The Long-Run Impact of the Great Recession on Student Debt *

SÉRGIO PINTO† MARSHALL STEINBAUM‡

September 8, 2023

Abstract

This paper investigates the effect of local labor market shocks during the Great Recession on subsequent student debt-related outcomes for a panel of 1 million student loan borrowers between the ages 17 and 34 in 2009, following that cohort's credit reports for the subsequent 10 years. We find that the Great Recession significantly increased student indebtedness, delinquency and default on student debt, and overall non-repayment of student loans. The Great Recession's effect on student indebtedness amplifies throughout the length of the panel, through 2019. We find that re-enrollment in higher education in response to the recession and declining state and local funding for public institutions are two likely mechanisms by which the recession exerted such a long-term impact on the financial status of student borrowers, beyond the initial economic shock and the sluggish labor market recovery.

*The authors thank their colleagues Eduard Nilaj and Laura Beamer, as well as the Jain Family Institute, for their advice and support on this project. Kevin Rinz, Doug Webber, Wilbert van der Klaauw, Dubravka Ritter, and two anonymous referees provided helpful comments. The authors also thank seminar participants at the Marriner Eccles Institute for Economics and Quantitative Analysis, American University, University of Nebraska Clark-England Labor Economics Conference, and the Society of Labor Economists.

†**University of Maryland at College Park** and **Instituto Universitário de Lisboa (ISCTE-IUL), DINAMIA'CET, Lisbon, Portugal**, stpinto@umd.edu.

‡**University of Utah Department of Economics**, marshall.steinbaum@utah.edu.

1 Introduction

In May 2009, at the height of job losses stemming from the Great Recession, the US Departments of Labor and Education issued joint guidance encouraging state employment agencies to notify unemployment insurance beneficiaries of opportunities to enroll in higher education and of the availability of Pell Grants. The same guidance requested university financial aid administrators to adjust aid awards to account for the effect of job loss on income.[1] The premise of that guidance was that higher education would mitigate the effect of labor market shocks both individually and in aggregate, by facilitating the sectoral reallocation of unemployed workers toward higher "skill" jobs.[2] Regardless of whether that guidance was the cause,[3] it nonetheless had the intended effect: enrollment in higher education increased substantially in the aftermath of the Great Recession, as did student indebtedness, most severely for individuals who lived in the local areas most adversely affected by the cyclical downturn.

The run-up in outstanding student debt pre-dates the Great Recession of 2008-2009. But this paper demonstrates that the Great Recession increased student loan balances, delinquency and default on student loans, and the non-repayment of student debt more broadly. As is typical during recessions, enrollment in higher education increased when the opportunity cost of foregoing labor force participation was low and employers demanded more educational credentials for scarce jobs.[4] Simultaneously, recession-driven contractionary fiscal policy at the state and local level forced public institutions to shift further toward tuition-dependent business models, likely inhibiting them from adjusting aid awards to fully account for students' and would-be students' loss of income during the recession, as envisioned by the aforementioned guidance.[5] Finally, deregulation in the higher education sector during the 2000s permitted for-

[1] Duncan (2009).

[2] Mitchell (2021*a*) contains an account of how this policy was arrived at by senior federal economic policy officials (pp 131-132).

[3] Barr and Turner (2018) conclude that the letters did in fact cause "an increase in enrollment of at least half a million UI recipients."

[4] Schmidt (2018), Modestino, Shoag and Ballance (2020).

[5] Chakrabarti, Gorton and Lovenheim (2020).

profit institutions to enter the market by offering credentials eligible for federal loans and to expand rapidly, responding to the recession-driven market opportunity more elastically than did traditional nonprofit institutions. For-profit higher education is associated with increased loan levels, non-completion, and default on student loans.[6]

This study has two aims: to quantify the Great Recession's effect on student debt and associated indicators of credit distress, including delinquency and non-repayment, and to identify, or alternatively rule out, the causal channels by which the recession affected student-debt-related outcomes. The causal channels we test are as follows:

1. Direct effect of the recession on the ability to repay pre-existing debt, e.g. due to recession-induced unemployment.

2. Increased educational attainment (and associated financial burdens) due to labor market credentialization. By credentialization, we mean that holding constant the occupation, job title, or pay, required educational attainment increases, as documented for the Great Recession and its aftermath by Modestino, Shoag and Ballance (2016, 2020).

3. Increased net tuition due to pro-cyclical state fiscal policy.

4. Entry and expansion of for-profit institutions (generally charging higher tuition than traditional nonprofits).

To address the first aim, we build on the literature that quantifies the long-run impact of the Great Recession on labor market outcomes using geographic variation in the recession's severity, most directly Yagan (2019). That study uses a panel of tax returns to estimate the effect of exposure to the Great Recession on subsequent earnings and employment rates through 2015. Exposure to the Great Recession is determined by where someone lived at the time it occurred, specifically the severity of the increase in the local unemployment rate. The inference is that living in a labor market in which the Great Recession was particularly severe exerted a

[6] Cellini and Turner (2019); Looney and Yannelis (2015).

negative influence on subsequent employment and earnings that was measurable at least six years later.

We document qualitatively similar effects for our panel of student loan borrowers: living in a commuting zone with a one-percentage-point-higher unemployment rate increase during 2007-2009 corresponds to a balance on outstanding debt that is $370 higher in 2019; it also leads to a probability of having fully repaid outstanding student loans that is 0.2 percentage points lower (this point estimate remains stable over time, although it becomes non-significant at the 5% level after 2014). Given that the average commuting-zone-level change in unemployment rate is approximately 3.8 percentage points, these estimates correspond to a "Great Recession effect" on student loan borrowing of about $1400.

To address the second aim, we undertake several different analyses. We divide the sample into borrowers who did or did not subsequently take out additional student loans associated with increased attainment after the Great Recession and examine repayment trajectories separately for those two groups of borrowers. Those who subsequently took on additional debt were much less likely to repay all their loans. We document that that decision to take on more debt is related to the severity of the Great Recession. We also show that the recession is related to declines in institution-level state funding, suggesting a relative shift toward tuition-dependent business models, as well as the entry and expansion of for-profit institutions. Collectively, these analyses provide support for all of the causal channels described above. Since they are not mutually exclusive, it's conceptually impossible to assign relative importance to each channel individually. The overall takeaway is that the recession was a temporary event that permanently shifted the political economy of higher education and its role in the labor market, resulting in growing student loan balances over time for a substantial minority of borrowers, with no reasonable probability of ever being repaid.

The significance of these findings is multifaceted: first, it validates in panel micro-data what has been apparent from more aggregated sources, that the Great Recession was a qualitative break with previous patterns of student indebtedness, heightening burdens on pre-existing

borrowers (as we document) and expanding the share of the population to have student debt in the first place (as others have shown, e.g. Morgan and Steinbaum (2018)). In doing so, it broadens our understanding of what causes students and their families to take on, and subsequently have difficulty paying off, student debt: namely, that it can be interpreted as a sign of disadvantage in the labor market, and specifically as an effect of hysteresis. Scholars have shown that employers facing slack labor markets require that workers have higher levels of education and skill, conditional on the job title being advertised.[7] Other work has shown that hiring rates decline disproportionately for young workers in recessions due to rationing job opportunities, but less so for workers with more educational credentials.[8] A natural outcome of that dynamic is that workers in need of jobs under slack labor market conditions would obtain more credentials by means of student debt, and also that given the labor market conditions that gave rise to that higher loan balance conditional on earnings, that they'd also have trouble repaying it. The expanding Income-Driven Repayment (IDR) programs feed this dynamic, by deferring otherwise-required payments if borrowers' incomes are too low. Life-cycle earnings progressions aren't steep enough to eventually enable IDR enrollees' balances to decrease and to eventually achieve full repayment. Even if the direct labor market impact of the Great Recession did eventually mitigate (as shown by Rinz (2022)), that would not necessarily halt the spiral of student debt if borrowers remain enrolled in IDR.

This narrative poses a severe challenge to the policy prescription embodied in the Obama administration guidance described above: encouraging unemployed or under-employed workers to have greater involvement in the higher education system did not remedy the recession's effect on their financial status as it relates to student debt, either individually or in aggregate.

This paper proceeds as follows: Section 2 reviews the literature on hysteresis following the Great Recession and labor market shocks in general, as well as the literature on the causes of student indebtedness. Section 3 describes the data sources used in this study and reports sum-

[7] Modestino, Shoag and Ballance (2016); Modestino, Shoag and Ballance (2020).
[8] Forsythe (2022).

mary statistics. Section 4 reports regression results for our main specification and addresses threats to identification. Section 5 undertakes supplemental analyses designed to prove or disprove the importance of the different mechanisms or causal channels for the Great Recession's effect named above. Section 6 discusses all our results in relation to the literature on student debt, higher education, and labor market dynamics. Section 7 concludes.

2 Literature Review

The basic purpose of this study is to link methodologies that investigate the medium- to long-term effects of labor market shocks with student debt-related outcomes. On the labor side, the closest predicate to this study is the aforementioned one by Yagan (2019), but others in broadly the same literature investigating the Great Recession in particular and using geographic variation in its severity as a proxy for the magnitude of the shock include Mian, Rao and Sufi (2013) and Mian and Sufi (2014). Charles, Hurst and Notowidigdo (2018*b*)'s findings relate to this study in that those authors document the negative correlation of labor market opportunities and college enrollment before and during the Great Recession, which may be a channel by which the student debt crisis is part of its legacy.

Rinz (2022) mostly replicates the findings in Yagan (2019) and extends them to 2017, by which time the Great Recession's effect on employment and earnings had begun to taper off and become statistically insignificant in some sub-groups. The former paper also disaggregates the Great Recession effect by "generation" and notes that the overall employment and earnings effects are different for each. Younger individuals experienced a smaller and more short-lived reduction in employment, but their Great Recession earnings effect was more severe (expressed as a percentage of pre-recession earnings) and longer-lasting, suggesting that they responded to the recession by taking lower-paying jobs, and that doing so exerted a negative effect on subsequent earnings levels and ascent up the job ladder, consistent with Haltiwanger et al. (2018). That job-ladder attenuation hypothesis in particular is consistent with our interpreta-

tion that ascending each 'rung' of the ladder requires more credentials and therefore more debt, while also making that debt more difficult to repay. By way of comparison to this paper, our panel consists of individuals who were between the ages of 17 and 34 in 2009, corresponding to birth cohorts 1974-1991. That includes the youngest members of "Generation X" and the oldest "Millennials."

Stuart (2022) uses geographic heterogeneity to estimate the effect of the 1980-82 recession on children's later-in-life education and earnings and finds it is substantial, and given the channel, extremely long-lasting. Using a similar geography-based proxy for exposure to every recession between 1973 and 2009, Hershbein and Stuart (2023) document long-run declines in employment-to-population ratios in local labor markets lasting through the present.

On the student debt side, studies have more recently focused on institutional segregation as an explanation for rising indebtedness, while treating macro labor dynamics as, at best, secondary.[9] Chakrabarti, Gorton and Lovenheim (2020) link declining state-level funding for higher education to increases in the probability of having student debt and to student loan balances among college students, a mechanism we consider in section 5.3. That study focuses more on state funding over the long run than on the business cycle, but the long-run decline in state funding consists of budget cuts in response to business cycle revenue reductions that are not restored when economic conditions improve. In this paper we show that state funding per FTE at public institutions declined substantially in the immediate aftermath of the recession before recovering in the late 2010s, but that notwithstanding the recovery, institutions did not meaningfully abandon the tuition-centric business models they moved toward during and soon after the Great Recession.

Chetty et al. (2020) does not discuss student debt per se, but does show intense sorting on parental income between institutions, as well as telling evidence that the sorting is becoming more pronounced as the children of relatively worse-off parents are decreasingly represented in selective public institutions wherein they might obtain financial aid, while there is no offset-

[9] e.g. Looney and Yannelis (2015), Cellini and Turner (2019).

ting increase in their representation among well-resourced private institutions. More intense sorting might give rise to more student indebtedness, although insofar as there is time variation in those authors' results, it is a secular trend rather than a business cycle effect.

Black et al. (2020) and Black, Turner and Denning (2023) study increases in federal student loan limits for undergraduate and graduate borrowers, respectively, and their effect on indebtedness, attainment, prices and cost of attendance, access, and subsequent labor market outcomes for borrowers. Those authors show that increasing undergraduate loan limits increases indebtedness but reduces delinquency and credit distress on student loans, with no apparent negative effect on other types of indebtedness. They also find that the resulting increase in attainment leads to better labor market outcomes in subsequent years. By contrast, their findings as to graduate student loans are far more negative: an increase in indebtedness and tuition without attendant student benefits in the form of improved attainment, access, or subsequent labor market outcomes. It is likely that the Great Recession effect we identify in this paper is preconditioned by the substantial increases in undergraduate and graduate student loan limits in the 2000s. Increased loan limits, the labor market credentialization effect of macroeconomic contraction, and pro-cyclical state fiscal policies resulting in declining public funding for academic institutions are likely the three main causes of the run-up in student debt during and after the Great Recession. The contribution of this paper is to establish the importance of the latter two causes using geographic variation in the recession's severity. But that is not to minimize the importance of increased loan limits, which may have invited the state-level policy response that shifted institutions toward tuition-centric business models.

The main interpretative difference between this study and Black et al. (2020) is that those authors find when student indebtedness increases due to increases in loan limits, difficulty with repayment declines, whereas we show that the Great Recession caused both greater indebtedness and greater difficulty with repayment. Those authors interpret the pattern they document as reflecting the higher likelihood of completion, and hence of improved labor market outcomes, the higher are the limits on the amount students can borrow during their undergraduate

education. The independent variation in this study is quite different: unemployment during the Great Recession. As such, it is not necessarily surprising that the increased indebtedness we find corresponds to greater difficulty with repayment, and although we don't directly observe labor market outcomes, our findings are consistent with a channel that leads from labor market slackness and unemployment to greater indebtedness to worse repayment outcomes, especially for borrowers who re-enroll in higher education (as that term is defined below). We also find that the Great Recession's effect amplifies over time throughout our panel. Black et al. (2020) also focus much more directly on traditional undergraduate students entering college full-time under the age of 20 and remaining continuously enrolled through graduation while borrowing at or below statutory loan limits in each year, and compare cohorts of that group differentially-affected by loan limit increases. Our sample is more diverse in the sense that we do not restrict it by age, enrollment status, or amount borrowed, aside from the overall sample frame of those aged 17-34 in 2009. It's likely that there are important differences in the ways that taking on student debt affects undergraduate and post-graduate students, and Black et al. (2020) is solely focused on the former. Black, Turner and Denning (2023), studying increases in graduate loan limits, do not report outcomes related to repayment or credit distress.

