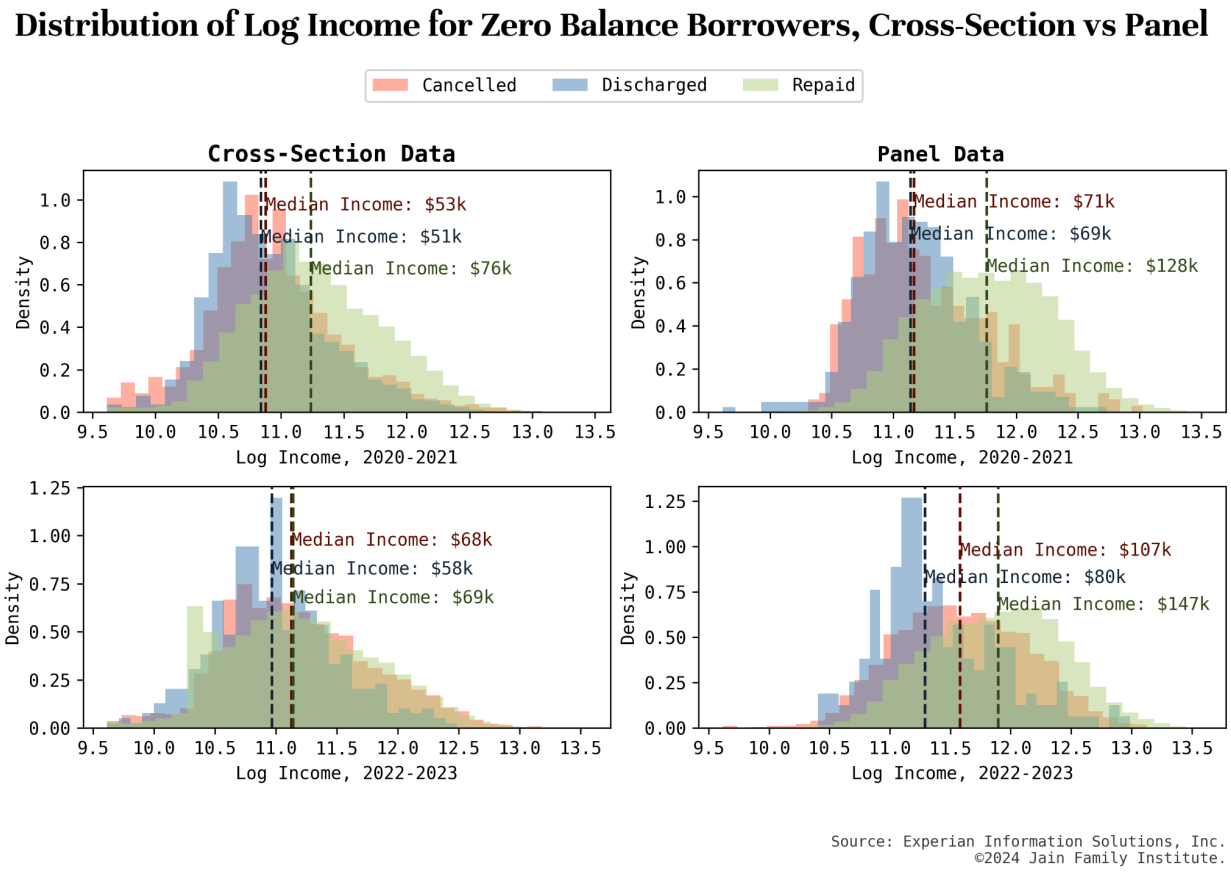

Figure 2.6 *- Income histograms for borrowers in each of the three zero-balance categories. In this case, we divide the sample(s) into 2020–2021 and 2022–2023 to visualize how the canceled borrower distribution has shifted as the cancellations that have been enacted shifted from closed-school discharge and Defense to Repayment to Public Service Loan Forgiveness.*

Finally, we report the breakdown by gender and by race. The patterns are unsurprising given what we know about who has student debt and why they take it on: women and racial minorities acquire more degrees (conditional on job or salary) to counteract labor market discrimination, and in the case of racial minorities, are segregated into higher-priced institutions with less generous financial aid policies. Black people and women are also more likely to work in the public sector, which both has greater credential requirements for a given annual salary and qualifies borrowers for PSLF, channeling them into eventual cancellation. Meanwhile, men and white people are better able to get a job that pays sufficiently to commence repayment on student loans directly after obtaining an undergraduate degree,

meaning that white men account for a disproportionate share of the repaid category of zero-balance borrowers.

