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

Daniel Herbst*
April 27, 2018

Please click here for most recent version

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. IDR reduces monthly loan payments to a fixed portion of earnings until debt is repaid or some forgiveness period has been reached. By aligning the repayment burden with the returns to college, IDR may prevent default and improve financial well-being among credit-constrained borrowers but carries a potential cost to social welfare through moral hazard. In this paper, I estimate the causal impact of IDR on repayment rates, balances, homeownership, and consumption proxies using a novel dataset linking the first administrative panel of federal student loan payments to credit bureau records for over one million student borrowers. My research design uses two complementary 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. I find evidence of liquidity benefits on both default and consumption margins. Within seven months of take up, IDR enrollees are 21 percentage points less likely to fall delinquent and pay down $90 more student debt each month compared to those who remain on standard repayment plans. IDR enrollees have credit scores that are 7.5 points higher, hold 0.1 more credit cards, and carry $240 higher credit card balances than non-enrollees one year after the servicing call, implying increased short-term consumption out of liquidity. IDR enrollees are also 2 percentage points more likely to hold a mortgage, an increase of 10 percent off of the pre-call mean, suggesting a positive effect of IDR on homeownership. Minimum monthly payments decrease by an average of $140 following IDR take up, but return to standard levels within one year of enrollment, minimizing the potential impact of moral hazard through loan forgiveness. 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, Princeton University, Louis A Simpson Building, Princeton, NJ 08544 (email: dherbst@princeton.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, Alexandre Mas, and Christopher Neilson, who provided invaluable advice throughout its development. I also benefited from the helpful comments of David Arnold, Barbara Biasi, Felipe Goncalves, Steve Mello, David Price, Maria Micaela Sviatschi, and other Princeton seminar participants. Financial support was provided by the Princeton Industrial Relations Section and the National Academy of Education Spencer Dissertation Fellowship.

1 Introduction

Each year, roughly one million borrowers default on their student loans, and millions more struggle to buy homes (Mezza et al. [2016]), accumulate wealth (Bleemer et al. [2017]), or choose their preferred career (Rothstein and Rouse [2011]). Many economists blame the rigid repayment terms on student loans, which require fixed payments averaging $350 per month early in borrowers' careers (Cortés et al. [2017]).[1] By far, the largest and most far-reaching policy response has been "Income-Driven Repayment" (IDR). First introduced in 2009, IDR allows students to repay a fixed portion of their yearly earnings until either their debt is paid off or some forgiveness period has been reached, offering an alternative to the flat, ten-year repayment schedule under traditional repayment plans. Enrollment in IDR has tripled since 2014, and today over $300 billion in debt is repaid through IDR (Department of Education [2017a]).

Even as IDR take-up continues to rise, its effects on borrowers 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. Evidence of high payment-to-income ratios (Looney and Yannelis [2015]) and limited credit (Gross and Souleles [2002]) among young borrowers suggests such liquidity may increase borrowers' welfare, allowing them to pay down their student debt while avoiding financial distress or sharp drops in consumption. However, these short-term liquidity benefits may be outweighed by the long-term costs of IDR's forgiveness provisions, which cancel any unpaid debt after twenty-five years of payments. If enrollees' incomes are *permanently* low, the repayment burden ultimately falls on the government, reducing social welfare through redistributive inefficiencies, administrative costs, and distorted incentives (i.e., moral hazard).

Assessing the costs and benefits of IDR requires causal estimates of its impact on repayment and consumption, but two obstacles have prevented researchers from identifying these effects. First, an empirical analysis of IDR requires high-frequency repayment data for many

[1]For 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."

borrowers, but until now these data have been unavailable.[2] Second, IDR enrollees are by design a very selected group, as only borrowers with a sufficiently high debt-to-earnings ratio would benefit from enrolling in IDR. Estimates which rely on cross-sectional comparisons of IDR enrollees to non-enrollees could therefore be biased in either direction, as low-income individuals typically have worse financial outcomes, but borrowers with high student debt balances are often highly educated and positively selected (Yannelis [2016]).

In this paper, I use novel administrative data from a large student loan servicing company to estimate the causal effect of IDR enrollment on defaults, credit scores, bankruptcies, and consumption proxies. 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 disbursals 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 U.S. federal student loan payments used in public research.

I use two complementary research designs to identify the effect of IDR on student borrowers, both of which exploit plausibly exogenous variation from loan servicing phone calls. 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 of IDR options. 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.

My second empirical strategy is an instrumental variables (IV) design exploiting the

[2]Economists 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]).

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 using the leave-one-out, mean IDR status among the borrowers who receive a call from that particular agent. The agent effectiveness measure is predictive of IDR take-up, but uncorrelated with borrower characteristics and pre-call outcomes. This research design resembles Kling [2006], which uses random judge assignment to estimate the effect of incarceration length on earnings, as well as subsequent research focusing on consumer bankruptcy (Dobbie and Song [2015]), juvenile incarceration (Aizer and Doyle [2015]), pre-trial detention (Dobbie et al. [2016a]), family welfare culture (Dahl et al. [2014]), disability insurance (Kostøl et al. [2017]), and foster care (Autor and Houseman [2010]). While the quasi-experimental nature of the instrumental variables design holds obvious advantages, the increased statistical precision of difference-in-differences estimates makes it my preferred design. Nonetheless, IV estimates provide complementary evidence to the difference-in-differences design, as results are generally consistent between the two strategies.

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 21.0 percentage points within seven months of take up. While IDR reduces monthly minimum payments by an average of $140 through a mechanical "first-stage" effect, the effect of reduced minimums on loan balances is dominated by more timely repayment; IDR borrowers pay down $90 more student debt each month, on average, than those on standard repayment plans. IDR enrollees have credit scores which are 7.5 points higher, hold 0.1 more credit cards, and carry $240 higher credit card balances than non-enrollees one year after the servicing call.[3] While I find no 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 10 percent off of the

[3] For reference, Dobbie et al. [2016b] 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.

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. These results suggest that short-term increases in cash-on-hand through IDR have long-term positive impacts on the financial health of credit-constrained borrowers.

In theory, my estimates combine three effects of IDR: (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. For two reasons, I argue my estimates capture a pure liquidity effect: First, minimum monthly payments decrease by an average of \$140 following IDR take up, but return to standard levels within one year of enrollment, a rate which makes future loan forgiveness unlikely under the terms of the IDR plan I study. 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.[4]

My analysis carries four 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 unravelling effect on a hypothetical "repayment-plan market."[5] 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.[6] My analysis therefore treats loan terms as policy pa-

[4]Note 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.