3 Data

The data used for the outcomes of interest and covariates in this study consist of a panel of borrowers provided by the credit bureau Experian. Relative to the population of Experian's master database (everyone with a credit report), the sample frame for the panel consists of individuals aged 17-34 in 2009 who had outstanding student debt at that time.[10] From there, 1 million individuals are randomly selected in 2009 and observed annually, including all of their credit records for student debt and otherwise, through 2019. We continue to observe

[10] The frame was slightly expanded to include individuals among the overall draw of 1 million who had had outstanding student loans and paid them off within a two-year window prior to 2009, but those individuals are excluded from this study.

sampled borrowers even if they fully repay their student loans during the 2009-2019 study period. Online Appendix A reports further details of sample construction.

The individual-level "consumer" panel includes both observed demographic characteristics (such as geographic location, derived from addresses) and imputed demographic characteristics (age and gender).[11] It also includes a variable highlighting the number of deferred student loans opened in the previous 12 months, which we call re-enrollment loans, i.e. new loans that we interpret to correspond to increased educational attainment. We use the appearance of new re-enrollment loans between two panel years as signifying that the borrower re-enrolled—another of our outcomes of interest—although we cannot tell whether an observed re-enrollment after 2009 is the continuation of a course of study started earlier or the beginning of a new course of study.

In addition to individual-level "consumer" data, we also have loan-level "trade" data, meaning each source of credit associated with an individual in the panel separately in each year. The trade data we use consist of all the student loans held by all of the individuals in the consumer-level panel. In addition, we use the presence of a bankruptcy flag on any loan associated with an individual, not just student loans, as one of our outcomes of interest. The loan-level data provide both the currently-outstanding balance and the origination date and origination amount of each loan, as well as the loan's current and past repayment status. We use these loan-level covariates in several ways:

- We construct "cohorts" of borrowers based on the earliest and latest origination dates of all the student loans in their portfolio. For example, one cohort of borrowers consists of all those borrowers whose earliest student loan origination date is in 2002 and whose latest loan origination date through 2009 is in 2006.

[11] The credit bureau possesses individuals' birthdates, which they use to identify the same individuals over multiple loan servicers and other of their input data sources. Hence, we believe the imputed age to be quite accurate after that starts to be reported in 2010, and of course given that plus our panel structure, we can infer age in 2009. The bureau also has full names, from which we believe gender is imputed, given that some records have different reported genders in different years and, for others, the gender is missing.

- For each year, we categorize each loan as in deferment, in forbearance, in repayment, repaid due to refinancing, and repaid due to being paid off.

For outcomes drawn from the trade data, we aggregate to the consumer level since our regressions are all conducted at the individual borrower level. For example, delinquency status is observed for loans, not consumers, but we specify delinquency-related outcomes like whether a consumer has any delinquent loan and the share of all outstanding non-deferred student debt that is delinquent at the consumer level.

Finally, in addition to the credit reporting data used for student-debt-related outcomes and covariates, we use labor data to construct the Great Recession shock. The basis of Yagan (2019) (and related literature) is that labor market dynamics in the local area where an individual lived between 2007 and 2009 are proxies for how shocked each individual was by the Great Recession. That paper uses the percentage-point change in the commuting zone unemployment rate between 2007 and 2009, and we follow that in our main specification. Commuting zone unemployment rates can be constructed from county-level unemployment rates in the Bureau of Labor Statistics Local Area Unemployment Statistics (LAUS) dataset, using the US Department of Agriculture's county-to-commuting-zone crosswalk.

3.1 Summary Statistics

To reiterate, the structure of our data is a panel of student loan borrowers first observed in 2009 and then annually through 2019, regardless of whether they continue to have positive student loan balances. Using the credit reporting data we define a total of seven outcomes of interest. They are:

1. **Total balance on outstanding student loans**.

2. A binary variable indicating **whether a borrower took out a re-enrollment loan in the previous year**.

3. A binary variable indicating whether **an individual has completed repayment on all their student loans**. Note that individuals remain in the panel (and thus are included in our regressions) after we code them as having completed repayment on their student loans. Thereafter, their total student loan balance is recorded as $0.

4. A binary variable indicating **whether a borrower has a bankruptcy flag on any loan in their credit record**.

5. **Total balance on delinquent or defaulted student loans**, for borrowers who have loans in repayment. Loans that are in deferment are assumed not to be delinquent or defaulted.

6. **Share of total loan balance in delinquency or default**, of total loan balance in repayment or forbearance, but not deferment.

7. A binary variable indicating **whether a borrower has any delinquent or defaulted student loan**.

Figure 1 plots the heterogeneous student debt repayment trajectories of the borrowers in the panel. In each year, each individual's total outstanding balance is compared to his or her 2009 balance. We plot the quantiles of the distribution of that ratio (computed year-by-year). The figure shows that the majority of borrowers had less student debt outstanding in 2019 than they did in 2009. However, a substantial minority of 2009 borrowers increased their loan balances in the intervening years, some by a great deal—likely signifying re-enrollment, enrollment in an income-driven repayment program (IDR), or both. The main empirical contribution of this paper is to relate these differences in borrower repayment trajectories to the severity of the Great Recession shock.

Table 1 reports summary statistics for the whole panel, regardless of age. The average age in 2009 is 28, and the average amount of outstanding student debt in that year is $22,806, already a right-skewed distribution since the median total student loan balance is $12,988. By 2019, the average was only slightly lower at $21,456, but the variance increased considerably. 46% of

the panel had fully repaid their student loans by 2019, corresponding to a balance of $0, while a significant minority increased their balances in the intervening decade. 38% of the panel is delinquent on at least one student loan at some point between 2009 and 2019.

We measure geographic variation in the Great Recession shock using the percentage point change in a commuting zone's unemployment rate between 2007 and 2009, following Yagan (2019). Figure 2(A) maps that by commuting zone, to get a sense of the geography of the Great Recession as we quantify it.

4 Results

4.1 Main Empirical Specification

The main regression equation is given by

$$y_{i,201t} = \alpha + \beta_{201t}\Delta Unemp_{c,2007-2009} + \gamma x_{i,2009} + \delta z_{i,201t} + \epsilon_{i,201t} \tag{4.1}$$

i indexes individuals in the panel and c refers to the commuting zone in which individual i resided in 2009. The outcome of interest is $y_{i,201t}$, where 201t is a panel year subsequent to 2009. $x_{i,2009}$ are covariates observable at the start of the panel, interpreted as individual-level characteristics that either pre-date or are independent of the effect of the Great Recession. $z_{i,201t}$ references covariates contemporaneous with the outcome of interest. For the continuous outcomes, $y_{i,201t}$ is expressed as the change from a 2009 baseline $y_{i,2009}$. We specify the regressions of continuous outcomes as differences from a 2009 baseline to take as full advantage of the panel structure of the data as possible, given that there is no variation within individuals over time in the severity of the Great Recession shock (i.e., each individual is exposed to a single shock, depending on the commuting zone where they reside in 2009). The total and delinquent balances are winsorized at the 99th percentile, which indirectly also affects the computation of the delinquency ratio. The remainder of the outcomes are 0-1 binaries, such as whether i had

completed student loan repayment by 201t or had any delinquent or defaulted loans in 201t.

β_{201t} is the coefficient of interest, interpreted as the causal effect of the Great Recession shock on the outcome $y_{i,201t}$. The identifying assumption underlying that interpretation is that the Great Recession shock is as good as randomly assigned to individuals, conditional on covariates $x_{i,2009}$ and $z_{i,201t}$. The magnitude of the estimated $\hat{\beta}$ is then interpreted as the effect of a one-unit-larger Great Recession shock (i.e., one-percentage-point-larger increase in the commuting zone unemployment rate between 2007 and 2009) on the outcome of interest in year 201t, expressed in units of the outcome of interest.

In all our estimates, the covariates are:

1. Gender.

2. 2009 student loan balance.

3. Contemporaneous local unemployment rate in the commuting zone where the borrower is observed in 201t (which may be different than the commuting zone where they are observed in 2009, which is used to assign the Great Recession shock).

4. Cohort-age fixed effects, which is a triple interaction of the year in which each borrower's oldest student loan was originated, the year in which each borrower's most recent student loan observed in 2009 was originated, and the borrower's age in 2009.

Hence, our coefficient $\hat{\beta}$ is estimated using variation in the Great Recession shock within cells defined by cohort and age, and net of controls for gender, "starting" student loan balance, and the subsequent state of the local labor market. We further address concerns about the assumption that the Great Recession shock is as good as randomly assigned net of these controls in Subsection 4.3 below.

4.2 Baseline Estimates

Table 2 reports estimates of equation 4.1 for all seven outcomes of interest for the year 2019, i.e. at the end of the panel. Figure 2(B) maps the change-in-student-loan-balance outcome (corresponding to column 1 of table 2) by commuting zone for 2019. We then plot the estimated coefficient $\hat{\beta}$ for each year between 2010-2019, for each outcome, in figures 3 and 4.[12]

Our key finding is that a one-percentage-point increase in the Great Recession shock corresponds to $370.25 more student debt in 2019, and the pattern of estimated coefficients shown in figure 3(A) means that this estimated "Great Recession effect" amplifies throughout the length of the panel, even after the actual Great Recession was long over and even after its lingering effect on labor market outcomes documented by Yagan (2019) and Rinz (2022) is shown to have dissipated. Its effect on student indebtedness remains, and indeed is still growing, at least as of 2019.

Time variation in the estimated coefficients for the other outcomes is further revealing about the dynamics of student indebtedness and delinquency, as well as economic distress more broadly, as our cohort of borrowers aged into what should have been their repayment years, and as the policy landscape of student borrowing and higher education finance also changed over those years. Figures 4(A), 4(B), and 4(C) depict the estimated coefficients for total delinquent or defaulted balance, the ratio of that delinquent balance to total outstanding student loan balance, and the indicator for having any student loan delinquent, respectively. They all show that the Great Recession corresponds to more delinquency, particularly in the first half of the sample. The coefficient for the indicator for any delinquency is negative in 2009 and 2010, which is probably related to the association between the recession shock and re-enrollment (and hence of deferment) discussed below. It is positive starting in 2011 for the rest of the sample.

[12] The coefficient plots begin in 2009 for all the binary outcomes, with the exception of having completed repayment, which given our empirical strategy - which relies on the borrower having outstanding student debt in 2009 - can only start in 2010.

The effect of the recession shock on the delinquency-related outcomes dissipates starting around 2016, which corresponds to the time period in which Income-Driven Repayment became more accessible and generous for distressed borrowers. By 2018 and 2019, the estimated coefficients are actually negative for the first two delinquency-related outcomes, total delinquent balance and delinquency ratio, meaning that borrowers who were exposed to a more severe Great Recession shock actually had lower delinquent balance on student loans than less-exposed borrowers by the end of the panel. Contrasted with the pattern in figure 3(A), the implication is that borrowers availed themselves of the option of expanded IDR by the second half of the 2010s, deferring currently-unaffordable loan payments and instead accumulating them into outstanding balances. Given that borrowers who were more exposed to the Great Recession were more likely to be delinquent in the first place, and likely suffered from long-term labor market hysteresis documented by Rinz (2022), they are exactly the borrower population expanded IDR was designed to relieve.

Figure 3(B) plots the coefficient estimates for the re-enrollment binary outcome. Between 2011-2014, the shock is associated with an increased probability of taking on additional student loans associated with re-enrollment in higher education, though we cannot distinguish between undertaking "new" degrees versus continuing toward a degree begun before 2009. From 2015-2019, the coefficient estimates for re-enrollment are near zero, reflecting both that this cohort is closer to having aged out of higher education altogether, and that for those still re-enrolling, that seems to no longer be related to the Great Recession. But as we discuss further in section 5, the recession-driven re-enrollment is a strong candidate for the lingering effect of the recession on student debt, since it appears to have fundamentally altered economic life cycles of affected borrowers thereafter.

Figure 3(C) plots the estimated coefficients for the completed-repayment outcome, which is in effect a mirror image of the total balance outcome coefficients shown in figure 3(A)—borrowers more affected by the recession are less likely to have completed repayment on their student loans. Aside from the direct labor market effect of the recession, re-enrollment is again

a strong candidate explanation.

Finally, figure 3(D) depicts the association between the recession shock and an indicator for whether the borrower has any loan with a bankruptcy flag in their credit record. That too shows an increasing positive association for 2009 through the mid-2010s, and thereafter the relationship diminishes, though it never becomes zero or statistically insignificant (even by 2019, it remains significant at the 10% level). That is consistent with the findings on hysteresis from the Great Recession, which eventually diminish, though the student-debt-related consequences of that hysteresis do not, at least within the first 10 years.

4.3 Identification

The empirical strategy employed in this paper is premised on the idea that it is possible to estimate the effect of macroeconomic fluctuations in individual microdata by making use of their geographically heterogeneous severity. Thus, once threat to identification that depends on that heterogeneity derives from the possibility that geographies experience macroeconomic fluctuations differently for reasons that pre-date their exposure to the fluctuation.

In the case of this study, the specific threat to identification is that labor market fluctuations during the Great Recession may be related to pre-existing heterogeneity in student-debt-related outcomes. That threat is heightened by the construction of our panel: to be included in it, borrowers need to have already had outstanding student debt in 2009. If it's the case that geographic heterogeneity in the Great Recession shock is related to pre-existing trends in student indebtedness in the population, it may also be the case that the subsequent indebtedness and difficulty with repayment that we ascribe to the Great Recession's effect may be due to an omitted variable driving both the dependent and independent variables in equation 4.1, and which affect the probability that individuals in our panel are selected into our sample in the first place.

For example, Charles, Hurst and Notowidigdo (2018*a*) and Charles, Hurst and Notowidigdo (2018*b*) show that individuals who lived in areas that experienced greater housing price infla-

tion in the 2000s, prior to the Great Recession, had lower educational attainment than would be predicted based solely on their demographics, and subsequently saw higher unemployment during the recession itself. Their interpretation of that pattern is that the housing boom drew workers into construction and related sectors who would otherwise have remained in higher education, and that these workers were then more prone to sector-specific shocks and had more difficulty re-allocating to less depressed sectors thanks to their lack of education. The implication of that finding that matters for assessing the exogeneity of the Great Recession shock with respect to student-debt-related outcomes is, if anything, that commuting zones that experienced a more severe shock had *less* student debt (and fewer student debtors for us to sample) than they otherwise would thanks to the preceding housing boom, and were less exposed to the run-up in student indebtedness that pre-dates the recession. If those authors are correct, then the association we document between the Great Recession and re-enrollment may reflect a kind of deferred attainment, i.e. higher education that the affected individuals would have obtained earlier, but for the pre-recession housing boom. It also means that we may be undersampling the geographies where the Great Recession was most severe, since individuals living there were less likely to have accumulated student debt in the 2000s.

