

Liquidity and Insurance in Student Loan Contracts: The Effects of Income-Driven Repayment on Borrower Outcomes

Daniel Herbst*

March 12, 2019

Abstract

Enrollment in income-driven repayment (IDR) plans for student debt has tripled in the past five years, yet little is known of its effects on borrower welfare. By pegging minimum monthly loan payments to borrowers' post-college income profiles, IDR may smooth consumption and improve financial well-being but carries a potential cost to social welfare through moral hazard. In this paper, I estimate the ex-post causal impact of IDR enrollment on student borrower outcomes using a novel dataset linking monthly federal student loan payments to credit bureau records for over one million student borrowers. My research design uses two identification strategies: a difference-in-differences design comparing borrowers with differential IDR take-up following delinquency calls from their loan servicer and an instrumental variables design exploiting variation in the tendency of randomly-assigned servicing agents to enroll borrowers in IDR. Within seven months of take up, IDR enrollees are 31 percentage points less likely to fall delinquent and pay down \$162 more student debt each month compared to those who remain on standard repayment plans. IDR enrollees have credit scores that are 8.6 points higher, hold 0.1 more credit cards, and carry \$293 higher credit card balances than non-enrollees one year after the servicing call, suggesting increased short-term consumption out of liquidity. IDR enrollees are also 1.7 percentage points more likely to hold a mortgage, an increase of 8 percent off of the pre-call mean, suggesting a positive effect of IDR on homeownership. Minimum monthly payments decrease by an average of \$169 following IDR take up, but return to standard levels within one year of enrollment, minimizing the potential impact of moral hazard through loan forgiveness. I find no evidence of labor supply effects through unemployment deferments or median zip-code level income. My results suggest IDR improves borrower welfare by correcting for a market failure in human capital financing, allowing financially distressed graduates to borrow against future income when they lack the credit or collateral to do so through private lending markets.

*Department of Economics, University of Arizona, McClelland Hall 401QQ, Tucson, AZ 85721 (email: dherbst@email.arizona.edu, website: www.danjherbst.com). I am extremely thankful to my advisers Will Dobbie and Ilyana Kuziemko for their guidance on this project, as well as Leah Boustan, Henry Farber, Alan Krueger, David Lee, Adam Looney, Alexandre Mas, Christopher Neilson, and Cecilia Rouse, who provided invaluable advice throughout its development. I also benefited from the helpful comments of David Arnold, Barbara Biasi, George Bulman, Felipe Goncalves, Steve Mello, David Price, Maria Micaela Sviatschi, and seminar participants at the APPAM Annual Research Conference, the CFPB Research Conference, the Federal Reserve Board, the IZA Economics of Education Workshop, the Jain Family Institute, Kansas State University, the MIT Golub Center, the National Academy of Education Research Conference, the National Tax Association Research Conference, Princeton University, the RAND Corporation, and the University of Arizona. Financial support was provided by the Princeton Industrial Relations Section, the MIT Golub Center for Finance and Policy, and the National Academy of Education Spencer Dissertation Fellowship.

1 Introduction

For many Americans, a college education is the largest investment they will make in their lifetimes. Yet, until recently, the only financing option for this investment has been short-term, fixed-payment debt through the federal government. Many economists blame these poorly structured debt contracts for the high default rates, low homeownership, and poor financial health observed among student borrowers.¹ Traditional student loan payments average over \$350 per month, fall on borrowers early in their careers, and provide no insurance against income shocks. Meanwhile, over one million student borrowers default each year, and millions more struggle to buy homes (Mezza et al., 2017), accumulate wealth (Bleemer et al., 2017), or choose their preferred career (Rothstein and Rouse, 2011), prompting policymakers to reconsider the structure of student loan contracts.

By far, the largest policy response to the student loans crisis has been Income-Driven Repayment (IDR). First introduced in the US in 2009, IDR sets minimum federal student loan payments to a fixed portion of borrowers' income until debt is repaid or some forgiveness period has been reached, offering an alternative to the flat, ten-year repayment schedule offered by traditional contracts. IDR enrollment has tripled since 2014, and today over \$300 billion in debt is repaid through IDR (Department of Education, 2017a). Similar programs have been adopted in the UK and Australia, where over 85 percent of students finance their higher education through income-contingent repayment schemes (Barr et al., 2017).

Even as IDR enrollment continues to rise, its effects on social welfare are largely unknown. By aligning debt repayment with the returns to college investment, IDR may provide the liquidity needed to help credit-constrained borrowers smooth consumption over the life-cycle and insure against transitory income shocks. However, these short-term liquidity benefits may be outweighed by the long-term social welfare costs of potentially higher debt balances. IDR promotes more timely repayment while reducing payment size, so its net effect on balance sheets is ambiguous. If enrollees' balances remain persistently high relative to traditional borrowers, debt forgiveness shifts the repayment burden to the taxpayer, reducing social welfare through redistributive inefficiencies, administrative costs, and distorted incentives.

¹For example, Barr et al. (2017) write, "The US student loan system is currently in crisis...mainly due to the fact that the US operates mortgage-type student loans: these are repaid over a set period of time, which places high repayment burdens on low earning graduates."

Assessing the costs and benefits of IDR requires causal estimates of its impact on borrowers' repayment, consumption, and financial health, but two obstacles have prevented researchers from identifying these effects. First, an empirical analysis of IDR requires high-frequency repayment data for many borrowers, but until now these data have been unavailable.² Second, IDR enrollees are, by design, a very selected group. Only borrowers with a sufficiently high debt-to-earnings ratio would benefit from enrolling in IDR. Low-income individuals typically have worse financial outcomes, but borrowers with high student debt balances are often highly educated and positively selected (Yannelis, 2016), so estimates that rely on cross-sectional comparisons of IDR enrollees to non-enrollees could therefore be biased in either direction.

In this paper, I use novel administrative data from a large student loan servicing company to estimate the causal effect of IDR enrollment on financial outcomes such as defaults, credit scores, and bankruptcies, as well as proxies for consumption, homeownership, employment, and income. The data I use link monthly loan records from a large loan servicer ("LLS") to administrative credit bureau information from TransUnion for over one million student borrowers. LLS manages disbursements and payments for over \$300 billion in federal student loans. The LLS data include detailed loan information (e.g., balances, payments, delinquency status, repayment plan) at a monthly frequency, as well as contact histories, zip code, age, institution attended, and college enrollment dates. These data are, to the best of my knowledge, the first panel of US federal student loan payments used in public research.

I use two complementary empirical strategies to identify the effect of IDR among borrowers who receive delinquency notice phone calls from LLS. First, I estimate the difference-in-differences between IDR enrollees (the "treatment group") and non-enrollees (the "control group") before and after receiving delinquency calls from LLS informing them about IDR.

²Economists have long bemoaned the dearth of available student loan data. In a 2015 meeting, William Dudley, President of the Federal Reserve Bank of New York, remarked, "Do borrowers who use programs like income-based repayment eventually succeed in paying off their debts? How do income-based repayment programs affect important decisions such as labor supply, consumption and household formation? These are important questions for the nation,...but it is very hard to answer these questions with existing data." (Dudley, 2015). In a column for the New York Times, Susan Dynarski echoes these sentiments, writing, "We are remarkably ignorant about student debt...But at the moment, the federal student loan data remains locked within the walls of the Education Department, with limited metrics trickling out," (Dynarski, 2015) and "Data on student loans are remarkably thin, given the size of this market. They are particularly inadequate for modeling and costing out income-based repayment plans," (Dynarski, 2014).

The identifying assumption for this design is that, in the absence of IDR, post-call outcomes for treatment and control groups would have exhibited parallel trends. To support this assumption, I show that trends in borrower outcomes are nearly identical across the two groups for several periods before the delinquency call. I also show that post-call responses between the two groups differ only for those calls that convert the treatment group to IDR. Responses to earlier “placebo” calls placed to *eventual* IDR borrowers closely track the response among non-IDR borrowers.

My second empirical strategy is an instrumental variables (IV) design exploiting the quasi-random assignment of delinquency calls to debt-servicing agents via an automatic dialing system. I use variability in agents’ tendencies to induce IDR take-up as a means of identifying the local average treatment effect (LATE) of IDR on individuals whose repayment plan decisions depend on the servicing agent to whom their delinquency calls are connected. I measure an agent’s IDR inducement rate (“agent score”) using the leave-one-out, post-call enrollment rate among borrowers who received a call from that particular agent. Agent score is predictive of IDR take-up, but uncorrelated with borrower characteristics and pre-call outcomes.³

Results suggest IDR increases loan repayment, credit scores, homeownership, and consumption among student borrowers. Relative to borrowers who remain on standard repayment plans, delinquency rates among IDR enrollees fall by 31.4 percentage points within seven months of take-up. While IDR reduces monthly minimum payments by an average of \$169 through a mechanical “first-stage” effect, the effect of reduced minimums on loan balances is dominated by more timely repayment; IDR borrowers pay down \$36 more student debt each month, on average, than those on standard repayment plans. IDR enrollees have credit scores which are 8.6 points higher, hold 0.1 more credit cards, and carry \$293 higher credit card balances than non-enrollees one year after the servicing call.⁴ While I find no

³My IV strategy closely resembles the “examiner design” used in a variety of other research settings, including incarceration Kling (2006); Aizer and Doyle (2015), pre-trial detention (Dobbie et al., 2018), consumer bankruptcy (Dobbie and Song, 2015), family welfare culture (Dahl et al., 2014), disability insurance (Maestas et al., 2013; Kostøl et al., 2017), and foster care (Autor and Houseman, 2010).

⁴For reference, Dobbie et al. (2016) finds that the removal of a flag designating Chapter 13 bankruptcy from one’s credit report is associated with a 9.8 point increase in credit scores, a \$143 increase in credit card balances, and a 2 percentage point increase in the likelihood of holding a mortgage one year after the flag removal.

effects on bankruptcies or auto loans, IDR enrollees are 2 percentage points more likely than non-enrollees to hold a mortgage after two years, an increase of 8 percent off of the pre-call mean. Estimated effects on financial outcomes outlast the “first-stage” effect of IDR on payment size: monthly minimum payments return to pre-call levels by fifteen months after the delinquency call. I also find no evidence of labor supply or income effects: unemployment deferments and median zip-code level income are not significantly different between IDR and standard borrowers within three years of the servicing call. These results suggest that IDR has long-term positive impacts on the financial health of borrowers through short-term increases in cash-on-hand.

In theory, IDR enrollment may influence borrower outcomes through three channels: (1) a liquidity effect through lower minimum payments, (2) a wealth effect through expected debt forgiveness, and (3) a moral hazard effect, as income-contingent forgiveness can distort labor supply or investment decisions.⁵ For two reasons, I argue my estimates capture a pure liquidity effect: First, minimum monthly payments decrease by an average of \$169 following IDR take-up, but return to standard levels within one year of enrollment. Borrowers facing this payment path should not expect loan forgiveness under the IDR plan I study, which only forgives outstanding balances after twenty-five years of successful payments. Second, because borrowers can opt-in to IDR at any time, any moral hazard from income-contingent loan forgiveness should apply to both treatment and control groups in my sample, as both groups are eligible for IDR and, by construction, aware of its existence. Estimates should therefore net out moral hazard effects, leaving a pure liquidity effect.

My analysis carries three important caveats. First, the effects I estimate hold interest rates and plan-specific repayment terms fixed. In a private market, one would expect the availability of IDR to have general-equilibrium effects on these parameters. For instance, adverse selection of individuals with low expected earnings into IDR could have an unraveling effect on a hypothetical “repayment-plan market” through adjustments to plan-specific rates. In the existing student loans environment, however, these types of effects are unlikely, as interest rates and repayment terms are held fixed by the federal government. My

⁵Note that increased liquidity may itself affect labor supply, even in the absence of forgiveness. I view such effects as conceptually distinct from moral hazard. See Chetty (2008) and Shimer and Werning (2008) for a discussion of liquidity versus moral hazard in the context of unemployment insurance.

analysis therefore treats loan terms as policy parameters rather than equilibrium objects. Nonetheless, even in the government-sponsored student loans market, equilibrium effects could potentially operate through political mechanisms like budgetary constraints or political pressure. IDR could also have general equilibrium effects on labor, marriage, or higher education markets. These potential market-level responses are excluded from the partial-equilibrium effects I estimate in this paper.

Second, because my sample includes only IDR-eligible borrowers while they are already in repayment, the treatment effects I identify exclude any ex-ante effects on decisions pertaining to college attendance, occupation choice, or principal borrowing amounts. In many ways, removing these effects is a theoretically attractive feature of my research design, as it allows me to isolate a single mechanism for observed effects—ex-post increases in liquidity. Nonetheless, these ex-ante responses to expected payment flexibility and loan forgiveness are an important component of IDR’s total impact on social welfare.⁶

Finally, the outcomes I measure carry several limitations. The credit data I use provide useful proxies for consumption and homeownership (Ganong and Noel, 2017; Di Maggio et al., 2016; Keys et al., 2014), but do not include direct measures of these variables. Proxies for labor market outcomes, measured by the receipt of unemployment deferments and median zip-code level income, are even more indirect.⁷ Likewise, while default and delinquency carry negative consequences for borrowers, the incidence and persistence of these consequences remains an open question. So, while improved repayment rates under IDR suggest increased borrower welfare, the magnitude of these welfare gains is somewhat ambiguous.

This paper complements a small but growing literature on student loan contracts and IDR. Beginning with Friedman (1962), many researchers have documented the theoretical benefits of publicly sponsored income-contingent student debt, citing both the failure of conventional credit markets and the need for more “equity-like” instruments for human capital investments (Chapman, 2006; Barr et al., 2017). A related stream of literature documents the revenue implications of various loan contracts by simulating repayment paths and loan forgiveness-incidence across different populations (Lucas and Moore, 2010; Johnston and

⁶See Abraham et al. (2018a) for an analysis of the ex-ante effects of IDR in a laboratory setting.

⁷Unemployment deferments allow borrowers working less than thirty hours per week to “pause” student loan payments for up to three years and forgive some or all of the accumulated interest.

Barr, 2013; Britton et al., 2018). Two studies, Abraham et al. (2018a) and Field (2009), investigate the *ex-ante* effect of student loan contract offerings on students' decisions.⁸ Finally, perhaps indicative of IDR's policy importance, several studies have aimed to evaluate measures which might improve take-up of IDR among eligible borrowers. Most notably, Abraham et al. (2018b) and Cox et al. (2018) document students' hypothetical plan choices under alternative framing, information, and income scenarios.

The contribution of this paper to the existing literature three-fold. First, I provide the first causal evidence of IDR's impact on loan repayment, consumption, homeownership, and employment. Second, I document several descriptive facts that are new and policy relevant. For example, this paper is the first to track federal student loan payments at a monthly frequency, providing insight into the precise pattern of delinquency and default. Third, this paper demonstrates the prevalence of liquidity constraints among student borrowers. While liquidity constraints and incomplete credit markets are well-studied phenomena, documenting them in the student borrowing context carries added importance given the contribution of human capital investments to economic growth. A failure of existing credit markets to finance higher education efficiently can carry significant macroeconomic consequences.

The remainder of this paper is organized as follows. Section 2 provides a brief overview of federal student loans and student loan servicing in the United States. Section 3 describes the student loan and credit bureau data and provides summary statistics. Section 4 motivates my empirical analysis with a model of consumption under alternative student loan repayment plans and discusses the potential welfare implications of IDR. Section 5 describes my empirical strategy. Section 6 presents results, and Section 7 provides interpretation. Section 8 concludes.

⁸Abraham et al. (2018a) use a lab experiment to measure the impact of hypothetical repayment options on stated career preferences. Their findings suggest IDR may encourage borrowers to take higher-paying careers with more variable earnings. Field (2009) documents evidence for framing effects (loan forgiveness versus tuition discounts) of career-contingent subsidies among NYU law students, suggesting student borrowers may suffer from debt-aversion.

2 Background

2.1 Overview of the Federal Student Loan System

Over 90% of student loans in the United States are federally subsidized and guaranteed.⁹ The government holds the liability on student loans, and interest rates are set by Congress.¹⁰ Student loans are not secured by collateral or subject to any credit check.¹¹ While the federal government holds the liability on student loans, the Department of Education contracts private debt servicing companies to handle disbursements, billing, and processing. Debt servicers disburse loans to colleges' financial aid offices, which apply disbursed funds directly to students' accounts. While the amount one can borrow from federal sources is capped by semester, virtually anyone attending an accredited institution is eligible to borrow at the same subsidized rate.¹²

The Department of Education sets repayment terms for student loans through repayment plans. Repayment plans specify the monthly minimum payments borrowers must make to their loan servicers, though borrowers can pay more than the minimum without penalty if they wish to pay down their debt early. The default repayment plan into which all borrowers are automatically enrolled is known as “standard repayment.” Under standard repayment, the minimum amount a borrower is required to pay follows a flat repayment schedule over ten years, so that minimum monthly payments are calculated as

$$\text{Monthly Standard Payment} = \frac{i * \text{Principal}}{1 - (1 + i)^{-(10*12)}}, \quad (1)$$

⁹A small private student loans market constitutes around ten percent of total student debt, mostly for creditworthy graduate students or borrowers who have exhausted their federal loan limits. In most cases, however, private lenders cannot compete with the subsidized rates offered by the government under the Federal Family Education Loan (FFEL) and Federal Direct Loan programs. Unless stated otherwise, I will use “student loans” to refer to loans originating from these federal programs.

¹⁰Congress has set rates on student loans since 1965, though it automated the process in 2013 with the Bipartisan Student Loan Certainty Act, which sets interest rates equal to the 10-year Treasury bond rate plus 205 basis points (360 bps for graduate students). Interest rates are fixed throughout the life of a loan and accrue as simple daily interest on principal only.

¹¹The exception is PLUS loans, for which parents can serve as cosigners, subject to a credit check.

¹²In the short term, borrowing costs can vary by financial need, as the “Subsidized Stafford Loan” program forgives interest accrued while the borrower is still in school, up to a means-tested limit.

where i denotes the monthly interest rate. Until 2010, the vast majority of borrowers in repayment were enrolled in standard repayment plans, with only a small fraction of borrowers choosing alternative financing options.

Income-driven repayment (IDR) plans first became available in 2009 as an alternative to standard repayment.¹³ Minimum payments under IDR are pegged to fifteen percent of borrowers' discretionary income, defined as the difference between adjusted gross income (AGI) and 150% of the federal poverty line (FPL).¹⁴ Specifically,

$$\text{Monthly IDR Payment} = 15\% * \left(\frac{\text{AGI} - 1.5 * \text{FPL}}{12} \right) \quad (2)$$

Monthly payments are capped at the standard minimum payment amount, and payments continue until the borrower's balance reaches zero. If a borrower successfully makes twenty-five years of payments under IDR, any remaining balance is forgiven, though any forgiven debt is treated as taxable income. Figure 1 provides a graphical comparison of IDR versus standard repayment plans under alternative income scenarios.

Borrowers can switch to IDR at any point in the repayment process. Opting-in requires completing an online form through the Department of Education, which verifies income and family size using information from a borrower's most recent federal tax return. Borrowers must recertify their income on a yearly basis, although they can adjust their payments more frequently by providing proof of income. If a borrower on IDR goes more than one year without recertifying income and family size, her payments automatically return to the standard payment amount, though her repayment plan is still classified as IDR.

Borrowers who fail to meet their monthly payments (i.e., "fall delinquent") under any repayment plan face penalties that increase in severity with the number of days past due.

¹³Since 2009, several versions of IDR plans have become available, including Income-Based Repayment (IBR), Pay-As-You-Earn (PAYE), and Revised-Pay-As-You-Earn (REPAYE). Eligibility criteria and repayment terms can vary across these plans, though they share the same general structure. For the purposes of this study, I focus on Income-Based Repayment (IBR), as it is the only IDR plan for which the borrowers in my sample are eligible. Any description or results pertaining to IDR are specific to the IBR plan, though the discussion and conclusion generalize to the broader concept of IDR. For more information on federal student loan repayment plans, see www.studentaid.gov.

¹⁴Payments for a married borrower who files jointly are prorated to her share of combined household student debt.

Between one and ten days past due, student loan servicers will contact delinquent borrowers through email or post. Between ten and ninety days past due, borrowers are charged late fees and contacted by phone at increasing frequency to encourage repayment and discuss repayment options. At 91, 181, and 271 days past due, borrowers are reported to credit bureaus, damaging their credit scores. Loans more the 270 days past due are considered eligible for default. Once in default, all remaining balance on student debt becomes due, and the Department of Education can garnish up to 15 percent of borrowers' wages or withhold their tax returns to collect on defaulted debt. In twenty states, the federal government can block the renewal of professional licenses for defaulted borrowers working in health care, education, and/or other licensed fields. Unlike other forms of consumer debt, student loans cannot be discharged by declaring Chapter 7 or Chapter 13 bankruptcy.¹⁵ Defaulted borrowers are ineligible for any future federal student aid (Department of Education, 2017a).

