

The rise in student loan defaults in the Great Recession*

Holger M. Mueller[†] Constantine Yannelis[‡]

June 2018

Abstract

We examine the rise in student loan defaults in the Great Recession by linking administrative student loan data at the individual borrower level to student loan borrowers' individual tax records. A Blinder-Oaxaca style decomposition shows that shifts in the composition of student loan borrowers and the massive collapse in home prices during the Great Recession can each account for approximately 30% of the rise in student loan defaults. Falling home prices affect student loan defaults by impairing individuals' labor earnings, especially for low income jobs. By contrast, when comparing the default sensitivities of homeowners and renters, we find no evidence that falling home prices affect student loan defaults through a home equity-based liquidity channel. The Income Based Repayment (IBR) program introduced by the federal government in the wake of the Great Recession reduced both student loan defaults and their sensitivity to home price fluctuations, thus providing student loan borrowers with valuable insurance against negative shocks.

*We thank Toni Whited (editor), an anonymous referee, David Berger, Jan Eberly, Caroline Hoxby, Amir Kermani, Theresa Kuchler, Andres Liberman, Adam Looney, Johannes Stroebl, Jeremy Tobacman, Eric Zwick, and seminar participants at NYU, Berkeley, and the Consumer Financial Protection Bureau for helpful comments. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the US Treasury or any other organization.

[†]NYU Stern School of Business, NBER, CEPR, and ECGI. Email: hmueller@stern.nyu.edu.

[‡]NYU Stern School of Business. Email: constantine.yannelis@stern.nyu.edu.

1. Introduction

Student loan default rates rose sharply in the Great Recession after having remained stable for many years. Given the significance of student loans for the financing of higher education, this sharp rise in student loan default rates is alarming. It has important consequences not only for the federal budget—more than 92% of all student loans are federal loans—but also for the defaulting student loan borrowers. Unlike other types of loans, student loans are not dischargeable in bankruptcy, and wages can be garnished for the rest of a borrower’s lifetime. Thus, besides the usual stigma associated with loan defaults—such as tainted credit scores and limited access to credit markets—the expectation of wages being garnished may affect student loan borrowers’ job search and incentives to work, while the fact that loan defaults can be observed by employers may affect their prospects of finding a job in the first place.¹

This paper examines the rise in student loan defaults during the Great Recession by linking administrative student loan data at the individual borrower level from the U.S. Department of the Treasury to individuals’ tax records from the Internal Revenue Service’s Compliance Data Warehouse. Our student loan data represent a 4% random sample of all federal direct and guaranteed student loans. Our final sample consists of over one million annual observations of student loan borrowers who are in repayment during the years of the Great Recession.

We begin by contrasting two potential candidate explanations for the surge in student loan defaults: shifts in the composition of student loan borrowers and the collapse in home prices during the Great Recession. Along with the rise in student loan defaults, the share of nontraditional borrowers attending for-profit institutions and community colleges

¹Using panel data from the National Longitudinal Survey of Youth 1997 (NLSY97), Ji (2016) finds that student loan borrowers spend 8.3% less time on their job search relative to nonborrowers. As a result, they earn 4.2% less annually in their first ten years after graduation. Similarly, survey evidence shows that 55% of student loan borrowers age 18 to 34 “accepted a job quicker to have income sooner” due to their student debt (Earnest, 2016). As for work incentives, Dobbie and Song (2015, 1300) conclude that debt relief “maintains the incentive to work by protecting future earnings from wage garnishment.” Whether loan default impairs job seekers’ ability to find jobs is less obvious. While survey evidence shows that 47% of US employers use credit checks to screen applicants (SHRM, 2010), the empirical evidence is mixed: Bos, Breza, and Liberman (2018) and Herkenhoff, Phillips, and Cohen-Cole (2016) find a negative effect of bad credit on employment, whereas Dobbie et al. (2017) find no significant effect.

increased by 16.9% from 2006 to 2009. These borrowers are riskier and have much higher default rates than traditional borrowers attending four-year colleges (Deming, Goldin, and Katz, 2012; Looney and Yannelis, 2015). Thus, compositional shifts may be able to explain some of the rise in student loan defaults. On the other hand, the massive collapse in home prices during the Great Recession—zip code level home prices declined by 14.4% from 2006 to 2009—caused a sharp drop in consumer spending, which adversely affected labor market outcomes (Mian, Rao and Sufi, 2013; Mian and Sufi, 2014; Stroebel and Vavra, 2017; Kaplan, Mitman, and Violante, 2016; Giroud and Mueller, 2017). In addition, declining home prices may have impaired student loan borrowers’ ability to borrow against home equity (Mian and Sufi, 2011; Bhutta and Keys, 2016), thereby limiting their access to liquidity and ability to make student loan repayments. Accordingly, the massive collapse in home prices—a highly salient and widely studied feature of the Great Recession—constitutes another potential candidate explanation for the rise in student loan defaults.

We find that compositional shifts and the collapse in home prices both matter. Cross-sectionally, changes in student loan default rates are positively related to changes in the share of nontraditional borrowers and negatively related to changes in home prices. Using a Blinder-Oaxaca style decomposition, we find that each of these two potential candidate explanations can account for approximately 30% of the rise in student loan defaults in the Great Recession. That being said, the time-series evidence appears to line up particularly well with the home price explanation. While student loan default rates began to increase at the onset of the Great Recession, along with the fall in home prices, the share of nontraditional borrowers had risen steadily throughout the 2000s. In fact, it had risen in every single year from 2000 to 2006.

In the remainder of the paper, we focus on the role of home prices for the rise in student loan defaults during the Great Recession—Looney and Yannelis (2015) provide an extensive discussion of the changing nature of borrower composition in the student loan market over the last 40 years. We begin by examining the role of home prices using aggregated data at the regional level. Regions are either zip codes, counties, or commuting zones. Regardless of whether we use a long difference specification in the spirit of Mian, Rao and Sufi (2013), Mian and Sufi (2014), and Giroud and Mueller

(2017), or a panel specification that includes region fixed effects, we consistently find a negative and highly significant sensitivity of student loan defaults to changes in home prices. We proceed by employing a panel specification using disaggregated student loan data at the individual borrower level. This panel specification also constitutes our main specification throughout the rest of the paper. Regardless of whether we include zip code, zip code \times cohort year, or individual borrower fixed effects, we find a negative and highly significant relation between changes in home prices and changes in student loan defaults. Our estimates are quantitatively similar to those from the long difference specification. Moreover, our results hold across all major institution types—four-year colleges, for-profit institutions, and community colleges.

Falling home prices can affect student loan defaults through various channels. Our data allow us to examine two of these channels in more detail: declining home prices can adversely affect labor market outcomes, and they can impair student loan borrowers' liquidity by limiting their ability to borrow against home equity. We find evidence in support of the labor market channel: for borrowers with annual earnings of \$60,000 or more, there is no significant association between home prices and student loan defaults. Furthermore, the point estimates are monotonically declining across labor income groups. In addition, we find that falling home prices are associated with large drops in labor earnings at the individual borrower level, and that this relation is only significant among low income borrowers. Altogether, these results are supportive of a labor market channel operating primarily through low income jobs.

On the other hand, we find no evidence in support of a home equity-based liquidity channel. Under this channel, homeowners should have a larger default sensitivity to changes in home prices than renters. We identify homeowners through Form 1098 submitted by mortgage lenders. Thus, we are able to identify homeowners as long as they have a mortgage, regardless of whether they file for the mortgage interest deduction. Maybe somewhat surprisingly, we find that homeowners and renters both respond similarly to changes in home prices. This remains true if we account for differences in labor earnings or family income. And it also remains true if we split our sample by age or repayment cohort to account for measurement error in homeownership.

We conclude with an evaluation of the Income Based Repayment (IBR) program introduced by the federal government in 2009, in the wake of the Great Recession. The purpose of IBR is to provide student loan borrowers with insurance against negative shocks by making their loan repayments contingent on discretionary income. Eligibility is based on a means test, which requires that the student debt be sufficiently large relative to discretionary income. To assess the efficacy of the IBR program, we conduct a triple difference analysis by examining the default responses of IBR eligible versus ineligible student loan borrowers to home price changes before and after the plan’s introduction. We find that the introduction of the IBR plan reduced both student loan defaults and their sensitivity to home price fluctuations. Importantly, this effect is entirely driven by IBR eligible student loan borrowers who actually took up the IBR repayment option. In contrast, IBR eligible student loan borrowers who did not take up the IBR repayment option continued to exhibit high default rates also after 2009.

This paper is part of a growing literature that focuses on household debt and defaults. Much of this literature focuses on mortgage defaults, emphasizing, for example, the role of screening by lenders (e.g., Keys et al., 2010; Keys, Seru, and Vig, 2012; Purnanandam, 2011) and specific “default triggers,” such as interest rate changes, negative equity, and employment losses (e.g., Elul et al., 2010; Gyourko and Tracy, 2014; Gerardi et al., 2018). Student loans constitute the largest source of nonmortgage household debt in the United States, with an outstanding balance of \$1.4 trillion. And yet, compared to the literature on mortgage defaults, there is relatively little systematic evidence on the determinants of student loan defaults and, especially, on what accounts for the sharp rise in student loan default rates during the Great Recession. Our study is an attempt to fill this void. We find that both shifts in the composition of student loan borrowers and the collapse in home prices matter for the rise in student loan defaults in the Great Recession. Home prices appear to operate primarily through a labor market channel by affecting student loan borrowers’ labor earnings, especially for low income jobs. By contrast, we find no evidence that falling home prices in the Great Recession affect student loan defaults through a home equity-based liquidity channel.

Our paper is also related to studies evaluating loan modification programs introduced

by the federal government in the wake of the Great Recession. As with the literature on mortgage defaults, many of these studies focus on mortgage modification programs, such as the Home Affordable Modification Program (HAMP) (e.g., Agarwal et al., 2017). Our paper is, to the best of our knowledge, the first systematic empirical study of the IBR program rolled out by the federal government in 2009. Similar to HAMP, monthly student loan payments are capped as a percentage of borrowers' discretionary income. We find that the IBR plan was successful at reducing student loan defaults and their sensitivity to home price fluctuations, thus providing student loan borrowers with valuable insurance against negative shocks.