Addressing the threat to identification in this study is hampered by the fact that our panel only begins in 2009. Hence, we do not observe our individuals prior to the onset of the Great Recession, and we are not able to run tests of the parallel trends assumption underlying the identification strategy directly in our data. Equation 4.1 is designed to mimic as closely as possible Yagan (2019)'s specifications that do rely on observing his panel pre-recession. Yagan (2019) is able to test the parallel trends assumption directly and concludes that it holds in relation to the outcomes that paper focuses on: employment and earnings, the indicators of labor market hysteresis. Given that the Great Recession shock is closely related to what those indicators measure, but not to student-debt-related ones, it would be strange if areas that experienced differentially-severe Great Recession shocks were selected according to ex-ante student-debt-related outcomes but not more directly labor-market-related ones. Further, it's worth pointing

out that the issue is not whether areas that experienced more severe Great Recession shocks also had worse ex-ante problems with student indebtedness and associated outcomes. We control for individuals' 2009 loan balance, contemporaneous with the conclusion of the shock period, so that would account for any level effect. Rather, the threat is differential growth paths for student debt across areas that are differentially-shocked during the Great Recession prior to its onset. If that were the case, then ex-post variation in student debt-related outcomes that we are attributing to the Great Recession shock would in fact be driven by some omitted variable driving both.

In the absence of individual-level longitudinal data from before the Great Recession that both reports student debt-related outcomes and is sufficiently geocoded to assign a Great Recession shock, we address that possibility to the best of our ability using College Scorecard data that begins prior to the Great Recession. The College Scorecard is a dataset that combines institution-level data from the Integrated Post-secondary Education Data System (IPEDS) with individual borrower student loan records from the National Student Loan Data System (NSLDS), reported at the institution level. For each institution in the College Scorecard, we utilize three outcome variables:

1. Median student loan debt of students who attended (and had student loans) upon entry into repayment (i.e. their loan principal at the time their student loans were no longer deferred);

2. Two-year cohort default rate (i.e. the share of each cohort of an institution's borrowers who defaulted within two years of entering repayment);

3. Number of borrowers in each median loan balance cohort (i.e., the number of individuals with positive loan balance over whom the median balance is computed).

We have the median loan amount and cohort size observable from 2001-2020, and the two-year cohort default rate observable from 2000-2012.[13] We map them to commuting zones based on

[13]College Scorecard switches from a two-year to a three-year cohort default rate after 2012, so since the main

the geographic information for each institution also reported in the College Scorecard. Figures 5(A) and 5(B) plot the time series for median loan balance and two-year cohort default rate. In both cases, they show a dramatic increase around the time of the Great Recession following little change earlier in the 2000s, validating the fact that motivates this study: the Great Recession was a turning point in the run-up in student debt.

Given this data, we run the following regression to relate these outcomes to the geographically heterogeneous Great Recession shock:

$$y_{j,t-2009} = \alpha + \beta_{t-2009}\Delta Unemp_{c,2007-2009} + X_{j,2009} + Z_{c,t} + \epsilon_{j,c,t} \tag{4.2}$$

$y_{j,t-2009}$ is one of the three outcomes: median loan balance, cohort default rate, or loan balance cohort size, for institution j in year $t-2009$, i.e. before or after the Great Recession. $\Delta Unemp_{c,2007-2009}$ is the same Great Recession shock as in equation 4.1. $X_{j,2009}$ is the "baseline" level of the median loan balance y in 2009, to make equation 4.2 as analogous as possible to the baseline specification outlined by equation 4.1. And $Z_{c,t}$ is the contemporaneous unemployment rate in commuting zone c in year t.

The results from estimating equation 4.2 are reported in figure 6. For all the outcomes, the pre-recession coefficient estimates are clustered around zero, consistent with the absence of pre-trends that would question the identifying assumption underlying the exogeneity of the Great Recession shock. After the Great Recession, the median loan balance coefficients are positive and significant, consistent with our earlier findings, e.g. in figure 3(A), albeit of smaller magnitude. We estimate the Great Recession effect on institution-level median loan balance is around $100. Likewise, the Great Recession shock causes the size of the loan balance cohorts to be significantly larger in the immediate aftermath of the recession, from 2010-2013, before

point of this analysis is to test whether there are pre-trends in estimating equation 4.1, we use the two-year cohort default rate variable we have during the pre-recession period, as well as a baseline value for 2009 since that is key to our regression specification. For this outcome, the value in year t corresponds to the default rate measured in FY_t, for the cohort that entered repayment during FY_{t-1}, where FY_t is a 12-month period that ends on September 30 of year t. For the median loan balance, the College Scorecard data pools data from two FY cohorts, so we rescaled the data so that the median balance in each year t is pooling those in FY_t and FY_{t-1}.

declining to a near zero effect by the end of the sample. The implication of the two findings combined is that the recession both increased the loan balances associated with borrowers at a specific institution, as well as the number of students at that institution who take out student loans, either because they are more likely to attend (enrollment increased), or because given they do attend, they are more likely to take out student loans. The difference between our estimated coefficient(s) on median institution-level loan balance in the College Scorecard data, versus the larger coefficient estimates from the credit reporting data, probably derives from several causes:

1. Increased institutional sorting and segregation (needier borrowers enrolled at less-resourced institutions) means that we would expect within-institution increases in loan balance to understate overall increases in loan balance, conditional on attainment.

2. Our main specification in equation 4.1 identifies the effect on average loan balance, rather than median, and our data on student loan balances shows an increasingly right-skewed institution over time (even despite Winsorizing student loan balances at the 99th percentile).

3. The median debt variable in the College Scorecard data only includes debt taken out for the degree for which the student actually enrolled. Hence, if a student takes out undergraduate loans at one institution and graduate loans at another (or, in fact, at the same one), and that borrower's loan balances are the same as that institution's median for each degree program, the College Scorecard median loan balance variable would not increase, even if the higher attainment corresponds to increased loan balance at the borrower level (i.e, if, in the absence of the Great Recession, this hypothetical borrower had not obtained a graduate degree).

We return to the effect of the Great Recession on enrollment in section 5 below.

Figure 6(B) shows no increase in cohort default rate due to the Great Recession in the three post-recession observations we have, but that too is consistent with earlier findings about

individual-level delinquency and default in the credit reporting data, since default is a drawn-out process, and it has to already have happened to be picked up in calculations of an institution-level cohort default rate. So we would not necessarily expect to see any Great Recession effect manifest for that outcome by 2012. Altogether, this analysis of institution-level College Scorecard data is both consistent with the quasi-exogeneity of the Great Recession shock (since the pre-treatment coefficient estimates are clustered around zero), and confirms the main findings of the analysis using student borrower credit reporting data about the Great Recession's effect.

The other major threat to identification is more subtle: that a randomly-selected sample of student borrowers in a given relatively wide age range, 18-34, reflects selection on our outcomes of interest that is differentially-severe across borrower age. Namely, having outstanding student debt into one's thirties bespeaks difficulty with repayment such that older borrowers in the panel are predisposed to have worse repayment outcomes relative to their contemporaries by age, whom we never observe because their debt was repaid prior to 2009. That threat to identification is independent of geographic variation in the Great Recession shock, but nonetheless may imply our estimates of the Great Recession's effect are skewed toward a subset of negatively-selected student borrowers, particularly at older ages: those who were having difficulty with repayment regardless.

5 Mechanisms

In this section, we consider in turn the mechanisms by which the Great Recession may be responsible for the increase in student indebtedness, delinquency, and non-repayment documented in section 4.[14]

[14] One potential mechanism we do not test is that the Great Recession may have brought students into higher education who were negatively selected for future earnings. We cannot test this mechanism since we only observe borrowers who already had student loans in 2009, and of those, we do not observe them prior to taking on student debt. Looney and Yannelis (2015) posit this as one contributor to declining student loan repayment and rising credit distress.

5.1 Direct effect of the recession on borrower financial status

The Great Recession dis-employed many workers, and aside from those who did not lose their jobs outright, many more endured long-running wage stagnation and declining labor market mobility associated with a scarcity of outside job offers. It's very likely that this impaired their ability to repay debt, including student loans.

We validate this using our bankruptcy flag outcome reported in section 4 and figure 3(D): borrowers who lived in areas more shocked during the Great Recession were more likely to later have one of their loans in bankruptcy. From an initial slightly negative relationship between the Great Recession shock and the bankruptcy flag outcome in 2009, the magnitude of the coefficient becomes positive and increases through the mid-2010s, when a one-percentage-point-larger value for the shock corresponds to a 0.4 percentage point-higher likelihood of having any loan in bankruptcy. This is broadly consistent with more in-depth studies of the Great Recession's effect on household financial status, e.g. Mian, Rao and Sufi (2013). Thereafter, the association diminishes but remains positive through the end of the panel. The implication for this paper is that there likely is a role for this direct channel on the other student-loan-related observables reported in section 4.

5.2 Increased educational attainment

As recounted at the start of this paper, policy-makers put a lot of faith in the idea that the way for workers to respond to the recession was to enroll in higher education, whether for the first time or after having entered the labor market already, then been displaced due to the macroeconomic upheaval. Figure 7 plots the share of our panel over age 23 (in each year) that re-enrolled (i.e., took out a new student loan that we categorize as a re-enrollment as opposed to re-financing/consolidation loan) between 2009 and 2019. That share peaks at just under 8% of the sample in 2010 and 2011, then declines to around 3% by 2019. Hence, a significant share of the borrowers we study seem to have taken that advice. Recall that figure 3(B) shows the

re-enrollment outcome is positively associated with the Great Recession shock from 2011-2014.

Figure B.1 shows the starkly different student loan repayment trajectories associated with the decision to re-enroll or not. Each subfigure is the equivalent of figure 1, but dividing borrowers based on whether or not they are ever observed to take out a re-enrollment loan during the length of the panel. The borrowers who don't are much more likely to repay their loans by the end of the panel in 2019, and only about 10% have a higher student loan balance at the end than they did at the beginning. By contrast, the borrowers who do re-enroll show almost exactly the opposite pattern: only a bit over 10% have fully repaid their loans by the end of the panel, and over 50% have more student debt outstanding at the end than they did at the beginning. The median borrower who re-enrolled had twice as much student debt in 2019 as in 2009, and the 90th percentile of borrowers who re-enrolled almost 10 times as much. We also restrict these figures to borrowers who were 23 years old or older in 2009 in order to filter out undergraduate students in the middle of a continuous spell of enrollment when they are first observed, such that continuing that spell after 2009 would be erroneously interpreted as re-enrollment.

In order to test this mechanism more fully, we re-run the regression whose estimated coefficients are plotted in figure 3(B) separately for borrowers who were under 23 versus 23 and over in 2009. Those estimates are depicted in figures 8(A) (for those under 23) and 8(B) (for those 23+). For borrowers under 23, the coefficient estimates are negative for the first few years following 2009, then become statistically indistinguishable from zero from 2012 onward. By contrast, the Great Recession shock effect for borrowers over 23 in 2009 is positive from 2010-2014, with the maximum point estimate in 2012. Hence, 8(B) looks similar to figure 3(B), but with estimates of larger magnitude in the years following the Great Recession. The implication is that the effect of the Great Recession on the youngest borrowers in the sample was to make them less likely to continue in higher education, whether because they finished more quickly, dropped out of undergraduate degree programs, or completed undergraduate programs but did not continue to graduate study. By contrast, for older borrowers, the Great

Recession caused them to re-enroll, whether in graduate or undergraduate degree programs.

We further break down the sample by whether borrowers re-enrolled or not (Table 3), by age group (Table 4), and by cohort (Table 5), where for simplicity we define cohort as origination-year-of-last-loan up through 2009, as opposed to the two-dimensional definition of cohort we use in our regressions. The borrowers who re-enrolled skew slightly younger than the overall sample, but a significant share from every age group re-enrolled. Those who did had a much higher rate of delinquency, and as stated earlier, eventually accumulated much more student debt (which they did not repay, at least by 2019). Recall that the "re-enrollment boom" in our sample peaked in 2011-2013 and then diminished. The fact that borrowers from every age group participated lends credence to the idea that re-enrollment was a specific outcome of the Great Recession, as opposed to the result of a single cohort of student borrowers aging out of higher education over time. If the latter were true, we would expect to see re-enrollment more concentrated among that cohort's youngest members than it in fact is. And as the different repayment trajectories for the two groups show, the decision to re-enroll or not turned out to be a decisive one in prefiguring whether a borrower would repay their loans within a ten-year window, or indeed ever see their loan balances decline.

To further illustrate these points, we split the sample based on whether borrowers re-enrolled or not between 2009 and 2019, and then estimate the baseline specification outlined in equation 4.1 on each sub-sample. Figures B.2(A) and B.2(B) show that it is not just that those who re-enrolled have higher loan balances, but also that among the borrowers who re-enrolled, those more negatively affected by the Great Recession shock accumulate vastly larger balances, and that disparity widens over time. The Great Recession shock effect on total loan balances in 2019 is about $100 for borrowers who did not re-enroll (diminished from an estimate of $150 in 2012, its maximum for that sub-group), whereas for borrowers who did re-enroll, the coefficient on the Great Recession shock is over $1000 in 2019, its maximum throughout the sample.

Considered from one point of view, it's hardly surprising that borrowers who took on more student debt took longer to pay off their accumulated balance than those who didn't. However,

the idea behind the guidance referred to at the start of this paper, and the larger imprecation that higher education attainment is the solution to labor market displacement, was that the attainment was supposed to pay off in the form of higher earnings, enabling the debt to be repaid. There's no indication in the data that borrowers who re-enrolled were later better able to make progress toward repayment than they would have been absent the increased debt—quite the reverse: their balances continued to increase, while borrowers who didn't re-enroll for the most part paid down their balances. What likely happened is that the increased debt plus relatively stagnant earnings qualified re-enrolled borrowers disproportionately for Income-Driven Repayment, and once enrolled in IDR, negative amortization is very likely, especially early in the repayment cycle when the preponderance of monthly payments consists of interest and hence reducing payments to a given portion of disposable income leaves a great deal of interest unpaid. Thus, the significant minority of the sample shown to have increased their balances between 2009 and 2019 in figure 1 is due to the combination of actual increased attainment and subsequent qualification for income-driven repayment, resulting in negative amortization and, implicitly, the later cancellation of outstanding balances at the end of the IDR repayment period. We view these findings as strong circumstantial evidence supporting the interpretation that the decision to re-enroll in the aftermath of the Great Recession permanently set otherwise-similar borrowers on a different life trajectory as it relates to their student debt, and that decision to re-enroll was related, at least in part, to the severity of the Great Recession where they happened to be located in 2009.

5.3 Declining state higher education funding

The remaining mechanisms pertain to the effect of the Great Recession on the political economy of higher education: did the recession lead to stagnant or even declining state funding for institutions, and thence to higher tuition and/or displacement of students from lower-priced public institutions to higher-price for-profit ones? We use institution-level data from IPEDS to answer these questions.

Figure B.3(A) shows that state and local funding per FTE at public institutions diminished substantially between 2009 and 2012, which would have both increased tuition and constrained seats available in classes and programs (potentially displacing students to higher-cost institutions), both mechanisms for increased student loan balances.

