

# The Student Debt Crisis is a Crisis of Non-Repayment

Marshall Steinbaum\*

November 17, 2020

## Abstract

This brief analyzes the increased prevalence of the non-repayment of student debt, primarily due to the increased takeup of the various Income-Driven Repayment (IDR) programs since the middle of the 2000s. It shows deteriorating repayment over time and across borrower cohorts, as well as suggestive evidence of a significant and growing repayment gap between minority and white borrowers. The implication of rising non-repayment is that the cancellation of a large (and increasing) share of outstanding student debt is inevitable.

Jain  
Family  
Institute

---

\*Assistant Professor, University of Utah Economics Department, and Senior Fellow in Higher Education Finance, Jain Family Institute. Thanks to Eddie Nilaj for research assistance and to Laura Beamer for coordinating this project.

# 1 Introduction

Think of the student debt crisis as an overflowing bathtub. On the one hand, too much water is pouring in: more borrowers are taking on more debt, thanks to increased demand for higher education in the face of rising tuition, stagnant wages, diminishing job opportunities for those with less than a college degree, and the power of employers to dictate that would-be hires have the necessary training in advance. On the other hand, the drain is clogged, so too little water is draining out: those who have taken on debt are increasingly unable to pay it off.

The last post in the Millennial Student Debt project used a new database of student debtors and their loan characteristics (matched to demographic and economic data in the American Community Survey) to document the former phenomenon, both in aggregate and particularly as it pertains to disadvantaged communities along multiple dimensions. Specifically, it showed the rapid growth of student debt levels and of debt-to-income ratios in the population at large, among people of all income levels. But this growth is concentrated among non-white borrowers, who have higher debt conditional on income and whose increased indebtedness over the past decade-plus is greater than for white borrowers. That racial disparity is particularly pronounced in the middle of the income distribution. It also showed that student-debt-to-income ratios have grown fastest in the poorest communities since 2008.

This post uses the same data to document the latter phenomenon: non-repayment by student loan borrowers is getting worse over time, especially so for non-white debtors. Over the last ten years, outstanding student loan debt has mounted and been assumed by a more diverse, less affluent group of students and their families than was the case for prior cohorts. A common policy response has been to wave away its impact on wealth, both individually and in aggregate, by saying that the debt finances its own repayment. First of all, so the claim goes, student debt finances college degrees that in turn pay off in the form of higher earnings, enabling debtors to repay. Second, expanded allowance for income-driven repayment (IDR), by capping debt service as a share of disposable income, eliminates the

worst forms of delinquency and default. The first claim says that repayment is inevitable, the second that it need not take place. Both claims together, however, serve to rationalize higher debt, higher tuition, higher attainment, and the forces driving all three.

IDR was designed to address a liquidity crunch: since students are graduating with more debt, they may not earn enough immediately upon entering the workforce to pay it down. That failure of earnings to align with debt service obligations means that a program to defer those obligations until earnings are realized would ameliorate delinquency and default, at the cost of capitalizing unpaid interest into a higher principal balance. The creation and expansion of IDR programs in the early 2010s did indeed serve to stop the growth of delinquency by the mid-2010s and reverse it, to the point that the share of accounts delinquent now is lower than it was before the Great Recession, despite the amount of debt and the number of debtors having increased continuously and diversified since then. For that reason, many higher education policy analysts have proposed further expanding IDR programs.

But IDR programs will never be successful as a solution to the student debt crisis, because they're designed to address a liquidity problem rather than the real problem: solvency. The problem with student debt is a problem of wealth. Students and their families are taking on debt because they don't have enough wealth to afford increasingly-costly, increasingly-mandatory higher education. The debt then itself exacerbates wealth disparities that the higher education it "paid for" doesn't rectify.

Moreover, we now have experience with steadily-increasing student indebtedness extending over two whole business cycles, through the expansion of the 2000s, the Great Recession, the long, slow recovery of the 2010s, and now, the current COVID recession. Any idea that temporary economic downturns were responsible for the crisis of non-repayment, and progress would be made up during expansions when labor markets are tight, has now been definitively disproved. What we've considered to be economic prosperity of the last ten years, prior to the pandemic, was in fact economically punishing to younger cohorts forced through

the wringer of increasingly costly higher education and into a labor market characterized by stagnant wages and deteriorating job ladders.

The fact that balances continue to grow and not be paid back shows that shifting the cost of higher education onto the backs of students, through state budget cuts, slack labor markets, and employer monopsony power, has resulted in a structural attenuation of the economic life cycle: balances carried further into mid-life, or taken on later in life to finance further education or a family member's education, impairing economic wellbeing for a widening and diversifying swathe of the population, inhibiting savings, increasing precarity, and draining the very incomes the student debt was supposed to increase.

Student debt cancellation is inevitable, in part because it's already happening: a larger and larger fraction of the loans originated in any given year are not going to be repaid at all, and many more will never be fully repaid. For example, as is shown below, 63.8% of the loans originated in 2018 already had higher balances than they did at origination by June 2019. That number is similar, around 60% and rising, for loans originated between 2013 and 2015, meaning that repayment should have begun following any deferment period, but in fact it has not. Rising IDR enrollment corresponds to deteriorating repayment trajectories and growing loan balances over time.