2.2 Study Setting: Student Loan Servicers

As one of ten federal student loan servicers, LLS manages disbursements and payments for over \$300 billion in student loans. Debt servicing is provided on behalf of the federal government, which hires the servicer through a series of contracts.¹⁶ As a part of its servicing operations, LLS makes frequent contact with delinquent borrowers to encourage repayment. When borrowers become fifteen or more days past due on their payments, their phone numbers are placed in a dialing queue. An automatic dialer then places calls to each of the numbers in this queue in rapid succession. If a call is unanswered, the dialer places it back at the bottom of the queue. Each answered call is immediately connected to a debt-servicing agent randomly selected from the pool of available agents. If no agents are available, the dialer places the borrower on hold until one becomes available. Such instances are extremely rare, however, as the dialer places calls at a rate to match agent availability, which is highly predictable over large numbers of agents.

¹⁵In rare circumstances, borrowers who demonstrate “undue financial hardship” can discharge their student debt in bankruptcy. Student loans can also be discharged if borrowers are disabled, deceased, or attended an institution which has since closed.

¹⁶Federal loan servicing contracts can be viewed online at www.studentaid.gov.

LLS employs over 300 servicing agents across four call centers. Agents are tasked with informing borrowers of their delinquent status, inquiring about their ability to repay, and informing them of repayment options. During a call session, the questions and responses of the agent are guided by a decision tree. The agent first asks if a borrower can make payments under their current plan. If not, the agent asks if the borrower is unemployed or a full-time student, as such borrowers can typically qualify for interest-free deferments. Finally, the agent “models-out” IDR payments for the borrower, eliciting information on annual income, marital status, and family size. Borrower responses are entered into the agent’s computer, which provides an estimate of monthly IDR payments according to Equation 1. If these “modeled-out” payments are lower than what the borrower is paying under the standard plan, the agent provides instructions for online IDR enrollment with the Department of Education. Agents are incentivized to bring delinquent accounts current, but face penalties if they fail to present borrowers with their best available options. Supervisors periodically monitor agents’ calls to ensure they meet federal compliance standards. If an agent does not offer IDR to a borrower deemed suitable for the option during a monitored call, the agent’s pay is reduced that month.

3 Data

The data I use to estimate the effect of IDR enrollment link administrative student loan repayment and contact data to credit bureau records for over one million borrowers. Data are drawn from LLS’s FFEL loan portfolio, which includes over \$90 billion in loans. The LLS loan data contain detailed repayment records for each borrower, including principal borrowing amounts, loan balances, minimum payments due, and dates of delinquency at a monthly frequency. They also include indicators for type of loan (e.g., Subsidized Stafford, PLUS), current repayment plan (e.g., Standard, IBR), and current loan status (e.g., forbearance, grace period, default). In addition to loan information, the LLS data contain borrower characteristics, including year of birth, 9-digit zip code, OPE ID for attended institutions, college attendance dates, and graduation status. Gender is inferred using first names.¹⁷

¹⁷The online appendix to Tang et al. (2011) provides a public-use list of common first names paired with the male-female proportions of New York City Facebook profiles with each name. LLS merged this list to

I merge demographic and loan information with LLS contact histories from 2011 onward. Contact history data provide a single observation for each point of contact, and include all incoming and outgoing calls in which the line was connected to a borrower in the sample. For each call in the data, I observe the date, time of day, incoming/outgoing status, and servicing agent identifier associated with the call. Agent identifiers are linked to a small set of agent characteristics, including work site location and work group (“claims aversion”, “skip tracing”, etc.).

Finally, borrowers in the LLS data are linked to yearly TransUnion credit bureau records from 2010 through 2017. The TransUnion data provide yearly balances, credit limits, delinquencies, and number of accounts for several categories of consumer debt, including mortgages, credit cards, and auto loans. They also include broader measures of financial health, like credit scores and bankruptcies.¹⁸ TransUnion data are merged to borrowers in the LLS data by SSN. 95 percent of borrowers are successfully matched to TransUnion records.

The analysis sample used in this study consists of 35,268 IDR-eligible borrowers drawn from LLS’s FFEL portfolio. To construct this sample, I begin with the universe of LLS’s FFEL borrowers with positive balances as of December 2011, excluding those who hold any private or Direct loans.¹⁹ From this population of 5.8 million borrowers, I select those borrowers who answered a delinquency call between 2013 and 2017, limiting the sample to 399,136 borrowers. I then construct a borrower-by-call panel, applying the following selection criteria to the 714,704 calls placed to borrowers in the sample: First, I remove borrowers who cannot be matched by zip code or first name to inferred measures of gender or income, leaving 629,359 calls. To remove calls which may be non-randomly assigned, I remove borrowers marked as non-native English speakers, as well as those more than 140 days delinquent at the time of the call, which leaves 528,771 calls.²⁰ Then, I remove

first names in their borrower records at my request.

¹⁸Additional details concerning TransUnion data can be found in Dobbie et al. (2017), Avery et al. (2003), and Finkelstein et al. (2012).

¹⁹While borrowers can hold loans from a mixture of FFEL, Direct, and private sources, the database I use only includes repayment information for FFEL borrowers. The analysis sample excludes borrowers with mixture of loans, so I can observe their complete repayment profile. Roughly fifteen percent of LLS’s 2012 FFEL borrowers also hold Direct loans, and fewer than ten percent hold private student loans.

²⁰LLS handles “Late Stage” delinquencies of 150 days or more through another division with different agent-assignment procedures.

borrowers who were already enrolled in IDR prior to their delinquency call, as they would not be “eligible” for call-induced IDR take-up. From the remaining group of 513,454 calls, I keep only those which reached the stage at which borrowers were provided information concerning their potential IDR payments (i.e., “modeled-out”), leaving 87,441 calls. I then exclude from the sample all calls made by agents with fewer than 100 total calls.²¹ From the resulting sample of 80,617 calls, I construct two balanced panels centered around call dates. For monthly repayment outcomes, I create a balanced panel of 30,359 calls with 30 leads and 10 lags. For yearly credit outcomes, I create a balanced panel of 16,504 calls with 3 leads and 3 lags.²²

Table 1 provides summary statistics for samples of interest. The “full sample” (column 1) is a random sample of 699,100 drawn from the population of LLS FFEL borrowers as of December 2011. The “analysis sample” (column 2) is the entire subpopulation of borrowers selected according to the criteria described in the previous paragraph. In the full sample, IDR has low take-up, with only 12 percent of borrowers enrolled in a plan. That share rises to 0 percent in the analysis sample, as it is constructed to include only borrowers who might benefit from the plan. Unsurprisingly, these borrowers are negatively selected: they have lower credit card limits, higher rates of bankruptcy and live in lower-income zip codes. While negative selection into the analysis sample is partly because borrowers must be delinquent in order to receive a phone call, it also reflects negative selection into IDR; borrowers with high income or low debt balances would not benefit from the plan, and hence would not reach the “modeled out” stage of the delinquency call. Modeled-out borrowers are therefore more representative of the “IDR-eligible” population than the larger sample of borrowers. Columns 3 and 4 of Table 1 break the analysis sample into control and treatment groups, where treatment is defined as IDR enrollment within four months of answering an LLS delinquency call.²³ Baseline variables for treated borrowers are largely comparable to

²¹Removing these calls reduces measurement error in the agent score instrument because estimates of the mean taken over a small number of calls are highly imprecise.

²²Note that the calendar date correspondence between monthly and yearly variables can vary dramatically depending on the relative timing of delinquency calls and data collection. If a borrower receives a call during the month in which yearly TransUnion data are collected, $t_{year} = 0$ corresponds to $t_{month} \in [-11, 0]$. If she receives the call one month later, $t_{year} = 0$ corresponds to $t_{month} \in [1, 12]$.

²³Note that treatment is defined at the call-level, not the borrower-level. For the borrower-level statistics reported in Table 1, the treatment group consists of all borrowers with *any* treated calls.

those for the control group. Note that 0 percent of the control group does eventually enroll in IDR, though never within four months of a delinquency call.

To assess the comparability of my study sample to the broader student borrowing population, I make use of a separate, nationally representative dataset from the 2008/2012 Baccalaureate and Beyond Longitudinal Study (B&B). Provided by the National Center for Education Statistics (NCES), the B&B data include restricted-use administrative loan and financial aid records linked to survey responses for a representative sample of four-year US college graduates in the spring of 2008, followed-up in 2011-2012. Table A1 provides summary statistics for the subsample of borrowers in the B&B data who took out federal student loans alongside the corresponding statistics for the full sample in the LLS data, restricted to include only 2008 graduates. Mean values for variables common to the two data sources are very similar between the two samples, suggesting the LLS data are largely representative of the general student borrowing population.

While the B&B comparison suggests LLS borrowers are representative of the larger federal borrowing population, there are two important caveats concerning external validity. First, individuals in my analysis sample are selected along two dimensions: (1) they have at least one instance of late payments, and (2) their loans originated prior to 2010. While the first restriction limits the sample to likely beneficiaries of IDR, the second restriction removes many borrowers for whom we would expect IDR to be most effective, as younger borrowers typically have higher debt-to-income ratios. Second, my sample is restricted to individuals eligible for the IDR plan known as Income-Based Repayment (IBR).²⁴ While this plan shares most features with other existing IDR plans like Pay-As-You-Earn (PAYE), results may not extend to hypothetical repayment schemes of policy relevance. In particular, my estimates will incorporate the effects of unattractive institutional features like staggered payment adjustments and complicated sign-up procedures that would likely be absent from an “ideal” system of income-contingent loan repayment.²⁵

²⁴Approximately twenty percent of the five million borrowers in IDR plans are enrolled in IBR (Department of Education, 2017b). For a detailed description of alternative IDR plans, see www.studentaid.gov.

²⁵For a discussion of potential improvements to the current structure of IDR in the US, see Dynarski and Kreisman (2013).

4 Model

In this section I develop a dynamic model of student loan repayment to motivate my empirical analysis. The model illustrates the various channels through which IDR may affect borrowers' consumption and repayment decisions.

4.1 Setup

Infinitely-lived borrowers are endowed with student debt D_0 and non-student debt assets $A_0 = 0$. Each period, borrowers draw income $y_t \sim N(\mu_t, \sigma)$ and make two decisions. First, they decide whether or not to pay down their student debt D_t by making their minimum monthly payment x_t . Second, they decide how much to save ($A_{t+1} > 0$) or borrow ($A_{t+1} < 0$) in non-student loan assets, but face a liquidity constraint L , imposing an upper bound on the amount they can borrow. If a borrower decides to repay ($q = 1$), she pays down x_t of her debt balance. If instead she decides to fall delinquent ($q = 0$), her debt balance does not change, and she pays a constant utility cost φ . To simplify exposition, assume no discounting or interest ($\beta = 1 + r = 1$).

The borrower's optimal repayment and consumption decision must satisfy the following:

$$V(y_t, D_t, A_t) = \max_{A_{t+1}, c_t, q \in \{0,1\}} \{u(c_t) - \varphi(1 - q) + E[V(y_{t+1}, D_{t+1}, A_{t+1})]\} \quad (3)$$

$$A_{t+1} \geq L \quad (4)$$

$$c_t = A_t - A_{t+1} + y_t - qx_t \quad (5)$$

$$D_{t+1} = D_t - qx_t. \quad (6)$$

To solve her problem, the borrower first decides whether to make her student loan payment by comparing the objective function on the right-hand side of Equation 3 under alternative repayment scenarios. In other words, she must compare the value of delaying debt payment x by one period against the utility cost φ . This decision can be characterized by a cutoff function for y_t ,

$$q_t = \mathbf{1} \{y_t > \psi(A_t, D_t)\}, \quad (7)$$

which is increasing in assets, A_t , and decreasing in student debt, D_t .²⁶ Once she has decided whether or not to make their payments, the borrower will save or borrow to the point where her marginal utility of current-period consumption is equal to the expected marginal value of assets next period, unless prevented by liquidity constraints. This consumption-savings decision is given by an Euler inequality,

$$u'(c_t) \geq E \left[\frac{\partial V}{\partial A}(y_{t+1}, D_{t+1}, A_{t+1}) \right], \quad (8)$$

which binds in period t if and only if $A_{t+1} > L$. The full solution to the model is given in Appendix B.

4.2 Model Predictions under IDR Payments

To evaluate the effects of IDR, I endogenize student loan payments in the model above to be a function of income and debt, $x_t = x(y_t, D_t)$, and consider how a change in repayment schemes affects optimal consumption and repayment decisions.²⁷ Let Standard (S) and IDR (I) repayment schemes be given by

$$x^S(y, D) = \min \left\{ \frac{D_0}{N}, D \right\} \quad (9)$$

$$x^I(y, D) = \min \left\{ \theta y, \frac{D_0}{N}, D \right\}. \quad (10)$$

For borrowers in repayment, the IDR plan is modeled as the minimum of some share θ of per-period income and the standard payment D_0/N . Borrowers under both plans never pay more than their remaining student debt balance D . Let m_t denote income net of loan payments, $m_t \equiv y_t - x_t$. The reduction in minimum payments offered by IDR can influence borrowers' consumption and repayment decisions through three channels.

²⁶Note that without liquidity constraints, the repayment decision is trivial, as non-repayment imposes a utility cost without improving in lifetime income. In this sense, non-repayment can be seen as a costly means of increasing liquidity: borrowers neglect to repay if and only if the benefit of x -higher cash-on-hand exceeds utility cost φ .

²⁷Here I restrict attention to IDR's liquidity effects through lower monthly minimum payments, deferring discussion of loan forgiveness to Section 4.2.

First, IDR weakly increases net income in the short term relative to the long term:

$$E [m_t^I] \geq E [m_t^S] \quad \forall t \leq N \quad (11)$$

$$E [m_t^I] \leq E [m_t^S] \quad \forall t > N. \quad (12)$$

The result is an intertemporal *smoothing effect*: pushing payments farther into the future flattens borrowers' expected net income profiles. As a result, financing consumption or student debt repayment in periods $t \leq N$ requires less borrowing, decreasing the likelihood of binding liquidity constraints. With less fear of binding constraints, borrowers place less value on precautionary savings: the marginal utility of assets in Equation 8 decreases and consumption increases.²⁸ Intuitively, individuals face less pressure to borrow because their reduced current-period payments are effectively “borrowed” through higher payments in future periods.

Second, in addition to smoothing income over time, IDR also smooths each period's income over states of the world, providing income insurance for student borrowers.²⁹ To see how, note that IDR reduces period- t variance in net income,

$$\text{Var}(m_t^I) \leq \text{Var}(m_t^S) \quad \forall t, \quad (13)$$

reducing income uncertainty increases consumption through an *insurance effect*. This insight comes from the literature on consumption dynamics: individuals facing higher income uncertainty are more likely to face binding liquidity constraints in all periods. A mean-preserving spread in the future income distribution will therefore increase the marginal value of assets and decrease current-period consumption out of choice (precautionary savings) or necessity (the liquidity constraint binds).³⁰

Finally, note that while the smoothing and insurance effects on repayment are unambiguously positive, their effects on consumption are only positive conditional on repayment. A

²⁸Proof in Appendix B.

²⁹I use the term “insurance” loosely, referring to state-contingent intertemporal transfers within individuals. Note that, ignoring loan forgiveness, income-contingent payments provide no insurance against *lifetime*-earnings risk.

³⁰See Ljungqvist and Sargent (2012). Proof in Appendix B.

borrower's consumption decision is made jointly with her repayment decision, and the binary nature of delinquency may produce negative crowd-out effects on consumption. In particular, if a borrower is failing to meet payments under the standard plan she may, in fact, *decrease* consumption under IDR if she uses all her increased cash-on-hand to avoid delinquency. In addition to the smoothing and insurance effects described above, observed consumption responses liquidity through reduced student loan payments will reflect this *repayment-crowd-out effect* among would-be defaulters shifting consumption spending towards student loan repayment.

The smoothing, insurance, and repayment-crowd-out effects of liquidity through lower minimum payments are summarized in the following expressions:

$$\Delta q_t \approx \underbrace{\sum_{k=0}^{\infty} -\frac{dq_t}{d\mu_k} \Delta E[x_k]}_{\text{smoothing } (-)} + \underbrace{\sum_{k=0}^{\infty} \frac{dq_t}{d\sigma_k^2} \Delta \sigma_k^2}_{\text{insurance } (-)} \quad (14)$$

$$\Delta c_t \approx \underbrace{\sum_{k=0}^{\infty} -\frac{dc_t^q}{d\mu_k} \Delta E[x_k]}_{\text{smoothing } (+)} + \underbrace{\sum_{k=0}^{\infty} \frac{dc_t^q}{d\sigma_k^2} \Delta \sigma_k^2}_{\text{insurance } (+)} + \underbrace{(c_t^1 - c_t^0) \Delta q_t}_{\text{repayment crowd-out } (-)} \quad , \quad (15)$$

where c_t^q denotes consumption holding the repayment decision constant at $q_t = q^S$. Note that assuming quadratic utility ($u''' = 0$) and no debt forgiveness ($\sum_{k=0}^{\infty} \Delta E[x_k] = 0$), removing liquidity constraints would reduce both terms above to zero, as individuals would simply borrow and save as needed to make their payments and achieve constant consumption over all periods. In this sense, IDR serves more like a market-correcting measure than a transfer program: Under standard repayment, borrowers who expect future earnings increases would ideally finance current period consumption through borrowing. Market failures may prevent such borrowing, however, as recent college graduates often lack credit or collateral necessary to overcome information asymmetries in lending. IDR alleviates this problem by pushing debt obligations further into the future during periods of low income. If borrowers face liquidity constraints, IDR should therefore decrease delinquency rates, increase consumption, or both. Conversely, if borrowers have alternative credit options, we would expect student loan repayment and consumption to remain unchanged under IDR. For these unconstrained

borrowers, IDR may instead have a deleveraging or savings effect, causing them to spend down higher interest credit cards or increase their savings.

Debt Forgiveness and Social Welfare

Liquidity effects may increase borrower welfare, but have no direct impact on lifetime income. Absent interest rate subsidies, reduced payments simply transfer wealth *within borrower* over time, and therefore carry no redistributive inefficiencies or moral hazard costs. By contrast, the forgiveness provisions of IDR effectively provide “wealth insurance” against lifetime-earnings risk and therefore hold very different implications for social welfare.³¹ First, because IDR borrowers pay no “premium” for wealth insurance, the government bears the expected cost of forgiven debts plus some risk premium.³² Such a transfer would impose a deadweight loss to social welfare through taxation while carrying ambiguous redistributive consequences: forgiveness is progressive conditional on student debt, but college graduates, especially those with high debt balances, have higher lifetime income than the general population. Second, insuring lifetime income for student borrowers can distort labor supply, occupation choice, or college attendance decisions through moral hazard, thereby reducing social welfare. For example, borrowers on IDR may take jobs with higher risk of low pay or unemployment as a result of being partially insured against income losses. By contrast, providing liquidity through reduced IDR payments is revenue neutral and non-distortionary. Income-contingent payments may affect labor supply, but such responses represent a corrective response to incomplete credit markets rather than moral hazard.³³

A full accounting of IDR’s social welfare effects would distinguish between liquidity and wealth effects and measure labor supply, college attendance, and principal borrowing responses through each channel. Such an analysis is beyond the scope of this paper. Nevertheless, the repayment path of IDR borrowers in my sample should provide some insight

³¹While my empirical results suggest few borrowers will ultimately qualify for debt forgiveness under IDR, I incorporate debt forgiveness into the model in Appendix B for completeness.

³²While insuring *idiosyncratic* risk between borrowers would be diversified away in large numbers, *systemic* risk could be costly to the government if, for instance, uncertainty over the business cycle makes it difficult to predict the amount of forgiven debt.

³³See Chetty (2008) for a more complete description of the distinction between moral hazard and liquidity in the context of unemployment insurance.

into the relative effects of liquidity versus expected debt forgiveness. The incidence of debt forgiveness depends on the likelihood of positive balances after twenty-five years of IDR payments. While I cannot directly test for moral hazard or measure realized loan forgiveness, the evolution of payment-to-debt ratios in the two years following IDR enrollment can inform forgiveness likelihood. If IDR payments return to the standard amount within one or two years, it suggests IDR holds little insurance benefit and moral hazard is unlikely to play a large role in borrower behavior.

5 Empirical Strategy

Consider the following empirical model of borrower i 's outcomes, t periods after receiving delinquency call c :

$$Y_{ict} = \beta_0 + \beta_1 IDR_{ic} + \beta_2 \mathbf{X}_{it} + \epsilon_{ict}, \quad (16)$$

where Y_{ict} denotes the outcome of interest, \mathbf{X}_{ict} is a vector of borrower control variables, IDR_{ic} is an indicator for IDR enrollment within four months of the call, and ϵ_{ict} is an error term.³⁴ Estimating β_1 in Equation 16 using OLS would likely yield biased estimates because preferences over repayment plan choices are correlated with unobserved borrower attributes. To overcome these biases and identify the causal effect of IDR on borrower outcomes, I employ two complementary empirical strategies. First, I estimate the difference-in-differences between IDR enrollees and non-enrollees before and after receiving delinquency calls. Second, I use an instrumental variables (IV) design which exploits the quasi-random assignment of servicing agents to these calls.