To test whether these trends are associated with the Great Recession shock, we run the following regression

$$y_{c,201t} = \alpha + \beta \Delta Unemp_{c,2007-2009} + X_{c,2009} + Z_{c,t} + \epsilon_{j,c,t} \tag{5.1}$$

where outcomes are aggregated across institutions within a commuting zone. $y_{c,201t}$ is the commuting-zone-level change in state and local funding per FTE across all institution types (although the amount of such funding for private institutions reported in IPEDS is negligible) in year $201t$, relative to 2009.[15] $X_{c,2009}$ is state and local funding per FTE in 2009. $Z_{c,t}$ is the commuting zone unemployment rate in year t.

Results are plotted in figure B.4(A). The Great Recession shock corresponds to a reduction in state and local funding per FTE of around $100-150 from 2010-2015. Thereafter the relationship starts to reverse and by the last two years of the panel, the coefficient estimates are positive. The latter finding is consistent with the rising trend in state funding in the latter part of the decade visible in figure B.3(A), and the timing of the coefficients suggests that recovery in state funding was delayed in areas more affected by the recession. Altogether, reductions in state funding are associated with (i) a more severe Great Recession shock, and (ii) increasing institution-level student indebtedness (figures 6(A) and 6(C)), consistent with a role for the rising-tuition channel.

[15] Since most public funding for institutions of higher education is state funding, it might seem more reasonable to run this regression at the state rather than the commuting zone level, using a state-level Great Recession shock rather than one constructed by commuting zones. We run it at the commuting zone level to ensure comparability with coefficients estimated from equations 4.1 and 4.2.

5.4 Entry and expansion of for-profit institutions

Finally, we consider whether the recession itself, plus the diminished funding for public institutions it brought about, re-allocated students from lower-cost public institutions to more expensive for-profit ones. Figure B.3(B) plots the total number of public, private nonprofit, and private for-profit institutions nationally between 2009 and 2019. The first two types show no meaningful change, while the number of private for-profit institutions increases from 2009 to 2013 and then declines. The question remains whether that inflation and subsequent deflation of the for-profit higher education bubble is associated with the geography of the Great Recession.

To answer that, we estimate equation 5.1 using the following commuting-zone-level outcomes:

1. Number of for-profit institutions (figure B.4(B)).

2. FTE enrollment at for-profit institutions (figure B.4(C)).

3. FTE enrollment at public institutions (figure B.4(D)).

The Great Recession shock does indeed correspond to increased number of and enrollment at for-profit institutions, although the timing is not quite aligned for the two outcomes: enrollment shows the largest effect in the first few years following the Great Recession, 2010-2013, and diminishes thereafter to be near zero. The effect on the number of for-profit institutions appears at a lag, peaking in 2014-2016. That pattern probably indicates that on the extensive margin, the for-profit sector responded to the business opportunity that the Great Recession seemed to offer, as enrollment at existing institutions swelled.

Enrollment at public institutions is immediately higher in more shocked commuting zones in 2010 and 2011, but thereafter the coefficient estimates on the Great Recession shock are negative. This may be due to the diminished funding and consequent tuition increases/capacity constraints highlighted in the previous subsection. The question that remains is whether the

for-profit boom drew solely from students who would have otherwise attended public institutions, or whether that was drawing on a population that otherwise wouldn't have attended at all. These findings provide some evidence for the former, but certainly don't rule out the latter.

6 Discussion

This paper has demonstrated that where the Great Recession was more severe, student loan borrowers subsequently took on more debt and had greater difficulty repaying it. The effect of the Great Recession on student loan balance and non-repayment has become more pronounced over time, even as other long-run hysteresis effects of the recession on employment and earnings have dissipated. The main implication of that pattern is that the Great Recession's effect will continue to be felt by those who were early in their working life at the time, due to the student debt they continue to carry with them.

In section 5 we test several mechanisms for how the Great Recession could have had that long-run effect. We find some support for all the mechanisms we test, but the two that appear notably significant are 1. re-enrollment in higher education, which permanently set borrowers on a course to non-repayment and likely negative amortization in IDR, and 2. declining state and local government funding relative to increased demand for higher education, burdening students with higher tuition and/or displacing them to the for-profit sector where they were more likely to take on more debt. The implication is that the guidance described at the start of this paper, to the effect that enrolling in higher education was a wise response to labor market displacement because it qualifies students for more job opportunities, did not turn out to be correct, at least as measured by student debt-related outcomes.

The aforementioned papers by Forsythe (2022), Modestino, Shoag and Ballance (2016) and Modestino, Shoag and Ballance (2020) point to our preferred interpretation of rising student indebtedness in response to the recession: employers hiring from slack labor markets increase the skills and credentials they demand of new hires with less experience, conditional on job

title and salary, giving the labor market a counter-cyclical "Last In, First Out" quality to the detriment of younger cohorts when jobs are scarce. In response, workers take on more degrees and therefore more debt to qualify themselves for a given job, and higher education institutions enter and expand to meet that recession- and slackness-generated demand. Because workers are competing for scarce jobs, their earnings don't increase as a result of the debt they take on. Rather, credentialization-driven re-enrollment functions as a way to limit earnings declines (or, in the extreme, non-employment) relative to the counterfactual where they don't take on debt. That rationalizes the decision to take it on despite facing slack labor demand, while simultaneously explaining why they'd have trouble paying it off.

The federal student loan program then operates as a semi-blank check to finance that process. ("Semi" because the program does impose loan limits on borrowers for undergraduate education, though those have been increased.[16] Since 2005, federal loans for graduate education have been uncapped, inviting the proliferation of high-tuition graduate programs despite costlier unsubsidized loans.) The system also includes solutions for the non-repayment that results from rising and more widespread indebtedness without rising earnings: refinancing old, defaulted debt with new, not-yet-defaulted loans[17] and enrolling borrowers who are otherwise likely to default into Income-Driven Repayment programs that reduce ongoing required payments at the expense of increasing nominal balances, until a later date in the far future when that nominal balance is eventually erased.[18]

Ameliorative measures on the back end such as IDR likely serve to make continued origination and rising indebtedness more palatable on the front end. By contrast, limiting the amount of student debt individuals are allowed to take on, at least within the federal student loan system, would likely worsen the effect of labor market slackness: without the ability to borrow to

[16] Black et al. (2020) show the increase in undergraduate loan limits increased student indebtedness.
[17] Mitchell (2021*b*) is a telling account of this.
[18] CBO (2020). Although it should be noted that the borrowers who enroll in Income-Driven Repayment tend be those with higher balances and (relatively) low income, for whom it is most worthwhile to overcome the bureaucratic burden. That population is not one that is necessarily likely to default in the counter-factual that IDR is not available (Collier, Fitzpatrick and Marsicano, 2022).

obtain jobs that require increased educational qualifications, would-be students would instead suffer earnings declines due to under-employment, or finance those credentials through other loans with less favorable terms.

7 Conclusion

This paper estimates the effect of the Great Recession on student indebtedness, delinquency, and non-repayment by applying the empirical approach adopted in the labor market hysteresis literature. Exposure to the Great Recession shock is on the basis of geography. Individuals who lived in more-shocked commuting zones in 2009 subsequently increased their student indebtedness and related outcomes, net of covariates. These effects continue to amplify throughout the length of the panel (through 2019), unlike prior findings on earnings and employment following the Great Recession that show its negative impact dissipating by the late 2010s. Our findings further suggest that increased attainment thanks to labor market credentialization and declining state funding for public institutions on a per-FTE basis are the likeliest mechanisms that would explain these results. Our conclusion is that the negative welfare impact of the Great Recession continues to be felt through its legacy of student indebtedness, particularly so for student borrowers who responded to credentialization by increasing their educational attainment.

References

Barr, Andrew, and Sarah Turner. 2018. "A Letter and Encouragement: Does Information Increase Postsecondary Enrollment of UI Recipients?" *American Economic Journal: Economic Policy*, 10(3): 42–68.

Black, Sandra E, Jeffrey T Denning, Lisa J Dettling, Sarena Goodman, and Lesley J Turner. 2020. "Taking It to the Limit: Effects of Increased Student Loan Availability on Attainment, Earnings, and Financial Well-Being." National Bureau of Economic Research Working Paper 27658.

Black, Sandra E., Lesley J. Turner, and Jeffrey T. Denning. 2023. "PLUS or Minus? The Effect of Graduate School Loans on Access, Attainment, and Prices." National Bureau of Economic Research Working Paper 31291.

CBO. 2020. "Income-Driven Repayment Plans for Student Loans: Budgetary Costs and Policy Options." Congressional Budget Office 55968.

Cellini, Stephanie Riegg, and Nicholas Turner. 2019. "Gainfully Employed?: Assessing the Employment and Earnings of For-Profit College Students Using Administrative Data." *Journal of Human Resources*, 54(2): 342–370.

Chakrabarti, Rajashri, Nicole Gorton, and Michael F. Lovenheim. 2020. "State Investment in Higher Education: Effects on Human Capital Formation, Student Debt, and Long-term Financial Outcomes of Students." National Bureau of Economic Research Working Paper 27885.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo. 2018*a*. "Housing Booms and Busts, Labor Market Opportunities, and College Attendance." *American Economic Review*, 108(10): 2947–2994.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo. 2018*b*. "Housing Booms, Manufacturing Decline and Labour Market Outcomes." *The Economic Journal*.

Chetty, Raj, John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan. 2020. "Income Segregation and Intergenerational Mobility Across Colleges in the United States." *The Quarterly Journal of Economics*, 135(3): 1567–1633.

Collier, Daniel A., Dan Fitzpatrick, and Christopher R. Marsicano. 2022. "Exploring the Relationship of Enrollment in Income-Driven Repayment to Borrower Demographics and Financial Outcomes." *Journal of Student Financial Aid*, 51(1): 2.

Duncan, Arne. 2009. "Key Policy Letters Signed by the Education Secretary and Deputy Secretary."

Forsythe, Eliza. 2022. "Why Don't Firms Hire Young Workers During Recessions?" *The Economic Journal*, 132(645): 1765–1789.

Haltiwanger, John C., Henry R. Hyatt, Lisa B. Kahn, and Erika McEntarfer. 2018. "Cyclical Job Ladders by Firm Size and Firm Wage." *American Economic Journal: Macroeconomics*, 10(2): 52–85.

Hershbein, Brad, and Bryan A. Stuart. 2023. "The Evolution of Local Labor Markets After Recessions." *American Economic Journal: Applied Economics*.

Looney, Adam, and Constantine Yannelis. 2015. "A Crisis in Student Loans? How changes in the characteristics of borrowers and in the institutions they attended contributed to rising loan defaults." *Brookings Papers on Economic Activity*, Fall.

Mian, Atif, and Amir Sufi. 2014. "What Explains the 2007-2009 Drop in Employment?" *Econometrica*, 82(6): 2197–2223.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics*, 128: 1687–1726.

Mitchell, Josh. 2021*a*. *The Debt Trap*. New York:Simon & Schuster.

Mitchell, Josh. 2021*b*. "Is the U.S. Student Loan Program Facing a $500 Billion Hole? One Banker Thinks So." *Wall Street Journal*.

Modestino, Alicia, Daniel Shoag, and Joshua Ballance. 2016. "Downskilling: Changes in Employer Skill Requirements Over the Business Cycle." *Labour Economics*, 41(C): 333–347.

Modestino, Alicia Sasser, Daniel Shoag, and Joshua Ballance. 2020. "Upskilling: Do Employers Demand Greater Skill When Workers Are Plentiful?" *The Review of Economics and Statistics*, 102(4): 793–805.

Morgan, Julie, and Marshall Steinbaum. 2018. "The Student Debt Crisis, Labor Market Credentialization, and Racial Inequality." Roosevelt Institute.

Rinz, Kevin. 2022. "Did Timing Matter? Life Cycle Differences in Effects of Exposure to the Great Recession." *Journal of Labor Economics*, 40(3): 703–735.

Schmidt, Erik. 2018. "Postsecondary Enrollment Before, During, and Since the Great Recession." U.S. Census Bureau Current Population Reports P20-580, Washington, DC. Section: Government.

Stuart, Bryan A. 2022. "The Long Run Effects of Recessions on Education and Income." *American Economic Journal: Applied Economics*, 14(1): 42–74.

Yagan, Danny. 2019. "Employment Hysteresis from the Great Recession." *Journal of Political Economy*, 127(5): 2505–2558.

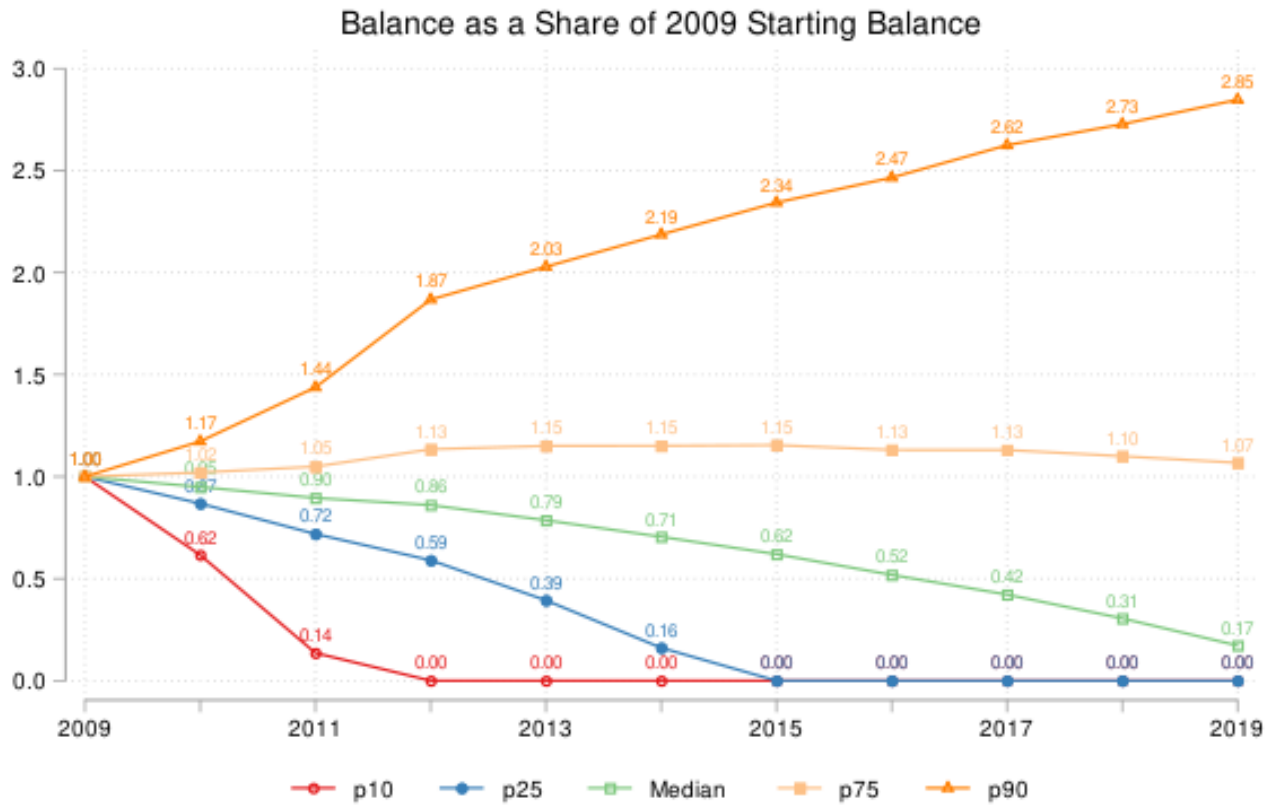

Figure 1. For our panel of individuals, all of whom have student debt when we first observe them in 2009, this chart shows the distribution of outstanding student loan balance as a share of initial 2009 balance. Individuals who have paid off all their student loans by a given year are zeroes. Individuals who have more student debt outstanding in a given year than they did when first observed in 2009 have a ratio greater than 1. Quantiles are computed separately year-by-year and plotted over time (thus each line does not represent any one individual).