The rest of this paper is organized as follows. Section 2 presents the data, variables, and summary statistics. Section 3 provides a Blinder-Oaxaca style decomposition to explain the rise in student loan defaults in the Great Recession. Section 4 studies the relation between home prices and student loan defaults at the individual borrower level. Section 5 examines two channels through which home prices can affect student loan defaults: through labor markets and through home equity-based borrowing. Section 6 provides an evaluation of the IBR program. Section 7 concludes.

2. Data, variables, and summary statistics

2.1. Data

Our student loan data come from the National Student Loan Data System (NSLDS), which is the main data source used by the US Department of Education to administer federal student loan programs. The NSLDS contains information on all federal direct and guaranteed student loans, accounting for more than 92% of the student loan market in the United States. Our analysis sample constitutes a 4% random sample of the NSLDS used by the US Department of the Treasury for policy analysis and budgeting purposes, drawn using permutations of the last three digits of an individual's social security number. The sample is constructed as a panel, tracking individual student loan borrowers over time.

Our sample includes student loans from both the Direct Loan program and the Federal Family Education Loan (FFEL) loan program. The two programs have similar rules and

offer products with identical limits and interest rates, which are set by Congress. The main difference between the two programs is that under the FFEL program, capital is provided by banks. Since 2010, all federal loans have been under the Direct Loan program, but during the time period studied here, loans were issued under both programs, with schools participating in either program. Our sample includes student loans for undergraduate as well as for postbaccalaureate graduate and professional degrees. While graduate PLUS loans are included in our sample, parent PLUS loans—which are for parents and not made to students directly—are not included. Also, our sample does not include private student loans. For the purpose of our analysis, we focus on student loan borrowers who are already in repayment. Student loan borrowers typically enter into repayment within six months after leaving their degree granting institution. Student loan borrowers who are in deferment or forbearance programs are treated as being in repayment.²

We have detailed information on loan disbursements, balances, and repayment. We also know the institutions that student loan borrowers attended, such as name and institution type. Our sample includes private not-for-profit, public not-for-profit, and four-year for-profit institutions. In addition, we have demographic information on the student loan borrowers and their parents from the Free Application for Federal Student Aid (FAFSA) form, which recipients of federal student loans are required to complete.

Our NSLDS student loan data are linked to deidentified tax data from the Internal Revenue Service’s Compliance Data Warehouse (CDW). The CDW sources data from W-2s and other tax returns. Besides individual labor earnings and total income, the tax data contain information on marital status, mortgage interest payments (Form 1098 filed by mortgage lenders), and number of individuals in a household. The latter information is used to calculate the poverty level of individuals when evaluating the IBR program. Earnings are defined as Medicare wages plus self-employment earnings. Total income additionally includes nonlabor income.

We match individual student loan borrowers to home prices at the zip code level using

²Federal student loan borrowers may be entitled to a loan deferment (if they are unemployed) or a forbearance (if the amount owed exceeds 20% of their gross income). These programs allow student loan borrowers to temporarily defer making payments, and interest may or may not accrue depending on the type of loan and specifics of the deferment or forbearance program.

home price data from Zillow.³ Home prices have been the focus of much of the empirical literature on the Great Recession, and Zillow home price data have been used by, e.g., Kaplan, Mitman, and Violante (2016), Bailey et al. (2017), Di Maggio et al. (2017), and Giroud and Mueller (2017), among others. We use home price data from 2006 to 2009. Changes in home prices from 2006 to 2009 based on Zillow data are highly correlated with the “housing net worth shock” in Mian, Rao and Sufi (2013) and Mian and Sufi (2014), “ Δ Housing Net Worth, 2006–2009.” The correlation at the Metropolitan Statistical Area (MSA) level is 86.3%. They are also highly correlated with changes in home prices from 2006 to 2009 using home price data from the Federal Housing Finance Agency (FHFA). The correlation at the MSA level is 96.4%. In line with prior research, we measure home prices in December of each year.

2.2. Variables and empirical specification

Our main outcome variable is an indicator of whether a student loan is in default for the first time (“new default”). New student loan defaults constitute a flow measure. By contrast, an indicator of whether a student loan is currently in default would be a stock measure. Using a flow measure allows us to relate the incidence of default to the underlying trigger event. With a stock measure, that is not possible, as the status of being in default is not informative about when the default was triggered. In fact, using new defaults as our measure, we know almost precisely when the default was triggered. A student loan goes into default within 270 days of a payment being missed. Once a student loan goes into default, the loan servicer has up to 90 days to report the default to the NSLDS. Thus, it takes about one year between when a payment is missed and when a new default is recorded in our administrative data. To account for this time lag, we always use student loan defaults in year $t + 1$. Accordingly, our focus is on home prices from 2006 to 2009 and student loan defaults from 2007 to 2010.

³We also aggregate our data at the county and commuting zone (CZ) level. We link zip codes to counties using the crosswalk from the US Department of Urban Development, and we link counties to CZs using the crosswalk from the US Department of Agriculture Economic Research Service. Our sample includes 12,749 zip codes, 1,234 counties, and 408 CZs with available home price data.

Our empirical specification is:

$$\pi_{i,t+1} = \alpha_t + \alpha_z + \beta \text{Home price}_{z,t} + \gamma X_{i,t} + \varepsilon_{i,t}, \quad (1)$$

where $\pi_{i,t+1}$ is an indicator of whether individual i defaults in year $t + 1$; $\text{Home price}_{z,t}$ is the home price (in logs) in zip code z in year t ; $X_{i,t}$ is a vector of controls, which includes loan balance, borrowing duration, family income, school type, and Pell grant aid; and α_t and α_z are year and zip code fixed effects, respectively. The year fixed effects capture any economy wide factors, such as aggregate economic conditions. The zip code fixed effects absorb any time invariant heterogeneity across zip codes, including any differences in borrower composition, college enrollment, or student loan volume, or any heterogeneity arising from different experiences during the preceding housing boom. In some of our specifications, we also include cohort year, zip code \times cohort year, or individual borrower fixed effects. Cohort year indicates the year in which a student loan borrower enters into repayment. Standard errors are clustered at the zip code level. In robustness checks, we alternatively cluster standard errors at the county level. Observations are weighted by individual student loan balances.

2.2. Summary statistics

TABLE 1 ABOUT HERE

Table 1 presents basic summary statistics. All variables are measured over the 2006 to 2009 period. The only exception is student loan defaults, which is measured over the 2007 to 2010 period. There are 1,071,049 annual borrower-level observations associated with 298,003 individual student loan borrowers. The average student loan borrower has \$23,757 in student debt and earns \$44,930 during our sample period. Total income, which includes nonlabor earnings, amounts to \$62,369. About 8% of student loan borrowers experience a drop in labor earnings of 50% or more relative to the previous year's earnings. Given the magnitude of these earnings drops, it is likely that they are associated with employment losses. By comparison, the average annual layoff rate during the Great Recession is about

7% (Fig. 1 in Davis, Faberman, and Haltiwanger, 2012). Student loan borrowers in our sample enter into repayment between 1970 and 2009. The average repayment cohort is 2002. In any given year, about 4% of all student loans default for the first time. When comparing this number to two- and three-year cohort default rates used by the US Department of Education, we note that these differ from our student loan default rates along two dimensions.⁴ First, our student loan default rates measure the annual flow of student loans that are in default for the first time. Second, cohort default rates measure student loan defaults in the first two or three years after borrowers enter into repayment, which is a period during which a disproportionately large fraction of student loan borrowers defaults. In contrast, our student loan default rates measure defaults across all repayment years.

About 39% of student loan borrowers own a home, which is much less than the national average of 68% during the sample period. This discrepancy is likely because student loan borrowers are younger and hence earlier in their life cycle. The average zip code level home price during the sample period is \$244,882. There is significant dispersion in home prices, though, ranging from \$26,800 in Youngstown, Ohio, to \$3,799,801 in Atherton, California. To reduce the sensitivity of our estimates to outliers, we use the natural logarithm of home prices in all our regressions.

Homeowners and renters live in neighborhoods with similar home prices. Also, they come from similar repayment cohorts. The typical homeowner enters into repayment in 2001, while the typical renter enters into repayment in 2002. That said, homeowners and renters differ along some important dimensions. In particular, homeowners have larger labor earnings, total income, and family income, and they are less likely to default on their student loans. In our empirical analysis, we confirm that homeowners have lower baseline default rates than renters. However, the question we are primarily interested in is not whether homeowners default less in general—which could be due to differences in labor earnings or access to financial resources—but rather whether they are less likely to default in response to declining home prices.

⁴Cohort default rates have been historically used by the US Department of Education at the cohort by school level to penalize schools with high student loan default rates.

FIGURE 1 ABOUT HERE

Fig. 1 shows the age distribution of student loan borrowers. As Panel A shows, most student loan borrowers enter into repayment in their early to mid-twenties. However, a large fraction of student loan borrowers enters into repayment in their late twenties, thirties, and even forties, reflecting the prominent role of nontraditional borrowers—those attending for-profit and other nonselective institutions—in our administrative data. Panel B shows the age distribution of all student loan borrowers in repayment. The average borrower in our sample is 37 years old. While the typical student debt repayment plan has a duration of ten years, student loan borrowers often have the choice among alternative repayment options, which can substantially increase the duration of their loans (Avery and Turner, 2012). For instance, by consolidating their loans, student loan borrowers may be able to extend their repayment terms to up to 30 years. This, in conjunction with the fact that some student loan borrowers enter into repayment in their thirties and even forties, explains why the age distribution in Panel B has a large right tail.