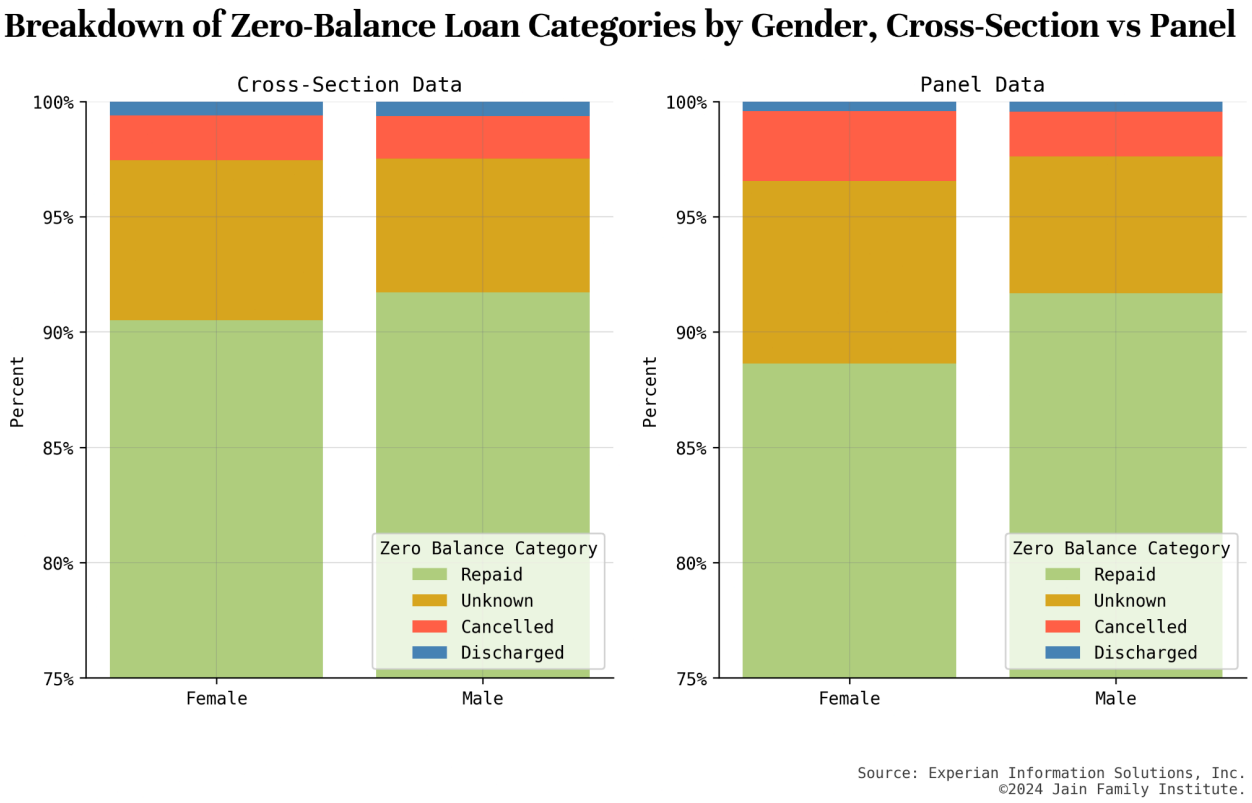

Figure 2.7 *- Zero-balance categories by gender. Since gender is assigned in the credit reporting data based on the borrower's name, the categorization is inexact; about seven percent of zero-balance borrowers in the cross-sectional samples and 4.5 percent in the panel samples were assigned to neither the Female nor Male category.*

Breakdown of Zero-Balance Loan Categories by Race, Cross-Section vs Panel

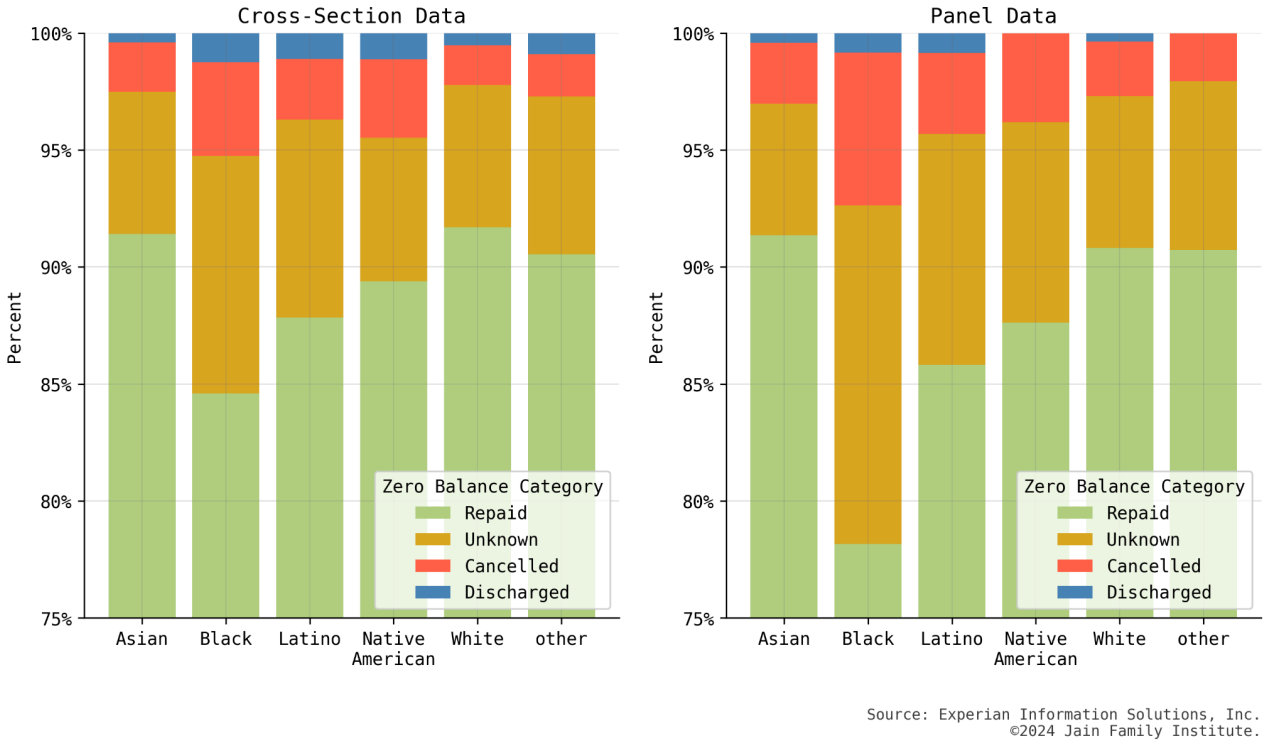

Figure 2.8 *- Zero-balance categories by race. A borrower's race is imputed from their residential location, using census tract-level race compositions from the American Community Survey.*