[5]In fact, the existence of adverse selection may explain why there exists virtually no private market in which to study the phenomenon. In concurrent work, I use private elicitations of expected future income among student borrowers to provide evidence in support of this hypothesis.

[6]A 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.

rameters rather than equilibrium objects. Nonetheless, 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 effect I identify excludes 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, removing ex-ante responses to expected payment flexibility or forgiveness leaves several unanswered questions important for the social welfare implications of IDR.

Third, my analysis is limited to the repayment and consumption-related outcomes available in my data. While these outcomes highlight the primary channel through which IDR is intended to benefit borrowers, IDR undoubtedly also influences borrower welfare through potential effects on labor supply and household formation. My analysis cannot directly speak to these channels, as my data do not include information on earnings, employment, or marital status.

Finally, for several reasons, the outcomes I measure can not perfectly capture welfare improvements through consumption or repayment: (1) The lack of IDR enrollment prior to 2013 makes it impossible to measure effects over long-term time horizons. (2) The credit card, auto loans, and mortgage data I use provide useful proxies for consumption and homeownership (Ganong and Noel [2017], Di Maggio et al. [2016], Keys et al. [2014]), but are not direct measures of these variables. (3) While default and delinquency carry negative consequences for borrowers, the welfare implication of these consequences remains an open question, so estimates of improved repayment rates under IDR are difficult to interpret.[7]

The remainder of this paper is organized as follows. Section 2 provides a brief overview of

[7]Different stages of non-repayment on student loans carry increasingly severe late fees and credit consequences, culminating in an official default classification after 270 days past due. "Default" is a legal distinction, defined as "Failure to repay a loan according to the terms agreed to in the promissory note." Defaulted loans can result in garnished wages, withheld tax returns, and revocation of professional licenses (Department of Education [2017a]).

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.

2 Background

2.1 Overview

Over 90% of student loans in the United States are federally subsidized and guaranteed. The government holds the liability on student loans, and interest rates are set by Congress.[8] Student loans are not secured by collateral or subject to any credit check.[9] While the federal government holds the liability on student loans, the Department of Education contracts private debt servicing companies to handle disbursals, 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.[10]

The Department of Education sets repayment terms for student loans through several repayment plans, each requiring borrowers make monthly payments to their loan servicer. The default repayment option 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

[8] Congress has set rates on student loans since 1965, though automated the process in 2013 with the Bipartisan Student Loan Certainty Act. The legislation mandates a rate-setting rule for student loans, setting 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.

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

[10] In 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.

payments are calculated as:

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

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.[11] 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).[12] Specifically,

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

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.[13]

A graphical comparison of repayment plans can be seen in Figure 1. The figure plots IDR and standard repayment paths under alternative income scenarios for a borrower leaving college with $18,000 in student loans. In panel A, the borrower's income is too high to qualify for reduced payments: her payments are identical under standard (dashed blue line) and IDR (dotted red line) plans. In panel B, the borrower's income is low enough to

[11] Since 2009, several IDR plans have become available, including Income-Based Repayment (IBR), Pay-As-You-Earn (PAYE), and Revised-Pay-As-You-Earn (REPAYE). The eligibility requirements and repayment terms can vary across these IDR 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 Federal Family Education Loan (FFEL) borrowers in my sample are eligible, though the discussion and conclusion generalize to the broader concept of IDR.

[12] Adjusted gross income (AGI) is an individual's total annual income minus specific tax deductions. The federal poverty line (FPL) is a government-specified income threshold determined by household size and state of residence. For 2017, the FPL for a family of four is $24,600 in every state but Alaska and Hawaii.

[13] Forgiveness periods can vary. While the specific IDR plan I study forgives loans only after twenty-five years, loans for more recent IDR borrowers are typically forgiven after twenty years, and non-profit or government employees can often qualify for a ten-year forgiveness period under the Public Service Loan Forgiveness program. Regardless of the period, any forgiven balance is taxed as income in the year of forgiveness.

reduce IDR payments in the first five years of repayment, but rises to the standard amount thereafter, extending the repayment period beyond the standard length of ten years but not long enough to qualify for loan forgiveness at twenty-five years. Panels C and D, by contrast, depict scenarios in which the borrower is granted partial loan forgiveness under IDR with panel D highlighting the "dynamic" response of IDR to a temporary earnings shock such as unemployment.

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.

Borrowers who fail to meet their monthly payments (i.e., "fall delinquent") under any repayment plan face penalties which increase in severity with the number of days past due. Between one and fifteen days past due, student loan servicers will contact delinquent borrowers through email or post. Between fifteen 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, and 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. Unlike other forms of consumer debt, student loans cannot be discharged by declaring Chapter 7 or Chapter 13 bankruptcy.[14] Defaulted borrowers are ineligible for any future federal student aid.

2.2 Setting: Student Loan Servicing

LLS manages disbursals and payments for over $300 billion in federal student loans. Debt

[14] In 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 (Department of Education [2017a]).

servicing is provided on behalf of the federal government, which hires the servicer through a series of contracts.

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.

LLS employs over 300 servicing agents across four call centers. Agents are tasked with informing borrowers of their delinquent status, inquiring about their intention 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 "models-out" IDR payments for the borrower, asking about their annual income, family size, and employment status. Borrower responses are entered into the agent's computer, which provides an estimate of monthly IDR payments according to Equation 1. The agent then provides the borrower with 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, the agent's pay is reduced that month.

3 Data

The data I use to estimate IDR link administrative student loan repayment data to credit bureau records for over one million borrowers. Borrowers in my samples are drawn from LLS's FFEL loan portfolio, which includes over $90 billion in loans. The LLS data contain detailed repayment records for each borrower, including principal borrowing amounts, loan balances, minimum payments due, and dates of delinquency at 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 some 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.[15]

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. Additional details concerning TransUnion data can be found in Dobbie et al. [forthcoming], Avery et al. [2003], and Finkelstein et al. [2012]. TransUnion data are merged to borrowers in the LLS data by SSN.

The loan data provided for this study consist of two samples of borrowers, each drawn from LLS's FFEL portfolio. The first, which I call the "full sample" ($N = 287,456$), is a random sample of FFEL borrowers with loans disbursed after 1995. While the random sample is not used in my main analysis, descriptive statistics from this sample are representative of the FFEL-borrowing population, and thus provide a useful benchmark to compare against my analysis sample. The second sample, which I call the "call sample" ($N = 164,090$), consists of the universe of LLS's FFEL borrowers who ever received a delinquency call from 2012 onward, excluding those who hold any private or Direct loans.[16]

Borrowers in the call sample are linked to loan servicer contact histories from 2009 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

[15] The 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 first names in their borrower records at my request.