The policy question is then not whether to cancel student debt, but how? Are we going to rely on increasing IDR and thus kick the ball further into the future, ruining more and more lives as we do so, while still not actually collecting on the government's loans and potentially imposing a large tax liability when the debt is finally canceled? Or can we honestly confront and solve a past policy failure in the present by cancelling debt now (and not taxing it), to prevent the debt spiral from getting any worse and rectifying the damage student debt does to household wealth?

IDR, and even modest cancellation proposals that still leave many borrowers with unpaid loans, underestimate the crisis of non-repayment and rest on the false assumption that entrance into the labor market allows borrowers to achieve solvency. The debate that is

trendy to have nowwhether cancellation at all is inequitable, how it can be structured to ensure it's equitable needs to be compared to a baseline in which the consequences of not cancelling debt are understood to be escalating debt peonage, exacerbating wealth inequality, followed by inevitable cancellation, because the debt won't be paid back.

## 2 Data

For this project, we have assembled a dataset that consists of individual credit reports linked to demographic and economic data reported from the American Community Survey using the subject's zip code. The credit reports are available in several formats: we have both loan-level data and individual consumer-level data. We also have both panel and cross-section data. For the panel data, one million consumers between the ages of 18 and 35 with a positive student loan balance were sampled in 2009, and then followed through 2019. For the cross-section data, one million consumers in that same age range are sampled, independently, in each year from 2009 to 2019. The loan-level data consists of the loans of all of these people, and the loans are linked to individuals with a unique consumer ID.

## 3 Analysis

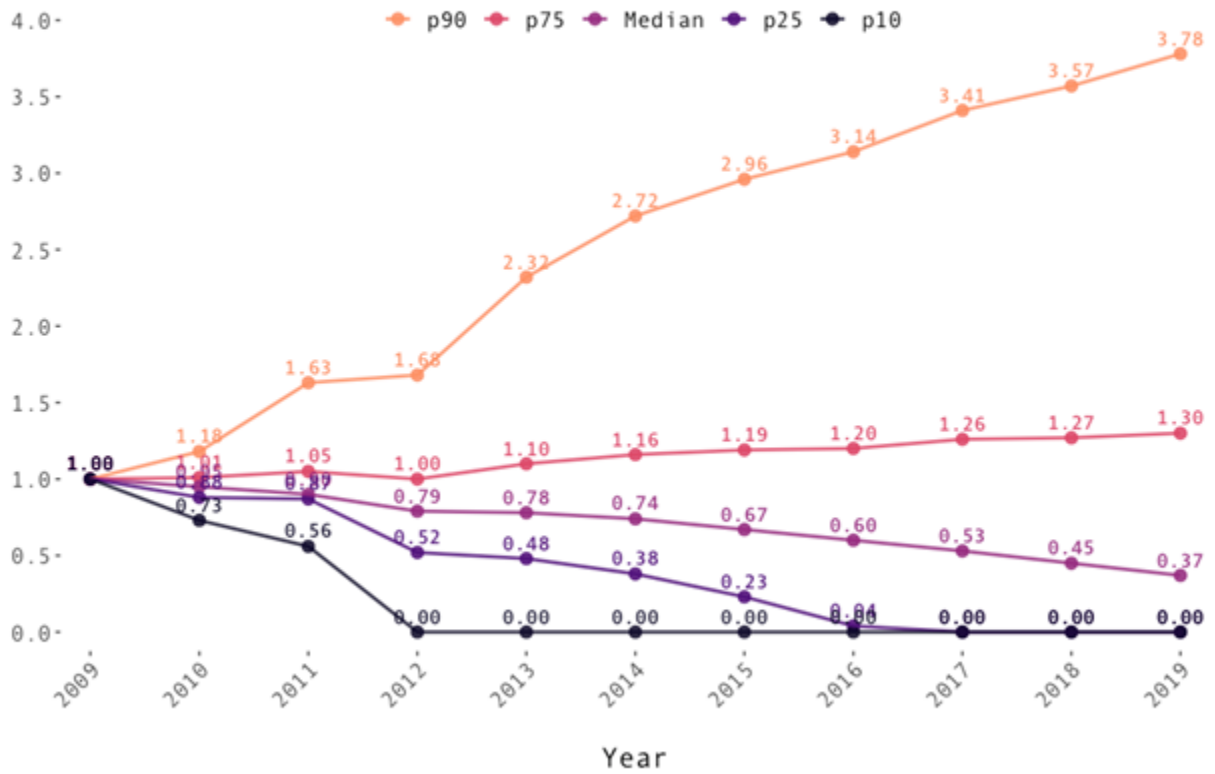
We look first at the panel data cohort, that is, 1 million individuals (between 18 and 35 years old) with student debt outstanding in 2009. It's important to keep in mind that the standard repayment schedule for student loans is a 10-year uniform payment plan, akin to a fixed-rate mortgage.

We compute the ratio of the amount of student debt outstanding in each year as a share of that individual's initial loan balance in 2009. We then compute the quantiles of the distribution of that ratio in each year. Figure 1 plots those quantiles over time. The darker lines (10th and 25th percentiles) show progress toward repayment for individuals who were probably already on the path to repayment when they were observed in 2009. Approximately

Figure 1

### Balance as a Share of 2009 Starting Balance

For the panel data cohort, this chart tracks outstanding student loan balance as a share of the "initial" 2009 balance, so a value of 1 indicates that an individual has the same amount of student debt outstanding in a given year as they did in 2009. Quantiles of the distribution of that ratio (computed year by year) are plotted over time.