III. The Impact of Canceling Student Debt

We perform a causal evaluation of student debt cancellation using just the panel dataset, which permits tracking a given borrower before and after we observe cancellation, and implementing an event-study specification that compares borrowers whose loans are canceled to a control group consisting of student borrowers whose loans are not canceled.

Specifically, we estimate a model of the following form:

$$y_{it} = \sum_{a=-n}^{N} \beta_a Cancellation_{i\tau} * post_t + \gamma_i + X_{it} + \epsilon_{it} \quad (1)$$

Where y_{it} is the outcome of interest for individual i at time t, $cancellation_{i\tau} * post_t$ indicates whether period t post-dates individual i'sloan cancellation in cohort τ. The cancellation cohorts are 2020, 2021, 2022, and 2023. In the small number of cases where the same individual has loans canceled in multiple cohorts, we assign the individual to the earliest cohort in which they have a loan canceled. γ_i represents individual fixed effects and X_{it} signifies individual-level, time-varying covariates.

In the implementation, we include one set of time-varying fixed effects: we interact the borrower's 2019 total student loan payment decile (computed across all borrowers in the treatment and control groups) with calendar time. The reason to include the 2019 payment-decile-by-calendar-time fixed effects is to ensure that borrowers with canceled loans are compared to borrowers at similar points in their economic (and natural) life cycles. Borrowers with and without canceled loans are likely to be on very different life cycle trajectories, depending on whether they enrolled in higher education during or after the Great Recession and, subsequently, their earnings in the labor market, as well as how much they had to rely on student loans to finance their education in the first place. If we simply compared borrowers whose loans were canceled between 2020–2023 to other borrowers in terms of the outcomes we study, the "effect of cancellation" we estimate would be dominated by those underlying differences between the two populations, and the parallel trends

assumption necessary for causal inference from a staggered difference-in-differences methodology likely would not hold. With that set of fixed effects, our treatment effect estimates are within cells of borrowers who had similar ex ante payment obligations, and whom we interpret to be comparable in terms of the pre-treatment burden of student debt. This payment decile is interacted with calendar time because there are likely to be specific time trends for borrowers at different stages of their economic life cycles.

β_a is the coefficient of interest, signifying the difference between borrowers whose loans were canceled and those who continued to make payments on outstanding loans, a periods before or after cancellation is observed. A causal interpretation of β_a requires parallel pre-treatment trends between the treatment and control group, and that the treatment has no effect on control units ("no spillovers").

The outcomes we consider are:

- Homeownership: whether individual i is observed to have a home mortgage at time t.
- "New" or "first-time" homeownership: whether individual i has a mortgage at time t and is not ever observed to have had one on their credit record prior to time t.
- Auto loan balance: total balance on individual $i's$ outstanding auto loans at time t.
- Credit score: individual $i's$ FICO score at time t.

We include all the borrowers who had canceled loans in 2020–2023 in our treatment group. Our control group is narrower in two respects than the not-canceled borrowers from the panel: we restrict it to those borrowers in repayment cohorts 2011 or earlier (meaning that their most recent loan had to have been originated in 2010 or earlier), and among whom we observe positive student loan payments during the pandemic repayment pause. The first restriction makes the control group more like the treatment group based on pre-2020 borrower characteristics. The second ensures that our comparison is with respect to borrowers who are to some degree constrained by their student loans during the study period, which we interpret to be the counterfactual for borrowers whose loans were canceled. The concern is that the repayment pause would mute the "effect" of cancellation, and hence attenuate the out-of-sample relevance of our findings now that the repayment pause has ended.

The table below reports summary statistics for the treatment cohorts and the control group based on 2019 observables, i.e. prior to treatment. This confirms the analysis we reported in Section II above: borrowers with canceled loans are economically worse-off than borrowers without canceled loans. The treatment group has higher loan balances, lower credit scores, and lower homeownership rates pre-treatment. We emphasize that the difference-in-differences methodology does not require the treatment and control groups to be identical in a level sense ex ante. Rather, the identifying assumption is simply that they follow parallel trends prior to treatment, so the control group constitutes a valid counterfactual for what would have happened to the treatment group in the absence of treatment.

Within the treatment group, the earlier-treated cohorts are relatively economically deprived and likely benefited from closed-school discharge and/or Defense to Repayment, since those programs are targeted at borrowers who attended for-profit institutions. The later-treated cohorts are disproportionately PSLF borrowers who are somewhat better-off, but not as well-off as the borrowers whose loans weren't canceled.

2019 Summary Statistics for Federal Cancellation Cohorts

	Cancellation Cohorts				Control Group
	2020	2021	2022	2023	
Number of Individuals	179	201	1,127	2,246	111,969
Average Age	37.65	37.86	38.15	38.36	38.30
Percentage Homeowners	20.67%	28.86%	49.42%	45.97%	58.00%
Average Student Debt Balance	$52,152	$58,749	$83,189	$98,983	$27,626
Average Credit Score	613.70	613.64	672.25	664.42	724.10

Source: Experian Information Solutions, Inc.