³⁴I fix my treatment variable IDR_{ic} to a specific month in order to capture the dynamic effects of IDR. It is important to note, however, that a borrower's repayment plan as of month four need not reflect her repayment plan in later months. As some IDR "treatment" borrowers leave the program, and gradually more "control" borrowers join it, my results will be attenuated relative to an "always IDR" versus "never IDR" comparison. Indeed, attrition from IDR after the one-year recertification period will play an important role in interpreting my results. See Figure 2.

5.1 Difference-in-Differences

My difference-in-differences design compares pre-/post-call differences in outcomes between borrowers who take up IDR (the “treatment” group) and borrowers who remain in standard repayment plans (the “control” group) to identify the causal impact of IDR on repayment and credit outcomes. Formally, the difference-in-differences specification takes the following form:

$$Y_{ict} = \gamma_i + \gamma_t + \left[\sum_{\tau \neq -1} \delta_\tau \cdot IDR_{ic} \cdot \mathbf{1}\{t = \tau\} \right] + \beta_1 IDR_{ic} + \beta_2 \mathbf{X}_{ict} + \epsilon_{it}, \quad (17)$$

where Y_{ict} denotes the outcome of interest, γ_i are individual fixed effects, γ_t are event-time fixed effects (months or years relative to call date), IDR_{ic} is an indicator for IDR enrollment within four months of the call, \mathbf{X}_{ict} is a vector of borrower control variables (including call date and time fixed effects), ϵ_{ict} is an error term, and δ_τ , the parameters of interest, are coefficients on IDR enrollment status which vary by event time. The specification omits γ_t and δ_τ terms at $t = -1$, so estimates can be interpreted relative to the baseline period of one month or year prior to the delinquency call.

Identification in the difference-in-differences specification comes from variation in the propensity to take up IDR following a delinquency call. Such variation includes differences in enrollment rates both between and within servicing agents. While between-agent variation is arguably random (see Section 5.2), within-agent variation is not. Internal validity of this specification therefore rests upon the assumption that, holding borrower-specific differences fixed, such within-agent variation in IDR take-up is exogenous with respect to outcomes. Stated differently, unbiased estimation in my difference-in-differences design assumes that post-call trends in outcomes would be the same in both groups of borrowers had neither taken up IDR.

Figures 2 through 7 provide graphical evidence in support of the common-trends assumption. The figures plot raw means for monthly repayment and annual credit outcomes, respectively, for IDR enrollees and non-enrollees, relative to call date and normalized by pre-call mean. Trends in pre-call outcomes appear similar between IDR and standard enrollees for several periods, diverging only after receiving the delinquency call. I also estimate IDR effects in an alternative differences-in-differences specification that controls for group-specific

linear trends in months or years prior to call.³⁵

Even if IDR and standard borrowers exhibit observably similar pre-trends, several potential violations to the exogeneity assumption exist. For example, estimates could be biased if treatment and control groups would have responded to delinquency calls differently in the absence of IDR. To address this concern, I develop a placebo test meant to simulate treatment and control responses absent IDR enrollment. Recall that a given borrower may receive several calls before eventually enrolling in IDR. In my main analysis, only the last of these calls is included in the treatment group. Previous calls made to the same borrower are designated as control calls, as they do not result in IDR take-up within four months. Such calls create a natural placebo group to test differential call response between borrowers; if, in the absence of IDR, treatment and control borrowers would have responded differently to their n th delinquency call, they would likely have had different responses to calls one through $n - 1$ as well.

Figure A1 plots raw pre- and post-call mean monthly outcomes for non-IDR control borrowers versus *eventual* IDR borrowers following earlier “placebo” calls that did not induce IDR take-up within the following twelve months. Difference-in-differences estimates for this placebo sample are reported in Tables A2, A3, and A4. Relative to main results, placebo estimates are either economically or statistically insignificant, suggesting the response to a delinquency call alone does not differ between IDR and non-IDR borrowers. The differential response estimated in my main difference-in-differences specification therefore reflects a pure IDR effect, not a “call effect.”

Estimates could also be biased if IDR enrollees experienced a shock at the time of a delinquency call that induced them into IDR take-up and influenced outcome variables. I argue that such instances are unlikely. Delinquency calls are *outgoing*, so their incidence is determined by LLS, not the borrower. While the timing of these calls are mechanically non-random *within-borrower*, they do not vary systematically *between* borrowers with ob-

³⁵Estimates from the specification including linear pre-trends can be interpreted as IDR’s impact on outcomes relative trend-predicted differences between groups. Formally, the model is given by

$$Y_{ict} = \gamma_i + \delta t \cdot IDR_{ic} \cdot \mathbf{1}\{t < 0\} + \left[\sum_{\tau \geq 0} \delta_{\tau} \cdot IDR_{ic} \cdot \mathbf{1}\{t = \tau\} \right] + \beta_1 IDR_{ic} + \beta_2 \mathbf{X}_{ict} + \epsilon_{it} \quad (18)$$

servably similar characteristics. If IDR borrowers were enrolling as a response to sudden shocks, outcomes should therefore vary from non-IDR borrowers in the months immediately preceding the call. It is possible that borrowers make IDR enrollment decisions based on expected *future* shocks to their financial well-being, but any bias created by forward-looking borrowers should be negative: the benefit of IDR is strictly decreasing in income and available credit, so borrowers who select into IDR based on expected shocks to their financial well-being should, all else equal, exhibit *lower* repayment, consumption, and credit scores relative to standard borrowers, attenuating any positive treatment effects of IDR.

5.2 Instrumental Variables

I complement the difference-in-differences design described above with an instrumental variables (IV) approach. The instrument used in my IV design is a measure of a quasi-randomly-assigned servicing agent’s tendency to induce IDR take-up among the borrowers she calls. The parameter I estimate in this specification is a local average treatment effect (LATE) identifying the impact of IDR enrollment among borrowers whose plan choice depends on agent-assignment.

I construct the instrument using a leave-one-out measure of agents’ ability to induce IDR enrollment, where enrollment is residualized to account for the timing and ordering of delinquency calls. This “agent score” is analogous to the measures of judge leniency used by Dahl et al. (2014) and Dobbie et al. (2018) to instrument for trial outcomes in court case settings. Residualizing enrollment removes two potential sources of endogeneity in the instrument. First, the start and end dates of an agent’s employment with LLS might coincide with time trends in IDR take-up unrelated to the agent’s ability. For instance, an agent employed before IDR became a popular repayment option might incorrectly be measured as “low-score.” Second, the shift choices of agents are non-random and may correlate with borrowers’ IDR enrollment if, for example, borrowers are more likely to hang up on calls received in the evening or on weekends.

I account for these potential selection effects by removing year, month, weekend, and hour-of-day fixed effects from observed IDR enrollments before calculating the leave-out

agent score. Specifically,

$$IDR_{ict}^* = IDR_{ict} - \gamma \mathbf{W}_{ict} \quad (19)$$

$$= Z_{icj} + \epsilon_{ict} \quad (20)$$

where \mathbf{W}_{ict} is a vector of call date and time controls and Z_{icj} is agent score. I calculate the residualized rate of IDR take-up, IDR_{ict}^* , using OLS estimates of γ in Equation 19. I then construct agent score Z_{icj} using the leave-one-out mean of this residualized rate:

$$Z_{icj} = \left(\frac{1}{n_j - 1} \right) \left(\sum_{k=0}^{n_j} IDR_{kcj}^* - IDR_{icj}^* \right) \quad (21)$$

where n_j denotes the number of calls made by agent j .³⁶

The residualized agent score distribution can be seen in Figure 8. Note that while both estimation strategies use balanced panels, the agent score instrument is calculated using the larger unbalanced panel of calls satisfying the all other sample selection criteria in Section 3. This sample includes calls from 286 different agents in four different call centers. Agents place 119 calls on average to borrowers in the sample, with a median of 92 calls.

Variation in agent score can be driven by several potential sources. Agents can vary in IDR-conversion through subtle variations in demeanor or tone, and borrowers often hang up or stop listening depending on the interaction. Conversations with agent supervisors suggest that factors like speech patterns and accents play a large role in keeping borrowers on the phone long enough to advise them about IDR. Agent score may also be influenced by agents' ability to provide clear details regarding plan payments and sign-up instructions. Borrowers must log into the Department of Education website using their social security number, authorize the IRS to transfer their tax return, correctly identify their loan program, and consent to change their payment plan. If an agent fails to properly explain these steps, a borrower may fail to enroll in IDR even if the agent convinces her to do so.

³⁶Removing borrower i 's outcomes from the calculation avoids reflection bias, the mechanical correlation between individual i 's take-up and average take-up calculated for any group containing i .

5.2.1 Identifying Assumptions

In order for two-stage least squares estimates to identify a local average treatment effect (LATE) for the causal impact of IDR take-up, the instrument must satisfy three conditions: (1) IDR take-up varies with agent assignment, (2) agent assignment predicts borrower outcomes only through its effect IDR take-up, and (3) agents’ tendency to induce IDR take-up is monotonic across borrowers.

To test the first identifying assumption, I estimate the first-stage relationship between the agent score instrument and observed IDR enrollment. Figure 8, which plots a local linear regression of IDR take-up against the agent score instrument, provides graphical evidence for this relationship. IDR take-up is monotonically increasing in agent score throughout the distribution of borrowers. The first-stage relationship seen in Figure 8 can be represented through following linear probability model:

$$IDR_{ict} = \alpha_0 + \alpha_1 Z_{ict} + \alpha_2 \mathbf{X}_{ict} + \epsilon_{ict} \quad (22)$$

In Table 3, I provide first-stage estimates from an OLS regression of Equation 22. Results imply a positive and highly significant relationship between agent score and IDR take-up. Point estimates range from 0.61 to 0.55, depending on the inclusion of borrower controls, and the F-statistic on a test of instrument significance equals 79.00.

The second identifying assumption requires that agent assignment be predictive of borrower outcomes only through its impact on repayment plan choice. One way this assumption could be violated would be if different types of borrowers were systematically assigned particular agents. Such violations are effectively ruled out by the automatic dialing mechanism—calls are mechanically assigned at random to the available agents working during a particular shift.³⁷

A more plausible threat to validity concerns the selection of calls into the study sample. In Section 3, I define the analysis sample to include only those calls during which agents have

³⁷Note that random assignment does not imply equal probability of assignment—an agent who makes shorter and more frequent phone calls will have a higher rate of availability during her shift. Any given delinquency call will therefore have a higher probability of being assigned to these “quicker” agents. The average call to which such agents are assigned, however, will nonetheless be no different from those calls assigned to relatively “slower” agents who make fewer calls per hour.

“modeled-out” the borrower’s hypothetical IDR payments. This selection criterion ensures all borrowers are aware of IDR and helps maintain comparability of treatment and control groups in the difference-in-differences design. However, recall from Section 2 that an agent informs a borrower of IDR and “models-out” her payments only after she has determined the borrower cannot adequately repay under her current plan. If agents exercise discretion in this determination, or if variation in their phone demeanor induces differential rates of borrower hangup before reaching this stage of the call, the sample would be selected based on agent-specific criteria which could potentially correlate with their IDR-inducement score and bias my estimates. To remedy this concern, I employ a strategy which combines the instrument construction described above with the selection correction techniques pioneered by Heckman (1979). Specifically, I construct a measure of agent-induced sample selection propensity by calculating the leave-one-out mean “modeled-out” rate, Z_{ict}^M , among all calls assigned to the agent on a given call. I perform this calculation on the larger, unconditional sample of 513,454 calls and follow the same procedure as Equations 19 through 21, replacing the treatment variable IDR_{ict} with $Modeled_{ict}$, an indicator for whether borrower i was “modeled-out” during phone call c . I then include the sample selection measure Z_{ict}^M in my instrumental-variables regressions to ensure that assignment of Z_{ict} is *conditionally* random.

Table 2 provides empirical evidence that, after correcting for agent modeling propensity and call date and time, borrowers do not vary systematically by agent score. Column 1 reports results from an OLS regression of realized IDR enrollment against several borrower characteristics and pre-call outcome variables, as well as call date and time fixed effects and modeling propensity Z_{ict}^M . Not surprisingly, estimates demonstrate non-random selection into IDR; holding date and time of call fixed, IDR enrollees are significantly more likely to be young, female, low-income, and hold lower balances across several types of debt. Column 2 reports results from an OLS regression of the agent score instrument on the same right-hand side variables. Estimated coefficients on borrower variables in this specification are statistically indistinguishable from zero, and the F-statistic on a test for whether all borrower variables can jointly predict agent score is 1.14.

Even if agents are randomly assigned to borrowers, the exclusion restriction may still be violated if agents can influence borrower outcomes through channels other than repayment plan choice. If, for example, agents who induce high IDR take-up also convince borrowers to

make timely payments, two-stage least squares estimates of IDR’s effects on repayment would be biased upwards. While it is impossible to rule out agent effects through non-IDR channels, loan servicing practices suggest that such threats to validity are unlikely. LLS’s delinquency calls are designed solely to provide borrowers with information on their repayment options. Agents provide no advice or counseling to borrowers, and follow a decision tree to present repayment alternatives.

The third identifying assumption requires monotonic agent effects across borrowers. To satisfy the monotonicity assumption, there can be no borrower for whom a higher (lower) score agent decreases (increases) the likelihood of IDR take-up. This assumption would be violated if certain agents “match” well with certain borrowers. For example, if some borrowers respond more favorably to female agents, their take-up may be higher under low-score female agents compared to high-score male agents. The presence of such “monotonicity violators” would generate a bias in my LATE estimation, the magnitude of which would increase with the number of violators and the difference in the marginal treatment effects between violators and non-violators (Angrist et al., 1996; Heckman and Vytlačil, 2005).

I implement two partial tests of the monotonicity assumption. First, I estimate the first stage relationship between my agent score instrument and IDR take-up within subgroups of my monthly analysis sample. If the monotonicity assumption is satisfied, these estimates should be non-negative for all subsamples. As Table A5 shows, estimated coefficients are positive across a variety of subgroups. Second, I calculate a variety of *group-specific* agent score instruments, capturing agents’ average IDR inducement rates within observably different subsamples.³⁸ Monotonicity requires a non-negative relationship between any of these subgroup-specific propensities. Figure A2 reports binned scatter plots and correlation coefficients for several pairwise comparisons of these group-specific instruments computed across the entire analysis sample. I find strongly positive correlations for each pair, suggesting agent inducement is similar across borrower characteristics.

³⁸Group-specific agent score instruments are calculated as

$$Z_{icj}^g = \left(\frac{1}{n_j^g - 1} \right) \left(\sum_{k=0}^{n_j^g} IDR_{kj}^* - IDR_{ij}^* \mathbf{1}_{\{i \in g\}} \right).$$

For example, Z_{icj}^{men} is the residualized, leave-one-out propensity of agent j to induce men into IDR.

6 Results

6.1 Repayment Outcomes

Table 4 reports difference-in-differences and IV estimates of IDR’s effect on minimum payments, changes in loan balances, and indicators for more than 10, more than 90, and more than 270 days delinquent. For each outcome listed on the left-hand side, Columns 2-5 provide estimated coefficients on the interaction between IDR take-up and consecutive seven-month periods from the pooled difference-in-differences specification given by Equation 17. Columns 6-9 provide estimated coefficients on IDR take-up from separate two-stage least squares regressions in each four-month period.³⁹ Figures 9 through A5 plot difference-in-differences coefficients on these outcomes, separated by each month following the date of the delinquency call.

6.1.1 Minimum Payments

The immediate effect of IDR enrollment on minimum payments is mechanical: given adjusted gross income, family size, and debt balance, one could directly calculate a borrower’s IDR payments using Equation 1. Estimating the change in payments for compliers in my sample, however, can provide insight into the mechanisms behind effects on other outcomes.⁴⁰ If we saw no effect on payments, any observed effects on delinquencies or consumption would be entirely attributable to expectations of future liquidity benefits or loan forgiveness, both of which are unlikely if income is too high for reduced payments in the short term. Moreover,

³⁹Results from the first two months following the call are omitted from Table 4 because it typically takes one or two months following contact to process and enroll borrowers in IDR, and eventual enrollees often forego making payments until IDR enrollment is complete. The relative timing of successful enrollment, next payment due date, and data collection date at the end of the calendar month adds further lag time before IDR effects can be realized.

⁴⁰Note that while the estimated effects of IDR on minimum payment help interpret estimated local average treatment effects on other outcomes, the more specific question of how IDR changes minimum payments can be directly measured, at least for the treatment group in my sample, because the counterfactual standard loan payments for IDR enrollees is a deterministic function of principal debt amounts and payment histories up until point they switch to IDR. More simply, the effect of IDR on payment size is approximately given by observed IDR payments minus payments in the month prior to receiving the delinquency call. Figure A4 provides a graphical illustration of this measured payment effect across the distribution of IDR enrollees in my analysis sample.

minimum payments for treated borrowers are a function of both payments under IDR as well as persistence in IDR. Treated borrowers who fail to recertify their incomes after one year of IDR payments will revert back to the standard repayment plan and face monthly minimums at least as high or higher than the payments they faced pre-enrollment.⁴¹ Minimum payment results therefore provide a useful “first stage” to measure the combined effect of payment size and persistence within the complier population.

In the difference-in-differences analysis, required payments for those on IDR fall by \$169 relative to those who remain on standard repayment three to nine months following the delinquency call. In the top panel of Figure 9, payments remain low until months twelve through fifteen, when they sharply increase to pre-call levels before gradually declining in subsequent months. IV estimates of IDR’s effect on payments are similar to difference-in-difference estimates. For those induced into IDR by their servicing agents, minimum payments lower by an average of \$187 three to nine months following IDR take-up and exhibit the same increase in later months. Payments are only \$98, \$62, and \$72 lower than non-enrollees in the following three seven-month periods.

6.1.2 Delinquencies

I measure IDR’s impact on delinquencies using three measures: the likelihood of falling more than 10 days delinquent, the likelihood of falling more than 90 days delinquent, and the likelihood of falling more than 270 days delinquent. The 10-, 90-, and 270-day benchmarks are points of increased delinquency penalties: At eleven days past due, borrowers begin to accrue late fees for delinquent loans. At 91 days past due, borrowers are reported to credit bureaus. At 271 days past due, a borrower becomes eligible for default. Defaulted loans can result in garnished wages, withheld tax returns, and revocation of professional licenses.

Monthly difference-in-differences estimates, shown in Figure 10, indicate a large negative effect of IDR enrollment for all three delinquency measures. IDR borrowers are 31 percentage points less likely to fall more than ten days delinquent relative to standard borrowers in the three to nine month period following the delinquency call, with a pre-call mean of 93pp. Corresponding estimates for months ten through sixteen, and months seventeen through

⁴¹Standard payments might be higher after a year on IDR because unpaid interest is, under some circumstances, recapitalized into the principal amount.

twenty-three, and months twenty-four through thirty are -14pp, -9pp, and -10pp, respectively. Estimates for more than 90 days delinquent (-8pp, -6pp, -2pp, and -4pp) and more than 270 days delinquent (-0.1pp, -0.5pp, -0.2pp, and -0.3pp) exhibit similar patterns, albeit mechanically smaller in magnitude. IV estimates for the effect of IDR on delinquencies are consistent with difference-in-differences estimates, though considerably less precise. Estimates of more-than-ten-day delinquencies are large and statistically significant for months three to nine (-35pp) and months ten to sixteen (-25pp), but attenuate in subsequent months. Corresponding results for longer delinquencies are too imprecise to be informative, though equality between IV and difference-in-differences estimates cannot be rejected for any delinquency outcomes.

6.1.3 Change in Balances

While minimum monthly payments measure the amount borrowers are required to pay to avoid late penalties, the high incidence of delinquency means that payments *due* provides a poor measure of payments *made*. I estimate effects on monthly changes in remaining balances to provide a sense of how IDR might impact borrowers' propensity to pay down their debts.⁴² Difference-in-differences estimates of IDR's effect on changes in debt balances suggest that, despite lowering required payments, IDR significantly increases the amount repaid, especially in the short term. Relative to standard borrowers, IDR enrollees pay down \$36 more debt in months three through nine than they did in the before their delinquency calls, when debt balances increased by an average of \$18. Some of these gains are lost in months ten through sixteen when attrition from IDR is high—point estimates suggest relative balance *increases* of \$25. However, these losses are more than recouped in months seventeen to twenty-three when balances decrease by \$28. The net effect implies IDR enrollees pay relatively more total debt over the entire thirty month period than non-enrollees.

It is important to emphasize that estimated effects on balances are relative, not absolute. On average, neither standard nor IDR borrowers are decreasing their total balances over this period. As the top-right panel of Figure 4 illustrates, change-in-balance estimates are driven

⁴²Note that, depending on the specific plan and minimum payment amount, IDR borrowers can sometimes receive partial forgiveness on accumulated interest. My measure of change in balances removes any interest forgiveness.

by a relatively slower absolute increase among IDR enrollees. Like long-term delinquency outcomes, IV estimates for monthly changes in debt balances are generally consistent with difference-in-difference estimates, though highly imprecise.