Source: Experian Information Solutions, Inc.  
© Jain Family Institute, 2020.

40% of the individuals with outstanding student debt in 2009 had paid off all of their student debt by 2019, within the standard 10-year repayment window (or sooner, given that they were probably already well into repayment in 2009).

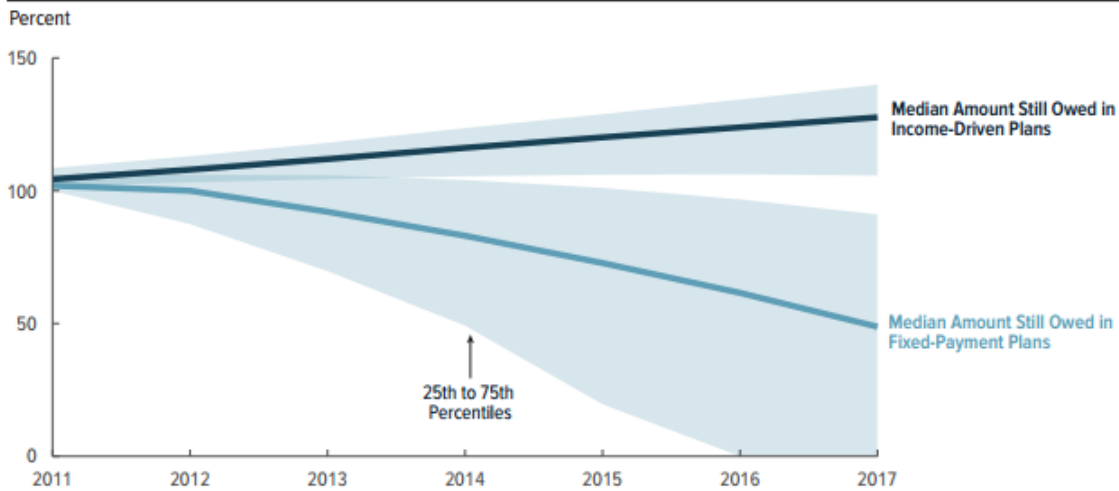
The purple line, the median of the distribution, by contrast, declines slowly and reaches 0.37 by 2019. That means that over 50% of the borrowers with outstanding debt in 2009 hadn't fully paid back their student loans ten years later. The 75th percentile stands at 1.3 in 2019, meaning that more than 25% of borrowers with outstanding student debt in 2009 had a larger student loan balance ten years later. The 90th percentile is much higher: 3.78, almost four times as much debt outstanding in 2019 as in 2009.

The lack of progress toward repayment can be caused by a number of different behaviors, as well as any combination of them: delinquency, enrollment in an IDR plan, and taking on more debt after these individuals are first observed in 2009 by re-enrolling, and thereby likely deferring payments on prior student loans. Given the age parameters of the dataset, it's likely that a number of borrowers in the 2009 cohort were not yet done with their education at that point, and thus took out more loans subsequently. It's common for student loans to increase in principal balance over their first few years after origination, due to deferral. On the other hand, the individuals in this cohort sample don't change over time. Their minimum age in 2009 was 18, so the minimum age of this group in 2019 is 28. Taking on more debt as they get older makes them less likely to pay off the debt they already had. And that pattern of behavior becomes more pronounced for subsequent cohorts, though this dataset does not directly track them over the full ten-year period that we have for the initial 2009 panel cohort.

Figure 2-5 from a recent Congressional Budget Office report on IDR shows a similar pattern. That chart plots progress to repayment by what CBO calls the "2012 cohort," which means individuals who began repayment in 2012 (and thus excludes the deferment period on those individuals' loans prior to that year). That report compares individuals who remained in the standard 10-year repayment plan throughout to those who utilized one of

Figure 2-5.

**Percentage of Original Balance Remaining for Borrowers Who Began Repaying Loans in 2012, by Type of Repayment Plan**



Source: Congressional Budget Office, using data from the Department of Education's National Student Loan Data System.

CBO examined borrowers who began repaying their loans in 2012 because the number of borrowers and the volume of loans in income-driven plans had begun increasing by that year and the borrowers' payments for the next several years could be observed. Borrowers who defaulted on their loans were excluded from the analysis.

the IDR programs. The former saw their balance decline over time, while, not surprisingly, the latter group's balance increased. That appears to confirm the interpretation of Figure 1 that a major contributor to the lack of progress toward repayment for many borrowers is enrollment in an IDR plan.