Figure 2. Commuting zone-level maps.

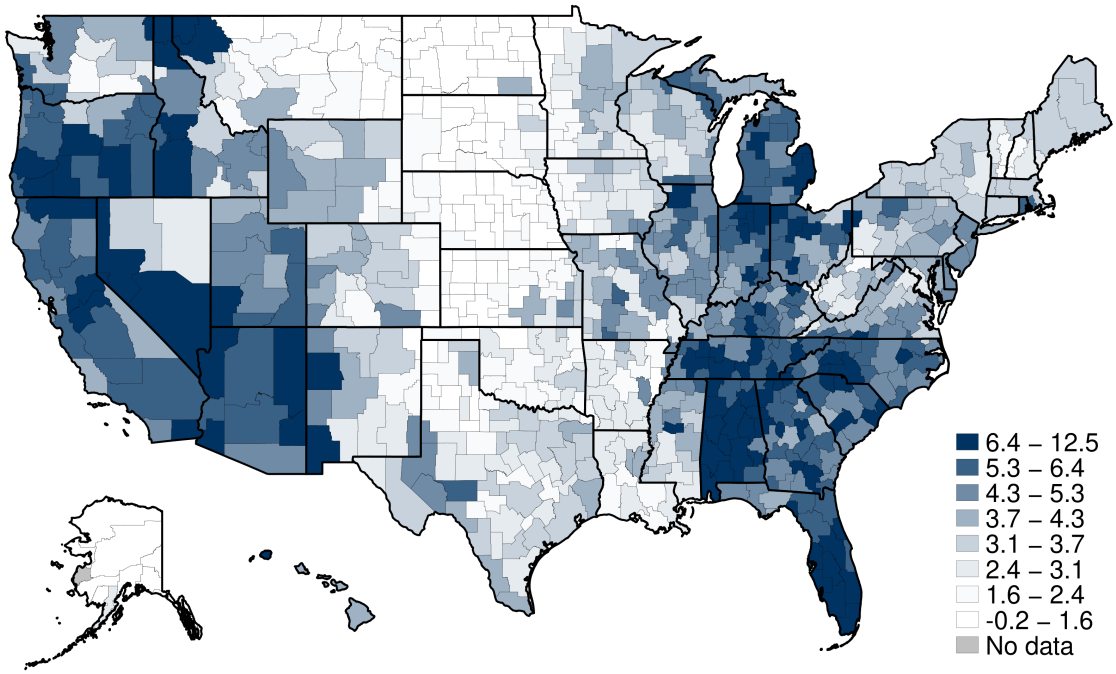

(A) This maps the Great Recession shock, as defined for the main specification in equation 4.1. It displays the percentage point change in unemployment rate by commuting zone, between 2007 and 2009, based on data from the Local Area Unemployment Statistics.

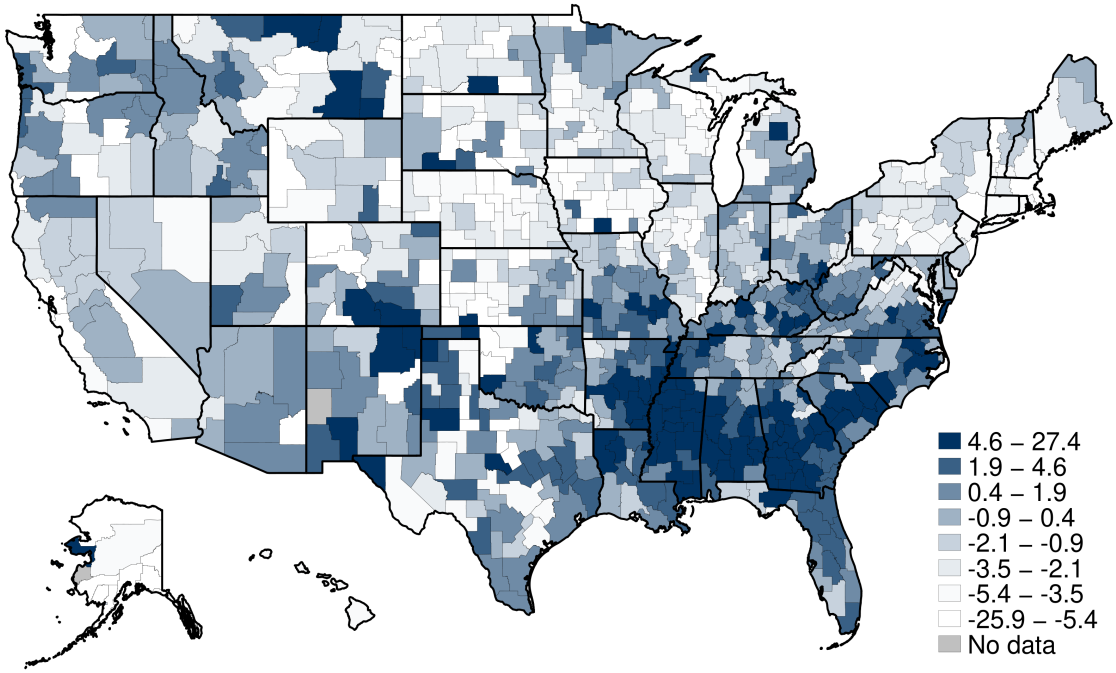

(B) This maps the change in average student loan balance by commuting zone between 2009 and 2019 for our sample of borrowers.

Figure 3. These figures visualize the estimates of the key coefficient $\hat{\beta}$ in equation 4.1 on the first four outcomes of interest listed in section 3.1.

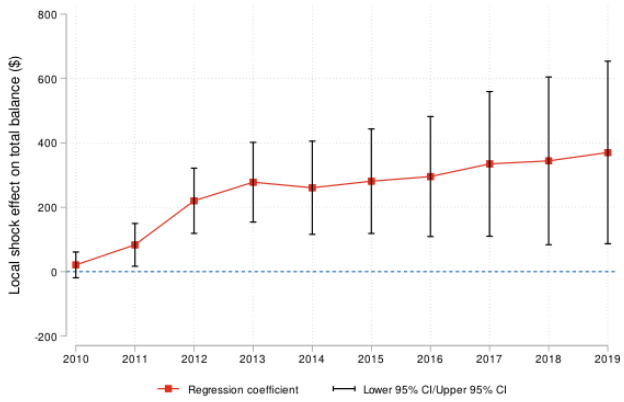

(A) Change in the total student loan balance, 2009-2019.

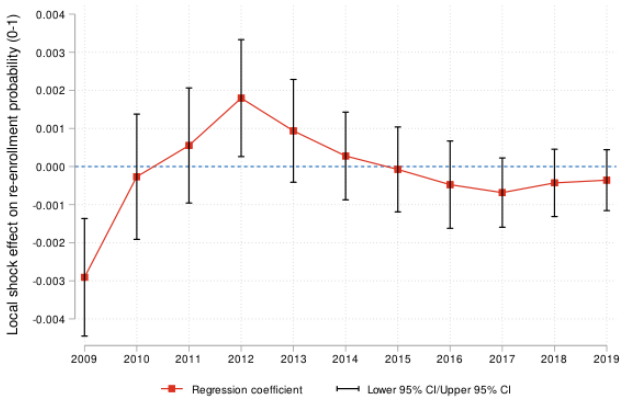

(B) Re-enrollment

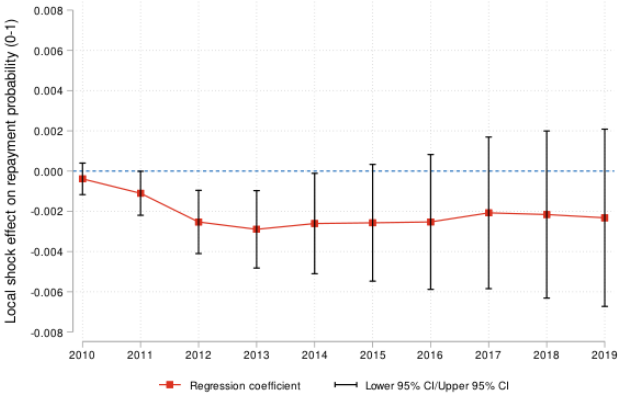

(C) Completed repayment.

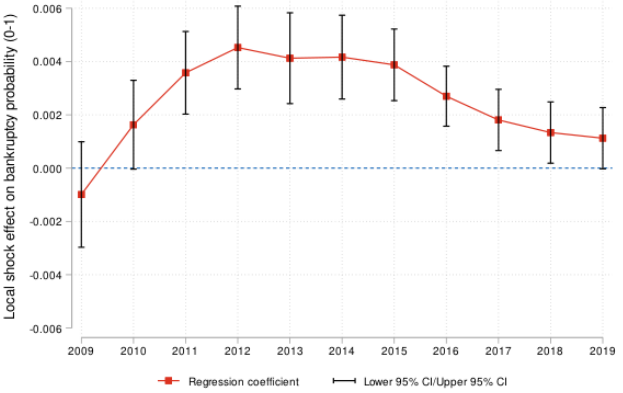

(D) Bankruptcy.

Figure 4. These figures visualize the estimates of the key coefficient $\hat{\beta}$ in equation 4.1 on the last three (delinquency-related) outcomes of interest listed in section 3.1.

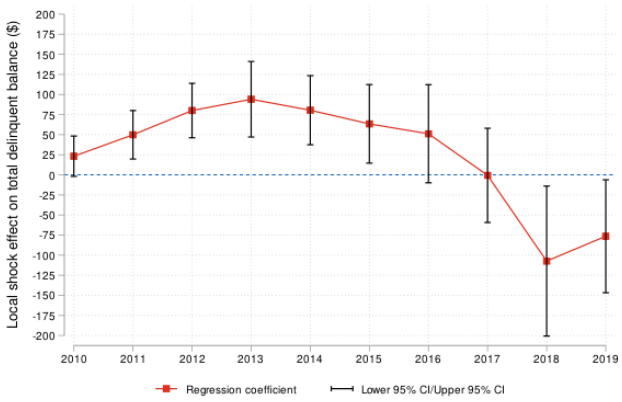

(A) Change in delinquent balance, 2009-2019.

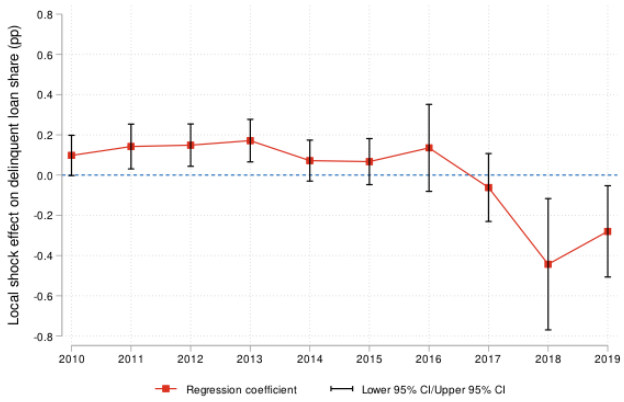

(B) Change in delinquency ratio, 2009-2019.

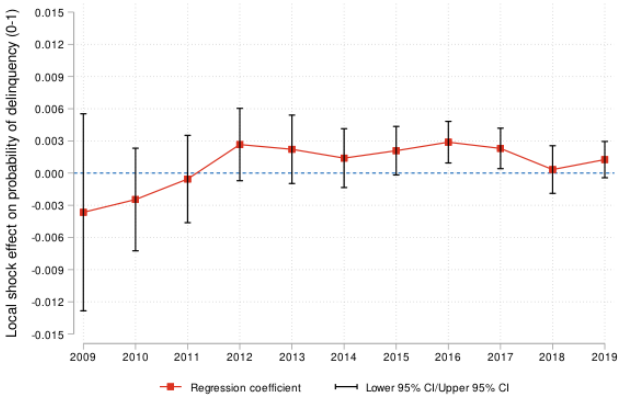

(C) Any delinquency.

Figure 5. These figures plot the median loan balance of students upon entering repayment (across institutions) and the two-year cohort default rate reported in College Scorecard. For the median loan balance, each year t pools students in two cohorts: those in FY_t and FY_{t-1}, where FY_t ends on September 30 of year t. For the default rate, the value in year t corresponds to the outcome measured in FY_t, for the cohort cohort that entered repayment during FY_{t-1}.

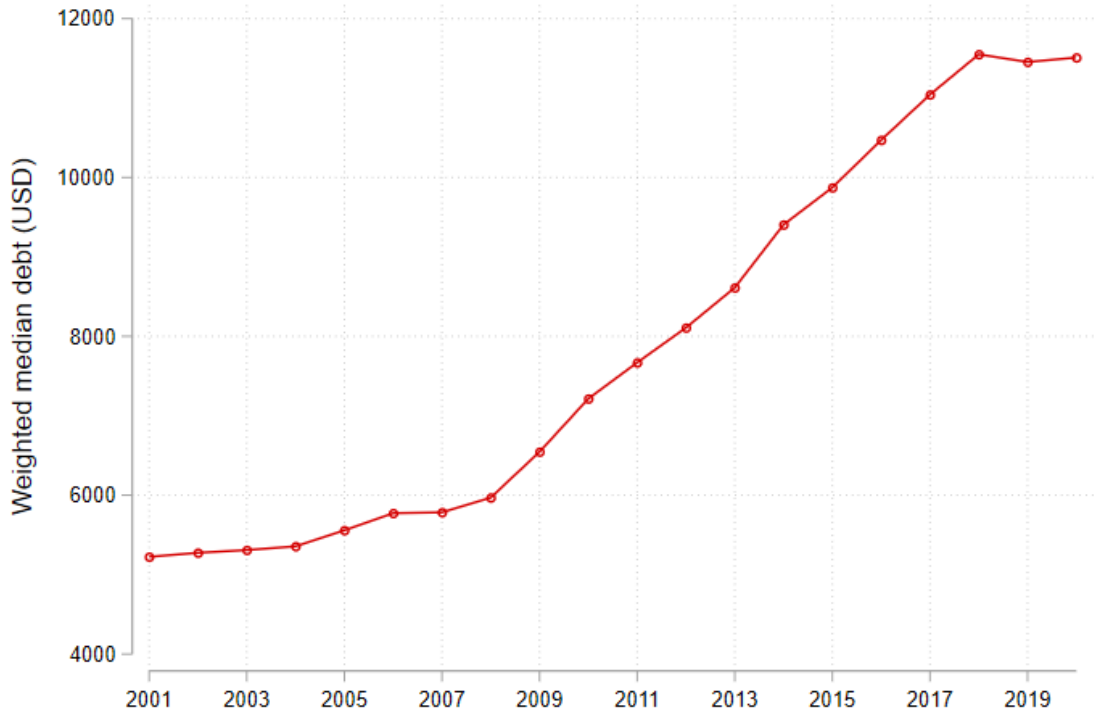

(A) Median Loan Balance.