[16] While 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 Call sample excludes borrowers with mixture of loans, so I can observe their complete repayment profile. Roughly 25% of the servicer's FFEL borrowers also hold Direct loans, and fewer than ten percent hold private student loans.

of agent characteristics, including work site location and work group ("claims aversion", "skip tracing", etc.).

I construct the instrument and analysis samples on the borrower-by-call level, applying the following selection criteria to all calls placed to borrowers in the call sample: First, I remove borrowers who cannot be matched by zip-code or first name to inferred measures of gender or income, leaving 581,148 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 401,088 calls.[17] Then, I remove 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 362,759 calls, I keep only those which reached the stage at which borrowers were provided information concerning IDR enrollment (i.e., "modeled-out"), leaving 66,248 calls.

While both estimation strategies use balanced panels, I use the larger unbalanced sample to calculate the instrument, excluding those calls made by agents with fewer than 40 total calls. Applying these criteria to the instrument sample reduce measurement error in my instrument because estimates of the mean taken over small cells have high variance. The instrument sample consists of 59,405 calls.

I use monthly and yearly balanced panels centered around delinquency call dates as my analysis samples for repayment and credit outcomes, respectively.[18] For monthly repayment outcomes, I create a balanced panel of 31,113 calls with 15 leads and 10 lags. For yearly credit outcomes, I create a balanced panel of 16,021 calls with 2 leads and 3 lags. The pooled sample of calls from either panel constitutes my analysis sample.

Table 1 provides summary statistics for the random sample, call sample, and analysis sample. In the full sample, IDR has low take up, with only 4 percent of borrowers enrolled in a plan. That share rises to 13 percent in the analysis sample, as it is constructed to include only borrowers who might benefit from the plan. Unsurprisingly, these borrowers

[17] LLS handles "Late Stage" delinquencies of 150 days or more through another division with different agent-assignment procedures.

[18] Note that the calendar-correspondence between monthly and yearly variables can vary dramatically depending on call dates and data collection dates. 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 $t_{month} \in [1, 12]$.

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. Looking within the analysis sample, we see that those who ultimately enroll in IDR are largely comparable to those who do not. IDR enrollees have slightly higher principal borrowing amounts and come from slightly lower-income zip codes, again reflecting characteristics of those who might gain more benefit from IDR.

While the analysis sample is broadly representative of the IDR-eligible population, there are two important caveats concerning external validity. First, while my dataset is nationally representative of borrowers attending college between 1995 and 2010, individuals in my 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. Estimated effects may therefore be small relative to population-wide effects.

Second, my sample is restricted to individuals eligible for one particular form of IDR. In the U.S., IDR exists through several repayment plans. While all IDR plans follow the same general structure, particulars of the payment calculation formula and forgiveness provisions under each plan can vary.[19] Furthermore, IDR plans include unattractive institutional features like staggered payment adjustments and complicated sign-up procedures. To the extent real-world IDR deviates from an "ideal" system of income-contingent loan repayment, estimates from this paper apply to the former and thus include any effects from these institutional features.[20] Generalizability of my results to alternative IDR plans, existing or

[19]Individuals in my analysis sample hold FFEL loans and not Direct loans, which means they are eligible for "Income-Based Repayment" (IBR), and not "Pay-As-You-Earn" (PAYE) or "Revised-Pay-As-You-Earn" (REPAYE) plans. Approximately twenty percent of the five million borrowers in IDR plans are enrolled in IBR (Department of Education [2017b]). For a detailed description of each plan, including eligibility criteria and repayment terms, see www.studentaid.gov.

[20]For a more detailed discussion of the limitations imposed by, and potential improvements upon, the current structure of IDR in the U.S., see Dynarski and Kreisman [2013].

hypothetical, are therefore limited.

4 Model

To motivate my empirical analysis, in this section I develop a dynamic model of student loan repayment. The model illustrates the various channels through which IDR may affect borrowers' consumption and default 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 student debt payments x_t. Individuals can save ($A_{t+1} > 0$) or borrow ($A_{t+1} < 0$) to finance consumption, but face a liquidity constraint L, imposing an upper bound on the amount they can borrow.[21] To simplify exposition, assume no discounting or interest ($\beta = (1 + r) = 1$). The borrower's problem is given by:

$$V(y_t, D_t, A_t) = \max_{A_{t+1}} \{u(c_t) + E[V(y_{t+1}, D_{t+1}, A_{t+1})]\} \tag{3}$$

$$\begin{aligned} A_{t+1} &\geq L && (4)\\ c_t &= A_t - A_{t+1} + y_t - x_t && (5)\\ D_{t+1} &= D_t - x_t && (6) \end{aligned}$$

Where minimum payments x_t are determined by a payment scheme modeled as a function of current-period income and remaining student debt:

$$x_t = x(y_t, D_t) \tag{7}$$

[21] While I do not explicitly model default in this model, a simple extension would be to incorporate a floor on consumption whereby borrowers make zero loan payments and incur a penalty to utility in period t if minimum payment x_t exceeds cash-on-hand by some amount, i.e. $x_t > (A_t - A_{t+1} + y_t) + \underline{c}$ In the baseline model, the incidence of binding liquidity constraints can be thought to capture delinquency or default in some loose sense.

First order conditions for the borrower's problem yield an Euler inequality:

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

which binds in period t if and only if $A_{t+1} > L$. The solution to the borrower's problem is characterized by Equations 4 through 8.[22] Unless prevented by liquidity constraints, individuals will save or borrow to so that the marginal utility of current-period consumption is equal to the expected marginal value of assets next period.

In the following sections, I analyze how a change from Standard ("S") to IDR ("I") repayment schemes affects optimal consumption in the model above, where

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

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

For borrowers in repayment, the IDR plan is modeled as the minimum of some share θ of per-period income and the standard payment amount D_0/N. Borrowers under both plans never pay more than their remaining balance D.