6.2 Credit and Consumption Outcomes

Because I restrict analysis of annual TransUnion data to a yearly panel of borrowers, credit and consumption outcomes are measured for a considerably smaller sample and at much lower frequency than loan repayment outcomes. I therefore rely on only the difference-in-difference strategy for these outcomes, as the IV strategy lacks the statistical power needed to draw meaningful inference.

Table 6 provides estimates of the effect of IDR on credit score, mortgage, and auto loan from the difference-in-differences analysis described in Section 5.1. For each outcome listed on the left-hand side, Columns 2-5 report coefficients on IDR in consecutive years from the pooled regression specified in Equation 17, beginning with the year of the delinquency call. Columns 6-9 report IDR coefficients from a regression which omits pre-call month dummies and includes a linear time trend.

Relative to those who remained in standard repayment, borrowers who enrolled in IDR experienced a statistically significant 7.73-point increase in credit scores within one year of the delinquency call, off of a pre-call mean of 595.78, followed by increases of 8.55 and 7.73 points one and two years after the call. Estimates of IDR's effect on bankruptcy filings and auto loans are statistically indistinguishable from zero for all three years following the delinquency call. IDR effects on the likelihood of holding a mortgage are also effectively zero in the year of the call, but increase gradually throughout the sample period, rising to 2 percentage points by year two, off of a pre-call mean of 22pp, an increase of 8 percent.

Table 7 provides difference-in-difference estimates for the effect of IDR on credit card balances, number of credit cards, and credit card limits. While I find no significant effects in the year of the delinquency call, IDR is associated with statistically significant increases in all three credit card measures one and two years following the call. Compared to standard borrowers, total balances on credit cards held by IDR enrollees increase by \$293 and \$304 one and two years after the delinquency call, corresponding to an increase of 19 percent off

of a pre-call mean of \$1569. Similarly, by the second year of enrollment, IDR borrowers hold 0.119 more credit cards (pre-call mean of 3.138) and have \$774 higher credit limits (pre-call mean of \$4926) compared to those who remained in standard repayment plans. Estimates of IDR’s effect of credit limits are modest and not statistically significant, suggesting liquidity effects are driven principally by the direct increase in cash-on-hand from lower loan payments rather than an increased ability to borrow on credit, at least within the first two years of IDR enrollment.

6.3 Employment Outcomes

My data do not include direct measures of employment or income. I can however, construct proxy measures for both variables using the LLS data. For employment, I use the incidence of unemployment deferment. Unemployment deferments allow borrowers working less than thirty hours per week to “pause” student loan payments for up to three years and forgive some or all of the accumulated interest. Unemployment deferment should be weakly preferred to either IDR or standard repayment, as it reduces payments to zero and forgives a portion of interest from accumulating for up to three years. For income, I use the median income among households in the borrower’s reported zip code, taken from the 2006-2010 American Community Survey (US Census Bureau, 2010).

Using these measures as outcome variables presents two complications. First, both unemployment deferment and median zip-code income are imperfect proxies. Moving to a higher income zip code may not reflect an actual increase in income, and not all unemployed borrowers file for unemployment deferments. Second, Once a borrower is on standard or IDR payment plans following the phone call, both deferment and zip code are determined, in part, by subsequent contact with LLS. Since such contact is endogenous to plan choice and repayment behavior, treatment effects estimated for these outcomes might be biased. For example, borrowers who remain on standard repayment plans are more likely to receive subsequent delinquency calls, and therefore more likely to update their income with LLS or enroll in unemployment deferment.⁴³ As a result, such borrowers will likely have

⁴³Figure A3 plots average number of additional points of contact for each month relative to the reference call. Recall that treated borrowers are defined as those who enroll in IDR within four months of the reference call, and thus qualify for reduced payments for some twelve-month period between month one and month

higher *reported* incomes and rates of unemployment, biasing income effects downward and employment effects upward.

To address this concern, I restrict estimation of IDR effects on labor market outcomes to eighteen and thirty months following the delinquency call. These windows correspond to periods in which both treatment and control borrowers are likely to have had recent contact with LLS. Since treatment borrowers will have been on IDR for just over one year or two years, their payments will have returned to the same levels as those for standard borrowers, so both groups face comparable probabilities of contact with LLS for purposes of (re-)enrolling in IDR or deferment. Results for employment and income proxies in month eighteen (“Year 1”) and month thirty (“Year 2”) are reported in Table 5. Estimates for both outcomes are not statistically different from zero in either period.

Judicious choice of time periods to measure employment outcomes mitigates some of the potential bias due to plan-dependent post-call contact. Nonetheless, there are still several reasons to believe employment and income proxies are confounded by such endogenous factors, even in these selected time periods. The strategy presumes that those who are unemployed or changed zip codes as of the designated time period would have that information reflected in the LLS data during the preceding months. While the benefits of deferment make that assumption plausible in the case of unemployment, it is almost certainly violated for zip code changes. Zip code information for those borrowers who violate this assumption will have been gathered in earlier periods, during which contact rates differed between standard and IDR borrowers. Results for labor market outcomes, especially income, should therefore be interpreted with caution.

7 Interpretation

In this section, I interpret my results by placing them in the context of the model laid out in Section 4. In short, results are consistent with a pure liquidity effect of IDR: reducing short-

fifteen. As expected, rates of contact are considerably lower for IDR borrowers than they are for standard borrowers during this window, but rise to approximate parity in months sixteen through eighteen. Treated borrowers are therefore equally likely as control borrowers to provide updated employment or zip code information to LLS during these months. Such information would be reflected in the month eighteen data.

term payments increases borrowers' cash on hand, allowing them to increase consumption and avoid default in periods of low income.

Minimum Payments

Ignoring any wealth effects through expected debt forgiveness, estimates of IDR's effect on minimum payments are analogous to a "first stage" in an estimation of the effects of increased liquidity, as reduced payments are a necessary condition for IDR to generate liquidity effects. While unlikely in my setting, it is possible for high-income borrowers who face no short-term benefit from IDR to sign up solely in anticipation of future income shocks years later. My results run counter to this hypothetical: the large and significant observed effect on minimum payments suggests increased liquidity drives at least part of results on other outcomes.

Around twelve months following their initial decline, minimum payments rise sharply and gradually decrease thereafter. To understand this pattern, recall that minimum monthly payments for borrowers in the treatment group are determined by their calculated IDR obligations as well as their persistence in the IDR plan, which requires income recertification after one year of enrollment. As the bottom panel of Figure 2 illustrates, the pattern of estimated payment effects between months twelve and thirty appears to be largely driven by the persistence margin. After the one-year recertification period expires, more than sixty percent of IDR enrollees in my sample do not extend their benefits. The reason behind this attrition cannot be directly determined. While some enrollees may simply decline re-enrollment because their incomes have risen enough to make them ineligible for reduced payments under IDR, the more likely explanation is a behavioral response to the burdensome recertification process required under IDR. Indeed, the gradual decrease in subsequent payments after the one-year mark suggests some enrollees *are* re-enrolling but are doing so after the one year deadline, resulting in a temporary rise in minimum monthly payments. This observed payment path for IDR borrowers provides evidence in support of proposals aimed at streamlining the recertification process (Cox et al., 2018), as large shocks to monthly debt obligations contradict the goals of the program.

Having established that IDR borrowers receive a large but short-lived decrease in their monthly debt obligations, results from delinquency and credit outcomes provide evidence for

how borrowers respond to this cash infusion. Recall that borrowers with reduced payments have three options: (1) keep consumption constant and add to their savings, (2) repay their student debt in a more timely manner, or (3) increase short-term consumption. While these options are not mutually exclusive, a repayment and/or consumption response is consistent with liquidity constraints and suggestive of social welfare improvement under IDR. My results provide evidence of both responses: the negative effect on delinquencies and positive effect on credit cards and mortgages suggest borrowers use the liquidity provided by IDR to prevent default, increase non-durable consumption, and purchase homes.

Delinquencies

Decreased incidence of delinquencies following IDR take-up holds two important implications. First, it provides a pathway for increased borrower welfare under IDR, as non-repayment can severely impact borrowers' credit and employment prospects.⁴⁴ Second, reduced delinquencies under IDR speaks to the long-standing debate over the determinants of default. Increased repayment following a reduction in minimum payments is suggestive of a liquidity motive for default rather than a strategic motive, as lower monthly payments should not influence strategic default decisions.⁴⁵

Note that over half of treated borrowers in my sample face IDR payments of zero dollars and thus cannot fall delinquent on their loans. While this mechanical result could still be characterized as a liquidity effect under a neo-classical model, it may be driven in part by psychological frictions or "hassle-costs" if borrowers facing payments of $\epsilon > 0$ dollars would face higher delinquency rates than zero-payment borrowers. To investigate the importance of this channel, I conduct my difference-in-differences analysis among a subsample of individuals with predicted nonzero IDR payments using the following procedure: First, I regress an indicator for positive IDR payments among IDR enrollees on a full set of demographic controls and pre-call student loan and credit variables. Second, I use these estimates to

⁴⁴Note that the magnitude of welfare effects through this channel are difficult to measure, as delinquencies serve only as a noisy indicator for the myriad of consequences associated with various stages of non-repayment described in Section 2.

⁴⁵I interpret "strategic default" as default decisions driven by total outstanding liability, which does not include decisions driven by expected future liquidity constraints. Also note that IDR may affect strategic default through expected loan forgiveness, though I argue forgiveness effects are unlikely later in the section.

predict the likelihood of having positive IDR payments among *all* borrowers in the analysis sample. Finally, I conduct my difference-in-differences analysis on the subgroup of individuals with greater than fifty percent predicted likelihood of positive IDR payments. Realized IDR payments are nonzero for more than seventy percent of treated individuals in this subsample, yet the repayment effects of IDR persist. Table A6 reports delinquency results for this subsample, and continues to find a large and significant effect on repayment rates.

Balances

In theory, IDR could affect balances on student loans in either direction. Lower required payments imply non-delinquent IDR borrowers pay down less of their debt than non-delinquent borrowers on standard repayment plans, increasing balances. However, the reduced default likelihood while on IDR has the opposite effect on balances. My results suggest the effect of reduced minimums on loan balances is dominated by more timely repayment: IDR borrowers pay down an average of \$36 more student debt each month relative to those on standard repayment plans. Importantly, however, this does not imply that IDR borrowers are closer to solvency than they were prior to enrollment, as balances are increasing for both treatment and control borrowers in my sample.

The negative balance result highlights the importance of considering counterfactual repayment behavior when evaluating the implications IDR has for government revenue. In the short term, net cash flows to the government depend not only on minimum payment schedules, but also on the incidence of IDR take-up and the distribution of relative repayment probabilities under each plan. Indeed, my results suggest that even though monthly minimums are small under IDR, total cash flows through completed payments may increase because so many borrowers would not have made timely payments under standard repayment.⁴⁶

In Figure A5, I use the full representative sample to plot actual total cash flows versus

⁴⁶While my analysis sample is selected to include those who are particularly prone to delinquency, so is the general population of likely IDR enrollees—in my representative sample of LLS borrowers, eventual IDR enrollees are 56 percent more likely to have fallen delinquent between July 2011 and June 2012 than those who never enroll. This finding is consistent with (Looney and Yannelis, 2015), who find that those who would benefit most from IDR (i.e. those with high debt-to-earnings ratios) are more likely to default under standard repayment plans.

predicted total cash flows in the counterfactual scenario in which all student borrowers enrolled in IDR in July 2012. I make these predictions using difference-in-difference estimates for the analysis sample re-weighted so that the joint distribution of pre-call observables matches that of the full representative sample. While extrapolating these estimates to a representative population carries a number of strong assumptions, the figure demonstrates how increased repayment likelihoods might ameliorate the potential fiscal cost associated with lower minimum payments, at least in the short term. Even in the long term, it is not clear whether IDR availability will impose a fiscal cost or benefit. While the government must eventually incur the cost of any forgiven loans, it will also avoid the costs of defaulted loans that are never repaid and the administrative cost of servicing serially delinquent borrowers.

Note that much like minimum payments, the pattern of both delinquency and balance effects appears driven by the recertification cycle. Treated borrowers show sustained decreases in delinquencies and balances relative to non-enrollees during the period when they qualify for reduced payments, but experience sharp increases in both measures as they fail to recertify their incomes twelve months after enrollment. In fact, enrollees pay down *less* average debt than non-enrollees for a few months immediately following initial certification period, likely because many wait until they have re-qualified for IDR before making any payments.

Credit Outcomes

The positive estimated effects of IDR on credit card balances, number of credit cards, and mortgage-holding rates are suggestive of a consumption and homeownership response to increased liquidity, but interpreting these results requires three important caveats. First, I have loosely referred to credit card balances as “proxies for consumption,” but balances reported in the credit bureau data capture both flows in credit card spending as well as the stock of unpaid debt. Increased credit card balances arising from more rolled-over debt, as opposed to higher transaction volumes, might not reflect a consumption response and could decrease borrower welfare. Given the positive estimated effects of IDR on credit scores, such increases in accumulated debt are unlikely to drive my credit card results. However, borrowers who were carrying debt before IDR enrollment may use their increased cash-on-

hand to *decrease* existing credit card debt. Such a deleveraging response would attenuate the estimated effect on credit card balances, understating the true consumption effect of IDR.

Second, even credit card responses driven entirely by changes in transaction volumes may not reflect changes in consumption. If credit card transactions reflect changes in short-term spending on durable goods, such responses would be more accurately interpreted as “expenditure effects.” For example, if IDR borrowers used their cash-on-hand to expedite purchases on automobiles, furniture, or appliances that they would have eventually bought anyway, their true consumption responses would depend upon the depreciation of these goods over the period in between counterfactual purchase dates. While this potential difference between expenditure and consumption effects could be large for results measured within short time horizons, the persistence of estimated credit card responses over three years suggests a true increase in combined durable/non-durable consumption relative to control borrowers.

Third, increased liquidity can have both direct and indirect effects on borrower outcomes, especially over long time periods. While the immediate increase in cash-on-hand borrowers receive through IDR is the primary channel through which we would expect increases in consumption or homeownership, the increased credit-access associated with higher credit scores provides a potential secondary channel. While this credit channel probably accounts for some of my results, especially over the two and three year time horizons, the impact on credit limits is modest relative to credit card balance and credit score effects in the short-term, suggesting liquidity effects are driven largely by the direct increase in cash-on-hand from lower loan payments rather than an increased ability to borrow on credit, at least within the first two years of IDR enrollment.

7.1 Alternative Mechanisms

The above discussion interprets estimates of IDR enrollment effects on repayment and credit outcomes as operating through increased short-term liquidity, but alternative potential channels exist. As discussed in Section 4, debt forgiveness provisions under IDR can also affect borrower behavior. If borrowers expect their loans to be forgiven, they may increase repayment to try and qualify for forgiveness or increase short-term consumption out of increases to their expected lifetime wealth. I argue this mechanism is unlikely. In order for the for-

givenness period to bind, a borrower must make 300 complete monthly loan payments on IDR and still hold an outstanding balance on her student loans. Minimum payments under standard plans would pay off balances after just 120 payments. The high rate of attrition from IDR therefore suggests forgiveness is very unlikely for many borrowers in my sample. Importantly, this rapid rise in payments does not necessarily imply an increase in income; borrowers may qualify for continued reduced payments but not recertify their income after twelve months of IDR. Even so, such borrowers should not expect loan forgiveness, unless they recertify later on. While I find some evidence of delayed recertification beyond fifteen months, the majority of initial enrollees have still not re-enrolled by month thirty. This high rate of attrition, whether driven by increased incomes, recertification “hassle-costs,” or behavioral phenomena like inattention or myopia, suggests few borrowers in my sample should expect their loans to be forgiven.

While my results do not appear to be driven by expected loan forgiveness, I cannot rule out behavioral mechanisms driving the positive effect on repayment and consumption. For example, myopia can drive effects on consumption upwards while reducing borrower welfare if increased cash-on-hand induces overspending in the short-run. Likewise, enrolling in a new repayment plan may make a borrower’s debt more psychologically salient, prompting her to keep current on payments even if the plan provides no real benefit.

8 Conclusion

In this paper, I estimate the ex-post causal impact of IDR enrollment on student debt repayment, financial health, and employment proxies using a novel dataset linking monthly federal student loan payments to credit bureau records for over one million student borrowers. My research design uses two identification strategies: a difference-in-differences design comparing borrowers with differential IDR take-up following delinquency calls from their loan servicer and an instrumental variables design exploiting variation in the tendency of randomly-assigned servicing agents to enroll borrowers in IDR. Within seven months of take-up, IDR enrollees are 31 percentage points less likely to fall delinquent and pay down \$162 more student debt each month compared to those who remain on standard repayment plans. IDR enrollees have credit scores that are 8.6 points higher, hold 0.1 more credit cards, and

carry \$293 higher credit card balances than non-enrollees one year after the servicing call, suggesting increased short-term consumption out of liquidity. IDR enrollees are also 1.7 percentage points more likely to hold a mortgage, an increase of 8 percent off of the pre-call mean, suggesting a positive effect of IDR on homeownership. Minimum monthly payments decrease by an average of \$169 following IDR take-up, but return to standard levels within one year of enrollment, as few borrowers recertify their incomes after a year on IDR. I find no evidence of labor supply effects through unemployment deferments or median zip-code level income. These results do not appear driven by borrower responses to expected loan forgiveness. Instead, they suggest IDR improves borrower welfare principally through a liquidity channel, providing short-term increases to cash-on-hand during periods of financial distress.

The results of my analysis raise several questions for future research. First, while this paper documents the ex-post liquidity benefits of IDR on several borrower outcomes, a full accounting of IDR's impact on social welfare would incorporate its ex-ante effects on college attendance, institution and major choice, career decisions, and principal borrowing amounts. Indeed, Abraham et al. (2018a) find survey evidence that IDR may influence borrowers' career paths, and a large literature on financial aid has shown that the amount and type of support students receive can affect their decisions both before and during college (Marx and Turner, 2015; Dynarski, 2003). Likewise, the proliferation of IDR may also affect the incentives of post-secondary institutions, which have been shown to strategically respond to students' financing options and labor market conditions (Armona et al., 2018; Turner, 2013).

Second, while the apparent lack of forgiveness-eligible borrowers in my sample provides clear evidence of liquidity benefits, it also leaves unanswered questions regarding the trade-off between insurance value and moral hazard inherent to income-contingent forgiveness policies. These questions hold significant policy importance, as a number of existing and proposed student loan programs offer more generous forgiveness provisions than those under the IDR plan studied in this paper. For instance, teachers and public service employees can apply for debt forgiveness after only ten years of IDR payments (Department of Education, 2017a), and President Trump has proposed a version of IDR with a fifteen-year forgiveness period (Douglas-Gabriel, 2015). Further empirical work is needed to assess the welfare benefits and potential distortionary effects of these debt-forgiveness policies.

Finally, my results highlight many shortcomings of the current structure of IDR plans in the United States and suggest potential improvements to the program. Most notably, hassle costs appear to be a large impediment for millions of potential IDR beneficiaries. Despite facing zero pecuniary costs to enrolling, initial take-up for IDR is remarkably low for my sample. As Cox et al. (2018) and Abraham et al. (2018b) demonstrate, this low take-up is likely driven by psychological frictions or behavioral constraints, considering IDR should be weakly preferred by any borrower behaving in accordance with classical economic theory.⁴⁷ Perhaps even more striking, less than half of those who do enroll manage to recertify their incomes within a year. These two observations call for two complementary reforms to IDR: First, take-up should be streamlined or made automatic. As the literature on retirement savings demonstrates (Duflo and Saez, 2004; Madrian and Shea, 2001), automating the enrollment process can have dramatic effects on take-up. Second, income verification should be less onerous. For instance, payments could be determined and deducted automatically through the tax system. These downstream improvements are crucial to ensure long term benefits under IDR, as my results illustrate how even those who are able to overcome initial barriers to enrollment are unlikely to stay in the program.

Even as important questions remain unanswered, this study provides a crucial step towards improving the system of financing for higher education in the US. IDR represents the largest change to this system in more than fifty years. Evaluating its overall impact on social welfare requires many considerations—the positive externalities of college, the redistributive impact of subsidies, the welfare gains from insuring human capital investments, and the distortionary costs of income-contingent benefits. This paper estimates one important component of this overall impact—the ex-post liquidity benefits of IDR enrollment. It therefore speaks not only to the impact of existing student loan policy, but also to larger question of how society can best finance investments in human capital.

⁴⁷Prior work has demonstrated the importance of such constraints in similar contexts, including college enrollment (Dynarski et al., 2018), college major choice (Hastings et al., 2015), and principal borrowing decisions (Marx and Turner, 2017).