Again, the purpose of IDR is to prevent a liquidity problem by deferring payments. The alternative is much higher delinquency and default rates than there would otherwise have been, so in that sense, the various IDR programs are a success on their own terms. But they are a failure insofar as they give rise to mounting balances over time, rather than facilitating repayment on a delayed schedule more favorable to borrowers. The structure of IDR, capped payments as a percentage of income in the present, possible cancellation of remaining balance in the far future, presumes that those entering into the program will eventually earn a high enough income to cover their full interest and principal loan payments before the date at which their loans would be cancelled. But many borrowers enter into IDR with no intention or prospect of ever leaving it. In fact, the program itself contains an incentive not to exit it:



the cancellation at the end, the balance of which the borrower would otherwise be responsible for. That is why a great deal of student debt cancellation has already been committed to and a rising share of new student loans will never be repaid. The policy question is essentially whether to continue to insist on futile interest payments in the meantime, or to recognize that the underlying debts are un-repayable now. Trying to push people out of IDR programs once they've structured their lives and careers to conform to their terms, in order to improve the prospects of repayment would likely lead to a flood of delinquencies, as better-paid jobs are hardly abundant and available for the asking in a monopsonized labor market. In any case it would be unfair to borrowers who've done as they were told by policy-makers, when in reality it was the policy-makers whose failures led to the crisis of non-repayment.

The lack of progress toward repayment means that many student loans increase in balance over their lifetime, the opposite of the standard repayment structure in which, following any deferment, principal is steadily drawn down by uniform loan repayments until the loan reaches a principal balance of zero upon full repayment. Thus, another window on the non-repayment of student loans is the share of loans that have a higher principal balance than they did when they were originated. The loan-level data in our credit reporting dataset enables us to compare the current balance on a loan when observed in the dataset to the balance of that loan at origination.

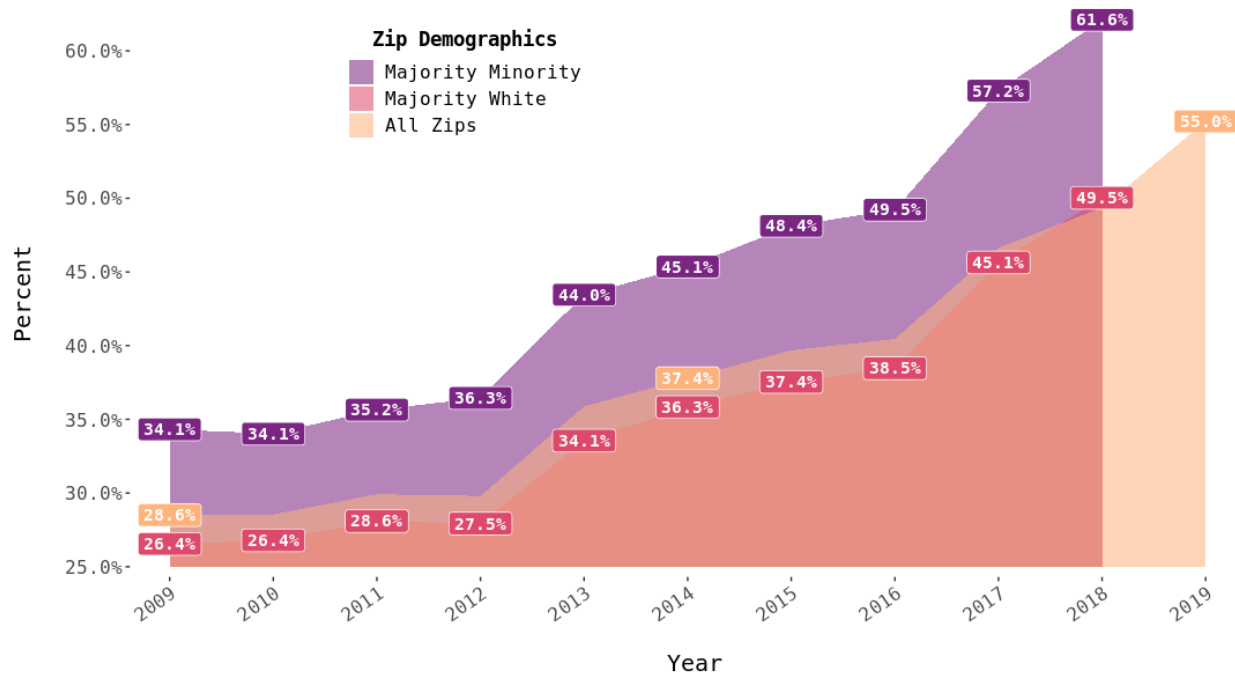
Figure 2 plots the percentage of loans at each observation date that had a higher balance at that point than they did at origination. This is taken from the loan-level cross section data, as opposed to the panel data used for Figure 1. In each year, one million individuals with student debt between the ages of 18-35 are observed, along with all of their student loans. The yellow series shows the share of student loans in each year of the cross section in which the current balance exceeds the starting balance.

The credit reporting data also includes each individual's zip code. We match that to zip-code-level demographic data from the American Community Survey to try to see how non-repayment differs by race. The pink series shows that same share of loans for individuals

Figure 2

### Share of Student Loans where Current Balance Exceeds Original

Each point references the percent of student loans that have a current balance that is higher than the loan's original balance, by year. This data is broken down to show the disparities by neighborhood racial demographics. Demographic data is from zip level ACS 5-year estimates (2009-2018). The 'All Zips' trend includes data from majority minority, majority white, as well as unknown/unreleased zips.