Table 1 *- Pre-treatment summary statistics for treatment cohorts and control group, observed in 2019..*

Figures 3.1 and 3.2 report the total student loan balance and actual payments made for the four treatment cohorts and the control group, respectively. Figure 3.1 shows that borrowers whose loans were canceled were accumulating balances pre-cancellation. After

cancellation, their balances drop to zero or near-zero and remain there, suggesting that our methodology for identifying canceled loans is sound. Figure 3.2 reveals several noteworthy patterns. First of all, even though the canceled borrowers have higher loan balances, the control group has higher payments—probably because the control group was less likely to be enrolled in IDR in the pre-treatment period. Once the repayment pause went into effect in 2020, the treatment group's payments dropped to zero or near-zero, whereas the control group's payments trend downward but are not nearly as affected by the repayment pause. This is by design: we select the control group on the basis that they were making payments during the pause, because we judge that is the better counterfactual for canceled borrowers once the pause has ended (as it now has). Finally, the significant decline in payments by the 2020 and 2021 cancellation cohorts in 2015 and 2016 likely indicates when those borrowers enrolled in IDR, which was expanded significantly in those years precisely to protect borrowers who had taken on student debt during the financial crisis and were having trouble with repayment. That pattern suggests that that population of borrowers (those whose debts were canceled in 2020 and 2021) were the onetime targets of that IDR expansion, which has been shown to improve credit outcomes while failing to alleviate student debt. The later-treated cohorts were more likely to be already enrolled in IDR when our analysis sample begins, as they would have had to have been in order to be eligible for PSLF by 2022 or 2023.

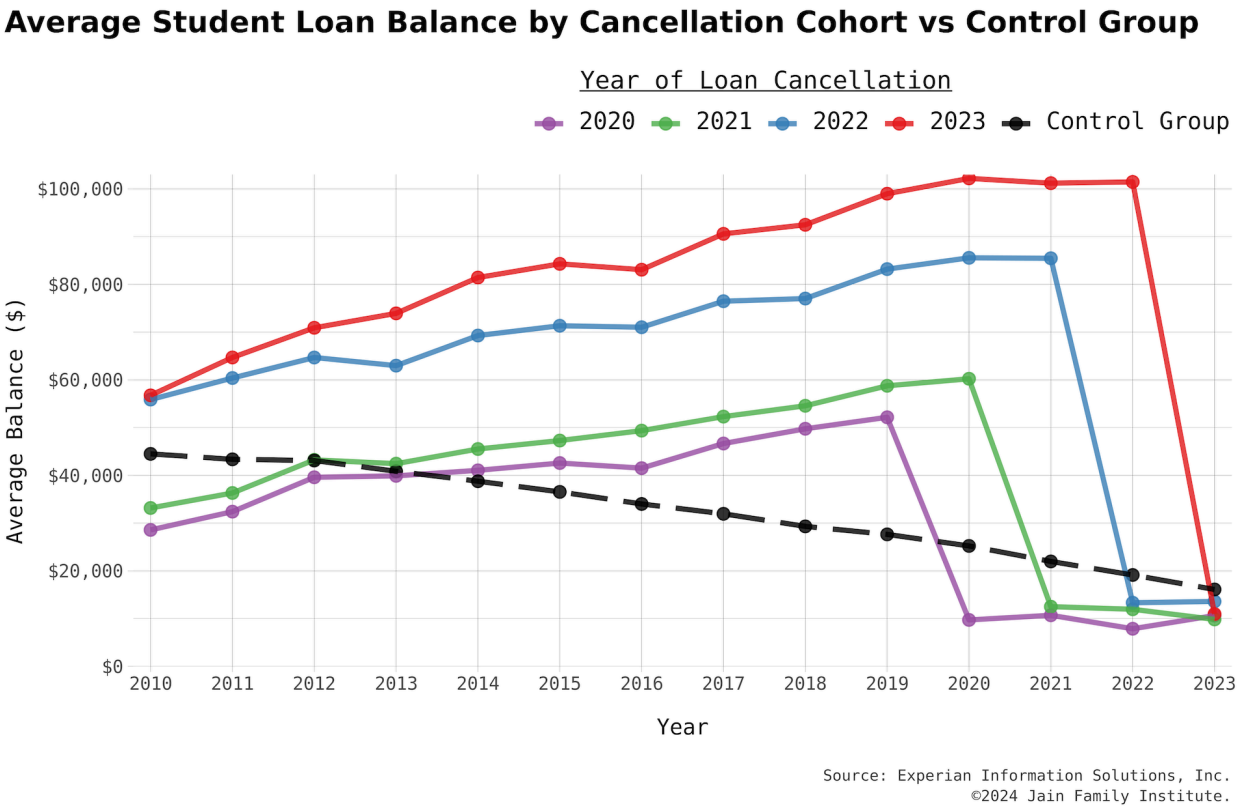

Figure 3.1 *- Average student loan balance over time for each cancellation cohort, and the control group.*