References

- Abraham, Filiz-Ozbay, Ozbay, and Turner (2018a, July). Behavioral effects of student loan repayment plan options on borrowers' career decisions: Theory and experimental evidence. National Bureau of Economic Research Working Paper 24804.
- Abraham, K. G., E. Filiz-Ozbay, E. Y. Ozbay, and L. J. Turner (2018b). Framing effects, earnings expectations, and the design of student loan repayment schemes. National Bureau of Economic Research Working Paper 24484.
- Aizer, A. and J. J. Doyle (2015). Juvenile incarceration, human capital, and future crime: Evidence from randomly assigned judges. *The Quarterly Journal of Economics* 130(2), 759–803.
- Angrist, J. D., G. W. Imbens, and D. B. Rubin (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association* 91(434), 444–455.
- Armona, L., R. Chakrabarti, and M. Lovenheim (2018). How does for-profit college attendance affect student loans, defaults and earnings? National Bureau of Economic Research Working Paper 25042.
- Autor, D. H. and S. N. Houseman (2010). Do temporary-help jobs improve labor market outcomes for low-skilled workers? evidence from "work first". *American Economic Journal: Applied Economics*, 96–128.
- Avery, R. B., P. S. Calem, G. B. Canner, and R. W. Bostic (2003). An overview of consumer data and credit reporting. *Fed. Res. Bull.* 89, 47.
- Barr, N., B. Chapman, L. Dearden, and S. Dynarski (2017). Getting student financing right in the US: lessons from Australia and England. Centre for Global Higher Education Working Paper.
- Bleemer, Z., M. Brown, D. Lee, K. Strair, and W. Van der Klaauw (2017). Echoes of rising tuition in students' borrowing, educational attainment, and homeownership in post-recession america. FRBNY Staff Report Number 820.

- Britton, J., L. van der Erve, and T. Higgins (2018). Income contingent student loan design: Lessons from around the world. *Economics of Education Review*.
- Chapman, B. (2006). Income contingent loans for higher education: International reforms. *Handbook of the Economics of Education 2*, 1435–1503.
- Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy 116*(2), 173–234.
- Cox, J. C., D. Kreisman, and S. Dynarski (2018). Designed to fail: Effects of the default option and information complexity on student loan repayment. National Bureau of Economic Research Working Paper 25258.
- Dahl, G. B., A. R. Kostøl, and M. Mogstad (2014). Family welfare cultures. *The Quarterly Journal of Economics 129*(4), 1711–1752.
- Department of Education (2017a). Federal student loan portfolio. Website. Accessed: 2017-10-19.
- Department of Education (2017b). New student loan report reveals promising repayment trends. Website. Accessed: 2017-10-19.
- Di Maggio, M., A. Kermani, B. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2016). Monetary policy pass-through: Mortgage rates, household consumption and voluntary deleveraging. *American Economic Review*.
- Dobbie, W., J. Goldin, and C. S. Yang (2018). The effects of pretrial detention on conviction, future crime, and employment: Evidence from randomly assigned judges. *American Economic Review 108*(2), 201–40.
- Dobbie, W., P. Goldsmith-Pinkham, N. Mahoney, and J. Song (2016). Bad credit, no problem? credit and labor market consequences of bad credit reports. National Bureau of Economic Research Working Paper 22711.
- Dobbie, W., P. Goldsmith-Pinkham, and C. S. Yang (2017). Consumer bankruptcy and financial health. *Review of Economics and Statistics 99*(5), 853–869.

- Dobbie, W. and J. Song (2015). Debt relief and debtor outcomes: Measuring the effects of consumer bankruptcy protection. *The American Economic Review* 105(3), 1272–1311.
- Douglas-Gabriel, D. (2015). Trump just laid out a pretty radical student debt plan. *The Washington Post*.
- Dudley, W. (2015). Opening remarks at the convening on student loan data conference. Federal Reserve Bank of New York.
- Dufló, E. and E. Saez (2004). Implications of pension plan features, information, and social interactions for retirement saving decisions. *Pension design and structure: New lessons from behavioral finance*, 137–153.
- Dynarski, S. (2014, September). An economist’s perspective on student loans in the United States. Gerald R. Ford School of Public Policy Working Paper.
- Dynarski, S. (2015). We’re frighteningly in the dark about student debt. *The New York Times*.
- Dynarski, S. and D. Kreisman (2013). Loans for educational opportunity: Making borrowing work for today’s students. *The Hamilton Project. Discussion Paper*.
- Dynarski, S., C. Libassi, K. Micheltmore, and S. Owen (2018). Closing the gap: The effect of a targeted, tuition-free promise on college choices of high-achieving, low-income students. National Bureau of Economic Research Working Paper 25349.
- Dynarski, S. M. (2003). Does aid matter? measuring the effect of student aid on college attendance and completion. *American Economic Review* 93(1), 279–288.
- Field, E. (2009). Educational debt burden and career choice: evidence from a financial aid experiment at nyu law school. *American Economic Journal: Applied Economics* 1(1), 1–21.
- Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, K. Baicker, and O. H. S. Group (2012). The Oregon health insurance experiment: evidence from the first year. *The Quarterly Journal of Economics* 127(3), 1057–1106.

- Friedman, M. (1962). *Capitalism and Freedom*. University of Chicago press.
- Ganong, P. and P. Noel (2017). The effect of debt on default and consumption: Evidence from housing policy in the great recession. Working Paper.
- Hastings, J., C. A. Neilson, and S. D. Zimmerman (2015). The effects of earnings disclosure on college enrollment decisions. National Bureau of Economic Research Working Paper 21300.
- Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica* 47(1), 153–161.
- Heckman, J. J. and E. Vytlacil (2005). Structural equations, treatment effects, and econometric policy evaluation 1. *Econometrica* 73(3), 669–738.
- Johnston, A. and N. Barr (2013). Student loan reform, interest subsidies and costly technicalities: lessons from the uk experience. *Journal of Higher Education Policy and Management* 35(2), 167–178.
- Keys, B. J., T. Piskorski, A. Seru, and V. Yao (2014). Mortgage rates, household balance sheets, and the real economy. National Bureau of Economic Research Working Paper 20561.
- Kling, J. R. (2006). Incarceration length, employment, and earnings. *The American Economic Review* 96(3), 863–876.
- Kostøl, A. R., M. Mogstad, B. Setzler, et al. (2017). Disability benefits, consumption insurance, and household labor supply. National Bureau of Economic Research Working Paper.
- Ljungqvist, L. and T. J. Sargent (2012). *Recursive macroeconomic theory*. MIT press.
- Looney, A. and C. Yannelis (2015). A crisis in student loans?: How changes in the characteristics of borrowers and in the institutions they attended contributed to rising loan defaults. *Brookings Papers on Economic Activity* 2015(2), 1–89.

- Lucas, D. and D. Moore (2010). Costs and policy options for federal student loan programs. a CBO study. *Congressional Budget Office*.
- Madrian, B. C. and D. F. Shea (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics* 116(4), 1149–1187.
- Maestas, N., K. J. Mullen, and A. Strand (2013). Does disability insurance receipt discourage work? using examiner assignment to estimate causal effects of ssdi receipt. *American economic review* 103(5), 1797–1829.
- Marx, B. M. and L. J. Turner (2015). Borrowing trouble? student loans, the cost of borrowing, and implications for the effectiveness of need-based grant aid. National Bureau of Economic Research Working Paper 20850.
- Marx, B. M. and L. J. Turner (2017, November). Student loan nudges: Experimental evidence on borrowing and educational attainment. National Bureau of Economic Research Working Paper 24060.
- Mezza, A., D. R. Ringo, S. Sherland, and K. Sommer (forthcoming). Student loans and homeownership. *Journal of Labor Economics*.
- Rothstein, J. and C. E. Rouse (2011). Constrained after college: Student loans and early-career occupational choices. *Journal of Public Economics* 95(1), 149–163.
- Shimer, R. and I. Werning (2008). Liquidity and insurance for the unemployed. *The American Economic Review* 98(5), 1922–1942.
- Tang, C., K. Ross, N. Saxena, and R. Chen (2011). What’s in a name: a study of names, gender inference, and gender behavior in Facebook. *Database Systems for Advanced Applications*, 344–356.
- Turner, L. J. (2013). The economic incidence of federal student grant aid. Unpublished manuscript. http://econweb.umd.edu/~turner/Turner_FedAidIncidence_Jan2017.pdf.

US Census Bureau (2010). Selected characteristics of the native and foreign-born populations: 2006-2010 American Community Survey 5-year estimates. Technical report, US Department of Commerce, Economics and Statistics Administration.

Yannelis, C. (2016). Asymmetric information in student loans. Working Paper.

Tables and Figures

Table 1: Summary Statistics

	Full Sample	Analysis Sample		
	(1) Pooled	(2) Pooled	(3) Control	(4) Treatment
<i>Panel A: LLS Data</i>				
IDR	0.117	0.348	0.231	1
Female	0.590	0.715	0.709	0.750
Zip Median Income	60.83	52.17	52.36	51.13
Age	36.97	39.26	39.28	39.10
Amount Borrowed	24.78	21.99	21.60	24.14
10+ Days Delinquent	0.382	0.765	0.775	0.711
90+ Days Delinquent	0.146	0.289	0.296	0.250
Days Delinquent	20.20	23.45	24.38	18.30
<i>Panel B: Credit Data</i>				
Credit Score	664.2	582.6	582.0	586.1
Bankruptcy	0.0862	0.162	0.159	0.178
Derogatory Rating	0.267	0.598	0.595	0.616
Number of Credit Cards	5.454	3.421	3.426	3.390
Credit Card Balances	3.975	1.479	1.488	1.429
Number of Mortgages	1.225	0.846	0.862	0.757
Mortgage Balances	67.36	31.49	32.84	24.00
Credit Card Limits	18.73	4.584	4.635	4.298
Number of Auto Trades	1.885	1.618	1.629	1.556
<i>N</i>	699100	34298	29051	5247

Note: This table reports summary statistics at the borrower level. The full sample is a random sample of the population of borrowers in LLS's FFEL portfolio who carried a positive loan balance as of December 31, 2011 and hold no private or Direct loans. The analysis sample is a subsample from the same population, selected according to the following criteria: borrowers received an IDR-modeled delinquency call from 2009 onward, have observable repayment histories ten months prior and thirty months following the phone call, are not recorded as non-english speakers, and were assigned to agents with at least 100 observed phone calls. Treated borrowers are those who enroll in IDR within four months of a delinquency call. IDR is an indicator for whether the borrower ever enrolled in IDR. Female is a measure of likelihood-female inferred from first name following Tang et al. (2011). Zip median income is the median 2010 income for the borrower's recorded 5-digit zip code. Days delinquent is the maximum number of days the borrower was ever past due on payments in the past year, and ever delinquent is an indicator for whether days delinquent is greater than 10. Number of calls is the total number of outgoing calls made to the borrower in the past year. IDR and treatment status reflect IDR enrollment histories through September, 2018. All other LLS variables are taken from administrative records as of June 30, 2012. Credit scores, bankruptcies, derogatory ratings, credit card, mortgage, and auto loan information are taken from TransUnion credit bureau data collected in August 2012.

Table 2: Balance Test

	(1)	(2)
	Agent Score*100	IDR*100
Female	0.058347 (0.043207)	1.656160*** (0.452140)
Amount Borrowed	-0.002080 (0.002340)	-0.104483*** (0.024767)
Age	0.002956* (0.001728)	-0.016085 (0.020268)
Lag Zip Median Income	-0.001667* (0.000991)	-0.026106*** (0.009598)
Lag Days Delinquent	-0.000075 (0.000739)	-0.003221 (0.007219)
Lag Minimum Payment	0.367353** (0.149353)	6.437119*** (1.873937)
Lag Remaining Balance	0.001935 (0.001639)	0.124716*** (0.019786)
Lag Credit Score	-0.000021 (0.000282)	0.035717*** (0.004283)
Lag Credit Card Balances	-0.007201 (0.006427)	-0.065190 (0.067098)
Lag Any Auto Trade	-0.019281* (0.011299)	-0.251060** (0.115486)
Lag Any Mortgage	-0.110630 (0.073337)	-0.398972 (0.722193)
Lag Mortgage Balances	0.000316 (0.000459)	-0.016237*** (0.003781)
Lag Number of Credit Cards	0.002617 (0.005706)	0.098107 (0.071924)
Lag Credit Card Limits	0.002149 (0.002942)	-0.087035*** (0.028729)
Mean Dep.	-0.406	11.160
F-stat	1.14	18.27
P-value	0.3287	0.0000
R-squared	0.161	0.018
N	30359	30359

Note: This table reports balance test results. The regressions are estimated on the call sample described in the notes to Table 1. Column 1 reports the estimated coefficients from an OLS regression of agent score multiplied by 100 against the variables listed, as well as agent modeling propensity and call year, month, and hour fixed effects. Agent score and modeling propensity are estimated using data from other phone calls placed by the same agent following the procedure described in Section 5. Column 2 reports estimates from an identical regression, except with the dependent variable equal to realized IDR take-up as of six months after the call, multiplied by 100. Robust standard errors two-way clustered at the borrower and agent level are reported in parentheses. The p-value reported at the bottom of columns 1-2 is for an F-test of the joint significance of the variables listed on the left. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 3: First Stage

	(1)	(2)
	IDR	IDR
Agent Score	0.6081489*** (0.0684207)	0.5515894*** (0.0776854)
Female		0.0162398*** (0.0044911)
Amount Borrowed		-0.0010334*** (0.0002477)
Age		-0.0001772 (0.0002025)
Lag Zip Median Income		-0.0002519*** (0.0000963)
Lag Days Delinquent		-0.0000318 (0.0000717)
Lag Minimum Payment		0.0623449*** (0.0186895)
Lag Remaining Balance		0.0012365*** (0.0001981)
Lag Credit Score		0.0003573*** (0.0000427)
Lag Credit Card Balances		-0.0006122 (0.0006734)
Lag Any Auto Trade		-0.0024043** (0.0011470)
Lag Any Mortgage		-0.0033795 (0.0071708)
Lag Mortgage Balances		-0.0001641*** (0.0000380)
Lag Number of Credit Cards		0.0009666 (0.0007212)
Lag Credit Card Limits		-0.0008822*** (0.0002880)
Mean Dep.	0.112	0.112
F-stat	79.00	50.41
P-value	0.0000	0.0000
R-squared	0.011	0.020
<i>N</i>	30359	30359

Note: This table reports first-stage results. The regressions are estimated on the call sample described in the notes to Table 1. Columns 1 and 2 report estimated coefficients from an OLS regression of IDR take-up within four months of a delinquency calls against the variables listed, as well as agent modeling propensity and call year, month, and hour fixed effects. Agent score and modeling propensity are estimated using data from other phone calls placed by the same agent following the procedure described in Section 5. Robust standard errors two-way clustered at the borrower and agent level are reported in parentheses. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 4: Difference-in-Differences and Instrumental Variables Estimates of the Effect of IDR Enrollment on Repayment Outcomes

<i>Dependent Variable</i>	Difference-in-Differences					Instrumental Variables			
	(1) Mean $t = -1$	(2) Mo. 3-9	(3) Mo. 10-16	(4) Mo. 17-23	(5) Mo. 24-30	(6) Mo. 3-9	(7) Mo. 10-16	(8) Mo. 17-23	(9) Mo. 24-30
Minimum Payment	0.200	-0.169*** (0.003)	-0.101*** (0.003)	-0.028*** (0.004)	-0.038*** (0.003)	-0.187*** (0.019)	-0.098*** (0.029)	-0.062 (0.040)	-0.072** (0.034)
Remaining Balance	24.963	-0.162*** (0.011)	-0.393*** (0.020)	-0.407*** (0.034)	-0.464*** (0.047)	0.045 (0.179)	-0.043 (0.329)	0.069 (0.537)	0.515 (0.715)
Δ Remaining Balance	0.018	-0.036*** (0.003)	0.025*** (0.003)	-0.028*** (0.003)	0.001 (0.003)	0.006 (0.035)	0.082** (0.033)	0.017 (0.037)	0.125*** (0.043)
10+ Days Delinquent	0.925	-0.314*** (0.006)	-0.143*** (0.006)	-0.086*** (0.007)	-0.096*** (0.007)	-0.349*** (0.102)	-0.252*** (0.096)	0.048 (0.091)	-0.112 (0.099)
90+ Days Delinquent	0.063	-0.076*** (0.003)	-0.061*** (0.003)	-0.021*** (0.004)	-0.041*** (0.004)	0.013 (0.073)	-0.106* (0.057)	0.094* (0.054)	-0.076 (0.064)
270+ Days Delnquent	0.000	-0.001 (0.000)	-0.005*** (0.001)	-0.002*** (0.001)	-0.003** (0.001)	0.013* (0.007)	0.015 (0.017)	-0.011 (0.014)	0.002 (0.021)
Call Time FE		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	30359	1244719	1244719	1244719	1244719	212513	212513	212513	212513

Note: This table reports difference-in-differences and two-stage least squares estimates of the effect of IDR enrollment on monthly loan repayment outcomes. Column 1 reports the dependent variable mean in the month prior to receiving a delinquency call. Columns 2-5 report coefficients on the effect of IDR enrollment in consecutive consecutive seven-month periods following the delinquency call from the pooled OLS regression specified in Equation 17. Each of Columns 6 - 9 report estimates from separate two-stage least squares regressions on outcomes in the same months. Regressions are estimated on the analysis sample as described in the notes to Table 1. Two-stage least squares models instrument for IDR enrollment with the agent score and control for agent modeling propensity following the procedure described in Section 5. All specifications include controls for call date and time, as well as amount borrowed, inferred gender, and zip-code median income. Robust standard errors, in parentheses, are clustered at the individual level. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 5: Difference-in-Differences and Instrumental Variables Estimates of the Effect of IDR Enrollment on Employment Outcomes

<i>Dependent Variable</i>	Difference-in-Differences			Instrumental Variables	
	(1) Mean $t = -1$	(2) Year 1	(3) Year 2	(4) Year 1	(5) Year 2
Unemployment Deferment	0.000	-0.001 (0.004)	0.003 (0.004)	-0.019 (0.041)	-0.012 (0.038)
Zip Median Income	52.240	0.028 (0.209)	-0.014 (0.225)	-2.340 (4.273)	-4.312 (4.691)
Call Time FE		Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes
N	30359	1244719	1244719	30359	30359

Note: This table reports difference-in-differences and two-stage least squares estimates of the effect of IDR enrollment on unemployment deferments and median zip-code income. Column 1 reports the dependent variable mean in the month prior to receiving a delinquency call. Columns 2-3 report coefficients on the effect of IDR enrollment in month 18 (“Year 1”) and month 30 (“Year 2”) from the pooled OLS regression specified in Equation 17. Each of Columns 3-4 report estimates from separate two-stage least squares regressions on outcomes in the same months. Regressions are estimated on the analysis sample as described in the notes to Table 1. Two-stage least squares models instrument for IDR enrollment with the agent score and control for agent modeling propensity following the procedure described in Section 5.2. All specifications include controls from call date and time, as well as amount borrowed and inferred gender. Robust standard errors, in parentheses, are clustered at the individual level. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 6: Difference-in-Differences Estimates of the Effect of IDR Enrollment on Financial Outcomes

<i>Dependent Variable</i>	Diff-in-Diff					Diff-in-Diff w/Trend			
	(1) Mean $t = -1$	(2) Year 0	(3) Year 1	(4) Year 2	(5) Year 3	(6) Year 0	(7) Year 1	(8) Year 2	(9) Year 3
Credit Score	595.780	7.734*** (1.463)	8.550*** (1.726)	7.726*** (1.900)	9.241*** (1.985)	6.712*** (1.015)	7.528*** (1.328)	6.703*** (1.559)	8.218*** (1.644)
Bankruptcy	0.188	-0.000 (0.003)	0.001 (0.004)	0.003 (0.005)	-0.002 (0.005)	0.000 (0.002)	0.001 (0.003)	0.004 (0.004)	-0.001 (0.004)
Any Mortgage	0.219	-0.004 (0.006)	0.011 (0.007)	0.017** (0.009)	0.018** (0.009)	-0.003 (0.005)	0.012** (0.006)	0.019** (0.007)	0.020** (0.008)
Any Auto Trade	0.717	0.004 (0.008)	0.007 (0.010)	0.009 (0.011)	0.018 (0.012)	-0.001 (0.006)	0.002 (0.009)	0.004 (0.010)	0.013 (0.011)