Source: Experian Information Solutions, Inc.  
© Jain Family Institute, 2020.

who reside in zip codes that are majority white, and the purple series shows that share for individuals who live in majority-minority zip codes. Even though there's no direct information about the race of the actual individuals in the credit reports, and plenty of non-white people live in majority-white zip codes, while a smaller but still significant share of white people live in majority-minority zip codes, we can be fairly certain from this data that non-repayment as measured by loans with a higher balance than their original balance is concentrated among non-white borrowers. That racial gap in non-repayment is significant throughout and rises slightly over time even while the overall rate of non-repayment also rises. (Note the kink point in 2016, after which non-repayment across all zip codes increases dramatically.) Thus, although borrowers with higher balances are likelier to enroll in IDR (because the savings from doing so is greater the higher are the payments you would otherwise have to make), and higher-balance borrowers are, in general, higher-income, it's also very likely that non-white borrowers are more likely to have enrolled in IDR, since they carry more debt conditional on income than white borrowers do. The pattern depicted in Figure 2, both over time and in the cross-section of zip codes, is consistent with that inference.

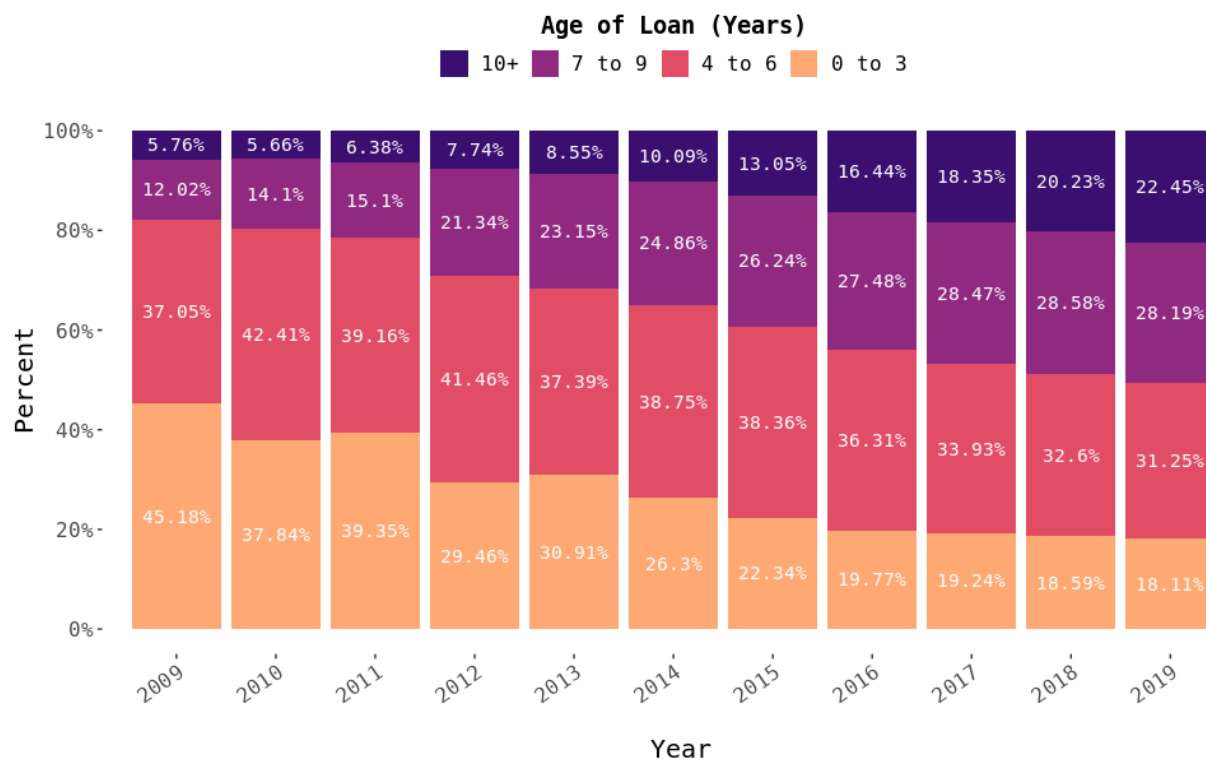
One concern with this metric as a measure of non-repayment is that loans in deferment increase their principal balance as a matter of course, and 'younger' loans are more likely to be in deferment. Therefore, it's theoretically possible that an increasing share of loans with a higher current balance than initial balance reflects the changing age distribution of loans: if more loans are taken out cohort-by-cohort, over time an increasing share of student loans would be younger and thus more likely to be in deferment. If that is true, then a greater share of student loans would appear by this metric to be in non-repayment status, even though that's because of rising share of loans in deferment, not IDR, delinquency, or some other reason pertaining to the non-repayment of loans over their term once repayment has begun.

Figure 3 addresses this concern by illustrating how the student loan age distribution has changed over time. It's getting older, not younger, which is itself evidence of declining

Figure 3

### Distribution of Student Loans, by Age of Loan (2009-2019)

The age distribution of the loans in each year of the cross-section (amounting to a million loan borrowers per year, including all of their loans) has been getting older over time.



Source: Experian Information Solutions, Inc.  
© Jain Family Institute, 2020.

repayment. It also suggests that the non-repayment or “level” effect outweighs the increased-origination or “composition” effect. Note especially the rising share of loans that are 10+ years old, which is to say, older than the term in which a traditional student loan is supposed to be paid off.