3. Blinder-Oaxaca decomposition

FIGURE 2 ABOUT HERE

Student loan default rates rose by 18.9% during the Great Recession.⁵ Fig. 2 contrasts two potential explanations for this increase. Panel A shows that, along with the rise in student loan defaults, the share of nontraditional borrowers attending for-profit institutions and community colleges rose steadily, from 39.0% in 2006 to 45.6% in 2009. These borrowers have much higher default rates than traditional borrowers attending four-year colleges. In 2006, for example, nontraditional borrowers were 2.8 times more likely to default than traditional borrowers (source: NSLDS). Thus, shifts in the composition of

⁵The 18.9% increase represents the change in new student loan defaults from 2007 to 2010. As discussed in Section 2.2, we focus on new defaults from 2007 to 2010, as it takes about one year between when a payment is missed and when a default is recorded in our administrative data.

student loan borrowers are a potential candidate explanation for the rise in student loan defaults during the Great Recession.⁶

Panel B shows a visually striking inverse relation between the rise in student loan defaults and the collapse in home prices during the Great Recession. Home prices may have affected student loan defaults through various channels, notably, through local labor markets (Mian and Sufi, 2014; Giroud and Mueller, 2017) or, more directly, by impairing student loan borrowers’ ability to borrow against home equity (Mian and Sufi, 2011; Bhutta and Keys, 2016). Accordingly, falling home prices may be able to explain some of the increase in student loan defaults during the Great Recession.

While we are careful not to interpret the time-series evidence as causal, we note that the graphs in Panel A depicting the rise in student loan defaults and the share of non-traditional borrowers do not line up particularly well. Indeed, student loan default rates began to increase only in 2007, whereas the share of nontraditional borrowers increased steadily throughout the 2000s, from 30.5% in 2000 to 39.0% in 2006. (It increased in every single year during this time period.) In stark contrast, the collapse in home prices in Panel B lines up almost perfectly with the rise in student loan defaults, especially if one accounts for the one-year time lag between when a payment is missed and when a loan default is recorded in administrative data.

FIGURE 3 ABOUT HERE

Fig. 3 examines the cross-sectional relation between changes in student loan defaults in the Great Recession, $\Delta \text{Log } default_{07-10}$, and either changes in the share of nontraditional borrowers, $\Delta \text{Log } NT \text{ share}_{06-09}$, or changes in home prices, $\Delta \text{Log } home \text{ price}_{06-09}$, at the zip code level. In Panel A, zip codes are sorted into percentile bins based on their value of $\Delta \text{Log } NT \text{ share}_{06-09}$. There are 12,749 zip codes in our sample. Accordingly, a percentile bin contains approximately 127 zip codes. For each percentile bin, the scatterplot shows

⁶“Repayment outcomes tend to be worse among borrowers who attend for-profit or community colleges [. . .]. Many of these types of borrowers accounted for a disproportionate share of the increase in student borrowing during the Great Recession” (Council of Economic Advisers, 2016, 4-5).

the mean of $\Delta \text{Log default}_{07-10}$ and $\Delta \text{Log NT share}_{06-09}$, respectively, where means are computed by weighting zip codes by total student loan balances. The scatterplot in Panel B, which depicts the relation between $\Delta \text{Log default}_{07-10}$ and $\Delta \text{Log home price}_{06-09}$, is constructed analogously, except that zip codes are sorted into percentile bins based on their value of $\Delta \text{Log home price}_{06-09}$. As is shown, there is a positive cross-sectional relation between changes in student loan defaults and changes in the share of nontraditional borrowers and a negative cross-sectional relation between changes in student loan defaults and changes in home prices. We confirm both of these relations in our regression analysis below.

TABLE 2 ABOUT HERE

Table 2 provides a Blinder-Oaxaca style decomposition to explain the rise in student loan defaults during the Great Recession. As column 1 shows, the coefficient from a zip code level regression of $\Delta \text{Log default}_{07-10}$ on $\Delta \text{Log NT share}_{06-09}$ is 0.3282. Since both variables are measured in logs, this coefficient represents an elasticity. Hence, a 1% increase in the share of nontraditional borrowers is associated with a 0.3282% increase in student loan default rates. Since the share of nontraditional borrowers increased by 16.9% between 2006 and 2009, this implies that shifts in the composition of student loan borrowers can explain approximately 29% ($16.9 \times 0.3282 / 18.9 = 0.293$) of the rise in student loan defaults in the Great Recession. As column 2 shows, the coefficient from a zip code level regression of $\Delta \text{Log default}_{07-10}$ on $\Delta \text{Log home price}_{06-09}$ is -0.4179 . Given that home prices fell by 14.4% between 2006 and 2009, this implies that the collapse in home prices during the Great Recession can explain approximately 32% ($14.4 \times -0.4179 / 18.9 = 0.318$) of the contemporaneous rise in student loan defaults.⁷ Accordingly, compositional shifts and the collapse in home prices together can explain roughly 60% of the increase in student loan defaults during the Great Recession.⁸

⁷The drop in home prices of 14.4% is very similar to the 14.5% drop reported by Giroud and Mueller (2017), also based on Zillow data, and the 14.9% drop reported by the St. Louis Fed based on FHFA data.

⁸As we do not force the shares to add up to one, this implies that about 40% of the rise in student

As we argue below, the collapse in home prices during the Great Recession affected student loan defaults primarily through a labor market channel. Hence, it might seem natural to use a more direct measure of labor market outcomes, such as employment changes. However, as column 3 shows, changes in employment from 2006 to 2009 at the zip code level have virtually zero explanatory power. While this may seem surprising, it is consistent with similar results in the mortgage default literature. As Gyourko and Tracy (2014, 87) point out, aggregate unemployment measures have been unsuccessful at explaining mortgage default due to “extreme” attenuation bias, resulting from the fact that these measures proxy for unobserved individual unemployment status. Based on simulations, the authors conclude that “the use of an aggregate unemployment rate in lieu of actual borrower unemployment status results in default risk from a borrower becoming unemployed being underestimated by a factor more than 100.”⁹

More generally, we focus on home prices as their massive collapse is a salient, and widely studied, feature of the Great Recession. Indeed, prior literature has argued that it underlies much of the rise in unemployment during the Great Recession—by causing large drops in local consumer spending by households (e.g., Mian, Rao and Sufi, 2013; Stroebl and Vavra, 2017; Kaplan, Mitman, and Violante, 2016; Mian and Sufi, 2014; Giroud and Mueller, 2017). In principle, falling home prices may have affected student loan defaults through a variety of channels, employment losses being just one of them. Other channels include drops in labor earnings while being employed (intensive margin) or lower propensities of finding a job while being unemployed, next to various nonlabor market channels. Much of this paper is devoted to studying some of these channels in more depth. As such, the elasticity reported in column 2 may be viewed as the total effect

loan defaults remains unexplained by either compositional shifts or changes in home prices. See Looney and Yannelis (2015, 51–60).

⁹Gerardi et al. (2018) overcome this measurement problem using household-level data from the Panel Study of Income Dynamics (PSID), which includes information on individuals’ employment status. We have verified the Gyourko-Tracy statement in our data using individual employment losses—measured as labor earnings drops of 50% or more relative to the previous year’s earnings—in lieu of regional employment changes. Consistent with their statement, we find that individual employment losses do better than regional employment changes and can explain about 20% of the rise in student loan defaults. Still, this is less than the 32% explained by home prices, consistent with home prices affecting student loan defaults through multiple labor (and nonlabor) market channels, employment losses being only one of them.

from all of these channels. With that being said, our study is first and foremost trying to understand the sharp rise in student loan defaults during the Great Recession. Thus, the home price mechanism emphasized in this paper may not readily apply to other time periods outside of the Great Recession.

There are two main takeaways from this section. One is that the collapse in home prices and the shift toward riskier nontraditional borrowers can each explain a sizable fraction of the rise in student loan defaults in the Great Recession, albeit the timing appears to line up much better with the fall in home prices. In the remainder of our paper, we focus on the role of home prices—Looney and Yannelis (2015) provide an in-depth discussion of the changing nature of borrower composition over the last 40 years. Second, given the strong empirical relation between student loan defaults and shifts toward non-traditional borrowers during the Great Recession (cf., Panel A of Fig. 3), we must be careful to isolate the effects of changes in home prices from those of changes in borrower composition. In our empirical analysis, we provide separate analyses by institution type, and we account for compositional differences by including cohort year, zip code \times cohort year, and individual borrower fixed effects.

4. Home prices and student loan defaults

TABLE 3 ABOUT HERE

Table 3 examines the role of home prices for the rise in student loan defaults during the Great Recession using aggregated data at the regional level. Panel A examines long differences in the spirit of Mian, Rao, and Sufi (2013), Mian and Sufi (2014), and Giroud and Mueller (2017). Thus, there is one observation per region. Regions are either zip codes, counties, or CZs. As is shown, the elasticity of student loan defaults with respect to home prices at the regional level is large and highly significant. This elasticity becomes larger as we broaden the level of regional aggregation—from zip codes to counties to CZs—suggesting that zip code (county) level home prices affect student loan defaults also in other zip codes (counties) within a given county (CZ). Panel B is similar to Panel A, except

that we exploit the panel dimension of our data. Accordingly, the dependent variable is student loan default in year $t + 1$, and the independent variable is home price (in logs) in year t . Since we measure home prices from 2006 to 2009 and student loan defaults from 2007 to 2010, this implies that there are (roughly) four observations per region. While this panel specification utilizes more observations than the long difference specification, serial correlation of the error term is a concern. To address this concern, we cluster standard errors at the region (zip code, county, CZ) level, allowing for arbitrary correlation of the residuals within a region over time. Further, to account for time invariant heterogeneity across regions, we include region fixed effects. Such region fixed effects were differenced out in our long difference specification. As can be seen, all results are similar to those in Panel A. Notably, the coefficient on home prices is again increasing in the level of regional aggregation.