4.2 Liquidity Effects through Reduced Payments

To isolate the potential channels through which IDR might influence consumption, I first consider alternative specifications for the payment scheme in Equation 7 and defer discussion of loan-forgiveness provisions to the following subsection. Let m_t denote income net of loan payments, $m_t \equiv y_t - x_t$. The reduction in minimum payments offered by IDR offers two distinct benefits to borrowers, both of which increase consumption through liquidity. First, IDR weakly increases net income in the short-term relative to the long-term:

$$E\left[m_t^I\right] \geq E\left[m_t^S\right] \quad \forall t \leq N \tag{11}$$

$$E\left[m_t^I\right] \leq E\left[m_t^S\right] \quad \forall t > N \tag{12}$$

[22] Proof in Appendix B.

The result is an intertemporal *smoothing effect*: pushing payments farther into the future flattens borrowers' expected net income profiles. As a result, financing consumption 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.[23] 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.[24] To see how, note that IDR reduces *ex-post* variance in net income:

$$\mathrm{Var}(m_t^I) \leq \mathrm{Var}(m_t^S) = \sigma^2 \quad \forall t \tag{13}$$

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).[25]

The complementary smoothing and insurance benefits of lower minimum payments are summarized in the following expression:

$$\Delta c_t = \underbrace{\sum_{k=0}^{\infty} -\frac{dc_t}{d\mu_k}\Delta E[x_k]}_{\text{smoothing}} + \underbrace{\sum_{k=0}^{\infty} \frac{dc_t}{d\sigma_k}[\sigma_k - \sigma]}_{\text{insurance}} \tag{14}$$

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 borrowers would simply borrow and save as needed to achieve constant consumption over all periods. In this

[23] Proof in Appendix B

[24] I use the term "insurance" loosely, referring to state-contingent liquidity benefits under IDR. Note that, ignoring loan forgiveness, income-contingent payments provide no insurance against *lifetime*-earnings risk.

[25] See Ljungqvist and Sargent [2012]. Proof in Appendix B.

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. IDR alleviates this problem by pushing debt obligations further into the future during periods of low income. If borrowers are liquidity constrained under standard repayment, 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.

Assessing IDR's potential liquidity effects is important for evaluating its welfare implications. If IDR provides liquidity to borrowers who would otherwise default or sharply reduce consumption, it corrects for an incomplete credit market and may increase total welfare: borrowers benefit from smoother consumption, improved credit, and fewer late fees, and the government saves on the considerable cost of collecting delinquent payments.[26] If, however, borrowers simply use their increased cash-on-hand to pay down existing debt or increase savings, flexible repayment amounts to little more than a redistributive policy transferring interest-savings to student borrowers through taxpayer-subsidized loans.

Observations of the student debt market provide support for potential liquidity benefits. Median monthly payments exceed 14% of median earnings for borrowers beginning standard repayment(Looney and Yannelis [2015]), and raising limits on credit card balances has twice the expenditure response for those under 44 as those over (Gross and Souleles [2002]). Together, these facts suggest college graduates might benefit from borrowing against future earnings to finance consumption and avoid delinquency in low-income periods.

[26]The Congressional Budget Office estimates annual administrative costs of the FFEL student loan program equal 0.67 percent of outstanding balances, or approximately $10 billion per year (Lucas and Moore [2010]).

4.3 Wealth Effects through Debt Forgiveness

Note that while the liquidity benefits described above increase borrower welfare, they do not increase lifetime income: reduced payments facilitate wealth transfers *within-borrower* over time, and therefore carry no cost to social welfare. By contrast, the forgiveness provisions of IDR, which forgive any outstanding debt after twenty-five years of successful payments, amount to a means-tested government transfer. [27]

To capture the loan forgiveness provisions of IDR, I make a simple modification to the model above. Let T denote the number of periods after which remaining balances are forgiven. In period $t = T + 1$, instead of Equation 6, we have:

$$D_{T+1} = 0 \tag{15}$$

Note that if borrowers expect loan forgiveness, $\sum_{k=0}^{\infty} \Delta E[x_k] < 0$ and the first term in Equation 14 includes a *wealth effect*. In this case, the expression can be decomposed to differentiate between liquidity and wealth effects:

$$\Delta c_t = \underbrace{-\frac{dc_t}{dD_{T+1}} E[D_T^I - x_T^I]}_{\text{wealth effect } (+)} + \underbrace{\frac{dc_t}{dD_{T+1}} E[D_T^I - x_T^I] + \sum_{k=0}^{T} \left(-\frac{dc_t}{d\mu_k} \Delta E[x_k] + \frac{dc_t}{d\sigma_k} [\sigma_k - \sigma] \right)}_{\text{liquidity effects } (+)} \tag{16}$$

4.4 Social Welfare Implications

While reduced payments and debt forgiveness provisions both increase consumption among IDR enrollees, they hold different implications for social welfare. By providing a ceiling on debt's share of lifetime income, debt forgiveness provides "wealth insurance" against lifetime-earnings risk. As with any social insurance program, however, providing state-contingent transfers carries costs. First, because IDR borrowers pay no "premium" for wealth insurance,

[27] Note that IDR may provide income transfers through other channels not included in this model. For instance, the subsidized interest rates on student loans provide a net-present-value benefit to borrowers whose minimum payments drop due to low income.

the government bears the expected cost of forgiven debts plus some risk premium.[28] Such a transfer would impose a deadweight loss to social welfare through taxation while carrying ambiguous redistributional consequences: forgiveness is progressive conditional on student debt, but college graduates, especially those with high debt balances, have higher lifetime income then 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.[29]

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. For several reasons, such an analysis is beyond the scope of this paper. First, IDR has not been available long enough for the twenty-five-year forgiveness period to bind, so assessing the incidence of forgiveness benefits necessarily would necessarily involve some degree of forecasting. Second, labor supply effects, like wages or hours worked, can only be detected insofar as they influence consumption decisions, and the post-graduate timing of IDR enrollment for individuals in my sample means I cannot estimate moral hazard effects on decisions like principal borrowing amounts or institution and college major choice. Third, some degree of wealth insurance is theoretically afforded to all IDR-*eligible* borrowers, not just current enrollees, as individuals on standard repayment plans usually can opt-in to IDR at any time. My research design cannot identify moral hazard through these market-wide effects, as all individuals in my sample are eligible for IDR.

Nevertheless, the repayment path of IDR borrowers in my sample should provide some insight 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

[28] While 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.

[29] See Chetty [2008] for a more complete description of the distinction between moral hazard and liquidity in the context of unemployment insurance.

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. If, instead, payments remain low for a prolonged period following enrollment, the likelihood of loan forgiveness is much higher and IDR could potential carry a high social welfare cost through moral hazard or redistributional inefficiencies.