To return to the bathtub metaphor, overall, the water in the tub is getting ‘older’ because the ‘old’ water that isn’t draining out and thus remains in the tub for longer counteracts the ‘young-ness’ of the ‘new’ water pouring in. The implication for Figure 2, then, is that even though you’d expect more loans to be out of deferment and in repayment due to the mix of loans getting older, we still see increasing balances and therefore not actual repayment. And that, in turn, is further evidence that if we had 10 years of panel data for subsequent cohorts following the 2009 cohort depicted in Figure 1, that data would show even worse progress toward repayment for those younger cohorts.

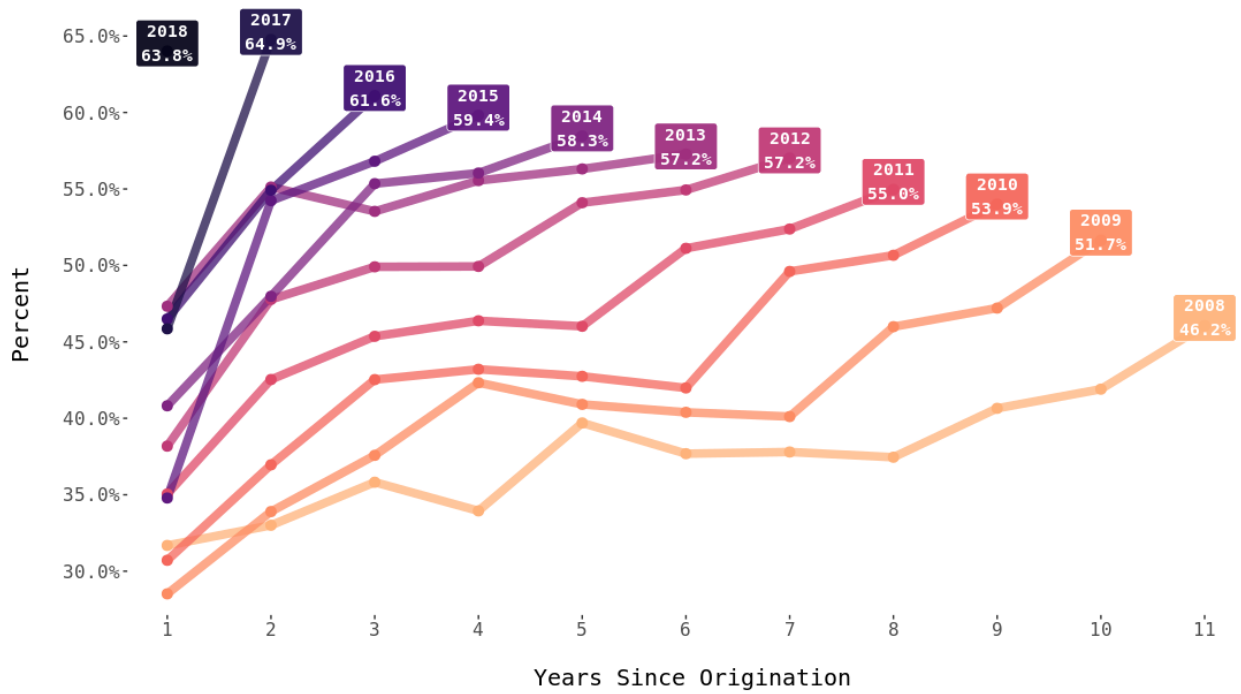
Figure 4 is a further illustration of the phenomenon of non-repayment similar to Figure 2. In this case, years refer to the loan’s origination year (its “vintage”), and the horizontal axis tracks the years since that initial origination year. Thus, we have longer series in the data covering 2009-2019 for loans originated near the start of that time period. As we can see, the share of loans with a higher current balance than initial balance rises vintage-by-vintage. All the vintages show an initial increase within the first few years, likely reflecting high prevalence of deferment for younger loans. They then level off, until the year 2016, which is a kink point for all vintages (thus reflecting the same pattern as in Figure 2). After that year, every vintage includes more loans with a higher current balance than initial balance, despite the fact that each vintage is getting older. And, most importantly, this metric shifts up with each vintage, suggesting that non-repayment is getting worse both over time and across vintages.

The kink point in 2016 likely reflects the expansion of IDR programs toward the end of the Obama administration, with the creation of the REPAYE program and aggressive marketing of other IDR programs as a solution to delinquency, as detailed in the aforementioned CBO

Figure 4

**Share of Current Balances that Exceed Original, by Origination Year**

Each line references student loans originating in a specific year. The y-axis is the share of student loans in the sample that have a balance above the original loan amount, including loans that were paid off over the period of analysis. The x-axis is years since the student loan's origination year.