TABLE 4 ABOUT HERE

In Table 4, as in all remaining tables of this paper, we employ a panel specification using highly disaggregated student loan data at the individual borrower level. Home prices are in logs and are lagged and measured at the zip code level. There are 1,071,049 annual borrower-level observations associated with 298,003 individual student loan borrowers. All regressions include year and zip code fixed effects. The year fixed effects capture any economy wide factors, such as aggregate economic conditions. The zip code fixed effects absorb any time invariant heterogeneity across zip codes, including any given differences in borrower composition, college enrollment, or student loan volume, or any differences arising from different experiences during the preceding housing boom. As column 1 of Table 4 shows, a 1% decline in home prices is associated with a 0.0113 percentage point increase in new student loan defaults. (Home prices at the zip code level declined by 14.4% between 2006 and 2009.) By comparison, the elasticity of -0.4179 in our long difference specification in Table 3 implies a 0.0150 percentage point increase in new student loan defaults, which is roughly of similar magnitude.¹⁰

¹⁰The rate of new student loan defaults is 3.6% in 2006. Hence, an elasticity of -0.4179 implies that

Columns 2 to 6 of Table 4 address various statistical and identification issues. In column 2, we cluster standard errors at the county level. As is shown, they become only slightly larger. In column 3, we include the full vector of controls from Eq. (1), which includes loan balance, borrowing duration, family income, school type, and Pell grant aid. As can be seen, our estimates remain virtually unchanged. Since it makes no difference whether these controls are included, we do not include them in our further analysis. While the inclusion of zip code fixed effects accounts for any fixed differences in borrower composition across zip codes, it is conceivable that the composition of student loan borrowers within a zip code has shifted over time. If such compositional shifts are correlated with home price changes, this could potentially confound our estimates. For example, student loan borrowers are more likely to default within the first few years after entering into repayment. If older repayment cohorts out-migrate in response to falling home prices, this could induce a negative correlation between changes in home prices and changes in default rates. In columns 4 and 5, we rule out such confounding factors by including either cohort year or zip code \times cohort year fixed effects, where a cohort is defined by the year in which a student loan borrower enters into repayment. As can be seen, our estimates remain very similar. In column 6, we include individual borrower fixed effects, thereby absorbing any time invariant heterogeneity across student loan borrowers, such as schools attended, major choice, family background, and credit history. Our estimates remain again similar.

TABLE 5 ABOUT HERE

Table 5 breaks down our main results by institution type. We have previously shown that shifts toward nontraditional borrowers are correlated with changes in student loan

a 1% decline in home prices translates into a $0.004179 \times 3.6 = 0.0150$ percentage point increase in new student loan defaults. New student loan default rates are lower than n -year cohort default rates used by the US Department of Education, for two reasons. First, by construction, n -year cohort default rates are larger than annual default rates. (For instance, the two-year (three-year) cohort default rate is 5.2% (9.1%) in 2006.) Second, n -year cohort default rates measure student loan defaults in the first n years immediately after student loan borrowers enter into repayment, which is a period during which a disproportionately large fraction of student loan borrowers defaults on their loans.

defaults. Accordingly, we now estimate our main specification within a given institution type. As columns 1 to 4 show, our main results hold across all institution types, albeit they are strongest (in a statistical sense) at for-profit institutions. These institutions—together with community colleges—also exhibit the largest point estimates. However, given that for-profit institutions and community colleges also tend to have much higher default rates (Deming, Goldin, and Katz, 2012; Looney and Yannelis, 2015), the implied elasticities are ultimately quite similar to those at (not-for-profit) public and private four-year colleges. Lastly, in columns 5 and 6, we break down our main results by institutional selectivity, as measured by Barrons, which is based on the fraction of applicants that institutions admit. As can be seen, our main results hold both across selective and nonselective institutions.

5. Home price channels

Home prices can affect student loan default through various channels. Our data allow us to examine two of these channels in more depth: changes in labor market outcomes and changes in homeowners' liquidity from borrowing against home equity.

5.1. Labor market channel

One of the main narratives of the Great Recession is that the collapse in home prices caused a drop in consumer spending, which adversely affected labor market outcomes. Employment losses and earnings declines, in turn, may have impaired student loan borrowers' ability to make repayments, especially if their labor income is low to begin with. To explore this labor market channel, we study the relation between home prices, labor earnings, and student loan defaults at the individual borrower level.

TABLE 6 ABOUT HERE

Table 6 breaks down our main results by borrowers' individual labor income. As can be seen, labor income matters. For borrowers with annual earnings of \$60,000 or more, there is no significant relation between home prices and student loan defaults. Also, the point estimates are declining across income groups—they are largest for low income borrowers

and smallest for high income borrowers. That being said, the results in Table 6 also raise questions. Are high income borrowers less likely to default in response to falling home prices because high income jobs are less affected? Or does the decline in home prices affect low and high income jobs alike, but high income borrowers have been able to build up savings in the past, allowing them to continue making repayments on their student loans? To address these questions, we now examine the relation between home prices and individual labor earnings.

FIGURE 4 ABOUT HERE

TABLE 7 ABOUT HERE

Fig. 4 shows the cross-sectional relation between changes in individual labor earnings during the Great Recession, $\Delta \text{Log earnings}_{06-09}$, and changes in home prices, $\Delta \text{Log home price}_{06-09}$, at the zip code level. The scatterplot is constructed the same way as the scatterplots in Fig. 3. As can be seen, larger declines in home prices are associated with larger declines in individual labor earnings. Table 7 examines this relation formally in a regression framework. The dependent variable is labor earnings at the individual borrower level, and the independent variable is home prices at the zip code level. All regressions include year and zip code fixed effects. Standard errors are clustered at the zip code level. To ensure that we capture variation in labor earnings that is likely to affect student loan repayment behavior, we focus on large drops in earnings of 50% or more relative to the previous year's earnings. Given the magnitude of these earnings drops, it is likely that they are associated with employment losses. As is shown, the relation between home prices and individual labor earnings is declining across income groups and only significant among low income borrowers. Hence, while high income borrowers may have been able to build up savings, their jobs also appear to be overall less affected by the fall in home prices.¹¹

¹¹The results in Table 7 complement results in prior literature showing that changes in home prices during the Great Recession affect aggregate employment (Mian and Sufi, 2014; Giroud and Mueller, 2017). Our results are obtained at the individual level, are based on labor earnings and, importantly, show that the “home price-labor market channel” operates primarily through low income jobs.

Altogether, the results in this section are consistent with a labor market channel, whereby declining home prices affect student loan repayment behavior through a decline in individual labor earnings. This labor market channel operates primarily through low income jobs. In contrast, high income borrowers do not face significant drops in their labor earnings in response to declining home prices and, as a result, are not more likely to default on their student loans.

5.2. Direct liquidity channel

Falling home prices during the Great Recession may have directly impacted student loan defaults through a liquidity channel. Precisely, they may have impaired student loan borrowers' ability to borrow against home equity (Mian and Sufi, 2011; Bhutta and Keys, 2016), thereby limiting their access to liquidity and ability to make repayments.

It is a priori unclear whether falling home prices should directly affect student loan defaults. Unlike mortgages, there are no strategic default incentives for student loans. Further, student loans are not dischargeable in bankruptcy, and wages—even social security benefits—can be garnished for the rest of a borrower's lifetime. Indeed, Mian and Sufi (2011) find that home equity-based borrowing during the housing boom is not used to pay down expensive credit card debt—even for households that heavily depend on credit card borrowing—suggesting that a direct impact on other types of household debt (i.e., other than mortgages) is anything but obvious.

We test for a home equity-based liquidity channel by comparing the default behavior of homeowners and renters in response to falling home prices. Intuitively, while falling home prices may have impaired homeowners' ability to borrow against home equity, there should be no corresponding liquidity effect for renters.¹² We identify homeowners through Form 1098 submitted by mortgage lenders, which includes mortgage interest payments, as well as the borrower's taxpayer identification number. Thus, we are able to identify homeowners as long as they have a mortgage, regardless of whether they file for the

¹²Renters' liquidity may have improved if falling home prices are passed through in the form of lower rents. However, this would only strengthen the argument that homeowners should be relatively more impaired than renters. See Rosen (1979) and Poterba (1984) for classic references.

mortgage interest deduction.¹³ The homeownership rate in our sample is 39%, which is considerably less than the national average of 68% during the sample period. Indeed, while our sample includes all student loan borrowers in repayment—many of which are in their thirties, forties, and even fifties (see Fig. 1)—they are still younger than the national average and hence earlier in their life cycle.

TABLE 8 ABOUT HERE

Table 8 examines whether homeowners respond more strongly to falling home prices than renters, as predicted by the liquidity channel. Maybe somewhat surprisingly, the interaction term *Home price* \times *Owner* is always small and insignificant. This is true regardless of whether we add controls, how we cluster standard errors, or which types of fixed effects we include. On the other hand, with the exception of column 5, the direct effect of homeownership is significant and has the predicted sign: absent any home price changes, homeowners are less likely to default than renters.¹⁴ Hence, while homeowners and renters differ in their baseline default likelihood, they respond similarly to changes in home prices, which is inconsistent with a direct liquidity channel.

TABLE 9 ABOUT HERE

A potential concern is that homeowners have access to more financial resources, and this could mask liquidity effects. Indeed, as is shown in Table 1, homeowners have higher labor earnings, higher total income, and higher family income than renters. We address this concern in Table 9 by including these variables and their respective interactions with home prices as controls in our regressions. While the controls have the predicted sign—individuals with higher labor earnings, total income, or family income default less and are

¹³Berger, Turner, and Zwick (2017) make a similar point.

¹⁴That the direct effect of homeownership is insignificant in column 5—which includes zip code \times cohort year fixed effects—suggests that it may be driven by cohort and regional effects, e.g., homeowners may be older and live in different neighborhoods than renters.

less sensitive to home price changes—the interaction term *Home price* \times *Owner* remains small and insignificant.