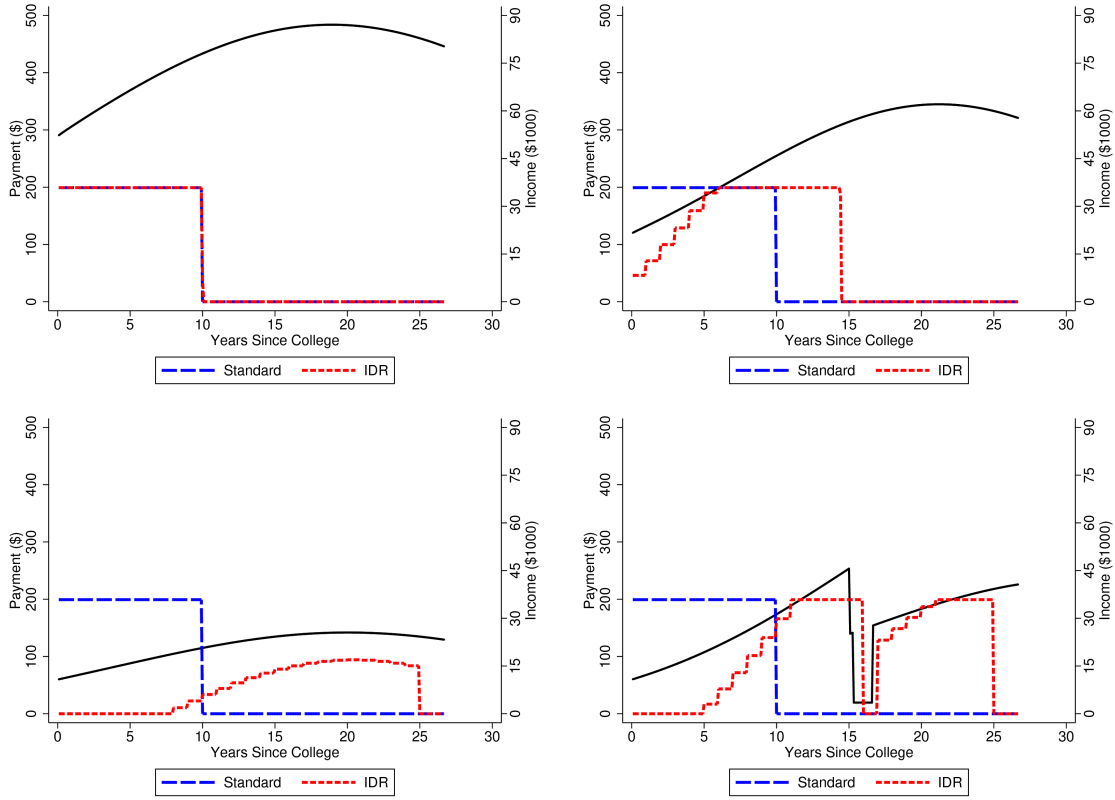
Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on yearly financial outcomes. Column 1 reports the dependent variable mean in the year prior to receiving a delinquency call. Columns 2-5 report coefficients on the effect of IDR enrollment in consecutive years following the delinquency call from the pooled OLS regression specified in Equation 17. Columns 6 - 9 report coefficients on the same yearly effect for a regression which omits pre-call year dummies and includes a linear time trend. The regressions are estimated on the analysis sample as described in the notes to Table 1. Both specifications include controls from call date and time, as well as amount borrowed, inferred gender, and zip-code median income. Sample size is observations from calls. Robust standard errors, in parentheses, are clustered at the individual level. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 7: Difference-in-Differences Estimates of the Effect of IDR Enrollment on Credit Cards

<i>Dependent Variable</i>	Diff-in-Diff					Diff-in-Diff w/Trend			
	(1) Mean $t = -1$	(2) Year 0	(3) Year 1	(4) Year 2	(5) Year 3	(6) Year 0	(7) Year 1	(8) Year 2	(9) Year 3
Credit Card Balances	1.569	0.108 (0.068)	0.293*** (0.093)	0.304*** (0.116)	0.407*** (0.129)	0.113** (0.056)	0.298*** (0.082)	0.309*** (0.105)	0.412*** (0.119)
Log Credit Card Balances	-2.360	0.033 (0.063)	0.281*** (0.078)	0.241*** (0.088)	0.274*** (0.090)	0.070 (0.050)	0.318*** (0.066)	0.278*** (0.076)	0.311*** (0.078)
Number of Credit Cards	3.138	-0.012 (0.034)	0.074 (0.048)	0.119** (0.058)	0.181*** (0.066)	0.025 (0.031)	0.111** (0.046)	0.155*** (0.057)	0.218*** (0.065)
Credit Card Limits	4.926	0.020 (0.162)	0.478** (0.215)	0.774*** (0.283)	0.980*** (0.323)	0.152 (0.135)	0.610*** (0.195)	0.906*** (0.264)	1.111*** (0.306)

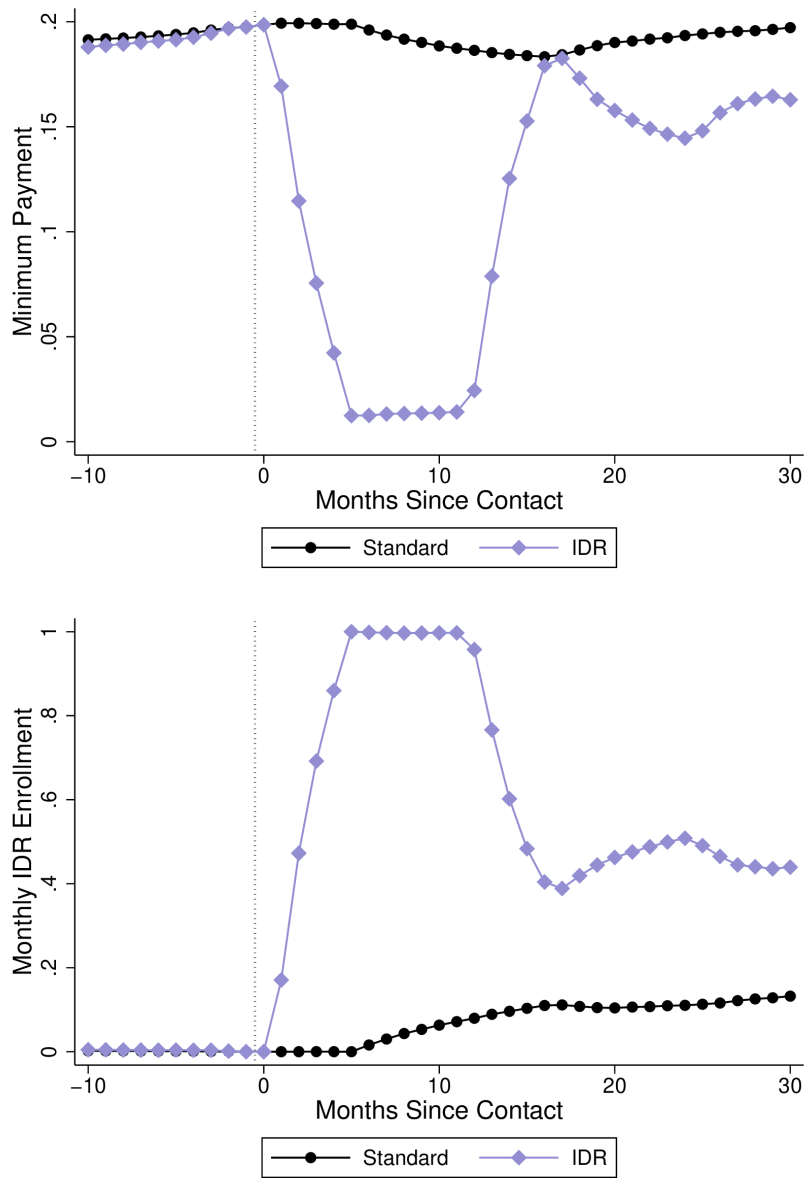
Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on yearly credit card outcomes. Column 1 reports the dependent variable mean in the year prior to receiving a delinquency call. Columns 2-5 report coefficients on the effect of IDR enrollment in consecutive years following the delinquency call from the pooled OLS regression specified in Equation 17. Columns 6 - 9 report coefficients on the same yearly effect for a regression which omits pre-call year dummies and includes a linear time trend. The regressions are estimated on the analysis sample as described in the notes to Table 1. Both specifications include controls from call date and time, as well as amount borrowed, inferred gender, and zip-code median income. Sample size is observations from calls. Robust standard errors, in parentheses, are clustered at the individual level. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Figure 1: Hypothetical Repayment Scenarios: IDR versus Standard Repayment



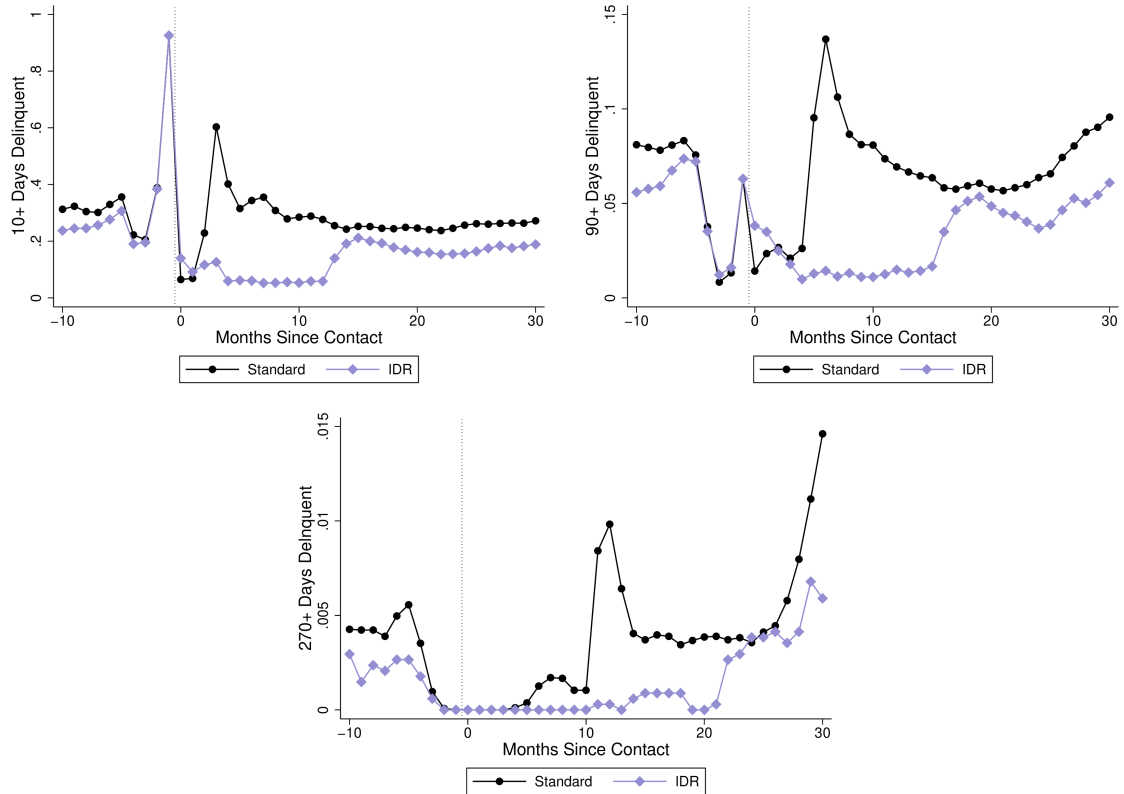
Note: This figure plots hypothetical repayment paths for standard and IDR plans under various income scenarios. Each panel represents an alternative income/repayment scenario for a borrower holding \$18,000 of student debt at the time she leaves college. The solid black line, plotted against the right axis, represents annual post-college income. The dashed blue and dotted red lines, plotted against the left axis, represent monthly payments under standard and IDR plans, respectively. The x-axis denotes years since leaving college. Repayment paths assume a 6.0 percent interest rate and no late payments.

Figure 2: Pre/Post-Call Trends in Minimum Payments and IDR Enrollment




Note: This figure plots the average monthly minimum payments and monthly IDR enrollment status for treatment and control borrowers in the analysis sample. The horizontal axis denotes time, in months, relative to the month of the loan servicing call. Outcomes are normalized to the average value for control borrowers in the month prior to the call. See Table 1 notes for additional details on the outcome measures and sample.

Figure 3: Pre/Post-Call Trends in Delinquencies



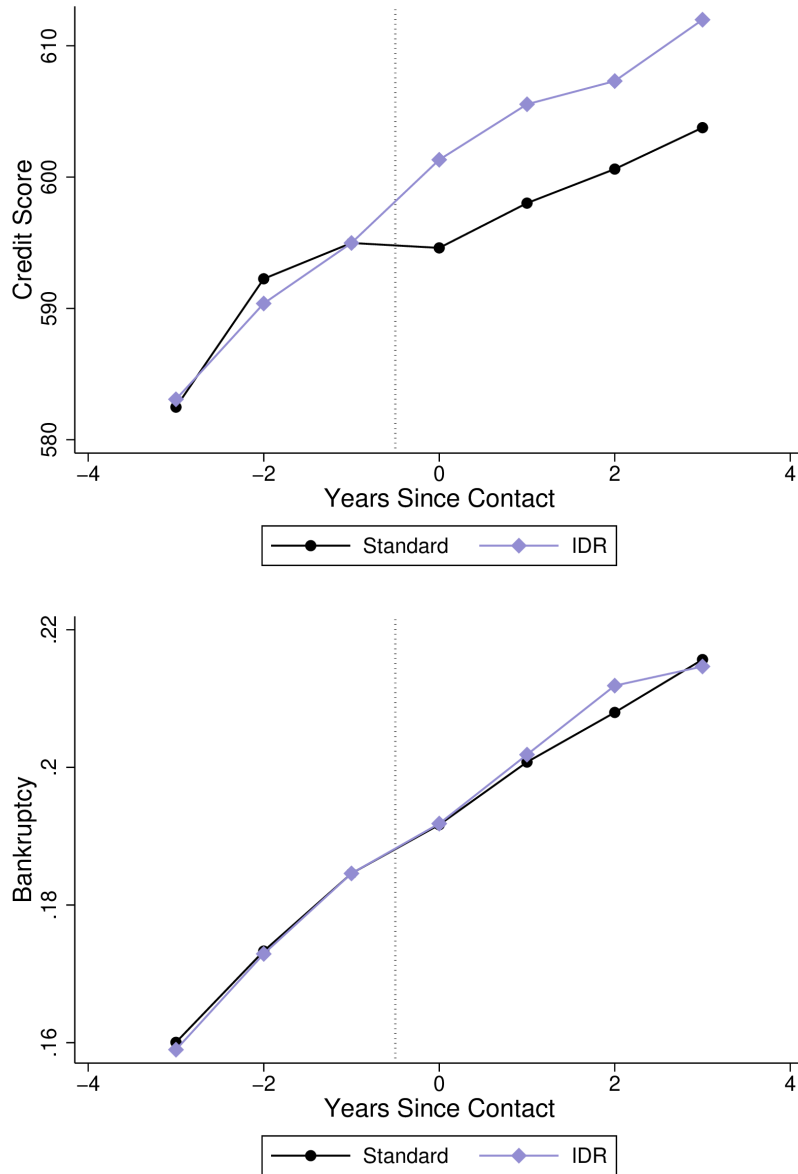
Note: This figure plots the shares of treatment and control borrowers more than 10, more than 90, and more than 270 days delinquent in the analysis sample. The horizontal axis denotes time, in months, relative to the month of the loan servicing call. Outcomes are normalized to the share of delinquent control borrowers in the month prior to the call. See Table 1 notes for additional details on the outcome measures and sample.

Figure 4: Pre/Post-Call Trends in Balances




Note: This figure plots the average total student loan balances and monthly changes in student loan balances, in thousands of dollars, for treatment and control borrowers in the analysis sample. The horizontal axis denotes time, in months, relative to the month of the loan servicing call. Outcomes are normalized to the average value for control borrowers in the month prior to the call. See Table 1 notes for additional details on the outcome measures and sample.

Figure 5: Pre/Post-Call Trends in Credit Scores and Bankruptcies



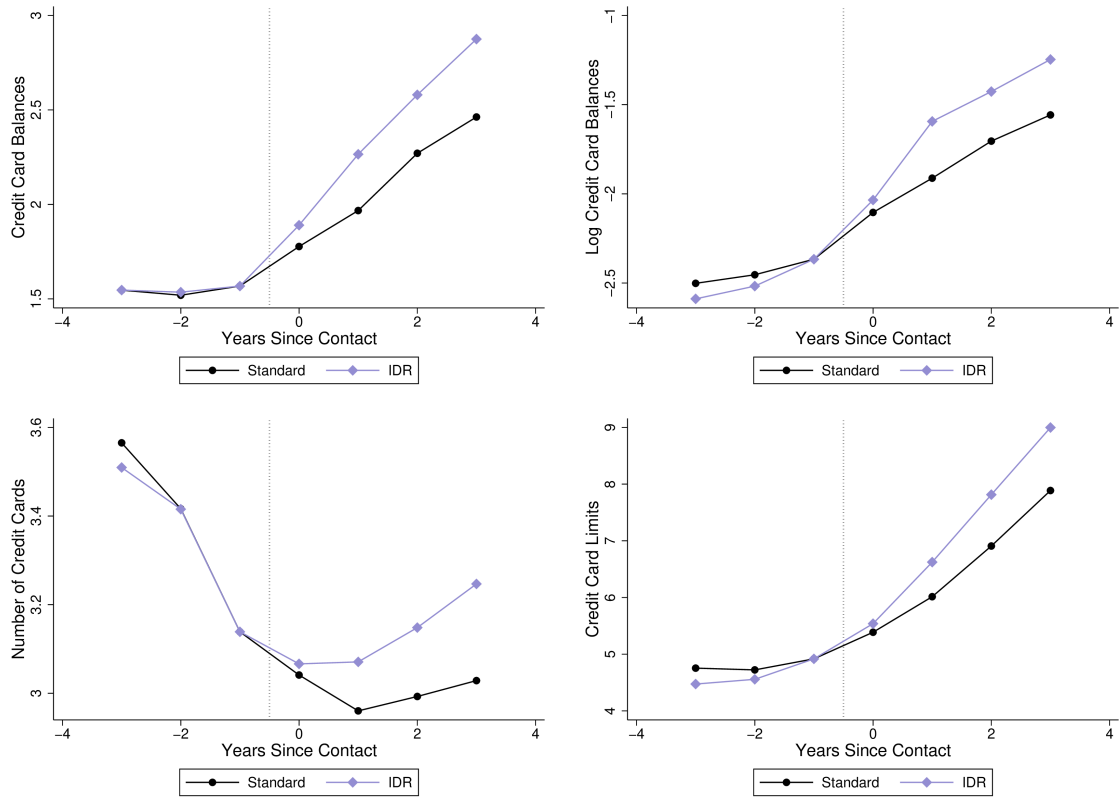
Note: This figure plots the average credit scores and bankruptcies for treatment and control borrowers in the analysis sample. The horizontal axis denotes time, in years, relative to the year of the loan servicing call. Outcomes are normalized to the average value for control borrowers in the year prior to the call. See Table 1 notes for additional details on the outcome measures and sample.

Figure 6: Pre/Post-Call Trends in Mortgages and Auto Loans




Note: This figure plots the shares of treatment and control borrowers holding mortgages and auto loans in the analysis sample. The horizontal axis denotes time, in years, relative to the year of the loan servicing call. Outcomes are normalized to the average value for control borrowers in the year prior to the call. See Table 1 notes for additional details on the outcome measures and sample.

Figure 7: Pre/Post-Call Trends in Credit Cards



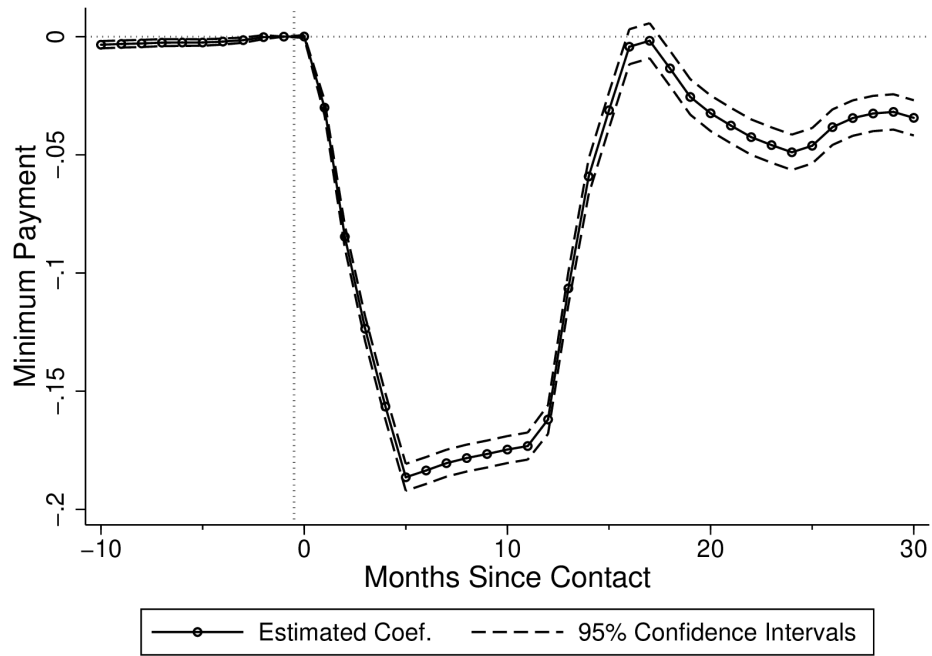
Note: This figure plots the average credit card balances, number of credit cards, and total credit card limits for treatment and control borrowers in the analysis sample. The horizontal axis denotes time, in years, relative to the year of the loan servicing call. Outcomes are normalized to the average value for control borrowers in the year prior to the call. See Table 1 notes for additional details on the outcome measures and sample.