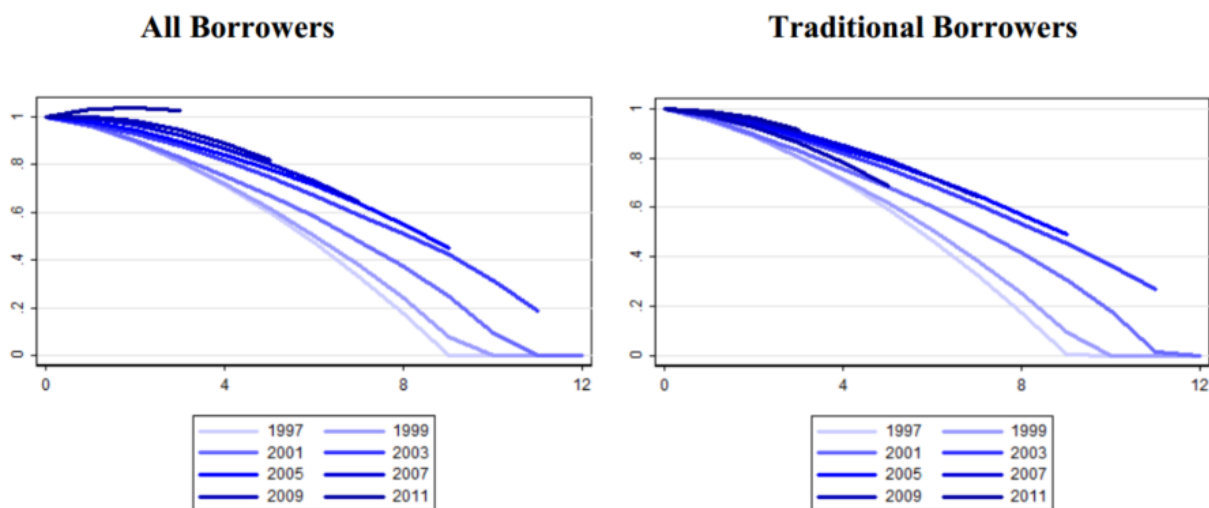


Source: Experian Information Solutions, Inc.  
© Jain Family Institute, 2020.

report.

Figure 4 is essentially the mirror image of Figure 15 from the 2015 paper “A crisis in student loans? How changes in the characteristics of borrowers and in the institutions they attended contributed to rising loan defaults” by Adam Looney and Constantine Yannelis. That figure is reprinted here. It tracks loan repayment for successive repayment cohorts (defined, as with the CBO, by year a borrower started repayment). That also shows repayment dwindling over time and across cohorts. Their data concludes in 2013. Since then, the economy has improved significantly (up until 2020), and yet repayment has only gotten worse.

**Figure 15: Slowing Rates of Repayment: Median Balance Remaining by Cohort and Year**



The subject of the paper by Looney and Yannelis is delinquency and default, primarily by what those authors call non-traditional borrowers, meaning they attended post-secondary education later in life and/or attended for-profit institutions. The authors rightly attribute the influx of non-traditional borrowers into the student loan system as a consequence of the Great Recession and the generally-poor labor market options available to workers in its aftermath, especially those with no post-secondary qualification. The purpose of their paper is to propose IDR as a solution to that rise in delinquency, which it is. In their conclusion,

the authors write:

Because of the ‘life cycle’ of borrowing, delinquencies and defaults are a lagging indicator, and the current [as of 2015] high rate of delinquency obscures some more favorable recent trends. In particular, the number of new borrowers at for-profit and 2-year institutions has dropped substantially, due to the end of the recession and to increased oversight of the for-profit sector, which is likely to improve the risk characteristics of future repayment cohorts. Moreover, borrowing by first-year borrowers and other enrolled students has declined, both in the number of borrowers and the amounts borrowed. These factors, coupled with efforts by the Department of Education to expand and encourage the use of income-based repayment programs are likely to put downward pressure on loan delinquency in the future, although with a lag.

This turned out to be an accurate prediction: delinquency has been in decline since that paper was released. But as we have shown, the result is rising balances and the crisis of non-repayment. The idea that deferring currently-due payments would make repayment easier when the labor market improved turned out not to be the case.