TABLE 10 ABOUT HERE

Lastly, we address concerns that the effects of homeownership may be attenuated by measurement error. Indeed, we only capture homeowners who have a mortgage. Those who own their home outright or have paid off their mortgage in full are misclassified as renters. Outright homeownership is infrequent, however. According to the National Association of Realtors (2006), 98% of first-time buyers and 89% of repeat buyers in 2005–2006 used a mortgage to finance their home. Accordingly, measurement error ought to be minimal especially among young homeowners who are likely to be first-time home buyers. Also, young homeowners are unlikely to have paid off their mortgage in full. In Table 10, we divide our sample into groups based on either age or repayment cohort. Consistent with attenuation bias, we find that the direct effect of homeownership is insignificant among older student loan borrowers (age $>$ 40 years or repayment cohort $<$ 2000). Importantly, however, the interaction term *Home price* \times *Owner* remains insignificant across all ages and repayment cohorts, except in column 1, where it is positive and (marginally) significant.¹⁵

6. Income Based Repayment program

Under the standard ten-year repayment plan, student loan borrowers can apply for a loan deferment (if they are unemployed) or a forbearance (if the amount owed exceeds 20% of their gross income).¹⁶ In the wake of the Great Recession, in 2009, the US Department of Education rolled out the IBR program. The purpose of income driven repayment

¹⁵The positive coefficient on the interaction term implies that, among younger student loan borrowers (age $<$ 30 years), homeowners are less likely to default than renters in response to falling home prices, which is inconsistent with the home equity-based liquidity channel.

¹⁶Other government-sponsored initiatives include grants and subsidized loans. By lowering monthly repayments, these initiatives reduce default risk. Surprisingly, students who are offered subsidized loans often turn them down, leaving money on the table (Cadena and Keys, 2013).

plans, such as IBR, is to provide student loan borrowers with additional insurance against negative shocks by making their loan repayments contingent on discretionary income. Under the IBR plan, repayments are capped at 15% of discretionary income, and repayment terms are extended to up to 25 years, after which all remaining student debt is forgiven.¹⁷ Eligibility is based on a means test, which requires that 15% of the borrower’s discretionary income be less than her payment under the standard ten-year repayment plan. Discretionary income is any income above 150% of the federal poverty level. Essentially, student loan borrowers are eligible for the IBR repayment option if their student debt is sufficiently high relative to their discretionary income.

To assess the efficacy of the IBR program, we classify student loan borrowers as IBR eligible and ineligible based on the means test. That is, we do not assign treatment status based on whether an individual actually enrolled in the IBR program, which is an endogenous choice, but based on whether she was eligible for enrollment. We later provide graphical evidence showing that changes in student loan defaults attributed to the IBR program come from (eligible) student loan borrowers who actually took up the IBR repayment option. We calculate IBR eligibility as $0.15 \times (E_{it} - \overline{E_{it}}) < P_{it}$, where E_{it} is individual i ’s labor earnings in year t , $\overline{E_{it}}$ is the federal poverty level—which varies from year to year and depends on household size—and P_{it} is the annual payment faced by individual i in year t under the standard ten-year repayment plan. Household size, including marital status and number of dependent children, is obtained from IRS records. Annual payments under the standard ten-year repayment plan, P_{it} , are computed using the amortization formula $P_{it} = L_{i0} \times (r_{it} + \frac{r_{it}}{(1+r_{it})^n - 1})$, where L_{i0} is the initial loan balance, r_{it} is the borrowing rate, and $n = 10$ is the number of years. IBR eligibility—which is based on the means test—is well defined for any given year, including years prior to the introduction of the IBR plan. Accordingly, we can compare student loan defaults by IBR eligible and ineligible borrowers before and after the plan’s introduction.

Given that the IBR program was introduced in 2009, we extend our sample period to include student loan defaults up to 2012.¹⁸ Thus, we consider home prices between

¹⁷The 15% cap was later reduced to 10% for new borrowers on or after July 1, 2014.

¹⁸We choose 2012 as the ending date year because a new insurance program, the Pay As You Earn

2006 and 2011 and student loan defaults between 2007 and 2012. Extending the sample period increases the number of annual observations to 1,556,296. To gauge the insurance value of the IBR plan, we conduct a triple difference analysis by comparing the default behavior of IBR eligible and ineligible borrowers in response to home price changes before and after the plan’s introduction. We estimate the following specification:

$$\begin{aligned} \pi_{i,t+1} = & \alpha_t + \alpha_z + \beta_1 \text{Home price}_{z,t} + \beta_2 \text{IBR eligible}_{i,t} + \beta_3 \text{Home price}_{z,t} \times \text{Post} \\ & + \beta_4 \text{IBR eligible}_{i,t} \times \text{Post} + \beta_5 \text{Home price}_{z,t} \times \text{IBR eligible}_{i,t} \\ & + \beta_6 \text{Home price}_{z,t} \times \text{IBR eligible}_{i,t} \times \text{Post} + \varepsilon_{i,t}, \end{aligned} \quad (2)$$

where $\pi_{i,t+1}$ is an indicator of whether individual i defaults in year $t + 1$; $\text{Home price}_{z,t}$ is the home price (in logs) in zip code z in year t ; $\text{IBR eligible}_{i,t}$ is a dummy indicating whether individual i passes the means test $0.15 \times (E_{it} - \overline{E_{it}}) < P_{it}$ in year t ; Post is a dummy that equals one beginning in 2009; and α_t and α_z are year and zip code fixed effects, respectively. Standard errors are clustered at the zip code level. Observations are weighted by individual loan balances.

The main coefficients of interest are β_2 , β_4 , β_5 , and β_6 . We would expect β_2 to be positive: IBR eligible borrowers—those with high ratios of student debt to income—should have higher default rates. The coefficient β_4 indicates the relative change in default rates of IBR eligible versus ineligible borrowers after the plan’s introduction. If the IBR plan is successful, we would expect β_4 to be negative. The coefficient β_5 indicates whether IBR eligible borrowers are more sensitive to home price changes. We would expect this coefficient to be negative: a decline in home prices increases student loan defaults, and this effect should be stronger for borrowers with high ratios of student debt to income. Lastly, the coefficient β_6 indicates whether the stronger default sensitivity of IBR eligible borrowers to home price changes is mitigated after 2009. If the IBR plan provides student loan borrowers with valuable insurance, then β_6 should be positive.

TABLE 11 ABOUT HERE

(PAYE) program, was introduced in December 2012.

Table 11 presents the results. Column 1 shows that home prices are negatively associated with loan defaults also during the extended sample period. In addition, IBR eligible borrowers—those with high ratios of student debt to income—are more likely to default on their loans. In column 2, we estimate the triple difference specification from Eq. (3). As is shown, the coefficient on *IBR eligible*, β_2 , is positive, and the coefficient on *IBR eligible* \times *Post*, β_4 , is negative. Together, these results imply that student loan borrowers with high ratios of student debt to income are more likely to default on their loans, and this effect is mitigated after the introduction of the IBR plan. Further, the coefficient on *Home price* \times *IBR eligible*, β_5 , is negative, while the coefficient on *Home price* \times *IBR eligible* \times *Post*, β_6 , is positive. Accordingly, while student loan borrowers with high ratios of student debt to income are more sensitive to home price changes, this effect is attenuated after the introduction of the IBR program. Altogether, the results in Table 11 show that the introduction of the IBR plan reduced student loan defaults in general, as well as their sensitivity to home price fluctuations.

The coefficients β_4 and β_6 indicate how student loan defaults and their sensitivity to home price fluctuations change after the introduction of the IBR plan. Both coefficients have the predicted sign but are only marginally significant. A potential concern is that student loan borrowers with high ratios of student debt to income may be unobservably different from borrowers with low ratios. We address this concern in column 3 by dropping borrowers with very low ratios from our sample. Effectively, we thus compare borrowers with a high degree of illiquidity, where some are eligible for the IBR plan and others are not.¹⁹ As can be seen, our results remain similar. Notably, the two main coefficients of interest, β_4 and β_6 , are now significant at the 5% level. In column 4, we account for unobserved heterogeneity among IBR eligible and ineligible borrowers by including borrower fixed effects. While some of the coefficients are insignificant due to lack of within borrower variation, the two main coefficients of interest, β_4 and β_6 , are significant at the 1% and 5% level, respectively.

¹⁹Under the means test, student loan borrowers are eligible for the IBR plan if $0.15 \times (E_{it} - \overline{E_{it}}) < P_{it}$. The sample restriction in column 3 requires that $0.75 \times (E_{it} - \overline{E_{it}}) < P_{it}$, thus eliminating all student loan borrowers with low ratios of student debt to discretionary income.

FIGURE 5 ABOUT HERE

The main assumption underlying our difference-in-differences analysis is that IBR eligible and ineligible borrowers exhibit parallel trends prior to the plan’s introduction. Fig. 5 provides evidence in support of this assumption. The white bars show student loan default rates of IBR ineligible borrowers. The gray bars show student loan default rates of IBR eligible borrowers. Eligibility is based on the means test, $0.15 \times (E_{it} - \overline{E_{it}}) < P_{it}$, which implies that it is well defined for any given year, including years prior to the introduction of the IBR plan. Given that there is a one-year time lag between when a payment is missed and when a default is recorded in the NSLDS, default rates in year $t + 1$ reflect eligibility (or take-up) status in year t . Beginning in 2009—showing up as 2010 due to the one-year time lag—we distinguish between IBR eligible borrowers who took up the IBR repayment option (black) and IBR eligible borrowers who did not take up the option (gray). Thus, prior to the introduction of the IBR plan, the gray bars pertain to IBR eligible borrowers in general, while after the plan’s introduction, they pertain to IBR eligible borrowers who did not take up the IBR repayment option.

Fig. 5 provides three main insights. First, and most important, IBR eligible and ineligible borrowers were on similar trends prior to 2009. Second, IBR eligible borrowers who did not take up the IBR repayment option (gray) continued on this trend after 2009. Thus, our results cannot be explained by IBR eligible borrowers suddenly experiencing a positive shock in 2009, which just happened to coincide with the introduction of the IBR program. Third, default rates of IBR eligible borrowers who took up the IBR repayment option (black) are very low, suggesting that the IBR program was successful at reducing student loan defaults.