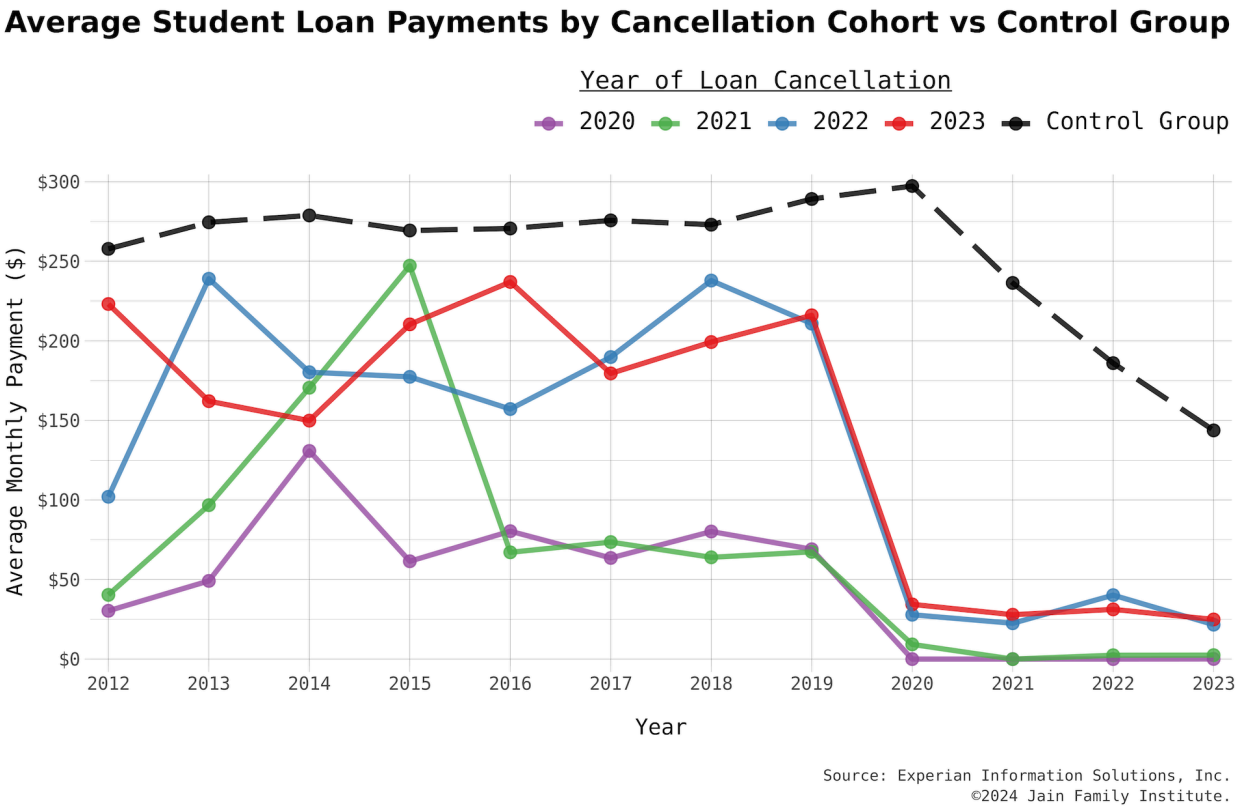

Figure 3.2 *- Total payments on student loans over time for four treatment cohorts and the control group.*

Results

We estimate equation (1) using the estimator of Sun and Abraham (2021), which is robust to staggered treatment timing with a treatment interpreted as absorbing. We report our results in the form of event study plots with five pre-treatment periods and three post-treatment ones. The 0th period in event time is the credit archive when a canceled loan is observed, and period -1 in event time is the reference period for all the estimates. Finally, we should be cautious about interpreting each relative-time coefficient, but especially the +3 coefficient (three years post-treatment), since that can only be estimated using the 2020 cancellation cohort, whereas the other coefficients are estimated off at least two different cancellation cohorts.