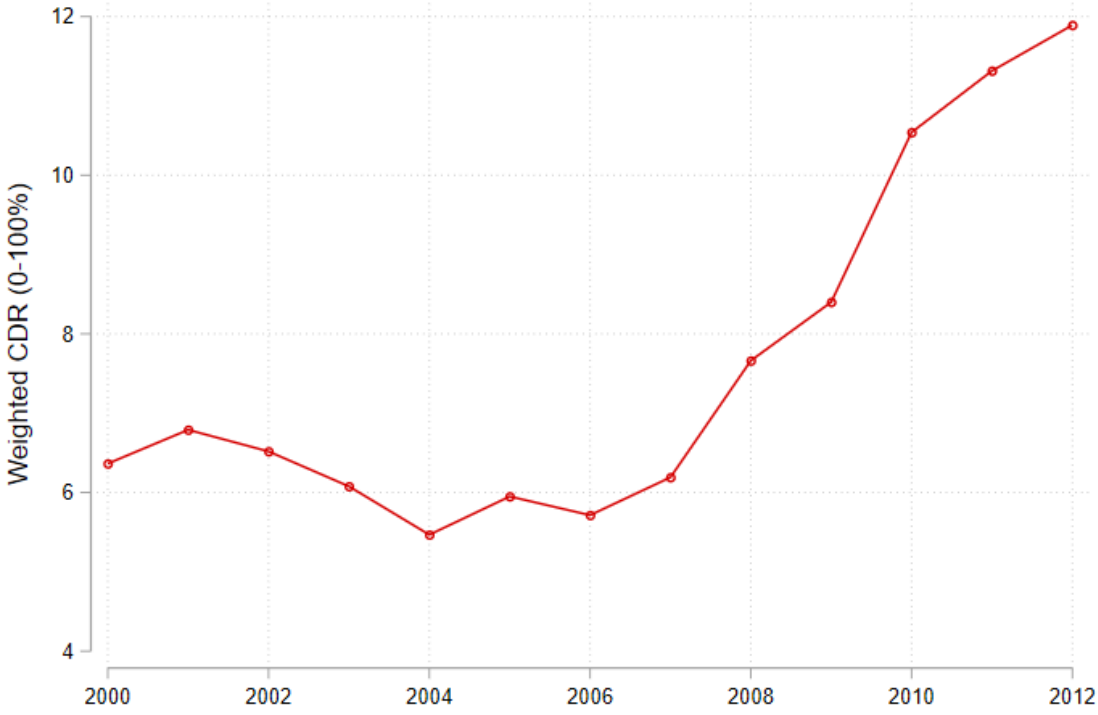

(B) Two-year cohort default rate.

Figure 6. These figures plot the estimated coefficient $\hat{\beta}$ from equation 4.2 using College Scorecard data on institution-level median student loan balance by cohort, two-year cohort default rate, and the number of students in each median loan balance cohort.

(A) Median loan balance.

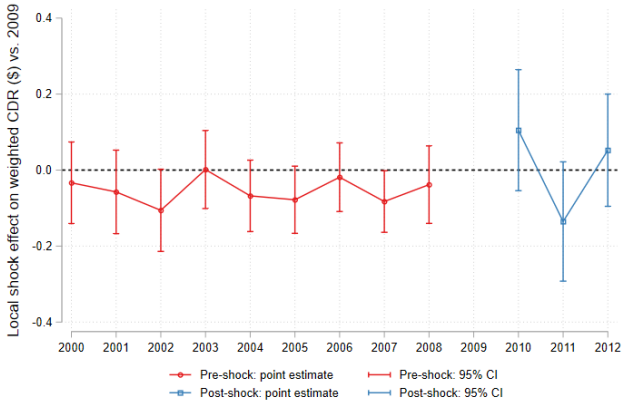

(B) Two-year cohort default rate.

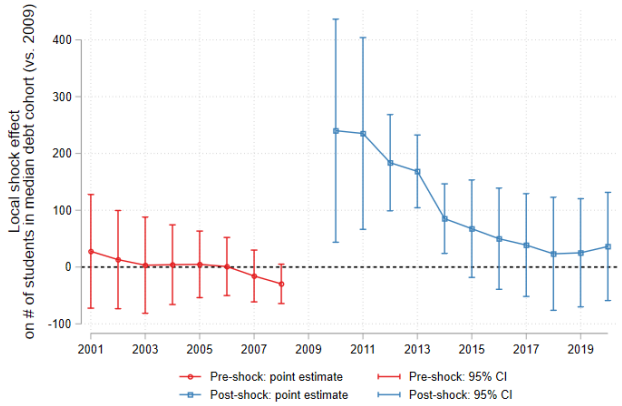

(C) Number of Students in Median Balance Cohort.

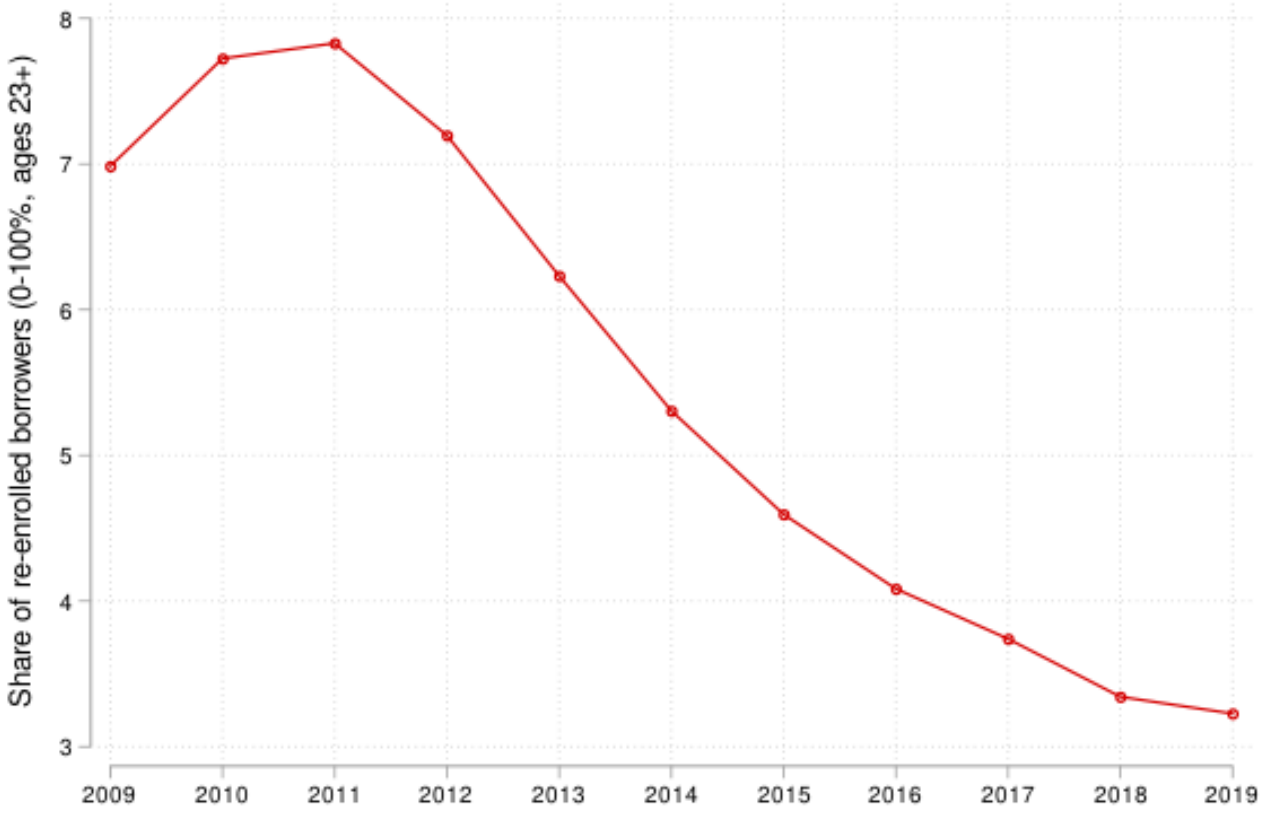

Figure 7. This figures displays, for those in the sample who are 23 or older in each year, the share that took out a re-enrollment loan over the previous year, for the period between 2009 and 2019.

Figure 8. This figure visualizes the estimation of the key coefficient of interest in equation 4.1, depicting the effect of the Great Recession shock on the probability of having re-enrolled, for each year, separately for borrowers who were less than 23 years old in 2009 and for borrowers 23 or older.

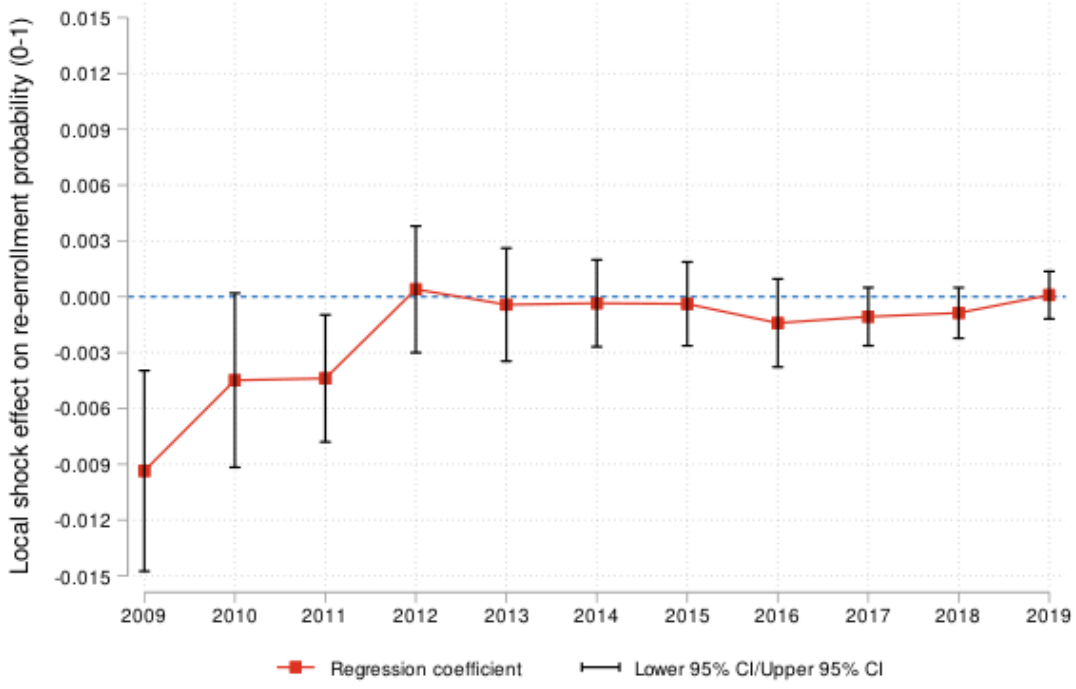

(A) Less than 23 in 2009

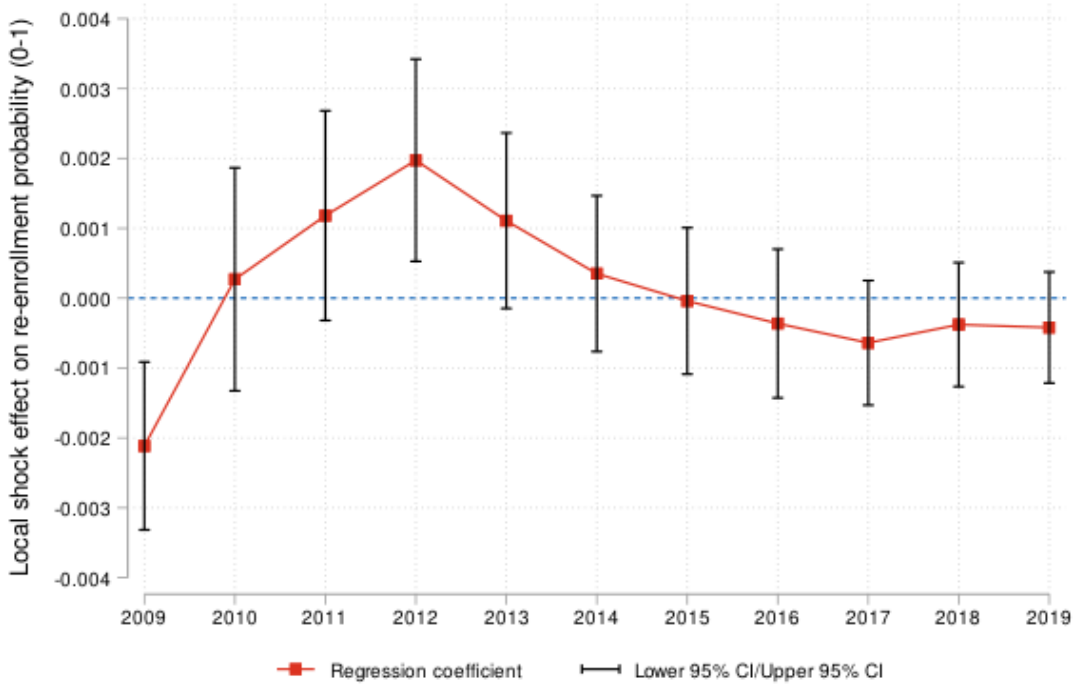

(B) 23+ in 2009

Table 1. Summary statistics for the full panel of student borrowers used for the main analysis.

Outcome	
Number of borrowers	787,154
Average age (2009)	28
Average loan balance (2009)	$22,805
Median loan balance (2009)	$12,988
Average loan balance (2019)	$21,456
Median loan balance (2019)	$2,440
% Finished repayment by 2019	46%
% Ever delinquent (2009-2019)	38%

Table 2. This table reports the result of estimating equation 4.1. We collect all seven outcomes of interest for the outcome year 2019 in one table. All specifications include controls for the borrower's gender, 2009 student loan balance, contemporaneous unemployment rate at the commuting zone of residence in 2019; they also include (Cohort)x(Age) fixed effects, defined as a triple interaction between (i) the year of origination for the oldest student loan, (ii) the year of origination of the most recent student loan up to 2009, and (iii) the borrower age in 2009.