FIGURE 6 ABOUT HERE

To provide further evidence in support of the parallel trend assumption, we estimate a variant of the specification in column 2 in which *Home price* \times *IBR eligible* \times *Post* is replaced by *Home price* \times *IBR eligible* \times t , where $t = 2007, \dots, 2011$. The yearly

coefficients $\beta_{6,t}$ indicate the extent to which the (higher) default sensitivity of IBR eligible borrowers to changes in home prices is mitigated in year t relative to the baseline year of 2006. The coefficients are plotted in Panel A of Fig. 6.²⁰ As is shown, IBR eligible and ineligible borrowers are on parallel trends prior to the introduction of the IBR program: the coefficients associated with 2007 and 2008 are statistically indistinguishable from the 2006 baseline coefficient. Second, the coefficient jumps in 2009, when the IBR plan is introduced. Third, the coefficient continues to increase gradually after 2009. To shed light on this gradual increase, Panel B shows take-up rates under the IBR plan. As can be seen, take-up is slow in the beginning but then gradually increases over time, consistent with the gradual increase of the coefficient in Panel A.

7. Conclusion

Student loan default rates increased sharply in the Great Recession. A Blinder-Oaxaca decomposition shows that shifts in the composition of student loan borrowers and the massive collapse in home prices during the Great Recession can each account for about 30% of the rise in student loan defaults. When exploring potential channels, we find that falling home prices affect student loan defaults primarily through a labor market channel by impairing student loan borrowers' labor earnings, especially for low income jobs. By contrast, we find no evidence that falling home prices affect student loan default behavior through a home equity-based liquidity channel.

In the wake of the Great Recession, in 2009, the federal government introduced the IBR program to reduce student loan defaults and insure student loan borrowers against negative shocks by making their loan repayments contingent on discretionary income. To assess the efficacy of the IBR program, we compare the default responses of IBR eligible versus ineligible student loan borrowers to home price changes before and after the program's introduction. We find that the introduction of the IBR program reduced both student loan defaults and their sensitivity to home price fluctuations, and that this result is entirely driven by IBR eligible borrowers who enrolled in the IBR program.

²⁰Given that there is a one-year time lag between when a payment is missed and when a default is recorded in the NSLDS, the coefficient associated with year t is plotted in year $t + 1$.

References


- Agarwal, S., Amromin, G., Ben-David, I., Chomsisengphet, S., Piskorski, T., Seru, A., 2017. Policy intervention in debt renegotiation: evidence from the Home Affordable Modification Program. *Journal of Political Economy* 125, 654–712.
- Avery, C., Turner, S., 2012. Student loans: do college students borrow too much—or not enough? *Journal of Economic Perspectives* 26, 165–192.
- Bailey, M., Cao, R., Kuchler, S., Stroebel, J., 2017. The economic effects of social networks: evidence from the housing market. *Journal of Political Economy*, forthcoming.
- Berger, D., Turner, N., Zwick, E., 2017. Stimulating housing markets. Unpublished working paper. University of Chicago.
- Bhutta, N., Keys, B., 2016. Interest rates and equity extraction during the housing boom. *American Economic Review* 106, 1742–1774.
- Bos, M., Breza, E., Liberman, A., 2018. The labor market effects of credit market information. *Review of Financial Studies* 31, 2005–2037.
- Cadena, B., Keys, B., 2013. Can self-control explain avoiding free money? Evidence from interest-free student loans. *Review of Economics and Statistics* 95, 1117–1129.
- Council of Economic Advisers, 2016. Investing in Higher Education: Benefits, Challenges, and the State of Student Debt. Penny Hill Press, Damascus, MD.
- Davis, S., Faberman R., Haltiwanger J., 2012. Labor market flows in the cross section and over time. *Journal of Monetary Economics* 59, 1–18.
- Deming, D., Goldin, C., Katz, L., 2012. The for-profit post secondary school sector. *Journal of Economic Perspectives* 26, 139–164.
- Di Maggio, M., Kermani, A., Keys, B., Piskorski, T., Ramcharan, R., Seru, A., Yao, V., 2017. Interest rate pass-through: mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review* 107, 3550–3588.
- Dobbie, W., Song, J., 2015. Debt relief and debtor outcomes: Measuring the effects of consumer bankruptcy protection. *American Economic Review* 105, 1272–1311.

- Dobbie, W., Goldsmith-Pinkham, P., Mahoney, N., Song, J., 2017. Bad credit, no problem? Credit and labor market consequences of bad credit reports. Unpublished working paper. National Bureau of Economic Research.
- Earnest Operations LLC, 2016. Student loans are changing the job hunt for recent grads.
- Elul, R., Souleles, N., Chomsisengphet, S., Glennon, D., Hunt, R., 2010. What “triggers” mortgage default? *American Economic Review* 100, 490–494.
- Gerardi, K., Herkenhoff, K., Ohanian, L., Willen, P., 2018. Can’t pay or won’t pay? Unemployment, negative equity, and strategic default. *Review of Financial Studies* 31, 1098–1131.
- Giroud, X., Mueller, H., 2017. Firm leverage, consumer demand, and employment losses during the great recession. *Quarterly Journal of Economics* 132, 271–316.
- Gyourko, J., Tracy, J., 2014. Reconciling theory and empirics on the role of unemployment in mortgage default. *Journal of Urban Economics* 80, 87–96.
- Herkenhoff, K., Phillips, G., Cohen-Cole, E., 2016. The impact of consumer credit access on employment, earnings and entrepreneurship. Unpublished working paper. National Bureau of Economic Research.
- Ji, Y., 2016. Job search under debt: aggregate implications of student loans. Unpublished working paper. Massachusetts Institute of Technology.
- Kaplan, G., Mitman, K., Violante, G., 2016. Non-durable consumption and housing net worth in the great recession: evidence from easily accessible data. Unpublished working paper. National Bureau of Economic Research.
- Keys, B., Mukherjee, T., Seru, A., Vig, V., 2010. Did securitization lead to lax screening? evidence from subprime loans. *Quarterly Journal of Economics* 125, 307–362.
- Keys, B., Seru, A., Vig, V., 2012. Lender screening and the role of securitization: Evidence from prime and subprime mortgage markets. *Review of Financial Studies* 25, 2071–2108.
- Looney, A., Yannelis, C., 2015. A crisis in student loans? How changes in the characteristics of borrowers and in the institutions they attended contributed to rising loan defaults. *Brookings Papers on Economic Activity*, Fall, 1–68.

- Mian, A., Sufi, A., 2011. House prices, home equity-based borrowing, and the US household leverage crisis. *American Economic Review* 101, 2132–2156.
- Mian, A., Sufi, A., 2014. What explains the 2007–2009 drop in employment? *Econometrica* 82, 2197–2223.
- Mian, A., Rao, K., Sufi, A., 2013. Household balance sheets, consumption, and the economic slump. *Quarterly Journal of Economics* 128, 1–40.
- National Association of Realtors, 2006. Profile of home buyers and sellers.
- Poterba, J., 1984. Tax subsidies to owner-occupied housing: an asset market approach. *Quarterly Journal of Economics* 99, 729–752.
- Purnanandam, A., 2011. Originate-to-distribute model and the subprime mortgage crisis. *Review of Financial Studies* 24, 1881–1915.
- Rosen, H., 1979. Housing decisions and the U.S. income tax: an econometric analysis. *Journal of Public Economics* 11, 1–23.
- Society for Human Resource Management (SHRM), 2010. Background checking: the implications of credit background checks on the decision to hire or not to hire.
- Stroebel, J., Vavra, J., 2017. House prices, local demand, and retail prices. *Journal of Political Economy*, forthcoming.

Fig. 1. Age distribution of student loan borrowers. Panel A shows the age of student loan borrowers in our sample at the time when they enter into repayment. Student loan borrowers typically enter into repayment within six months after leaving their degree granting institution. Panel B shows the age of student loan borrowers in our sample based on all borrower-year observations.

Panel A: Age of student loan borrowers when they enter into repayment



Panel B: Age of student loan borrowers in repayment


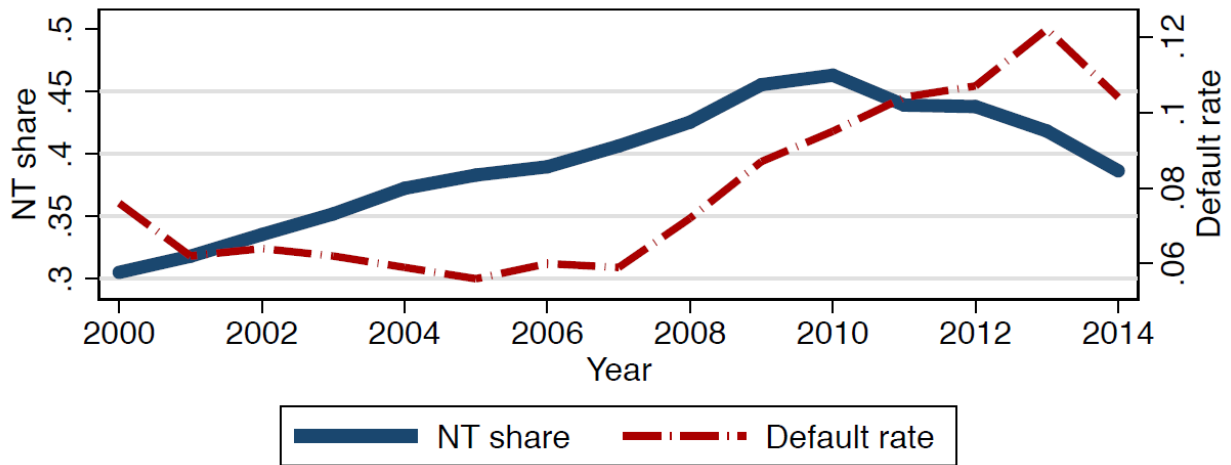


Fig. 2. Time-series evidence. Panel A shows the relation between student loan defaults and the share of nontraditional borrowers attending for-profit institutions and community colleges (“NT share”). Default rate is the two-year cohort default rate, defined by the last year in which the cohort has been in repayment for two years. A student loan goes into default if it is more than 270 days past due. When a loan goes into default, the loan servicer has up to 90 days to report the default to the NSLDS. Thus, there is approximately a one-year time lag between when a payment is missed and when a default is recorded in the NSLDS. Panel B shows the relationship between student loan defaults and home prices. Home price index is the Zillow Home Value Index, which is normalized to one in 1996.