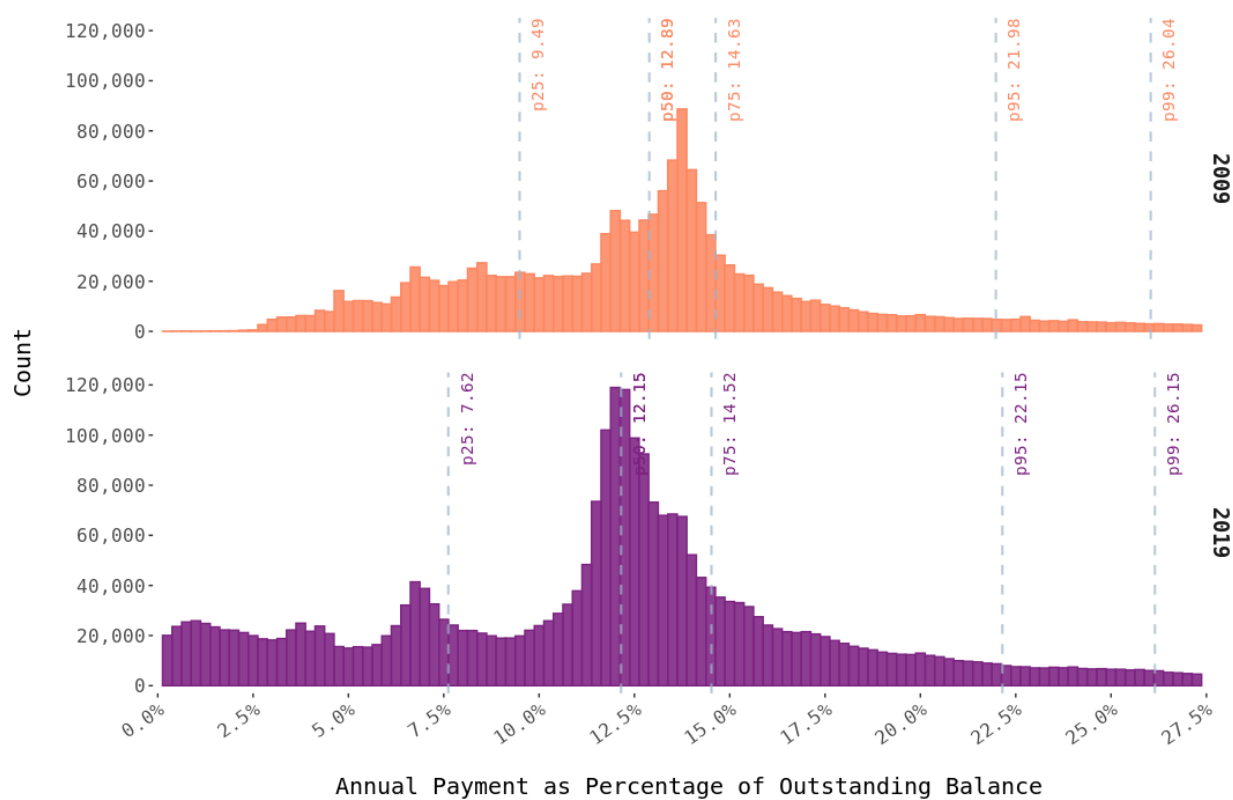
We performed one final exercise to identify the effect of non-repayment, which is depicted in Figure 5. For each loan, we calculate the required annual payment amount (which is either the standard repayment schedule for a given initial balance and interest rate or an adjusted amount under IDR) as a share of the total outstanding balance, to see how the cost of carrying student loans has changed over time. Figure 5 plots two histograms for that number, one for all the loans in the cross-section data for 2009, the other for 2019. For reference, a standard 10-year uniform repayment plan would have that number in the neighborhood of 12-13% initially (i.e. when the total balance is highest, before progress toward repayment), depending on the interest rate. Over the life of a loan that is actually being repaid, the number increases since the required annual payment stays the same and the balance decreases. For that reason, as the distribution of outstanding loans gets older (recall



Figure 5

### Distribution of Annual Payments Due as Share of Outstanding Balance

Distribution of student loans, held by 18-35 year-olds, showing the loan's estimated annual payment amount (monthly payment due \* 12) as a percentage of the outstanding balance. Data conditional on loans with an outstanding balance of \$100 or more. The vertical lines show the 25, 50, 75, 95, and 99 percentiles for both years. Scale capped at 27.5% for viewing purposes.



Source: Experian Information Solutions, Inc.  
© Jain Family Institute, 2020.

Figure 3), we should expect to see the distribution of this annual-payment-to-outstanding-balance ratio shift to the right.

Instead, we see a much larger share of loans have required payments of zero or near-zero, which likely reflects the advent of IDR between 2009 and 2019. For loans in IDR, this number would decline over time, as the required annual payment amount is capped (conditional on income) and the balance is growing. There are also more significant mass points in the 2019 distribution around 7.5% and 12.5%, which may also reflect IDR required payment levels, although without income data it's hard to be sure of that. Finally, there are, overall, higher payment ratios in the right half of the distribution, so the variance in the whole distribution has increased a good deal. That is probably for the same basic reason as we see rising variance in progress to repayment in Figure 1 and elsewhere in this analysis: as the federal student loan program has grown in size, more borrowers are coming from low-income and minority communities. As borrowers have diversified, more are in a worse position repay and therefore suffer from rising balances over time.

## 4 Conclusion

In light of rising student loan balances and their non-repayment, higher education policy analysts have tended to favor either expanding IDR and streamlining enrollment therein or limiting the reach of IDR and capping loan amounts, in the hope of re-directing students into less expensive institutions and into careers that will, ostensibly, enable them to repay their loans as scheduled. The effect of that is likely to be credit rationing: less advantaged students taking on more expensive private loans or not receiving higher education at all. The effect of expanding IDR is what we've already seen: balances mounting over time, with little or no progress toward repayment, even when the economy is as near to full employment as has been achieved in the United States in recent memory. Both of these alternative diagnoses decline to face reality head-on: we already have a great deal of student debt outstanding that

isn't being repaid and isn't going to be repaid, and that debt resides disproportionately with minority borrowers victimized both by labor market credentialization and a lack of family wealth.

The implications are dire when it comes to household formation and asset accumulation over the life cycle, which is attenuated for millennials and will be for subsequent generations as well. It also likely impairs marriage and fertility, labor market mobility, and entrepreneurship, not to mention retirement security for families who take on student debt on behalf of their children. The grand experiment in shifting the cost of higher education to individual students and workers, at the same time as we closed off access to the labor market for anyone without post-secondary credentials, driving a broader and more diverse swathe of the population into the increasingly-expensive higher education system, has resulted in greater educational attainment, but not a better-paid workforce. Instead the debt is a lifetime drag on social mobility, widening wealth disparities between people whose families could pay for their education, or who needed less education in order to qualify themselves for professional careers, and those who had to finance it themselves and who need more of it to access opportunities that are rationed on the basis of race and class background.