1. Homeownership

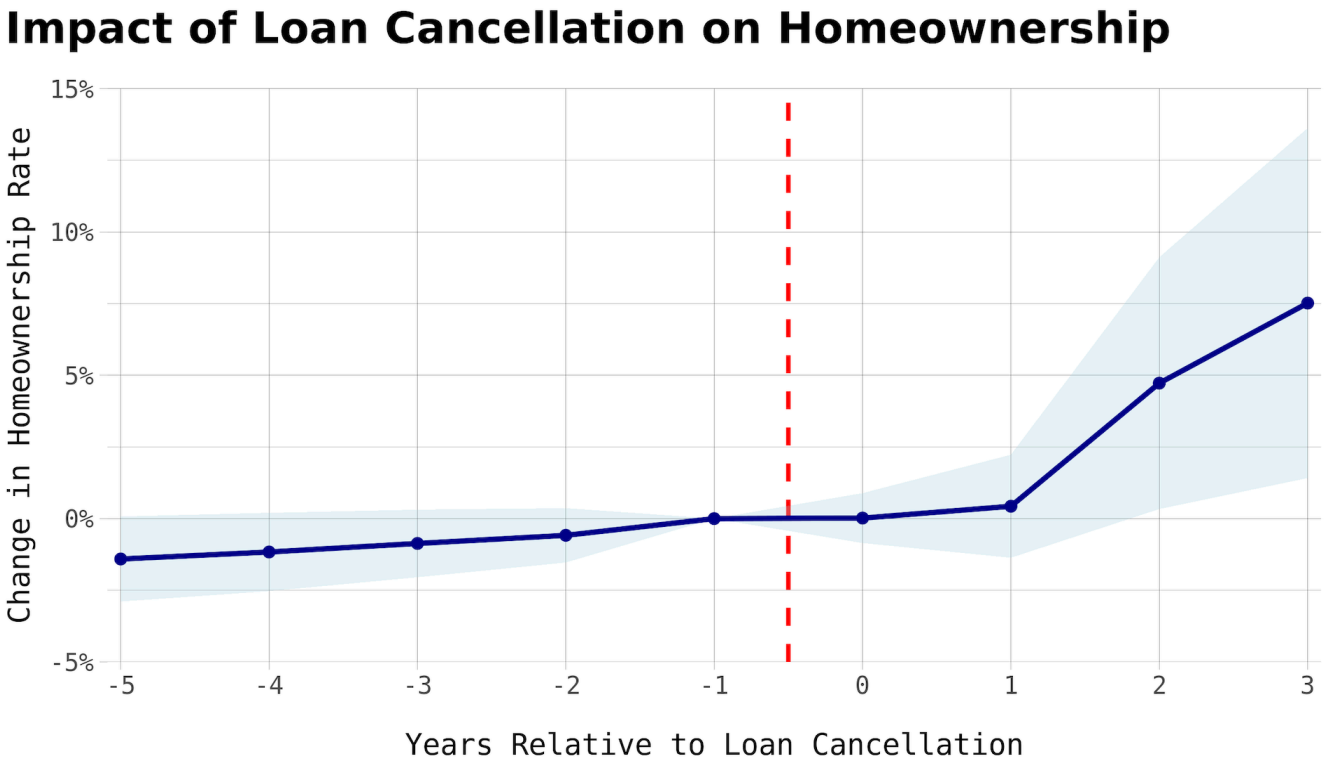

Note: Sun and Abraham estimator with individual and 2019 student loan payment decile by year fixed effects.

Source: Experian Information Solutions, Inc.

Figure 3.3 *- Event study estimate of the effect of student debt cancellation on homeownership, as measured by whether the borrower is observed to have a home mortgage on their credit report.*

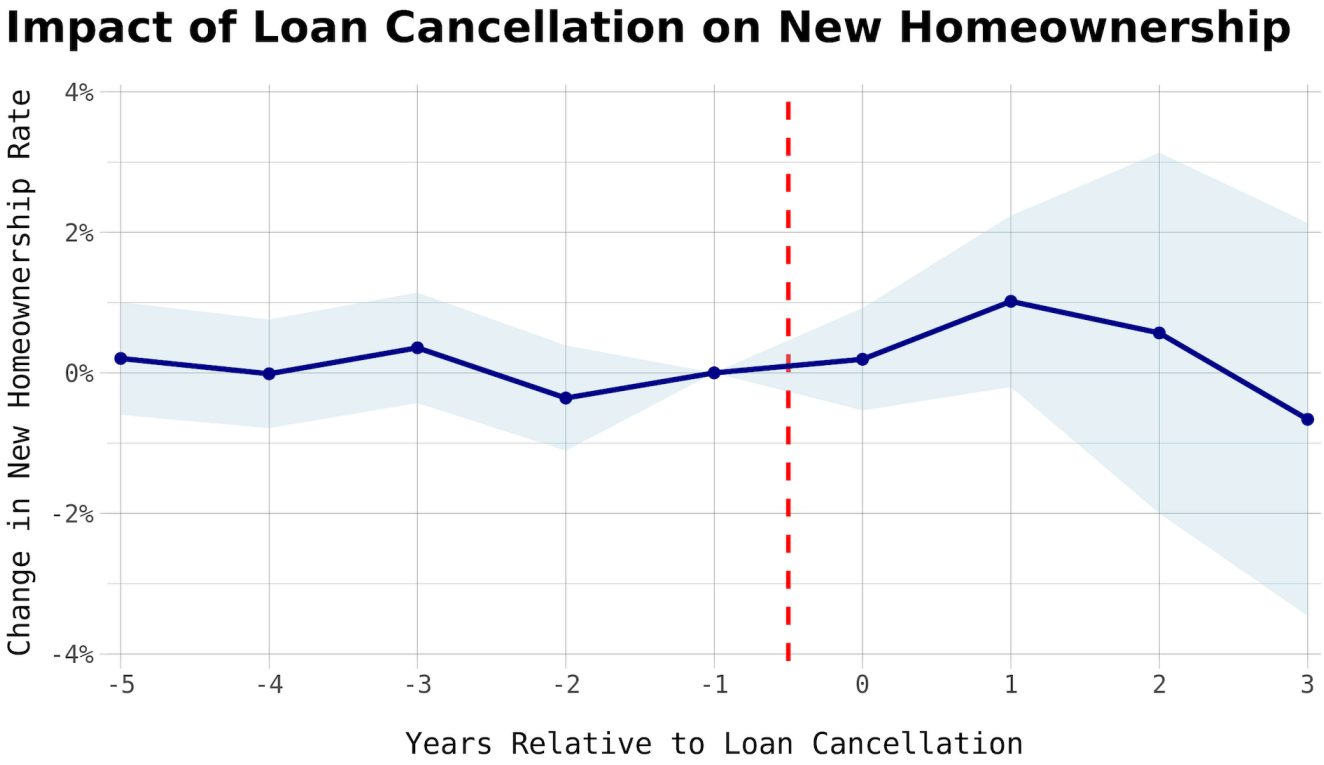

Note: Sun and Abraham estimator with individual and 2019 student loan payment decile by year fixed effects.

Source: Experian Information Solutions, Inc.