Panel A: Student loan defaults and share of nontraditional borrowers



Panel B: Student loan defaults and home prices


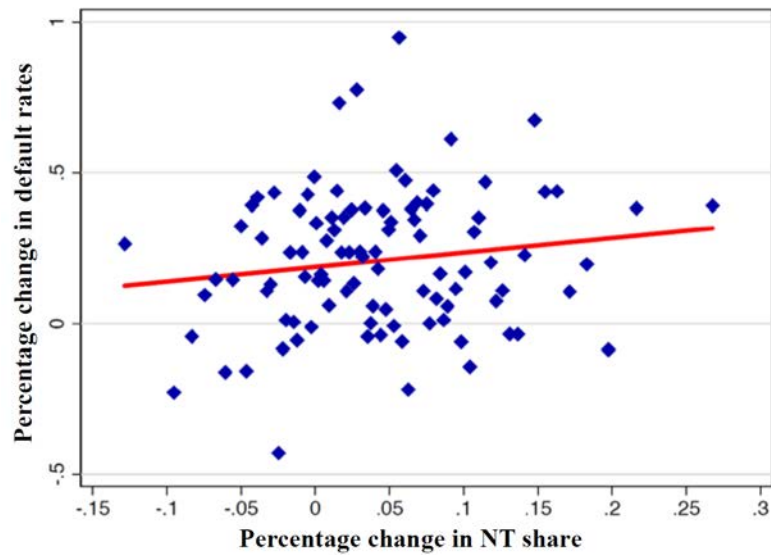


Fig. 3. Cross-sectional evidence. Panel A shows the relation between the percentage change in student loan defaults from 2007 to 2010, $\Delta \text{Log default}_{07-10}$, and the percentage change in the share of nontraditional borrowers attending for-profit institutions and community colleges from 2006 to 2009, $\Delta \text{Log NT share}_{06-09}$, at the zip code level. Zip codes are sorted into percentile bins based on their value of $\Delta \text{Log NT share}_{06-09}$. For each percentile bin, the scatterplot shows the mean of $\Delta \text{Log default}_{07-10}$ and $\Delta \text{Log NT share}_{06-09}$, respectively, where means are computed by weighting zip codes by total student loan balances. Panel B shows the relation between the percentage change in student loan defaults from 2007 to 2010, $\Delta \text{Log default}_{07-10}$, and the percentage change in home prices from 2006 to 2009, $\Delta \text{Log home price}_{06-09}$, at the zip code level. The scatterplot is constructed analogously to that in Panel A, except that zip codes are sorted into percentile bins based on their value of $\Delta \text{Log home price}_{06-09}$.

Panel A: Student loan defaults and share of nontraditional borrowers



Panel B: Student loan defaults and home prices

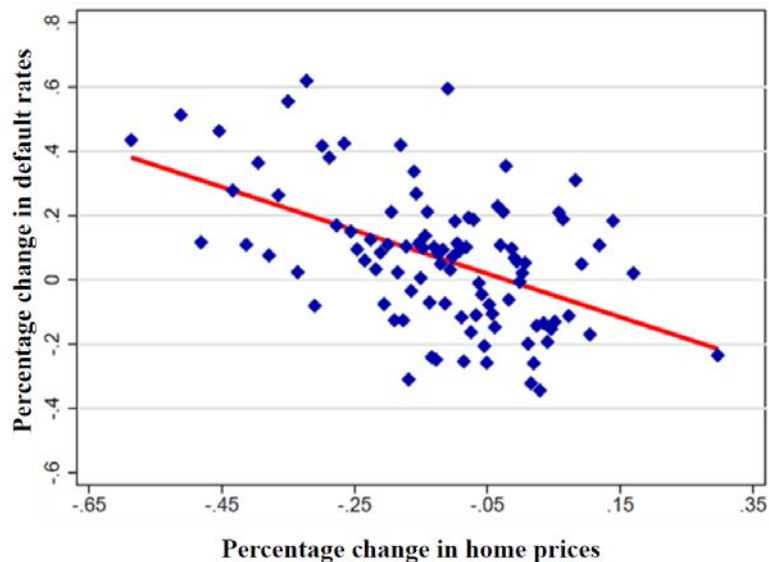


Fig. 4. Home prices and individual labor earnings. The scatterplot shows the relation between the percentage change in student loan borrowers' individual labor earnings from 2006 to 2009, $\Delta \text{Log } earnings_{06-09}$, and the percentage change in home prices from 2006 to 2009, $\Delta \text{Log } home\ price_{06-09}$, at the zip code level. The scatterplot is constructed analogously to that in Panel B of Fig. 3.



Fig. 5. IBR eligibility, take-up, and student loan defaults. This figure shows student loan default rates of IBR eligible and ineligible student loan borrowers. The white bars show student loan default rates of IBR ineligible student loan borrowers. The gray bars show student loan default rates of IBR eligible student loan borrowers (before the introduction of the IBR plan) and IBR eligible student loan borrowers who did not take up the IBR repayment option (after the introduction of the IBR plan), respectively. The black bars show student loan default rates of IBR eligible student loan borrowers who took up the IBR repayment option. Given that there is a one-year time lag between when a payment is missed and when a default is recorded in the NSLDS, student loan default rates in year $t+1$ reflect eligibility (or take-up) status in year t . The IBR plan was introduced in 2009, which implies that its impact on student loan defaults shows up for the first time in 2010. IBR eligibility is based on the means test, which is described in Section 6.





Fig. 6. IBR insurance value and take-up rate. Panel A plots the yearly coefficients $\beta_{6,t}$ from a variant of Eq. (2) in which *Home price* \times *IBR eligible* \times *Post* is replaced with *Home price* \times *IBR eligible* \times *t*, where $t = 2007, \dots, 2011$. The yearly coefficients indicate the extent to which the default sensitivity of IBR eligible student loan borrowers to home price changes is mitigated in a given year relative to the baseline year of 2006. Given that there is a one-year time lag between when a payment is missed and when a default is recorded in the NSLDS, the coefficient associated with a given year t is plotted in the following year, $t+1$. The IBR program was introduced in 2009, which implies that its impact on student loan defaults shows up for the first time in 2010. IBR eligibility is based on the means test, which is described in Section 6. The specification includes year and zip code fixed effects. Standard errors are clustered at the zip code level. The dashed lines represent a 95% confidence interval. Panel B plots the take-up rate of the IBR program as a percentage of all student loan borrowers in repayment.

Panel A: IBR insurance value



Panel B: IBR take-up rate




Table 1

Summary statistics.

The table shows basic summary statistics. Means and standard deviations (SD) are based on annual observations at the individual borrower level. Observations are weighted by individual student loan balances. Labor earnings are Medicare wages plus self-employment earnings. Total income additionally includes nonlabor income. Employment loss is a drop in labor earnings of 50% or more relative to the previous year's earnings. Repayment cohort is the year in which a student loan borrower enters into repayment. Default is an indicator of whether a student loan borrower defaults on her student loan(s) in a given year. A student loan goes into default if it is more than 270 days past due. When a loan goes into default, the loan servicer has up to 90 days to report the default to the NSLDS. Thus, there is approximately a one-year time lag between when a payment is missed and when a default is recorded in the NSLDS. Home prices are measured at the zip code level. All variables are measured over the 2006 to 2009 period, except default, which is measured over the 2007 to 2010 period. Home price data are from Zillow. All other data are from a 4% percent random sample of the NSLDS matched to de-identified IRS tax data.

	All borrowers		Homeowners		Renters	
	Mean	SD	Mean	SD	Mean	SD
Loan balance	23,757	31,520	27,009	34,066	21,862	29,765
Labor earnings	44,930	54,254	62,422	69,688	33,470	36,765
Family income	42,675	54,394	50,972	55,844	37,290	52,388
Total income	62,369	98,345	95,588	126,794	40,755	65,583
Employment loss	0.08	0.27	0.06	0.23	0.09	0.29
Repayment cohort	2002	6	2001	6	2002	6
Default	0.04	0.19	0.02	0.12	0.05	0.23
Homeowner	0.39	0.49	1	0	0	0
Home price	244,882	171,694	243,653	171,286	245,633	172,019

Table 2

Blinder-Oaxaca decomposition.

The table presents a Blinder-Oaxaca style decomposition to explain the rise in student loan defaults during the Great Recession. Elasticity is the coefficient from a regression at the zip code level of the percentage change in student loan defaults from 2007 to 2010, $\Delta \text{Log default}_{07-10}$, on either the percentage change in the share of nontraditional borrowers attending for-profit institutions and community colleges from 2006 to 2009, $\Delta \text{Log NT share}_{06-09}$ (column 1), the percentage change in home prices from 2006 to 2009, $\Delta \text{Log home price}_{06-09}$ (column 2), or the percentage change in employment from 2006 to 2009, $\Delta \text{Log employment}_{06-09}$ (column 3). Zip codes are weighted by total student loan balances. A student loan goes into default if it is more than 270 days past due. When a loan goes into default, the loan servicer has up to 90 days to report the default to the NSLDS. Thus, there is approximately a one-year time lag between when a payment is missed and when a default is recorded in the NSLDS. Change is the value (in percent) of either $\Delta \text{Log NT share}_{06-09}$ (column 1), $\Delta \text{Log home price}_{06-09}$ (column 2), or $\Delta \text{Log employment}_{06-09}$ (column 3). Share explained is $\text{Elasticity} \times \text{Change} / \Delta \text{Log default}_{07-10}$, where $\Delta \text{Log default}_{07-10}$ is 18.9%. Home price data are from Zillow. Employment data are from the US Census Bureau's zip code business patterns. All other data are from a 4% random sample of the NSLDS matched to de-identified IRS tax data. Standard errors (in parentheses) are clustered at the zip code level. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	$\Delta \text{Log default}_{07-10}$		
	NT share (1)	Home price (2)	Employment (3)
Elasticity	0.3282	-0.4179	-0.0004
Change	16.9%	-14.4%	-4.7%
Share explained	29.3%	31.8%	0.0%

Table 3

Home prices and student loan default at the regional level.