	(1) Change in total balance (vs. 2009)	(2) Change in delinquent balance (vs. 2009)	(3) Change in delinquent ratio (vs. 2009)	(4) Any delinquency (2019)	(5) Completed repayment (2019)	(6) Re-enrolled over last year (2019)	(7) Any bankruptcy (2019)
Unemp. rate shock	370.25**	-76.45**	-0.28**	0.0013	-0.0023	-0.0004	0.0011*
	(144.43)	(35.81)	(0.12)	(0.0009)	(0.0022)	(0.0004)	(0.0006)
Contemporaneous CZ unemp. rate	268.89	65.81	0.26*	0.0086***	-0.0096***	0.0016***	0.0010*
	(188.02)	(41.13)	(0.15)	(0.0011)	(0.0024)	(0.0006)	(0.0006)
Female	4,900.01***	916.36***	-0.71***	0.0061***	-0.0813***	0.0192***	0.0082***
	(165.79)	(47.23)	(0.18)	(0.0008)	(0.0018)	(0.0007)	(0.0006)
2009 loan balance (000s USD)	-313.89***	14.54***	0.06***	-0.0000	-0.0040***	-0.0002***	-0.0001***
	(11.57)	(1.35)	(0.00)	(0.0000)	(0.0001)	(0.0000)	(0.0000)
Observations	787,154	441,243	352,311	787,154	787,154	787,154	787,154
R-squared	0.102	0.008	0.013	0.026	0.134	0.009	0.005
Cohort x Age FEs	Y	Y	Y	Y	Y	Y	Y

*** p<0.01; ** p<0.05; * p<0.10
Notes: Robust standard errors clustered at the level of the 2009 CZ of residence.

Table 3. Summary statistics for the student borrowers who never re-enrolled after 2009 versus those who did.

Outcome	Never re-enrolled	Re-enrolled
Number of borrowers	566,699	220,455
Share of sample	72%	28%
Average age (2009)	28	27
Average loan balance (2009)	$21,976	$24,938
Median loan balance (2009)	$12,417	$14,715
Average loan balance (2019)	$10,116	$50,605
Median loan balance (2019)	$0	$33,541
% Finished repayment by 2019	59%	13%
% Ever delinquent (2009-2019)	34%	49%

Table 4. Summary statistics by student borrower age group (in 2009).

Outcome	17-24	25-29	30-34
Number of borrowers	195,268	328,873	263,013
Share of sample	25%	42%	33%
Share ever re-enrolled	39%	27%	22%
Average loan balance (2009)	$16,449	$23,810	$26,269
Median loan balance (2009)	$9,422	$13,934	$14,860
Average loan balance (2019)	$19,477	$21,359	$23,046
Median loan balance (2019)	$2,092	$2,250	$3,082
% Finished repayment by 2019	46%	46%	46%
% Ever delinquent (2009-2019)	47%	36%	35%

Table 5. Summary statistics by the origination year of each borrower's most recent student loan prior to 2009 when they are first observed.

Outcome	1992-2003	2004-2006	2007-2009
Number of borrowers	159,502	308,588	319,064
Average age (2009)	30	28	26
Average loan balance (2009)	$13,226	$21,811	$28,556
Median loan balance (2009)	$7,249	$13,322	$17,013
Average loan balance (2019)	$9,863	$16,946	$31,613
Median loan balance (2019)	$0	$1,358	$10,507
% Finished repayment by 2019	65%	47%	35%
% Ever delinquent (2009-2019)	24%	30%	53%

Online Appendix

A Sample Construction

Credit Reporting Data

The main sample for this study is a panel of one million student loan borrowers between the ages of 17 and 34 in 2009, sampled from the Experian credit bureau's master database. Those individuals are then followed through 2019 and observed annually on the same date, June 30th.[19]

We see both individual-level "consumer" data and loan-level "trade" data for each of the million borrowers in the panel. Observations at each level have a stable identifier across years, and loans are linked to the same consumer-level ID. Experian computes consumer-level variables itself based on the trade data, which it sources mostly from loan servicers. We follow that logical ordering here, i.e., we start by explaining our interpretation of observations in the trade data, then explain how we aggregate to consumers and use the consumer-level demographic variables Experian reports, e.g. location, age, and gender.

Loan-level Data

Categorization of Loans: We categorize student loans by repayment status as follows

1. In deferment: *enhanced special comment* code 29, "payment deferred."
2. In forbearance: *enhanced special comment* code CP, "account in forbearance."
3. Repaid due to refinancing/consolidation: *balance amount* = 0 and *account condition code* 05 ("account transferred to another office") or 10 ("account renewed or refinanced")

[19] Even though the sampling date from Experian's database is always June 30th, each observation also has a "reporting date" which is the date on which that record was last updated in Experian's database. Hence we define observations as "outdated" if their reporting date is more than a year before their sampling date.

4. Repaid due to being paid off: *balance amount* = 0 and any *account condition code* other than 05 and 10.

5. In repayment: All remaining loans.

Delinquency and default: We categorize a loan as being in delinquency or default if *enhanced payment status* is in the ranges 22-29, 42, 67, 69, 71-84, or 87-98. But any loan with a zero balance is assumed not to be delinquent, regardless of *enhanced payment status*.

Missing loan balance: Here we explain how we impute loan balances for observations in which that outcome is missing.

- If *account condition code* = A2, "paid account/zero balance," the balance amount is set to 0.
- If *account condition code* = 10, "account renewed or refinanced," the balance amount is set to 0.
- If *account condition code* = 05, "account transferred to another office," the balance amount is set to 0. (We think this code also corresponds to a refinancing, i.e. a new loan with a different identifier.)
- If *account condition code* = A3, "closed account," the balance amount is set to 0.
- *account condition code* = 88 ("claim filed with government for insured portion of balance on defaulted loan"), 92 ("Claim filed for insured portion of balance"), 67 ("Debt included in or discharged through Chapter 7, 11 or 12 Bankruptcy"), 68 ("Account legally paid in full for less than the full balance"), 69 ("Debt included in or discharged through Chapter 13 Bankruptcy"), and 97 ("Unpaid balance reported as a loss"), the balance amount is set to zero. The implication is that all of these codes signify that the borrower's responsibility to repay the loan has been extinguished, and hence it is properly removed from that borrower's balance sheet.

- If *account condition code* = A1, "open account," multiple rules are used to impute positive balances where the reported balance is missing:
  - If balance amount in $t+1$ is positive and non-missing, and the reporting date exceeds the origination date for the loan, assign the balance amount the corresponding value from $t+1$; loop over multiple times to finish this iterative process.
  - If balance amount in $t+1$ is positive and non-missing, and the reporting date does not exceed the origination date for the loan, set the balance amount to zero; loop over multiple times to finish this iterative process.
  - If balance amount in $t+1$ is zero, then set the balance amount in t to zero.
  - Repeat this process in reverse chronological order from 2019 to 2010.
  - For loans that continue to have a missing balance but are observed with a positive, non-missing balance in $t-1$, that $t-1$ balance is imputed forward in chronological order.
  - Open loans for which we only ever observe a missing balance were dropped from the sample, since it's not possible to tell the balance.
- If *account condition code* = 93 ("account seriously past due/account assigned to attorney, collection agency or credit grantor's internal collections department") multiple rules are used:
  - If the loan is coded as 44 under *enhanced special comment* ("Student loan permanently assigned to governmen"), they are set to 0.
  - If the loan has a non-missing balance amount in $t-1$, that balance is extended until 2021 at the value available in $t-1$.

Consumer-level Data

Cohort: We assign borrowers to cohorts based on the origination year of their earliest and latest student loan observed in 2009.

Total Balance: We compute a borrower's total outstanding balance in a given year as the sum of the outstanding balance on each individual loan. Where an individual borrower has a zero total balance in year t and positive balances in year $t-1$ and $t+k$, and we do not observe a re-enrollment loan to account for that subsequent increase in balance, we interpolate the total balance at the consumer level on the assumption that the relevant loans are unobserved in the zero-balance year(s) due to failure to track loans, for example following transfers from one servicer to another.

Delinquent Balance: We compute a borrower's total delinquent balance as the sum of the balance on each loan coded as delinquent, unless the loan is also coded as in deferment. The share of loans coded as both delinquent and in deferment is very small. If a borrower's loans are all in deferment, this variable is set to missing.

Any Delinquency: Coded "1" if the borrower has any loans that are coded as delinquent and not in deferment.

Re-enrollment: We aim to distinguish borrowers who incur new student loans after 2009 due to "re-enrolling" in higher education. We determine whether an individual borrower re-enrolled after 2009 based on the following: we first calculate the gap between the reference year and the year of the borrower's most-recently-opened student loan. We then use the consumer-level variable *stu0807*, "number of deferred student trades opened in the last 12 months." If that number is greater than 0 and the gap between the reference year and the year of the most recently-opened student loan is 0 or 1, we determine that that borrower re-enrolled in the reference year. This procedure is based on the assumption that a "new" loan is a refinancing loan

if it is not deferred. We also define a cumulative variable for re-enrollment, whether a given borrower ever re-enrolled according to the above definition during the entire sample 2010-2019.

The share of borrowers who re-enroll (on a non-cumulative basis) peaks in 2010 at about 9% of the sample, and steadily diminishes to 3% thereafter.

Age: We drop borrowers with inconsistent age profiles across years of the sample, since in those cases it appears that the borrower identifier is in fact associated with more than one individual (and hence other demographic information, such as location, may also not be accurate).

Gender: We believe that Experian imputes its gender assignment based on an individual's first name. To address the instances where individuals' gender was not consistent over time, everyone was assigned their modal gender. With this correction, 52% of our sample is coded as female, 41% as male, and 7% as unknown.

Location: The implementation of our empirical strategy requires assigning borrowers to commuting zones in each year—the 2009 commuting zone to be able to assign the value of the Great Recession shock, and every subsequent year in order to be able to assign a local unemployment rate based on the borrower's current location. Assigning commuting zones requires knowing each individual's county of residence in 2009. The credit reporting data does not make county information directly available, but it includes geographic identifiers for core-based statistical areas (CBSA), states, zip codes, and Census tracts. This allows for a precise match at the county level in 98% of the observations for the full 2009-2019 period.

An additional limitation of the original data is that the Census tract and CBSA information for 2009 is incorrect, as it was defined as identical to that of 2010, although the state and zip code data are correct for every year. For 2009, the approach was as follows: (i) for the cases where the zip code and state are identical in 2009 and in 2010 (approximately 70% of the 2009 sample), we assume the remaining geographic identifiers are also identical to those of 2010; (ii)

for the cases where the zip code or the state change, we assign the individual to the county with the highest share of residences for that particular zip code, based on the US Housing and Urban Development (HUD) 2010 zip code-to-county crosswalk.

LAUS data

LAUS reports employment, unemployment, and labor force counts at the county level, which we aggregate to commuting zones using the 2000 crosswalk made available by the US Department of Agriculture's Economic Research Service.

The commuting-zone-level change in the unemployment rate during the Great Recession (the "Great Recession shock") can be computed in a very straightforward way using the county-level variables for labor force, employed, and unemployed population that are available for every year.

B Supplemental Tables and Figures

Figure B.1. For our panel of individuals, all of whom have student debt when we first observe them in 2009, these figures show the distribution of outstanding student loan balance as a share of initial 2009 balance for those who are 23 or above in 2009, who 1. never re-enrolled during the 2009-2019 period, and 2. re-enrolled at some point after 2009. Individuals who have paid off all their student loans by a given year are zeroes. Individuals who have more student debt outstanding in a given year than they did when first observed in 2009 have a ratio greater than 1. Quantiles are computed separately year-by-year and plotted over time (thus each line does not represent any one individual).

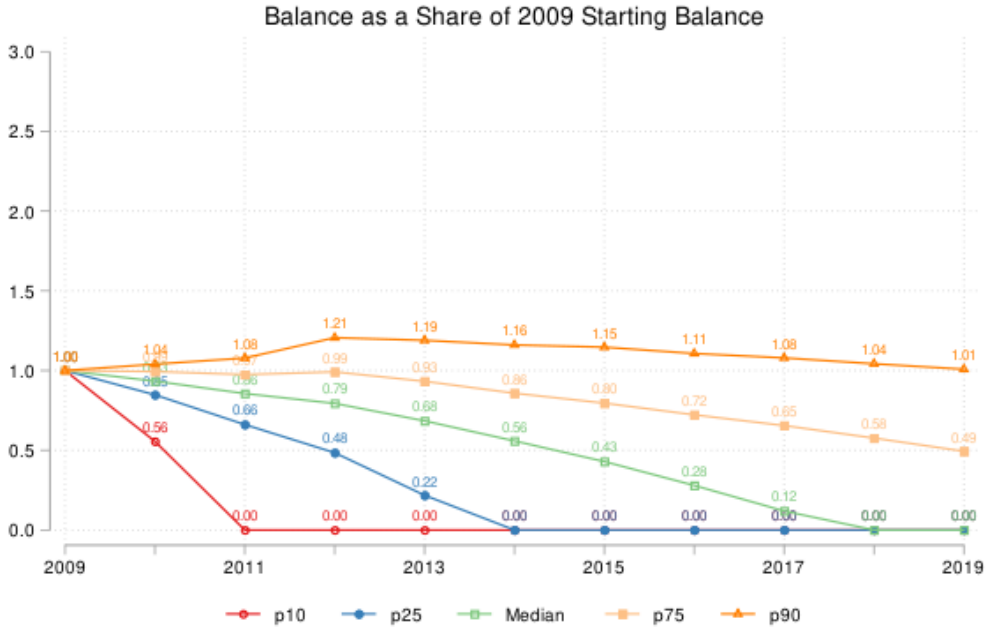

Never re-enrolled

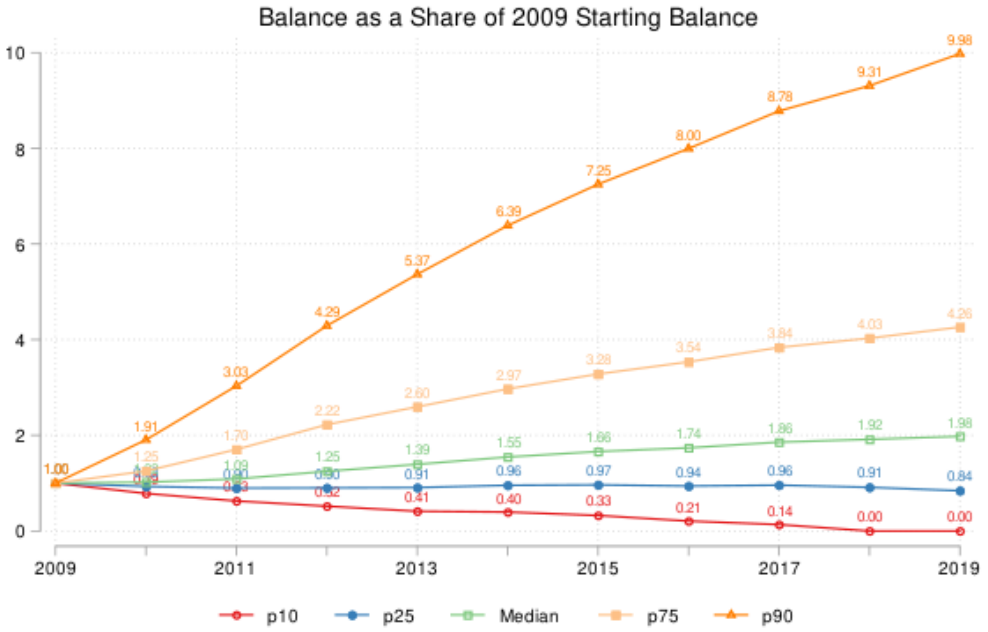

Re-enrolled after 2009

Figure B.2. This figure visualizes the estimation of the key coefficient of interest in equation 4.1, depicting the effect of the Great Recession shock on the change in total balance since 2009, separately for (i) borrowers who did not re-enroll in 2009 or later, and (ii) for those who did re-enroll at any point betwen 2009 and 2019. The specification that generates these coefficients includes the same controls and fixed effects as those mentioned in Table 2.