Figure 3.4 *- Event study estimate of the effect of student debt cancellation on new homeownership, as measured by whether the borrower is observed to have a home mortgage on their credit report in year t and did not have one in any previous year.*

We find that having had a student loan canceled increases the likelihood that a borrower is a homeowner by five percentage points two years post-treatment, increasing to 7.5 percentage points three years post-treatment, though we repeat our note of caution about interpreting the three-years-post estimate. The overall homeownership indicator can be interpreted loosely as cumulating the first-time-homeownership outcome year-by-year; hence, the latter shows cancellation increases the likelihood of being a first-time homeowner by just under one percentage point within one to three years of treatment.

2. Total auto loan balance

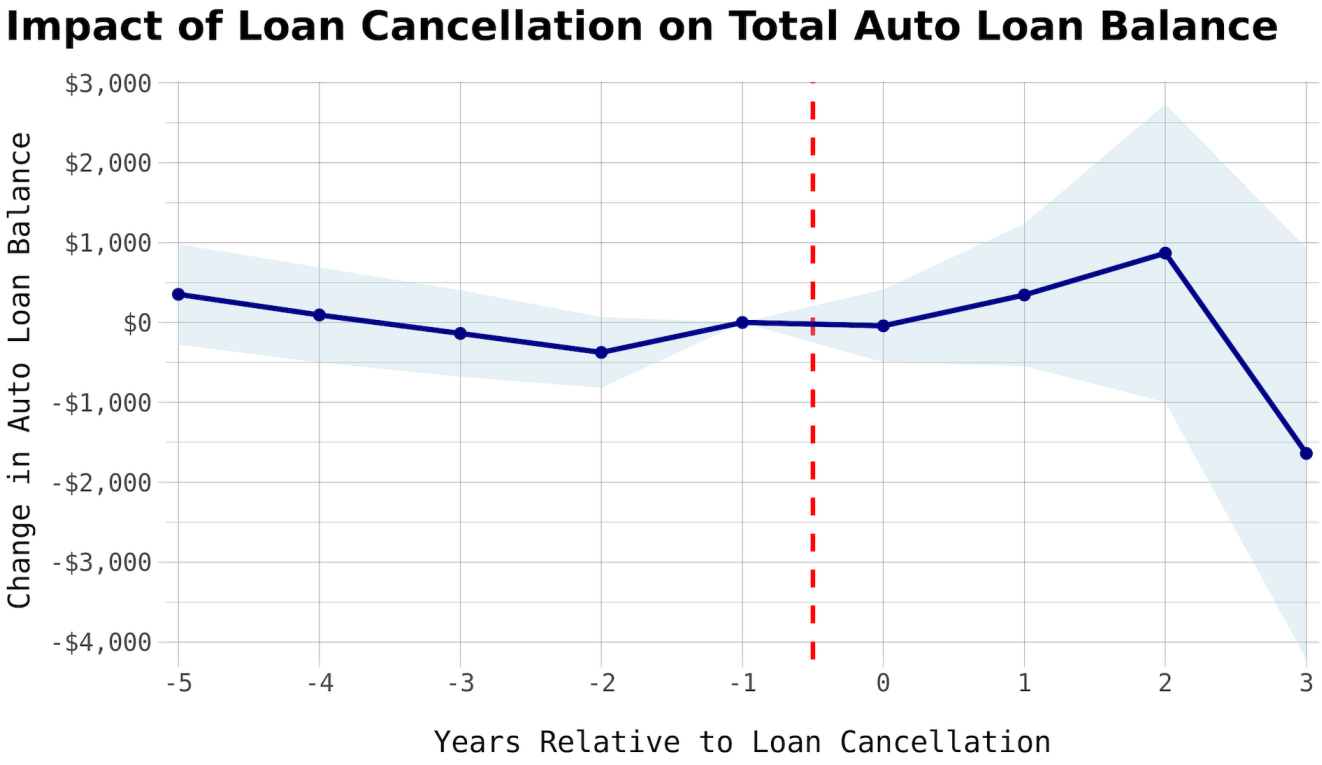

Note: Sun and Abraham estimator with individual and 2019 student loan payment decile by year fixed effects.

Source: Experian Information Solutions, Inc.
©2024 Jain Family Institute.

Figure 3.5 *- Event study estimate of the effect of student debt cancellation on total auto loan balance.*

We find cancellation increases total balance on all outstanding auto loans by about $1,000 two years post-treatment, before the trend reverses in year three. As with the homeownership outcome, we think it's likely the three-years-post coefficient would be different if we had additional observations on canceled borrowers that far out.