In Panel A, the dependent variable is the percentage change in student loan defaults from 2007 to 2010, $\Delta \text{Log default}_{07-10}$, at the regional level. A student loan goes into default if it is more than 270 days past due. When a loan goes into default, the loan servicer has up to 90 days to report the default to the NSLDS. Thus, there is approximately a one-year time lag between when a payment is missed and when a default is recorded in the NSLDS. $\Delta \text{Log home price}_{06-09}$ is the percentage change in home prices from 2006 to 2009 at the regional level. In Panel B, the dependent variable is the rate of student loan defaults in year $t+1$ at the regional level. Home price_t is the home price (in logs) in year t at the regional level. All columns include region and year fixed effects, and standard errors (in parentheses) are clustered at the regional level. In both panels, regions are weighted by total student loan balances. Regions are either zip codes (column 1), counties (column 2), or commuting zones (column 3). Home price data are from Zillow. All other data are from a 4% random sample of the NSLDS matched to de-identified IRS tax data. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

Panel A: Long difference specification

	$\Delta \text{Log default}_{07-10}$		
	Zip code (1)	County (2)	Comm. zone (3)
$\Delta \text{Log home price}_{06-09}$	-0.4179*** (0.117)	-0.674*** (0.214)	-0.712*** (0.222)
Observations	12,749	1,234	408

Panel B: Panel specification

	Default _{t+1}		
	Zip code (1)	County (2)	Comm. zone (3)
Home price _t	-0.00614*** (0.0022)	-0.00639*** (0.00244)	-0.00728*** (0.00260)
Year fixed effects	Yes	Yes	Yes
Region fixed effects	Yes	Yes	Yes
Observations	48,252	4,994	1,644

Table 5

Main results by institution type.

The table presents variants of the specification in column 1 of Table 4 in which the sample is divided into subsamples based on institution type (columns 1 to 4) and institutional selectivity (columns 5 and 6), respectively. Public refers to public not-for-profit four-year institutions. Private refers to private not-for-profit four-year institutions. For-profit refers to for-profit institutions. Comm. college refers to public and private not-for-profit two-year institutions. Institutional selectivity, as measured by Barron's, is based on the fraction of applicants that institutions admit. Observations are weighted by individual student loan balances. Home price data are from Zillow. All other data are from a 4% random sample of the NSLDS matched to de-identified IRS tax data. Standard errors (in parentheses) are clustered at the zip code level. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Default _{t+1}					
	Public (1)	Private (2)	For-profit (3)	Comm. college (4)	Nonselective (5)	Selective (6)
Home price _t	-0.00764** (0.00380)	-0.00787* (0.00453)	-0.0377*** (0.00969)	-0.0220* (0.0114)	-0.0205*** (0.00563)	-0.00668** (0.00317)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Zip code fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	462,777	344,497	136,517	127,258	458,141	612,908

Table 6

Main results by individual labor earnings.

The table presents variants of the specification in column 1 of Table 4 in which the sample is divided into subsamples based on student loan borrowers' individual labor earnings. Observations are weighted by individual student loan balances. Home price data are from Zillow. All other data are from a 4% random sample of the NSLDS matched to de-identified IRS tax data. Standard errors (in parentheses) are clustered at the zip code level. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Default _{t+1}			
	< \$20,000 (1)	\$20,000– \$40,000 (2)	\$40,000– \$60,000 (3)	> \$60,000 (4)
Home price _i	-0.0176** (0.00687)	-0.0129** (0.00536)	-0.0113** (0.00538)	-0.00498 (0.00396)
Year fixed effects	Yes	Yes	Yes	Yes
Zip code fixed effects	Yes	Yes	Yes	Yes
Observations	353,771	341,299	197,294	178,685

Table 7

Home prices and individual labor earnings.

The table presents variants of the specification in column 1 of Table 4 in which the sample is divided into subsamples based on borrowers' individual labor earnings, and in which the dependent variable, *Earnings drop_t*, is an indicator of whether a student loan borrower's individual labor earnings drop by 50% or more relative to the previous year's earnings. Observations are weighted by individual student loan balances. Home price data are from Zillow. All other data are from a 4% random sample of the NSLDS matched to de-identified IRS tax data. Standard errors (in parentheses) are clustered at the zip code level. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Earnings drop _t			
	< \$20,000 (1)	\$20,000– \$40,000 (2)	\$40,000– \$60,000 (3)	> \$60,000 (4)
Home price _t	-0.0135** (0.00662)	-0.0108** (0.00466)	0.00682 (0.00423)	-0.000485 (0.00397)
Year fixed effects	Yes	Yes	Yes	Yes
Zip code fixed effects	Yes	Yes	Yes	Yes
Observations	353,771	341,299	197,294	178,685

Table 9

Homeownership and financial resources.

The table presents variants of the specification in column 1 of Table 8 in which either *Labor earnings_t*, and *Home price_t × Labor earnings_t* (column 1), *Total income_t*, and *Home price_t × Total income_t* (column 2), or *Family income_t*, and *Home price_t × Family income_t* (column 3) are included as additional regressors. Labor earnings and Total income are described in Table 1. Observations are weighted by individual student loan balances. Home price data are from Zillow. All other data are from a 4% random sample of the NSLDS matched to de-identified IRS tax data. Standard errors (in parentheses) are clustered at the zip code level. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

Resources:	Default _{t+1}		
	Labor earnings (1)	Total income (2)	Family income (3)
Home price _t	-0.0112*** (0.00285)	-0.0112*** (0.00297)	-0.0113*** (0.00285)
Home price _t × Owner _t	0.000972 (0.000882)	0.000972 (0.000896)	0.00101 (0.000885)
Owner _t	-0.0439*** (0.0108)	-0.0439*** (0.0110)	-0.0431*** (0.0109)
Home price _t × Resources _t	0.000211*** (0.0000181)	0.0000889*** (0.0000171)	0.0000628*** (0.00000679)
Resources _t	-0.00325*** (0.000219)	-0.00142*** (0.000202)	-0.000977*** (0.0000700)
Year fixed effects	Yes	Yes	Yes
Zip code fixed effects	Yes	Yes	Yes
Observations	1,062,914	1,062,914	1,062,914

Table 10

Homeownership results by age and repayment cohort.

The table presents variants of the specification in column 1 of Table 8 in which the sample is divided into subsamples based on either age (columns 1 to 3) or repayment cohort (columns 4 to 6). Repayment cohort is the year in which a student loan borrower enters into repayment. Observations are weighted by individual student loan balances. Home price data are from Zillow. All other data are from a 4% random sample of the NSLDS matched to de-identified IRS tax data. Standard errors (in parentheses) are clustered at the zip code level. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Default _{t+1}					
	Age			Repayment cohort		
	< 30 (1)	30–40 (2)	> 40 (3)	< 2000 (4)	2000–2005 (5)	> 2005 (6)
Home price _t	-0.0114*** (0.00441)	-0.0110*** (0.00407)	-0.0171*** (0.00638)	-0.0142** (0.00606)	-0.00898** (0.00397)	-0.0107** (0.00449)
Home price _t × Owner _t	0.00247* (0.00132)	0.000786 (0.00137)	-0.000855 (0.00205)	-0.00167 (0.00208)	0.000619 (0.00118)	0.000887 (0.00149)
Owner _t	-0.0680*** (0.0162)	-0.0423** (0.0168)	-0.0154 (0.0253)	0.00370 (0.0257)	-0.0350** (0.0144)	-0.0513*** (0.0182)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Zip code fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	362,166	427,006	298,313	283,557	470,889	423,021

Table 11

Income based repayment program.

The table presents variants of the specification in column 1 of Table 4 in which the sample period is extended to include student loan defaults up to 2012, and in which $IBR\ eligible_t$, $Home\ price_t \times Post$, $IBR\ eligible_t \times Post$, $Home\ price_t \times IBR\ eligible_t$, and $Home\ price_t \times IBR\ eligible_t \times Post$ are included as additional regressors. $IBR\ eligible_t$ is a dummy indicating whether a student loan borrower passes the means test in year t , as described in Section 6. $Post$ is a dummy that equals one beginning in 2009. In column 3, the sample is restricted to student loan borrowers with relatively high ratios of student debt to income, as described in Section 6. Columns 1 to 3 include zip code fixed effects. Column 4 includes individual borrower fixed effects. All columns include year fixed effects. Observations are weighted by individual student loan balances. Home price data are from Zillow. All other data are from a 4% random sample of the NSLDS matched to de-identified IRS tax data. Standard errors (in parentheses) are clustered at the zip code level. *, **, and *** denotes significance at the 10%, 5%, and 1% level, respectively.

	Default _{t+1}			
	Full (1)	Full (2)	Restricted (3)	Full (4)
Home price _t	-0.00419** (0.00209)	-0.00338 (0.00209)	-0.00439 (0.00326)	-0.00441* (0.00231)
IBR eligible _t	0.0256*** (0.000531)	0.0425** (0.0183)	0.0510** (0.0235)	0.0274 (0.0229)
Home price _t × Post		-0.000366 (0.000816)	-0.00168 (0.00138)	-0.000967 (0.000928)
IBR eligible _t × Post		-0.0404* (0.0215)	-0.0594** (0.0269)	-0.0707*** (0.0258)
Home price _t × IBR eligible _t		-0.00259* (0.00146)	-0.00345* (0.00188)	-0.00157 (0.00184)
Home price _t × IBR eligible _t × Post		0.00314* (0.00173)	0.00467** (0.00217)	0.00513** (0.00209)
Year fixed effects	Yes	Yes	Yes	Yes
Unit fixed effects	Zip code	Zip code	Zip code	Individual
Observations	1,556,296	1,556,296	658,504	1,556,296