Figure 8: First Stage



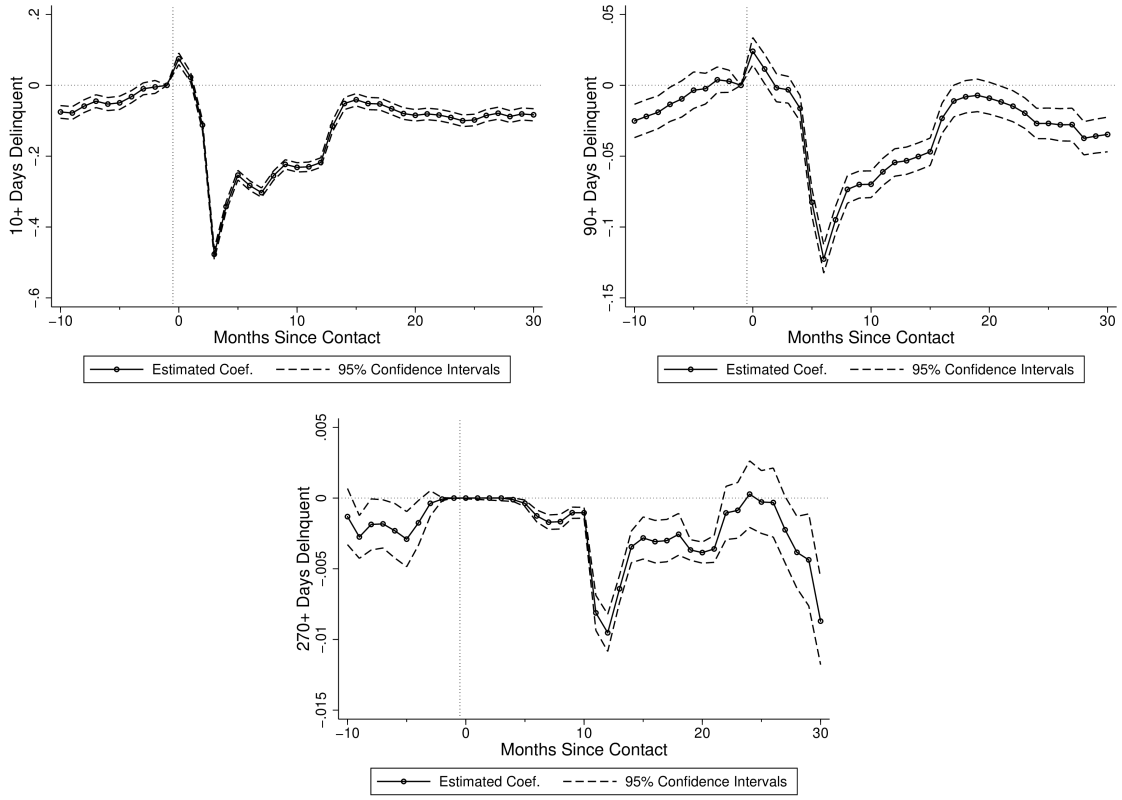
Note: This figure reports first-stage effects and distribution of agent scores across delinquency calls, where agent score is the leave-out mean IDR take-up calculated using data from other calls made by the agent following the procedure described in Section 5. The solid and dashed lines, plotted against the right axis, represent predicted means with 95% confidence intervals from a local linear regression of residualized IDR take-up on agent score. The histogram, plotted against the left axis, provides the distribution of agent scores across all delinquency calls in my analysis sample. All regressions include the full set of call date and time fixed effects.

Figure 9: Difference-in-Differences: Minimum Payments and IDR Enrollment



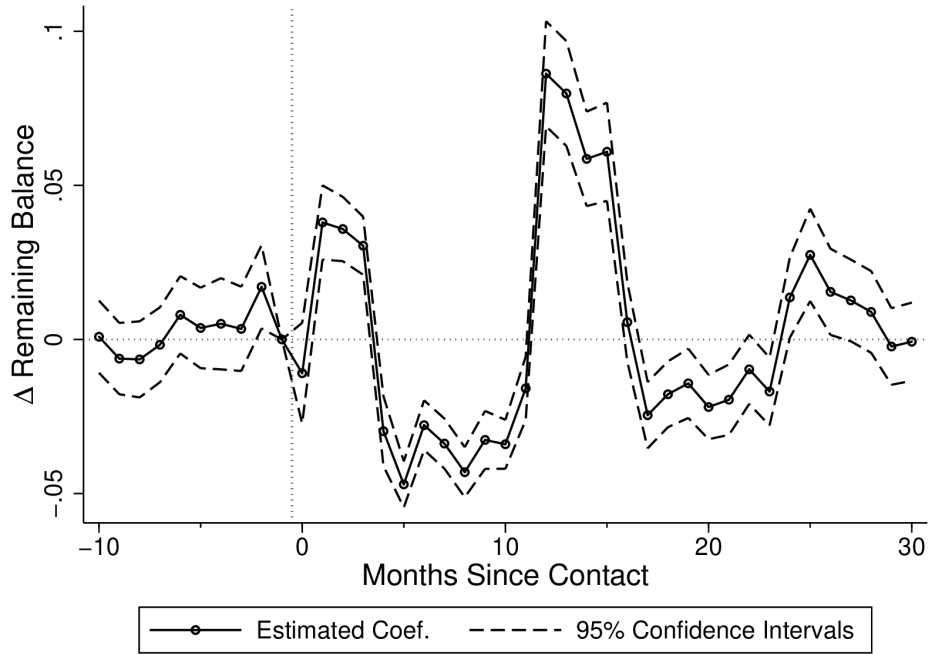
Note: This figure reports estimated coefficients from the non-parametric difference-in-differences specification described in equation 17. Each point represents the estimated effect of post-call IDR status on minimum monthly payment at a given time period relative to the date of delinquency call. Relative months are plotted along the x-axis. Dashed lines represent 95% confidence intervals using robust standard errors clustered at the borrower level. All regressions include individual and call date/time fixed effects.

Figure 10: Difference-in-Differences: Delinquencies



Note: This figure reports estimated coefficients from the non-parametric difference-in-differences specification described in equation 17. Each point represents the estimated effect of post-call IDR status on the likelihood of being more than 10, more than 90, and more than 270 days delinquent at a given time period relative to the date of delinquency call. Relative months are plotted along the x-axis. Dashed lines represent 95% confidence intervals using robust standard errors clustered at the borrower level. All regressions include individual and call date/time fixed effects.

Figure 11: Difference-in-Differences: Balances



Note: This figure reports estimated coefficients from the non-parametric difference-in-differences specification described in equation 17. Each point represents the estimated effect of post-call IDR status on borrowers' month-to-month change in debt balances at a given time period relative to the date of delinquency call. Relative months are plotted along the x-axis. Dashed lines represent 95% confidence intervals using robust standard errors clustered at the borrower level. All regressions include individual and call date/time fixed effects.

Appendix A Additional Tables and Figures

Table A1: Summary Statistics: LLS & Nationally Representative Sample

	2008 Bacchalaureate Recipients	
	(1) LLS	(2) B&B
Female	0.598	0.601
Zip Median Income	60.00	60.42
Age	31.99	29.42
Amount Borrowed	19.17	19.79
Minimum Payment	0.170	0.184
Any Mortgage	0.252	0.327
Mortgage Balances	47.29	47.14
<i>N</i>	318828	8730

Note: This table reports summary statistics at the borrower level. The LLS sample (Column 1) is a random sample of the population of borrowers in LLS's FFEL portfolio who graduated in 2008 and made any loan payments from 2010 onward. The B&B sample (Column 2) consists of all student borrowers in the 2008/2012 Baccalaureate and Beyond Longitudinal Study—a separate, nationally representative dataset of four-year college graduates in 2008. B&B data are derived from FAFSA records, the National Student Loan Database System (NSLDS), and survey responses. Variable definitions follow those from Table 1. Number of observations for the B&B sample are rounded to the nearest ten.

Table A2: Placebo Difference-in-Differences Estimates: Repayment Outcomes

<i>Dependent Variable</i>	(1) Mean $t = -1$	Diff-in-Diff	Diff-in-Diff w/Trend
		(2) Mo. 3-9	(3) Mo. 3-9
Minimum Payment	0.216	0.001 (0.001)	0.002** (0.001)
Remaining Balance	26.461	0.105*** (0.029)	0.128*** (0.032)
Δ Remaining Balance	0.016	-0.013 (0.009)	-0.001 (0.007)
10+ Days Delinquent	0.811	-0.078*** (0.014)	-0.061*** (0.011)
90+ Days Delinquent	0.049	-0.013* (0.007)	-0.014** (0.006)
270+ Days Delnquent	0.000	0.001 (0.001)	-0.001*** (0.000)

Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on monthly loan repayment outcomes for only non-IDR calls. Column 1 reports the dependent variable mean in the month prior to receiving a delinquency call. Column 2 reports coefficients on the effect of IDR enrollment in months three through nine from the pooled OLS regression specified in Equation 17. Column 3 reports coefficients for the same time period for a regression which omits pre-call month dummies and includes a linear time trend. The regressions are estimated on the analysis sample as described in the notes to Table 1. Both specifications include controls from call date and time, as well as amount borrowed, inferred gender, and zip-code median income. Sample size is 37,226 observations from 50,033 calls. Robust standard errors, in parentheses, are clustered at the individual level. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table A3: Placebo Difference-in-Differences Estimates: Financial Outcomes

<i>Dependent Variable</i>	Diff-in-Diff	
	(1) Mean $t = -1$	(2) Year 0
Credit Score	593.434	1.178 (1.693)
Bankruptcy	0.185	0.006 (0.005)
Any Mortgage	0.207	0.003 (0.009)
Any Auto Trade	0.711	0.019* (0.010)

Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on yearly financial outcomes for only non-IDR calls. Column 1 reports the dependent variable mean in the year prior to receiving a delinquency call. Column 2 reports coefficients on the effect of IDR enrollment the year of the delinquency call from the pooled OLS regression specified in Equation 17. The regressions are estimated on the subsample of calls for which no borrowers enrolled in IDR during the call-panel window, but some borrowers eventually enrolled following later delinquency calls. Regression includes controls from call date and time, as well as amount borrowed, inferred gender, and zip-code median income. Sample size is 37,226 observations from 18,613 calls. Robust standard errors, in parentheses, are clustered at the individual level. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table A4: Placebo Difference-in-Differences Estimates: Credit Cards

<i>Dependent Variable</i>	Diff-in-Diff	
	(1) Mean $t = -1$	(2) Year 0
Credit Card Balances	1.568	0.177** (0.072)
Log Credit Card Balances	-2.382	0.083 (0.076)
Number of Credit Cards	2.948	0.015 (0.049)
Credit Card Limits	4.874	0.239 (0.147)

Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on yearly financial outcomes for only non-IDR calls. Column 1 reports the dependent variable mean in the year prior to receiving a delinquency call. Column 2 reports coefficients on the effect of IDR enrollment the year of the delinquency call from the pooled OLS regression specified in Equation 17. The regressions are estimated on the subsample of calls for which no borrowers enrolled in IDR during the call-panel window, but some borrowers eventually enrolled following later delinquency calls. Regression includes controls from call date and time, as well as amount borrowed, inferred gender, and zip-code median income. Sample size is 37,226 observations from 18,613 calls. Robust standard errors, in parentheses, are clustered at the individual level. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table A5: First Stage by Subgroup

	Gender		Age		Amount Borrowed		Credit Score	
	(1) IDR	(2) IDR	(3) IDR	(4) IDR	(5) IDR	(6) IDR	(7) IDR	(8) IDR
Agent Score	0.672*** (0.102)	0.230** (0.107)	0.648*** (0.127)	0.457*** (0.095)	0.619** (0.257)	0.548*** (0.076)	0.434*** (0.122)	0.635*** (0.088)
Subsample	Women	Men	> 40	≤ 40	> 50K	≤ 50K	> 600	≤ 600
Mean Dep.	0.116	0.100	0.114	0.109	0.128	0.110	0.125	0.102
Controls?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-stat	43.64	4.60	26.15	23.35	5.81	51.56	12.65	51.96
P-value	0.0000	0.0332	0.0000	0.0000	0.0169	0.0000	0.0005	0.0000
R-squared	0.020	0.029	0.026	0.021	0.055	0.020	0.026	0.017
N	21841	8518	14674	15685	2187	28171	12616	17743

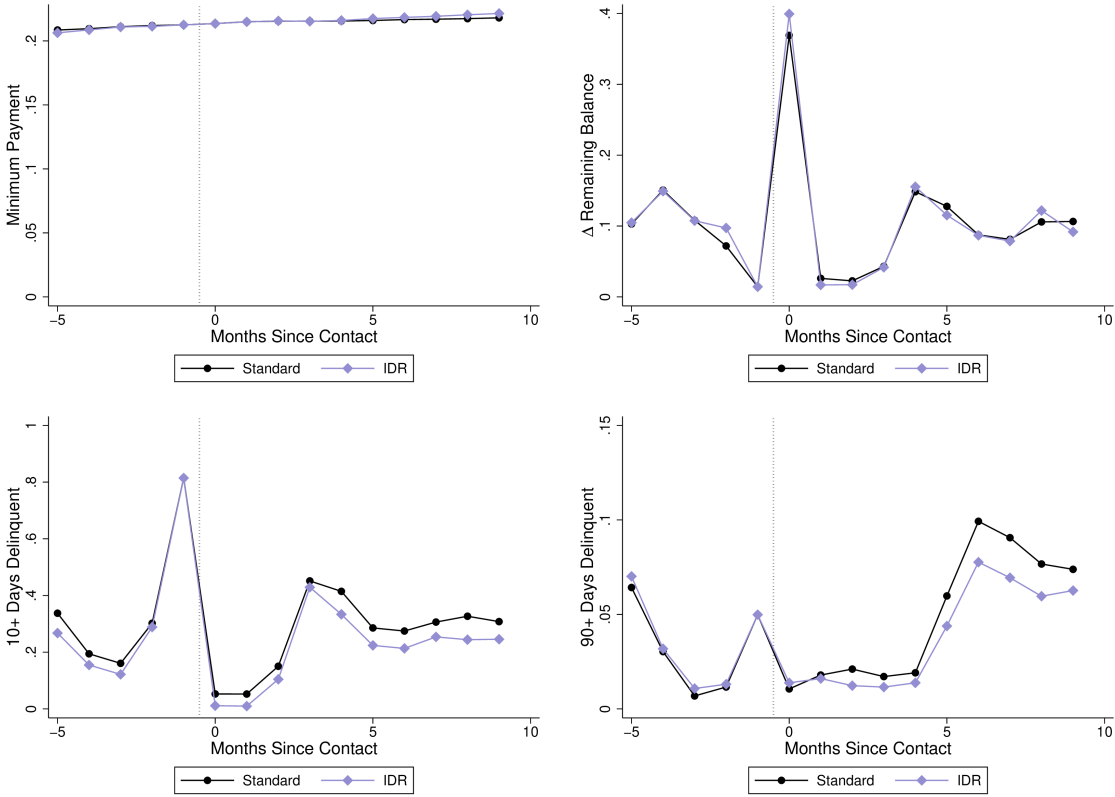
Note: This table reports first-stage results by subgroup. The regressions are estimated on subsamples defined by applying the criteria in the “Subsample” row to the analysis sample described in the notes to Table 1. Agent score is estimated using data from all other phone calls placed by the same agent following the procedure described in Section 5. IDR is an indicator for IDR take-up as of six months after the call. Robust standard errors two-way clustered at the borrower and agent level are reported in parentheses. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table A6: Difference-in-Differences Estimates of the Effect of IDR Enrollment on Repayment Outcomes: Predicted Non-Zero Payments

<i>Dependent Variable</i>	(1) Mean $t = -1$	Diff-in-Diff				Diff-in-Diff w/Trend			
		(2) Mo. 3-9	(3) Mo. 10-16	(4) Mo. 17-23	(5) Mo. 24-30	(6) Mo. 3-9	(7) Mo. 10-16	(8) Mo. 17-23	(9) Mo. 24-30
Minimum Payment	0.570	-0.346*** (0.015)	-0.211*** (0.017)	-0.073*** (0.023)	-0.089*** (0.023)	-0.348*** (0.015)	-0.213*** (0.017)	-0.075*** (0.023)	-0.091*** (0.023)
Remaining Balance	58.163	-0.432*** (0.062)	-0.992*** (0.104)	-1.209*** (0.170)	-1.485*** (0.232)	-0.423*** (0.069)	-0.984*** (0.113)	-1.200*** (0.180)	-1.476*** (0.243)
Δ Remaining Balance	0.035	-0.129*** (0.017)	-0.000 (0.020)	-0.100*** (0.019)	-0.037* (0.020)	-0.091*** (0.017)	0.038* (0.021)	-0.061*** (0.019)	0.001 (0.021)
10+ Days Delinquent	0.925	-0.263*** (0.022)	-0.143*** (0.022)	-0.107*** (0.023)	-0.094*** (0.023)	-0.251*** (0.021)	-0.132*** (0.020)	-0.095*** (0.021)	-0.082*** (0.023)
90+ Days Delinquent	0.060	-0.052*** (0.011)	-0.043*** (0.011)	-0.022* (0.012)	-0.024* (0.014)	-0.036** (0.014)	-0.027* (0.015)	-0.006 (0.015)	-0.008 (0.017)
270+ Days Delnquent	0.000	-0.002 (0.002)	-0.003 (0.002)	-0.003 (0.002)	-0.001 (0.004)	-0.001** (0.000)	-0.001* (0.001)	-0.001 (0.001)	0.000 (0.003)

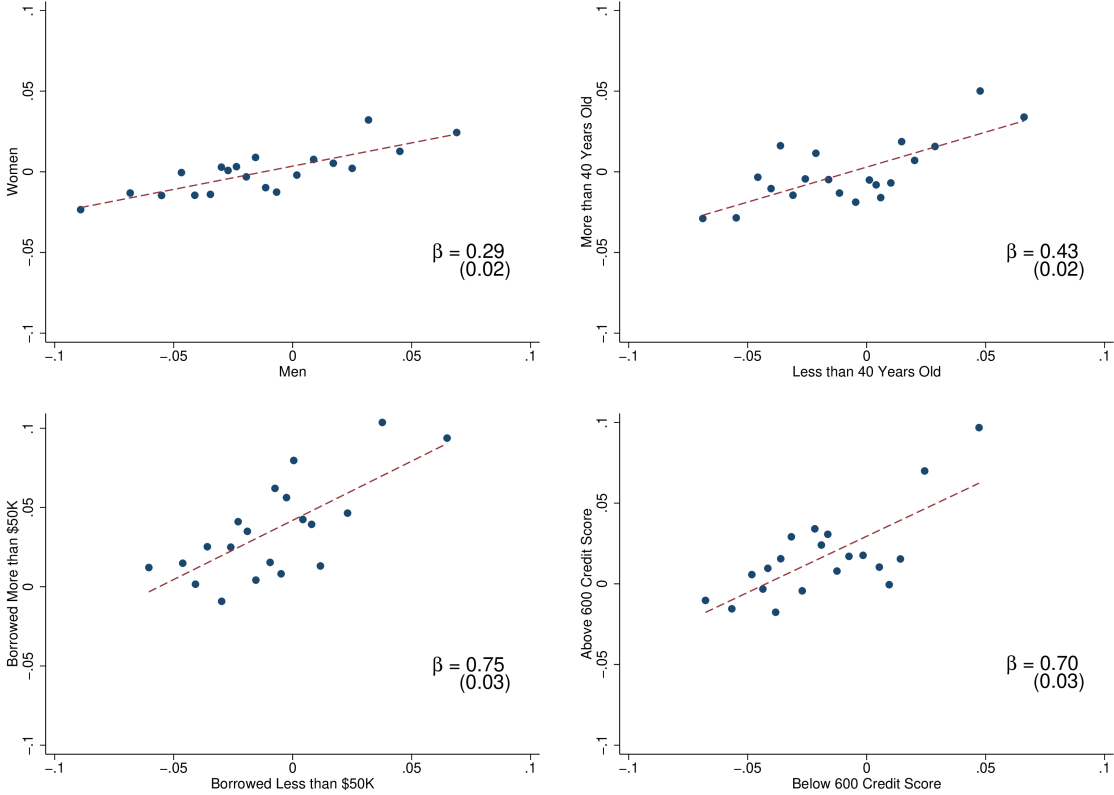
Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on monthly loan repayment outcomes for those predicted to have non-zero IDR payments. Column 1 reports the dependent variable mean in the month prior to receiving a delinquency call. Columns 2-5 report coefficients on the effect of IDR enrollment in consecutive consecutive seven-month periods following the delinquency call from the pooled OLS regression specified in Equation 17. Each of Columns 6 - 9 report coefficients on the same monthly effect for a regression which omits pre-call month dummies and includes a linear time trend. The regressions are estimated on the analysis sample as described in the notes to Table 1. Both specifications include controls from call date and time, as well as amount borrowed, inferred gender, and zip-code median income. Sample size is 12,355 observations from 2,316 calls. Robust standard errors, in parentheses, are clustered at the individual level. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Figure A1: Pre/Post-Call Trends in Repayment Outcomes: Placebo Test



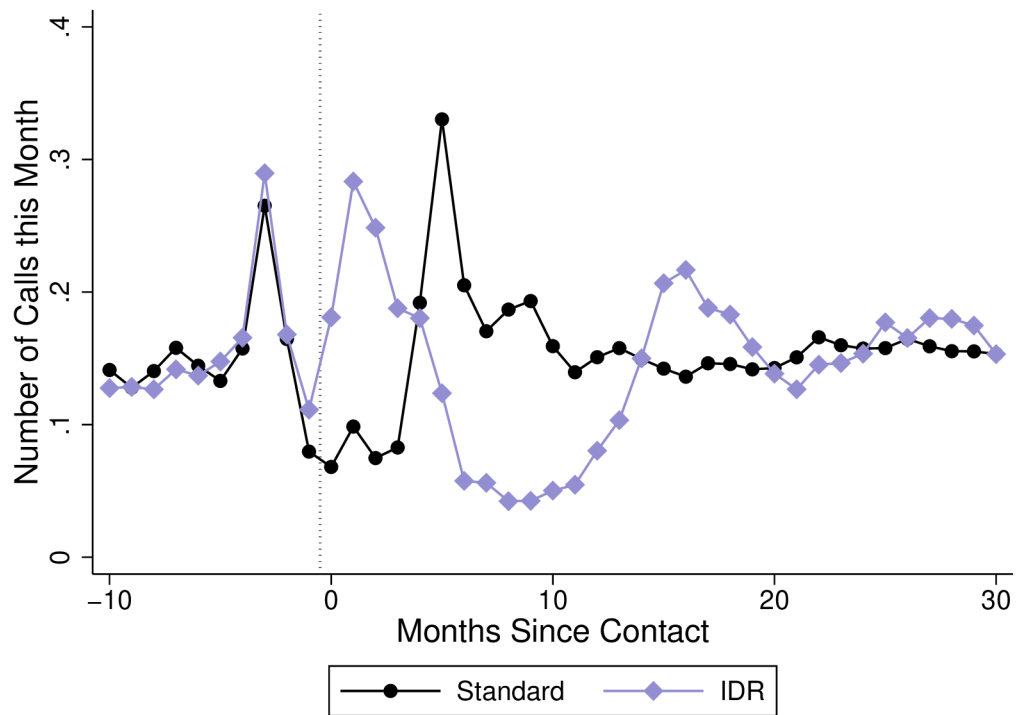
Note: This figure plots the average loan repayment outcomes for *eventual* IDR enrollees following previous delinquency calls that did not end in enrollment versus non-enrollees in the analysis sample. The horizontal axis denotes time, in months, relative to the month of the loan servicing call. Outcomes are normalized to the average value of the outcome for non-enrollees in the month prior to the call. See Table 1 notes for additional details on the outcome measures and sample.