5 Empirical Strategy

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

$$Y_{ict} = \beta_0 + \beta_1 IDR_{ic} + \beta_2 \boldsymbol{X}_{it} + \epsilon_{ict} \tag{17}$$

where Y_{ict} denotes the outcome of interest, IDR_{ic} is an indicator for IDR enrollment within three months of the call, $\boldsymbol{X}_{ict}$ is a vector of borrower control variables, and ϵ_{ict} is an error term. Estimating β_1 in Equation 17 using OLS would likely yield biased estimates, because preferences over repayment plan choices are correlated with unobserved borrower attributes. For example, borrowers expecting an increase in earnings would likely have a weaker preference for IDR than those with lower earnings expectations. In this case, OLS estimates of IDR's effect on measures of financial well-being would be biased *downwards*, reflecting the selection of individuals low earnings potential into the plan.

The overcome the biases associated with OLS and identify the causal effect of IDR on repayment and credit 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.

5.1 Difference-in-Differences Specification

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:

$$Y_{ict} = \gamma_i + \gamma_t + \left[\sum_{\tau \neq -1} \delta_\tau \cdot IDR_i \cdot \mathbf{1}\{t = \tau\} \right] + \gamma_1 IDR_{ic} + \gamma_2 \boldsymbol{X}_{ict} + \epsilon_{it} \tag{18}$$

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 three months of the call, $\boldsymbol{X}_{ict}$ is a vector of borrower control variables (including call date and time fixed effects), ϵ_{ict} is an error term, and δ_t are coefficients on IDR enrollment status which vary non-parametrically by event time. The specification omits α_t and δ_t terms at $t = -1$ to prevent perfect multicollinearity. Estimates can therefore 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 can come from a variety of sources, including both between-agent differences captured by the instrumental variables strategy I describe in Section 5.2 and *within*-agent variation over calls. While the following section argues that between-agent variation is as-good-as-random, 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 is exogenous with respect to outcomes. Stated differently, unbiased estimation in my difference-in-differences design assumes that trends in outcomes would be the same in both groups of borrowers had neither taken up IDR.

Figures 2 - 5 provide graphical evidence in support of the common-trends assumption. The figures plot raw means for monthly repayment and credit bureau outcomes, respectively, for IDR enrollees and non-enrollees, normalized by month- or year-of-call and 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.[30]

Even if IDR and standard borrowers exhibit observably similar pre-trends, other potential violations to the exclusion restriction exist. For example, IDR-enrollees may have experienced a shock at the time of their delinquency call that both induced them into IDR take-up while also influencing 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 observably similar debt 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 exhibit the same or similar trends before delinquency calls but make IDR enrollment decisions based on expected *future* shocks to their financial well-being. While my difference-in-differences design cannot rule out this scenario, 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.

Finally, the quasi-experimental nature of my instrumental variables design provides an informal robustness check to my differences-in-differences design. While differences between the two approaches can be partly attributable to heterogeneous treatment effects between their respective complier populations, similar IV and difference-in-differences estimates would suggest that the variation driving the latter is plausibly exogenous. I make this comparison for monthly repayment outcomes, where IV estimates are sufficiently precise, to lend further credibility to difference-in-differences estimates of IDR's effect on credit outcomes, where my

[30] Estimates 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

$$Y_{ict} = \gamma_i + \delta t \cdot IDR_i \cdot \mathbf{1}\{t < 0\} + \left[\sum_{\tau \geq 0} \delta_\tau \cdot IDR_i \cdot \mathbf{1}\{t = \tau\}\right] + \gamma_1 IDR_i + \gamma_2 \boldsymbol{X}_{ict} + \epsilon_{it} \tag{19}$$

IV approach is too underpowered to conduct meaningful inference.

5.2 Instrumental Variables

I complement the difference-in-differences design described above with an instrumental variables (IV) approach which isolates quasi-random variation in IDR enrollment. 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. I interpret observed differences in outcomes between borrowers assigned to different agents as the causal impact of the change in IDR take-up associated with those agents. The parameter I estimate in this specification is a local average treatment effect (LATE), which identifies the impact of enrolling in IDR among borrowers whose plan choice depends on agent-assignment.

Instrument Calculation I construct my instrument using a residualized, leave-one-out measure of agents' ability to induce IDR take-up that accounts 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. [2016a] to instrument for trial outcomes in court case settings. The approach solves two problems one might see in a non-residualized 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 take-up 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 as well as a linear time trend from observed IDR enrollments before calculating the leave-out agent score. Specifically,

$$\begin{aligned} IDR^{*}_{ict} &= IDR_{ict} - \gamma \boldsymbol{W}_{ict} && (20) \\ &= Z_{ijc} + \epsilon_{ict} && (21) \end{aligned}$$

where $\boldsymbol{W}_{ict}$ is a vector of call date and time controls and Z_{ijc} is agent score. I calculate residualized rate of IDR take-up, IDR^*_{ict}, using OLS estimates of γ in Equation 20. I then construct agent score Z_{ijc} using the leave-one-out mean of this residualized rate:

$$Z_{icj} = \left(\frac{1}{n_j - 1}\right)\left(\sum_{k=0}^{n_j} IDR^*_{kj} - IDR^*_{ij}\right) \tag{22}$$

where n_j denotes the number of calls made by agent j. Removing borrower i's outcomes from the calculated mean is important to avoid reflection bias: including one's own outcome in her assigned agent score would introduce the same estimation errors on both left- and right-hand side of the regression, incorrectly attributing unobserved between-borrower differences in IDR take-up to agent ability. My two-stage least-squares estimation avoids this problem by using the leave-one-out agent score, calculated as in Equation 22, as an instrumental variable for whether a borrower enrolls in IDR.

The residualized agent score distribution can be seen in Figure 6. The analysis sample includes calls from 144 different agents in four different call centers. Agents place 216.06 calls on average to borrowers in the analysis sample, with a median of 183 calls.

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 empirically evaluate the first-stage relationship between the agent score instrument and observed IDR take-up. Figure 6, 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 6 can be represented through following linear probability model:

$$IDR_{ict} = \alpha_0 + \alpha_1 Z_{ict} + \alpha_2 \boldsymbol{X}_{ict} + \epsilon_{ict} \tag{23}$$

In Table 3, I provide first-stage estimates from an OLS regression of Equation 23. Results imply a positive and highly significant relationship between agent score and IDR take-up. Point estimates range from 0.62 to 0.60, depending on the inclusion of borrower controls, and the F-statistic on a test of instrument significance equals 21.89.