3. Credit Score

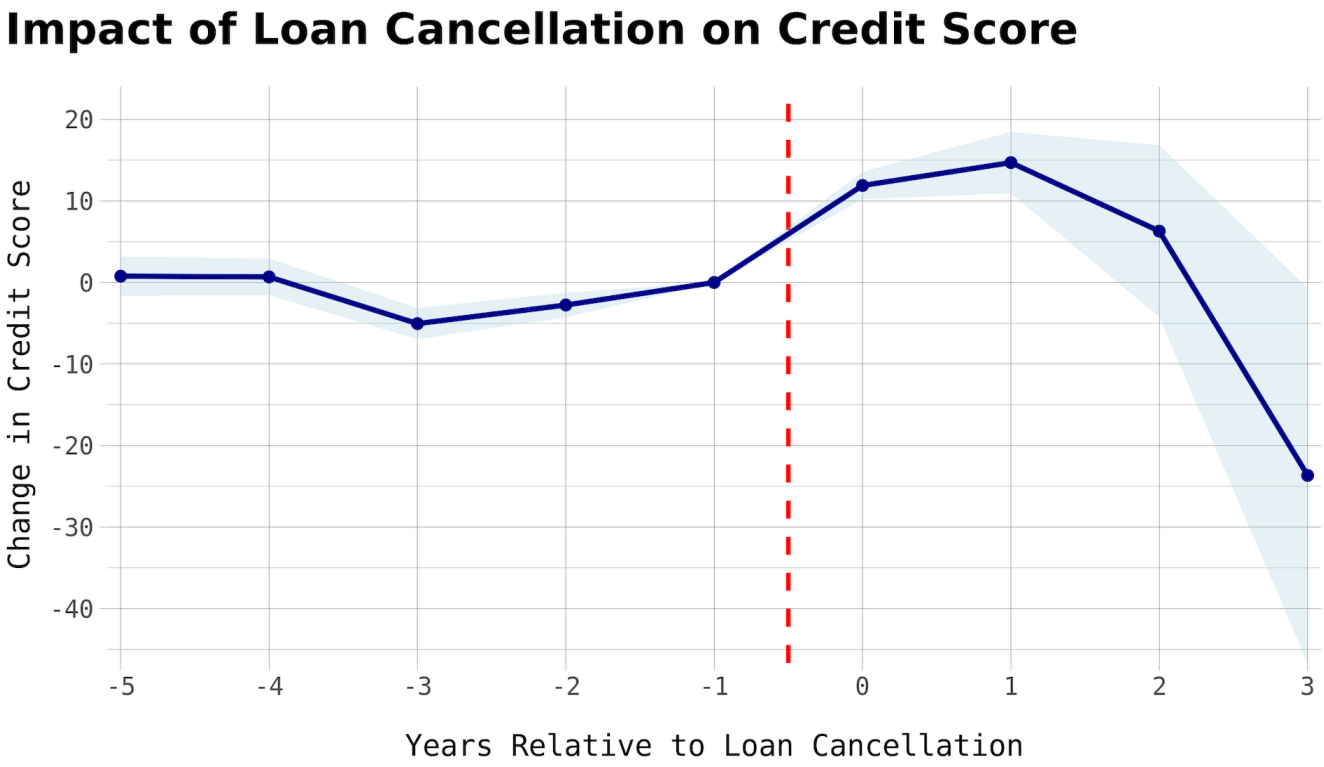

Note: Sun and Abraham estimator with individual and 2019 student loan payment decile by year fixed effects.

Source: Experian Information Solutions, Inc.

Figure 3.6 *- Event study estimate of the effect of student debt cancellation on credit score.*

We find that cancellation increases credit score by around 10 points on average immediately post-cancellation. In this case, we observe the trend reversing both two-years and three-years-post, which is stronger evidence for a dynamic effect, perhaps by borrowers increasing/replacing their debt load once their student debt obligations are diminished. That is consistent with what other scholars have found regarding the student loan repayment pause.

Conclusion

We examine the short-to-medium-run impact of having one's student loans canceled on individual financial outcomes and economic status. We report suggestive evidence that once-encumbered student borrowers are able to substitute secured loans, and by implication, asset ownership, for the unsecured student loans that are removed from their household balance sheet. Nothing about that finding is remotely surprising, other than how it stands relative to misinformation about the identity and financial status of student borrowers as is widely circulated in the popular press and even among some researchers. The implication is that student borrowers are substantially encumbered by their debt while it remains outstanding; the ones whose debt is canceled are those who are otherwise unable to repay, and when cancellation does occur, it enables previously-encumbered borrowers to get on with their economic life cycles, potentially leading to other changes in household composition, career choice, and asset accumulation. In other words, canceling student debt encourages social reproduction, while leaving it outstanding and irremediable causes economic stasis.

These findings also imply a significant difference between reductions in payment obligations, such as IDR or the student loan repayment pause of 2020–2023, and actual principal balance reductions, on which this analysis is focused. These two policy paradigms are often contrasted, with the former proffered as a more "targeted" alternative that would alleviate student debt burdens for the borrowers most in need of help. We show that the two sets of policies accomplish two different sets of outcomes. Even when payments are reduced on outstanding student debt, borrowers are not any more able to make asset purchases relative to the counterfactual of full monthly payments. If IDR and the repayment pause actually remedied the economic life-cycle-impairing effects of student debt, we would observe their effect beginning one or more years pre-treatment in figures 3.3–3.6, because the control group is selected based on having made payments during the repayment pause, while the treatment group largely did not (as shown in figure 3.2). The absence of such pre-trends indicates that neither IDR nor the repayment pause addresses this aspect of the relationship between student debt burdens and the accumulation of household wealth.

In addition to simply examining a longer post-period once it occurs in the real world, particularly following the end of the repayment pause in late 2023 (when we expect the differences between canceled and not-canceled borrowers to be, if anything, more dramatic), another extension to the results reported here would be to adopt a treatment-intensity specification that leverages variation between borrowers whose loans are canceled, with respect to either the dollar amount of debt removed or its relative share of their total assets or liabilities. From figure 3.1, we can see that, to some extent, our treatment cohort-level breakdown is proxying for two separate dimensions of heterogeneity: which cancellation pathway borrowers benefited from, and how much student debt was canceled in dollar terms. We leave distinguishing more completely between these two different dimensions of heterogeneity, as well as others, to future work.