Figure A2: Group-Specific Instrument Correlations



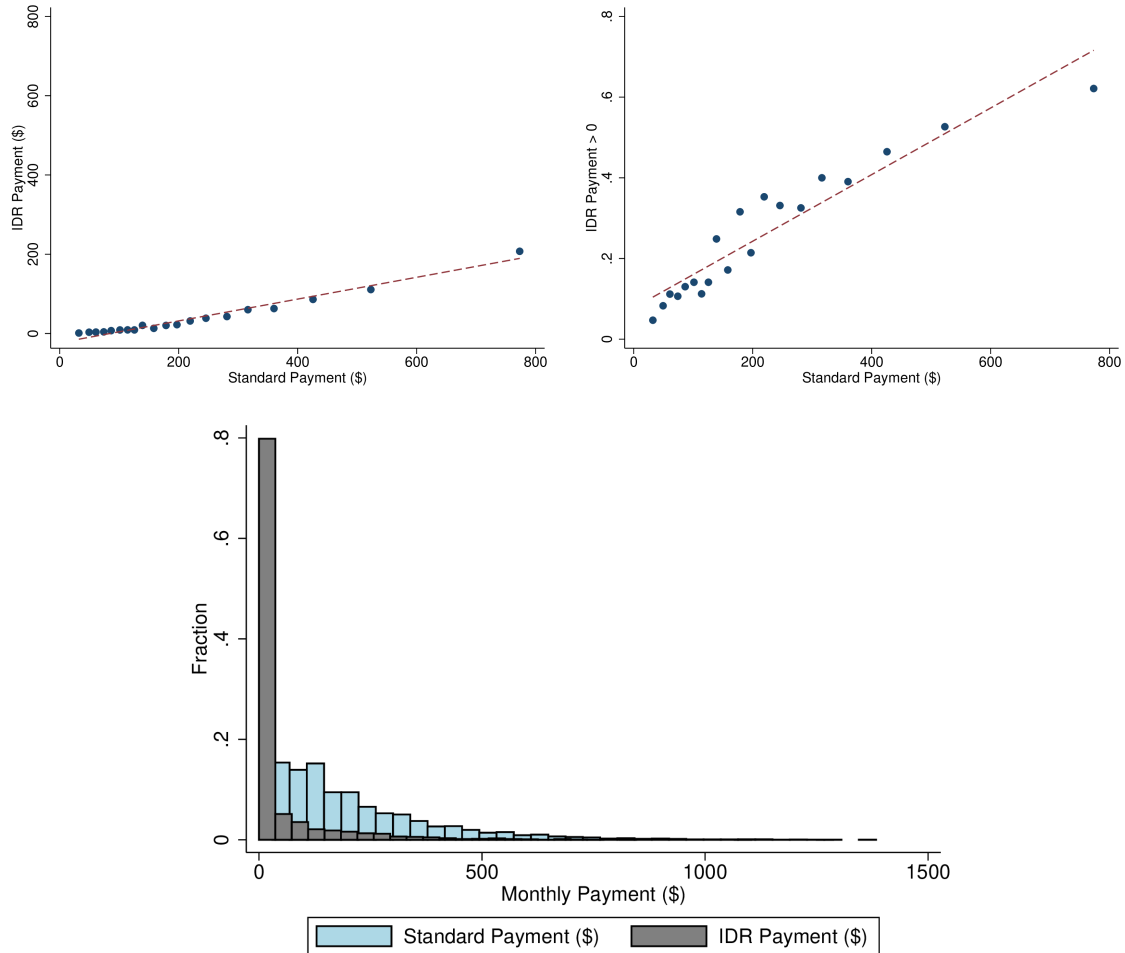
Note: This figure plots binned correlations between group-specific instruments Z_{icj}^g . Each axis measures the residualized, leave-one-out propensity of every call’s assigned agent to induce IDR take-up among individuals in the group specified by the axis label. I also plot the linear best fit line estimated using OLS and report the associated coefficients and standard errors in the upper left corner of each panel.

Figure A3: Pre/Post-Call Points of Contact



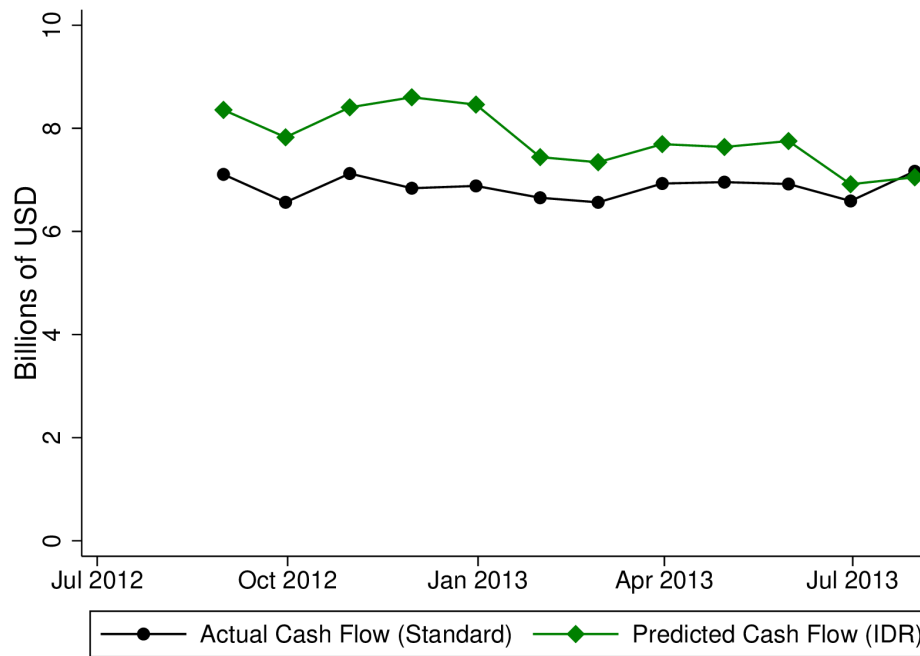
Note: This figure plots the average monthly points of contact (incoming calls, outgoing calls, and web chats) between borrowers and LLS for IDR enrollees and non-enrollees in the analysis sample. The horizontal axis denotes time, in months, relative to the month of the loan servicing call. See Table 1 notes for additional details on the outcome measures and sample.

Figure A4: Standard versus IDR Payments among IDR Enrollees



Note: This figure plots the relationship between pre-call standard payments and post-call IDR payments. The binned scatter plot is constructed using payment amounts one month before and six months following the delinquency call for borrowers in the analysis sample who take up IDR. The top-left panel plots average standard payment size against average IDR payment size. The top-right panel plots average standard payment size against the share of individuals with IDR payments greater than zero. The bottom panel plots histograms for standard and IDR payments. See Table 1 notes for additional details on the sample.

Figure A5: Predicted Cash Flows to Government Under IDR



Note: This figure plots actual total cash flows versus predicted total cash flows for the counterfactual scenario in which all student borrowers enrolled in IDR in July 2012. Predictions are generated using monthly difference-in-difference estimates for the analysis sample re-weighted so that the joint distribution of pre-call observables matches that of the full representative sample from Table 1. Values are scaled to reflect total national student loan balances as of June 2012.

Appendix B Mathematical Appendix

Model Solution

The agent's problem is given by:

$$V(y_t, D_t, A_t) = \max_{A_{t+1}, c_t, q \in \{0,1\}} \{u(c_t) - \varphi(1 - q) + E[V(y_{t+1}, D_{t+1}, A_{t+1})]\} \quad (23)$$

$$A_{t+1} \geq L \quad (24)$$

$$c_t = A_t - A_{t+1} + y_t - q_t x_t \quad (25)$$

$$D_{t+1} = D_t - q_t x_t \quad (26)$$

$$A_0 = \bar{A}_0 \quad (27)$$

$$D_0 = \bar{D}_0 \quad (28)$$

$$\lim_{t \rightarrow \infty} \prod_{i=0}^{t-1} A_i \geq 0 \quad (29)$$

I begin by reformulating the Bellman Equation 23. Let $V^r(y_t, D_t, A_t)$ denote the value function in the case where the borrower repays ($q = 1$), and let $V^n(y_t, D_t, A_t)$ denote the value function when she does not repay ($q = 0$):

$$V^r(y_t, D_t, A_t) = \max_{A_{t+1} \geq L} \{u(A_t - A_{t+1} + y_t - x_t) + E[V(y_{t+1}, D_t - x_t, A_{t+1})]\} \quad (30)$$

$$V^n(y_t, D_t, A_t) = \max_{A_{t+1} \geq L} \{u(A_t - A_{t+1} + y_t) - \varphi + E[V(y_{t+1}, D_t, A_{t+1})]\}. \quad (31)$$

I can then rewrite Equation 23 as

$$V(y_t, D_t, A_t) = \max \{V^r(y_t, D_t, A_t), V^n(y_t, D_t, A_t)\}. \quad (32)$$

Repayment Decision The borrower's optimal repayment decision is determined by comparing value of repayment versus non-repayment:

$$q_t = \mathbf{1} \{V^r(y_t, D_t, A_t) > V^n(y_t, D_t, A_t)\} \quad (33)$$

Note from Equation 31 that non-repayment simply moves today's payment by one period at cost φ , so that Equation 33 can be written as

$$q_t = \mathbf{1} \{V^r(y_t, D_t, A_t) > V^r(y_t + x, D_t + x, A_t) - \varphi\} \quad (34)$$

Since $u'(\cdot) > 0$ and $u''(\cdot) < 0$, the envelope theorem implies both left and right sides of the inequality in Equation 33 are strictly increasing and concave in y_t , and thus satisfy the single-crossing property. The repayment decision can therefore be characterized with the cutoff value ψ that satisfies the following:

$$V^r(\psi, D_t, A_t) = V^r(\psi + x, D_t + x, A_t) - \varphi, \quad (35)$$

so that

$$q_t = \mathbf{1} \{y_t > \psi(D_t, A_t)\}, \quad (36)$$

Consumption Decision Solving Equation 30 for optimal assets under repayment, A_{t+1}^r , the Lagrangian can be written as:

$$\mathcal{L} = u(A_t - A_{t+1} + y_t - x_t) + E[V(y_{t+1}, D_t - x_t, A_{t+1})] - \lambda(L - A_{t+1})$$

where $\lambda \geq 0$. The first-order condition for next period's assets A_{t+1}^r is:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial A_{t+1}}(y_{t+1}, D_{t+1}, A_{t+1}^r, A_t) &= 0 \\ \frac{dE[V(y_{t+1}, D_{t+1}, A_{t+1}^r)]}{dA_{t+1}^r} &= -\lambda + u'(A_t - A_{t+1}^r + y_t - x_t) \end{aligned} \quad (37)$$

Applying the envelope theorem to Equation 37 yields:

$$u'(A_t - A_{t+1}^r + y_t - x_t) = E \left[\frac{\partial V}{\partial A}(y_{t+1}, D_{t+1}, A_{t+1}^r) \right] + \lambda_1 \quad (38)$$

Since $\lambda > 0$ if and only if $A_{t+1}^r = L$, Equation 38 implies:

$$u'(A_t - A_{t+1}^r + y_t - x_t) \geq E \left[\frac{\partial V}{\partial A}(y_{t+1}, D_{t+1}, A_{t+1}^r) \right], \text{ with "=" if } A_{t+1}^r > L \quad (39)$$

Following the same process for the $q = 0$ case, solving for optimal saving in Equation 31 yields

$$u'(A_t - A_{t+1}^n + y_t) \geq E \left[\frac{\partial V}{\partial A}(y_{t+1}, D_t, A_{t+1}^n) \right], \text{ with "=" if } A_{t+1}^n > L, \quad (40)$$

so that, in general,

$$u'(c_t) \geq E \left[\frac{\partial V}{\partial A}(y_{t+1}, A_{t+1}, D_{t+1}) \right] \text{ with "=" if } A_{t+1} > L, \quad (41)$$

$$(42)$$

which, along with a conditions 24, 25, 26, and 36 characterize the solution.⁴⁸

Proof of Liquidity Effects

Note that the solution from the previous section implies:

$$c_t = \begin{cases} E[c_{t+1}] & \text{if } A_{t+1} > L \\ A_t - L + y_t - q_t x_t & \text{if } A_{t+1} = L \end{cases} \quad (43)$$

Consider the solution $c_t^* = E[c_{t+1}]$ when the borrowing constraint does not bind. In this case $q_t = 1$ and $A_{t+1}^* > L$, so

$$A_t + y_t - x_t - c_t^* > L \quad (44)$$

$$c_t^* < A_t - L + y_t - x_t \quad (45)$$

$$E[c_{t+1}] < A_t - L + y_t - x_t \quad (46)$$

⁴⁸Technically, the borrower must also satisfy a transversality condition, given by $\lim_{t \rightarrow \infty} u'(c_{t-1}) A_t = 0$. See Ljungqvist and Sargent (2012) for a proof of the condition.

We can therefore rewrite (43) as:

$$c_t = \min \{A_t - L + y_t - q_t x_t, E[c_{t+1}]\} \quad (47)$$

Iterating forward, we have

$$c_t = \min \{A_t - L + y_t - q_t x_t, E[\min \{A_{t+1} - L + y_{t+1} - q_{t+1} x_{t+1}, E[c_{t+2}]\}]\} \quad (48)$$

Smoothing Effect

First I analyze how IDR affects current consumption through a smoothing of the expected net-income profile. Note that for all t , IDR “backloads” minimum payments relative to standard plans:

$$E[x_t^I] = E\left[\min\left\{\theta y_t, \frac{D_0}{N}\right\}\right] \leq \frac{D_0}{N} = E[x_t^S] \quad \forall t \leq N \quad (49)$$

$$E[x_t^I] = E\left[\min\left\{\theta y_t, \frac{D_0}{N}, D_t\right\}\right] \geq 0 = E[x_t^S] \quad \forall t > N \quad (50)$$

Now consider the effects of a decrease in period- k minimum payments x_k on period- t consumption c_t , for all $k \in [t, N]$. Without loss of generality, let $N = t+1$. Period- t consumption increases through two channels: first, lower x_t means the ceiling on present-day consumption is mechanically lifted by Δx_t . Second, borrowers *expect* a higher ceiling on consumption in period $t+1$. As long as period $t+1$ liquidity constraints held some positive probability of binding under the original payment scheme, reducing minimum payments will increase current-period consumption by lowering precautionary savings; borrowers spend down their assets today because they expect less need to supplement their potentially constrained consumption tomorrow. Formally, suppose that $\frac{dc_t}{dx_t} > 0$, which would require the right-hand side of 47 to decrease following a drop in x_t . The budget constraint implies this scenario would only be possible if A_{t+1} increased more than $-\Delta x$, which would, in turn, increase $E[c_{t+1}]$ which would violate Equation 43.

Now consider a commensurate *increase* in period- q minimum payments x_q on period- t consumption c_t , for all $q > N$. Assume that the period- t borrowing constraint does not bind

and expand Equation 48 an additional period, so:

$$c_t = E [\min \{A_{t+1} - L + y_{t+1} - x_{t+1}, E [\min \{A_{t+2} - L + y_{t+2} - x_{t+2}, E [c_{t+3}]\}]\}] \quad (51)$$

By the same logic as above, an increase in period $t + 2$ payments x_{t+2} *decreases* period t consumption. However, the net effect of a revenue-neutral “backloading” of payments must be positive. To see why, recall that $E[y_q] \leq E[y_{q'}]$ for all $q < q'$, and compare the effects of payment changes in different periods on period- t consumption in Equation 51:

$$\frac{dc_t}{dx_q} \geq \frac{dc_t}{dx_{q'}} \quad \forall q' > q \quad (52)$$

Any change in payments such that $\Delta x_q < \Delta x_{q'}$ but $\sum_t \Delta x_q = 0$ implies a net *increase* in c_t . Intuitively, the precautionary response to net income changes in the near future is greater than the response to net income changes in the distant future, as borrowers can gradually accumulate precautionary savings for the latter over many periods.

Insurance Effect

In addition to flattening the net-income profiles, IDR also reduces the per-period variance of net income. Let m denote income net of loan payments, $m \equiv y - x$.

$$\text{Var}(m^I) - \text{Var}(m^S) = \int_{-\infty}^{\infty} (y - x^I)^2 dF_y - \int_{-\infty}^{\infty} (y - x^S)^2 dF_y \quad (53)$$

$$= \int_{-\infty}^{\infty} (y - \max \{(1 - \theta)y, y - x_t^S\})^2 dF_y - \int_{-\infty}^{\infty} (y - x^S)^2 dF_y \quad (54)$$

$$= \int_{-\infty}^{\frac{x^S}{\theta}} [(y - \theta y)^2 - (y - x^I)^2] dF_y < 0 \quad (55)$$

Rewriting Equation 48 in terms of net income, we have:

$$c_t = \min \{A_t - L + m_t, E [\min \{A_{t+1} - L + m_{t+1}, E [c_{t+2}]\}]\} \quad (56)$$

Note that for any period $k > t$, a mean-preserving contraction in m_t will decrease the likelihood of low realizations of net income and, hence, a binding liquidity constraint; a decrease in $\text{Var}(m_{t+1})$ will increase the value of $E[\min\{A_{t+1} - L + m_{t+1}, E[c_{t+2}]\}]$, inducing less precautionary savings and more consumption in period t .

Loan Forgiveness and Wealth Effects

To capture the loan forgiveness provisions of IDR, I make two simple modifications to the model. First, since debt forgiveness is contingent on timely payments, I assume the borrower repays in every period, $q = 1$, and focus on the consumption decision.⁴⁹ Second, I modify the evolution of student debt, D_t , so that borrowers expect zero loan payments after some forgiveness period T . Specifically, in period $t = T + 1$, instead of Equation 6, we have:

$$D_{T+1} = 0. \quad (57)$$

Note that if borrowers expect loan forgiveness, then $\sum_{k=0}^{\infty} \Delta E[x_k] < 0$ and the first term in Equation 15 includes a *wealth effect*. In this case, the expression can be decomposed to differentiate between liquidity and wealth effects:

$$\Delta c_t = \underbrace{-\frac{dc_t}{dD_{T+1}} E[D_T^I - x_T^I]}_{\text{wealth effect (+)}} + \underbrace{\frac{dc_t}{dD_{T+1}} E[D_T^I - x_T^I] + \sum_{k=0}^T \left(-\frac{dc_t}{d\mu_k} \Delta E[x_k] + \frac{dc_t}{d\sigma_k} [\sigma_k - \sigma] \right)}_{\text{liquidity effects (+)}}. \quad (58)$$

⁴⁹Technically, while defaulted loans cannot qualify for forgiveness, it is possible to miss payments and still qualify for IDR forgiveness if the borrower makes up those payments later on.