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. Table 2 provides empirical evidence that, conditional on call date and time, borrowers do not vary systematically by agent score. Column 1 reports results from an OLS regression of realized IDR take-up against several borrower characteristics and pre-call outcome variables, as well as call date and time fixed effects. Not surprisingly, estimates demonstrate considerable 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 1.17.

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, the institutional details in my research setting provide suggestive evidence that such threats to validity are minimal. 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; any borrower who declines the IDR option under a high agent must not be induced into IDR under a comparatively low score agent. 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.

Agent Variation Variation in agent score can be driven by several potential sources. First, while agents ask questions and provide responses according to a standardized decision tree, several nodes of this tree require agent judgment. For example, before informing borrowers of alternative repayment options, an agent must first ask if a borrower is "capable of making payments under her current plan." This question seldom elicits an unambiguous "yes" or "no" response, so agents must make this determination based on their own judgment. To the extent agents are afforded discretion in their responses to borrowers, their decisions are undoubtedly influenced by the incentive scheme discussed in Section 2.2. For instance, more risk averse agents may present the IDR option to a broader swath of borrowers to avoid potential compliance infractions. Second, variation may exist in agents' perceived geniality. Even agents reciting the same words can vary in IDR-conversion through subtle variations in demeanor or tone, and borrowers often hang up or stop listening depending on the interaction. Similarly, while call data are not detailed enough to verify, 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. Third, agent score may 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.

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 10 or more and 90 or more days delinquent. For each outcome listed on the left-hand side, Columns 2-4 provide estimated coefficients on the interaction between IDR take-up and consecutive three-month periods from the pooled difference-in-differences specification given by Equation 18. Columns 5-7 provide estimated coefficients on IDR take-up from separate two-stage least squares regressions in each three-month period. [31] Figure 8 plots coefficients on these outcomes using the same specification, but separated by each month following the date of the delinquency call.

Minimum Payments IDR's effect on minimum payment is mechanical: given 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. If we saw no effect on payments, any observed effects on delinquencies or consumption would be entirely attributable to the behavioral effects of expected loan forgiveness.[32]

In the difference-in-differences analysis, required payments for those on IDR fall by $140 relative to those who remain on standard repayment. In the top-left panel of Figure 7, payments remain relatively stable in the twelve months following the delinquency call, then sharply increase to their pre-call levels, so that IDR and standard borrowers face similar minimum payments fourteen to eighteen months following their delinquency calls. The convergence in payments corresponds to the one-year recertification period for IDR, suggesting that IDR-enrollees in my sample do not extend benefits beyond one year. 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 $200 four to seven months following IDR take-up and exhibit the same rapid return to standard levels one year following the delinquency call, as seen in the top left panel of Figure 8.

Delinquency I measure IDR's impact on two measures of delinquency: the likelihood of falling ten or more days delinquent, and the likelihood of falling ninety or more days

[31] Results from the first three 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 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.

[32] The expectation of future liquidity benefits could technically have a behavioral response as well, though upward-sloping wage profiles, and the ability to opt-into IDR at any point in the future make "expected-liquidity" channels unlikely.

delinquent. The ten and ninety day benchmarks are points of increased penalties: at ten days past due, borrowers begin to accrue late fees for delinquent loans, and at ninety days past due, borrowers are reported to credit bureaus. Figures 7 and 8 show a large negative effect of IDR enrollment for both delinquency measures. In the difference-in-differences analysis, IDR borrowers are 21 percentage points less likely to fall ten or more days delinquent relative to standard borrowers in the four to seven months following the delinquency call, with a pre-call mean of 22pp. Corresponding estimates for months seven through nine and ten through twelve are -16pp and -4pp, respectively. Estimates for ninety or more days delinquent exhibit a similar pattern: the likelihood of falling ninety or more days delinquent falls by 8pp, 8pp, and 7pp in months four through six, seven through nine, and ten through twelve after the delinquency call, respectively, compared to a pre-call mean of 2pp. IV estimates for the effect of IDR on delinquencies are qualitatively similar to difference-in-differences estimates, though less precise and larger in magnitude.

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 provides a sense of how IDR might impact borrowers' propensity to pay down their debts. Both IV and 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. While not statistically significant, point estimates for the instrumental variables analysis suggest IDR decreases monthly balances by $90 in months four through six and $80 in months seven through nine after the delinquency call, off of a pre-call mean of $130. Difference-in-difference estimates are qualitatively similar and statistically significant: relative to standard borrowers, IDR enrollees pay down $70 to $90 more debt each month than they did before their delinquency calls.

6.2 Credit and Consumption Outcomes

For evaluating IDR's impact on credit and consumption outcomes in the TransUnion data, I rely principally on difference-in-difference estimates, as instrumental variables estimates are too low-powered to draw meaningful inference.[33]

Table 5 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 2-4 report coefficients on IDR in consecutive years from the pooled regression specified in Equation 18, beginning with the year of the delinquency call. Columns 5-7 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 5.01-point increase in credit scores within one year of the delinquency call, off of a pre-call mean of 593.26, followed by increases of 7.53 and 5.93 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 20pp, an increase of 10 percent.

Table 6 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 \$240 (0.33 log points) and \$400 (0.38 log points) one and two years after the delinquency call off of a pre-call mean of \$1420, corresponding to an increase of 28 percent. Similarly, by the second year of enrollment, IDR borrowers hold 0.14 more credit cards (pre-call mean of 3.00) and have \$340 higher credit limits (pre-call mean of \$2920) 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.

[33] For completeness, instrumental variables estimates for IDR's effect on credit outcomes are reported in Appendix A.

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

Minimum Payments

Ignoring the potential effects of 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: a necessary condition for IDR to generate observable liquidity effects is that it measurably reduces payments during my sample period. While the alternative scenario seems unlikely, it is possible for high-income borrowers to sign up for IDR solely in anticipation of future income shocks. 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.

Note that while the estimated effects of IDR on minimum payment provide valuable insight into the mechanisms driving IDR's 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 A1 provides a graphical illustration of this measured payment effect across the distribution of IDR enrollees in my analysis sample.

Having established that IDR borrowers receive an increase in short-term cash-on-hand, 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 and Balances

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 carries negative consequences described in Section 2. Note, however, 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. 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 "stragetic" motive, as lower payments should not influence strategic default decisions.[34]

In theory, IDR could affect balances on student loans in either direction. Lower required payments mean non-delinquent IDR borrowers pay down less of their debt than non-delinquent borrowers on standard repayment plans, increasing balances. However, the reduced likelihood to default 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 $90 more student debt each month, on average, than those on standard repayment plans.