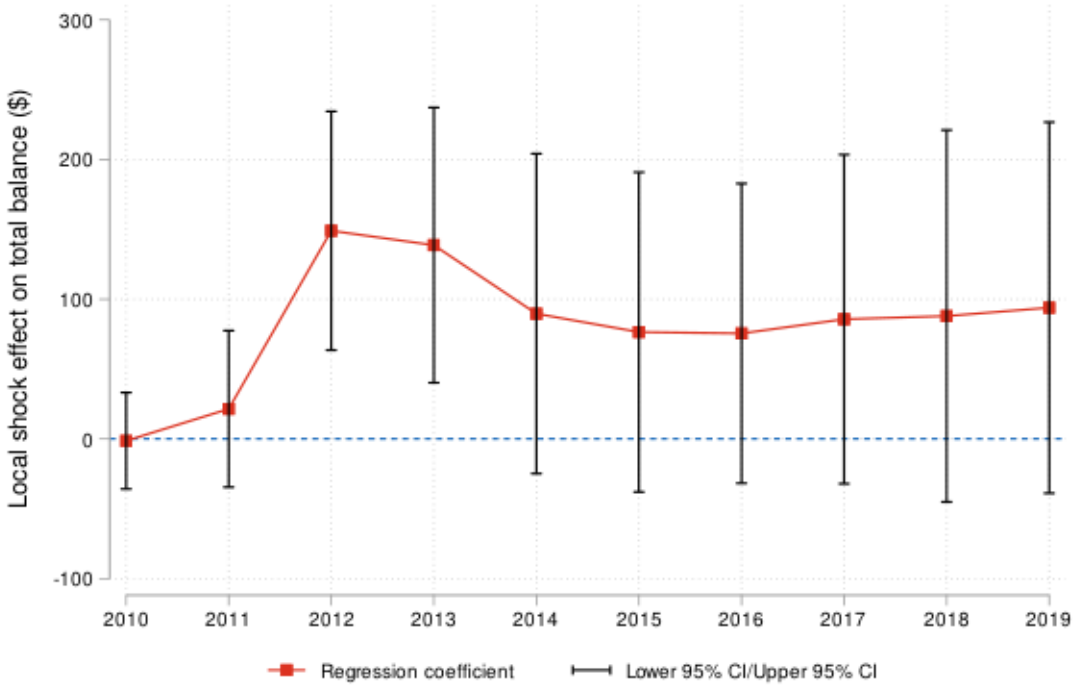

(A) Never re-enrolled between 2009 and 2019

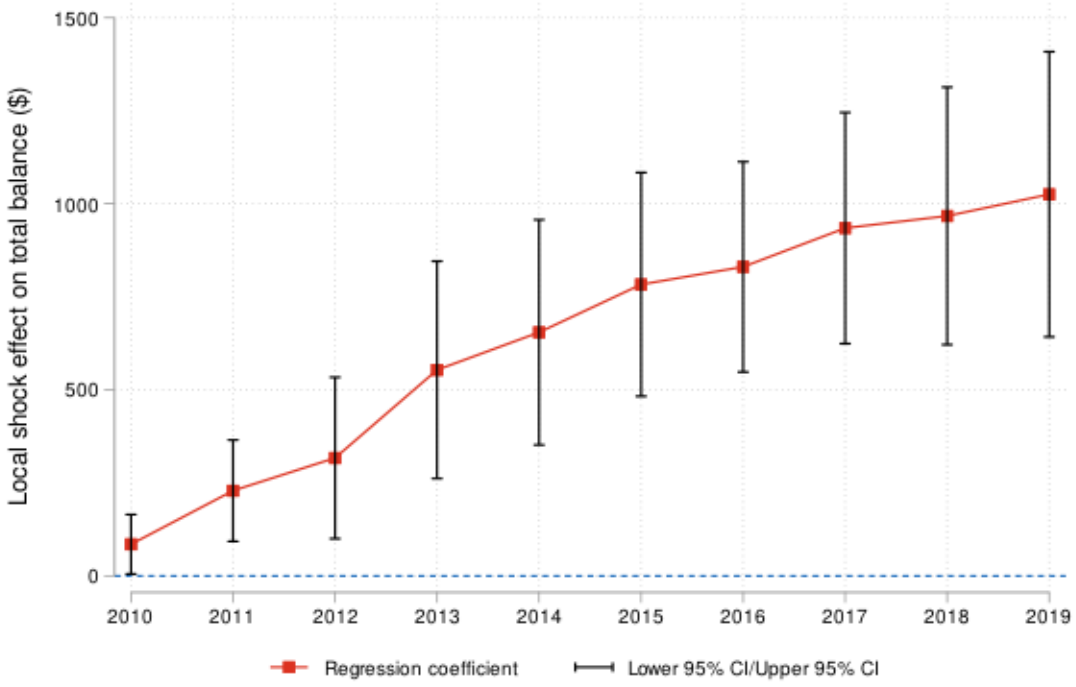

(B) Re-enrolled between 2009 and 2019

Figure B.3. This figure reports 1. the total state and local funding per FTE at public institutions, and 2. the total number of public, private nonprofit, and private for-profit institutions nationally, both for 2009-2019.

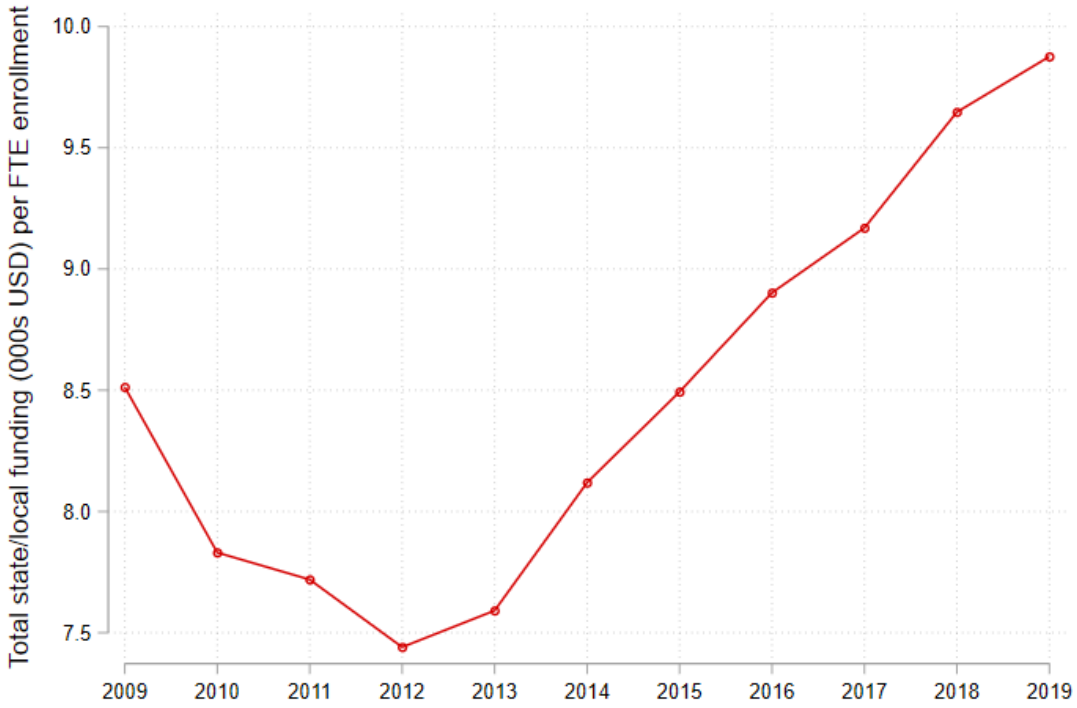

(A) Total state and local funding per FTE at public institutions.

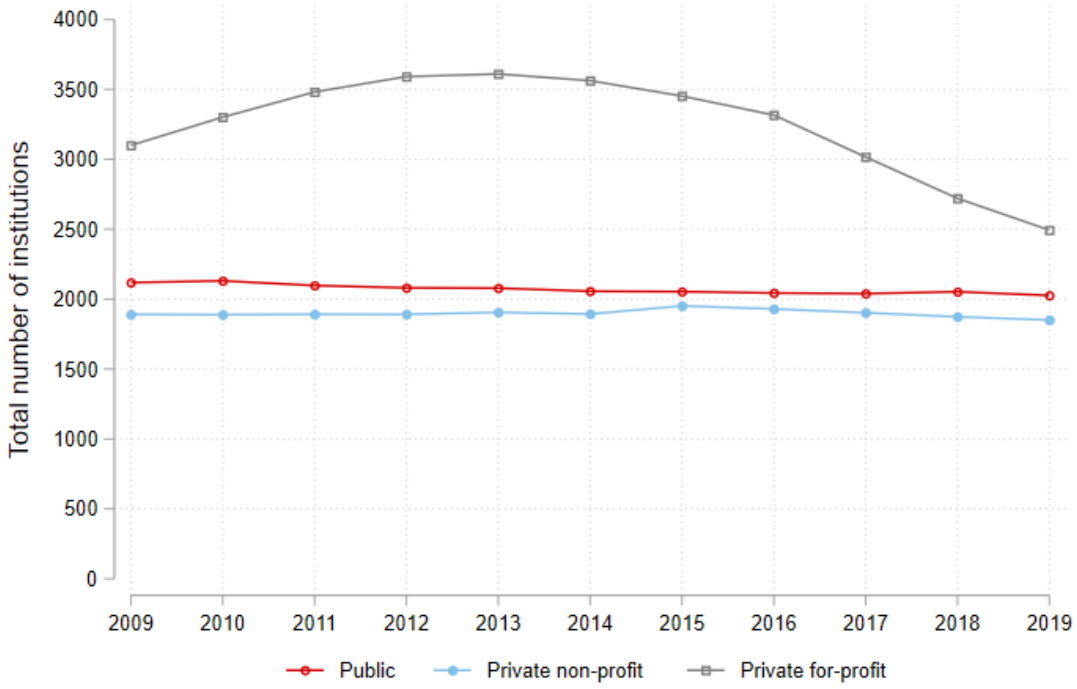

(B) Total number of institutions by type.

Figure B.4. This figure visualizes the estimation of the key coefficient $\hat{\beta}$ in equation 5.1, depicting the effect of the Great Recession shock on state funding per FTE, entry of for-profit institutions, and enrollment at for-profit and public institutions. The regressions from which these key coefficients of interest are estimated include controls for the contemporaneous CZ unemployment rate and for the level of the outcome of interest in 2009.

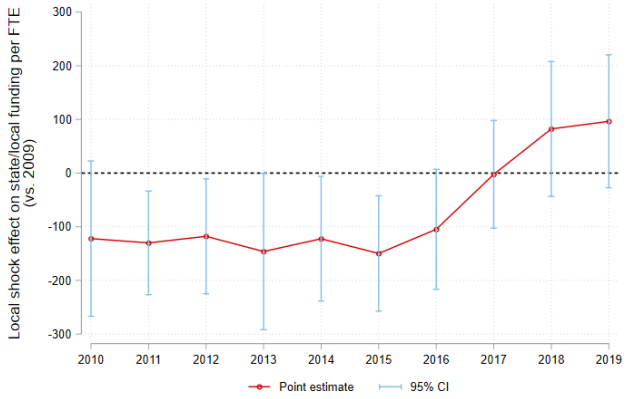

(A) State and local funding per FTE.

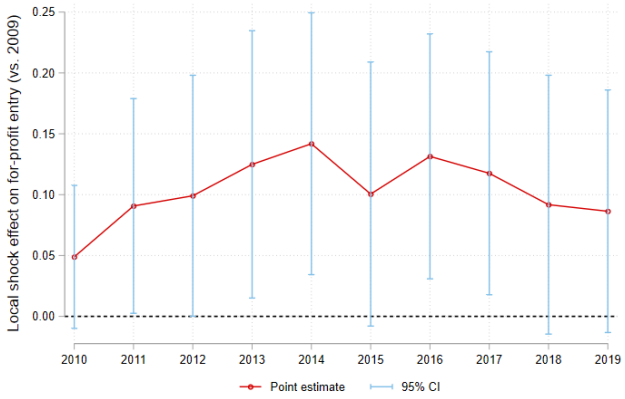

(B) Number of For-Profits.

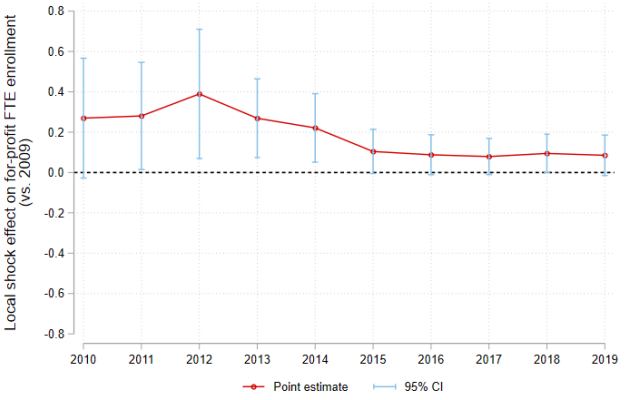

(C) FTE Enrollment at For-Profits.

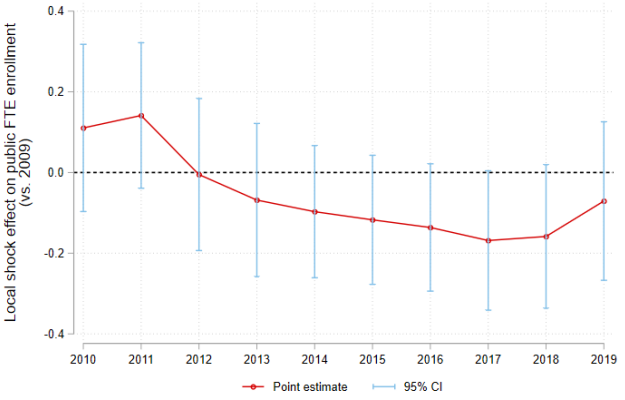

(D) FTE Enrollment at Publics.