Credit Outcomes

The positive estimated effects of IDR on credit card balances and number of credit cards are suggestive of a consumption response to increased liquidity, but interpreting these results requires two 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

[34] I 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 loan forgiveness. I rule out forgiveness effects later in the section.

flows in credit card spending as well as the stock of unpaid debt. For this reason, the parameter I identify is not a marginal propensity to consume (MPC) out of liquidity. To the extent lower required payments have a deleveraging effect allowing borrowers to pay down existing debt on credit cards, I likely underestimate the true consumption response to IDR.

Second, the liquidity effects I estimate may be driven by multiple mechanisms. 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 through which individuals might be affected. While this credit channel probably accounts for some of my results, especially over the two and three year time horizons, the estimated impact on credit limits is modest or null, 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.

7.1 Loan Forgiveness

By increasing expected lifetime income, debt forgiveness provisions under IDR can also affect borrower behavior. If borrowers expect their loans to be forgiven, they may increase short-term consumption regardless of liquidity constraints. While expected loan forgiveness could, in theory, drive effects of IDR on repayment or consumption, my results are more consistent with a pure liquidity effect. In order for the forgiveness 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 rapid return of minimum payments to standard levels therefore suggests forgiveness is very unlikely for 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. It is possible that payments may drop again in the future, making debt forgiveness more likely. Forgiveness effects on consumption under these circumstances would require one of two scenarios, both of which I argue are unlikely. First, if rising payments reflect increased incomes, payments might fall again in the future if income declines more permanently. While I cannot rule out this possibility, it seems unlikely, as income profiles do not generally exhibit sharp, transitory *upward* shocks followed by permanent decreases. Second, if rising payments reflect a failure to recertify income, borrowers may have permanently low incomes through the end of the sample period and choose to reenroll in IDR at a later date. These borrowers may qualify for forgiveness if they successfully make enough reduced payments following their reenrollment. However, if consumption were effected through the expectation of future forgiveness via potential future reenrollment, one would expect those effects to exist for both treatment and control borrowers in my sample, as those who remain in standard repayment plans following the delinquency call are aware of the forgiveness benefit they would receive from IDR should they decide to enroll at a later date.

7.2 Alternative Mechanisms

While my results suggest that IDR improves welfare by providing liquidity to credit-constrained borrowers, I cannot rule out all alternative mechanisms driving the positive effect on repayment and consumption. Behavioral mechanisms, in particular, may be consistent with my results while holding different welfare implications. 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.

Sixty-nine percent of IDR-enrollees in my sample face IDR payments of zero dollars and thus cannot fall delinquent on their loans. While this mechanical result would still be characterize as a "liquidity effect" under a neo-classical model, to the extent borrowers facing payments of $\epsilon > 0$ dollars would face higher delinquency rates than zero-payment borrowers, the result may be driven in part by pyschological frictions or "hassle-costs." To attenuate this mechanical effect, I conduct my difference-in-differences analysis for repayment 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 predict the likelihood of having positive IDR payments among *all* borrowers in my 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. Appendix Table A1 reports delinquency results for this subsample, and continues to find a large and significant effect on repayment rates.

8 Conclusion

In this paper, I use new administrative data from a large student loan servicing company and exploit two sources of plausibly exogenous variation from servicing phone calls to isolate the causal effect of IDR enrollment on default, bankruptcy, consumption proxies, and household balance sheets. I find evidence of increased liquidity on both default and consumption margins. IDR decreases delinquency rates by 16 percentage points within seven months of take up. Relative to borrowers who remain on standard repayment plans, IDR enrollees have credit scores which are 7.53 points higher, hold 0.1 more credit cards, and carry \$240 higher credit card balances one year after the servicing call, implying increased short-term consumption out of liquidity. Minimum monthly payments decrease by an average of \$140 following IDR take up, but return to standard levels within one year of enrollment, suggesting limited forgiveness incidence and minimal moral hazard effects. My results suggest IDR improves borrower welfare principally through a liquidity channel, providing short-term increases to cash-on-hand in periods of financial distress.

The results of my analysis raise several questions for future research. First, while this paper establishes the benefits of IDR on consumption, the same benefits may carry over into other outcomes. Income, employment, and occupation choice might all be affected by reduced loan payments, especially since liquidity benefits are concentrated in the years after college when labor market decisions are most consequential. Likewise, IDR may have ex-ante effects on borrowers' college attendance decisions, major choices, or principal borrowing amounts. In both cases, the relative effects of liquidity and forgiveness hold important implications for social welfare: a borrower may make "riskier" career decisions because she can borrow against future income, or because her future income is insured.

Second, while the apparent lack of forgiveness-eligible borrowers in my sample provides clear evidence of liquidity benefits, it also leaves unanswered questions regarding moral hazard and student debt forgiveness more generally. 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 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 on loan forgiveness is needed to assess the welfare benefits and potential distortionary effects of these policies.

Third, evidence for IDR's benefits raises the natural question of why so few students enroll in the plan. According to classic economic theory, IDR should be weakly preferred to standard repayment for any borrower. There are no pecuniary costs to enrolling, yet my results suggest many individuals who would benefit from IDR do not take it up. There are a number of potential reasons for this low take-up, such as "hassle-costs" of enrollment, inattention, or incomplete information.[35] Further research into these psychological frictions may reveal policy interventions to help borrowers take advantage of the benefits offered by IDR.

[35] Prior work has demonstrated the importance of information in similar contexts, including principal borrowing decisions for student loans (Marx and Turner [2017]) and college major choice (Hastings et al. [2015]).

References

Anna Aizer and Joseph J Doyle. Juvenile incarceration, human capital, and future crime: Evidence from randomly assigned judges. *The Quarterly Journal of Economics*, 130(2): 759–803, 2015.

David H Autor and Susan N Houseman. Do temporary-help jobs improve labor market outcomes for low-skilled workers? evidence from "work first". *American Economic Journal: Applied Economics*, pages 96–128, 2010.

Robert B Avery, Paul S Calem, Glenn B Canner, and Raphael W Bostic. An overview of consumer data and credit reporting. *Fed. Res. Bull.*, 89:47, 2003.

Nicholas Barr, Bruce Chapman, Lorraine Dearden, and Susan Dynarski. Getting student financing right in the us: lessons from australia and england. *Centre for Global Higher Education*, 2017.

Zachary Bleemer, Meta Brown, Donghoon Lee, Katherine Strair, and Wilbert Van der Klaauw. Echoes of rising tuition in students' borrowing, educational attainment, and homeownership in post-recession america. Working Paper, 2017.

Raj Chetty. Moral hazard versus liquidity and optimal unemployment insurance. *Journal of political Economy*, 116(2):173–234, 2008.

Kristle Cortés, Gary Wagner, Christopher Vecchio, and Guhan Venkatu. Is there a student loan crisis? not in payments. *Policy*, 6:17, 2017.

Gordon B Dahl, Andreas Ravndal Kostøl, and Magne Mogstad. Family welfare cultures. *The Quarterly Journal of Economics*, 129(4):1711–1752, 2014.

Department of Education. Federal student loan portfolio. Website, 2017a. URL `https://studentaid.ed.gov/sa/repay-loans/default/get-out`. Accessed: 2017-10-19.

Department of Education. New student loan report reveals promising repayment trends. Website, 2017b. URL `https://www.ed.gov/news/press-releases/`

`new-student-loan-report-reveals-promising-repayment-trends`. Accessed: 2017-10-19.

Marco Di Maggio, Amir Kermani, Benjamin Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao. Monetary policy pass-through: Mortgage rates, household consumption and voluntary deleveraging. *American Economic Review*, 2016.

Will Dobbie and Jae Song. Debt relief and debtor outcomes: Measuring the effects of consumer bankruptcy protection. *The American Economic Review*, 105(3):1272–1311, 2015.

Will Dobbie, Jacob Goldin, and Crystal Yang. The effects of pre-trial detention on conviction, future crime, and employment: Evidence from randomly assigned judges. 2016a. National Bureau of Economic Research Working Paper.

Will Dobbie, Paul Goldsmith-Pinkham, Neale Mahoney, and Jae Song. Bad credit, no problem? credit and labor market consequences of bad credit reports. National Bureau of Economic Research Working Paper 22711, 2016b.

Will Dobbie, Paul Goldsmith-Pinkham, and Crystal Yang. Consumer bankruptcy and financial health. *The Review of Economics and Statistics*, forthcoming.

Danielle Douglas-Gabriel. Trump just laid out a pretty radical student debt plan. *The Washington Post*, 2015.

William Dudley. Opening remarks at the convening on student loan data conference. 2015.

Susan Dynarski. An economist's perspective on student loans in the united states. Working paper, Education Policy Initiative: Gerald R. Ford School of Public Policy, September 2014.

Susan Dynarski. We're frighteningly in the dark about student debt. *The New York Times*, 2015.

Susan Dynarski and Daniel Kreisman. Loans for educational opportunity: Making borrowing work for today's students. *The Hamilton Project. Discussion Paper*, 2013.

Amy Finkelstein, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P Newhouse, Heidi Allen, Katherine Baicker, and Oregon Health Study Group. The oregon health insurance experiment: evidence from the first year. *The Quarterly journal of economics*, 127(3):1057–1106, 2012.

Peter Ganong and Pascal Noel. The effect of debt on default and consumption: Evidence from housing policy in the great recession. Working Paper, 2017.

David B Gross and Nicholas S Souleles. Do liquidity constraints and interest rates matter for consumer behavior? evidence from credit card data. *The Quarterly journal of economics*, 117(1):149–185, 2002.

Justine Hastings, Christopher A. Neilson, and Seth D. Zimmerman. The effects of earnings disclosure on college enrollment decisions. June 2015. National Bureau of Economic Research Working Paper 21300.

Benjamin J Keys, Tomasz Piskorski, Amit Seru, and Vincent Yao. Mortgage rates, household balance sheets, and the real economy. National Bureau of Economic Research Working Paper 20561, 2014.

Jeffrey R Kling. Incarceration length, employment, and earnings. *The American Economic Review*, 96(3):863–876, 2006.

Andreas Ravndal Kostøl, Magne Mogstad, Bradley Setzler, et al. Disability benefits, consumption insurance, and household labor supply. National Bureau of Economic Research Working Paper, 2017.

Lars Ljungqvist and Thomas J Sargent. *Recursive macroeconomic theory*. MIT press, 2012.

Adam Looney and Constantine Yannelis. 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, 2015.

Deborah Lucas and Damien Moore. Costs and policy options for federal student loan programs. a cbo study. *Congressional Budget Office*, 2010.

Benjamin M Marx and Lesley J Turner. Student loan nudges: Experimental evidence on borrowing and educational attainment. National Bureau of Economic Research Working Paper 24060, November 2017.

Alvaro Mezza, Daniel R Ringo, Shane Sherland, and Kamila Sommer. On the effect of student loans on access to homeownership. Working paper, 2016.

Jesse Rothstein and Cecilia Elena Rouse. Constrained after college: Student loans and early-career occupational choices. *Journal of Public Economics*, 95(1):149–163, 2011.

Robert Shimer and Iván Werning. Liquidity and insurance for the unemployed. *The American Economic Review*, 98(5):1922–1942, 2008.

Cong Tang, Keith Ross, Nitesh Saxena, and Ruichuan Chen. What's in a name: a study of names, gender inference, and gender behavior in facebook. *Database Systems for Adanced Applications*, pages 344–356, 2011.

Constantine Yannelis. Asymmetric information in student loans. Working paper, 2016.

Table 1: Summary Statistics

	Full Sample	Analysis Sample		
	(1)	(2)	(3)	(4)
	Pooled	Pooled	non-IDR	IDR
Panel A: LLS Data				
IDR	0.0405	0.141	0	1
Female	0.570	0.713	0.708	0.742
Zip Median Income	59.23	52.25	52.42	51.24
Age	32.92	43.26	43.28	43.14
Amount Borrowed	19.11	27.48	27.09	29.84
10+ Days Delinquent	0.500	0.998	0.999	0.995
90+ Days Delinquent	0.287	0.737	0.754	0.635
Days Delinquent	93.99	167.3	171.6	141.0
Number of Calls		5.183	5.328	4.298
Panel B: Credit Data				
Credit Score	676.3	588.9	587.6	596.4
Bankruptcy	0.127	0.196	0.194	0.207
Derogatory Rating	0.331	0.553	0.551	0.568
Number of Credit Cards	5.020	3.081	3.066	3.169
Credit Card Balances	4.104	1.581	1.582	1.575
Number of Mortgages	1.156	0.682	0.691	0.631
Mortgage Balances	67.25	25.87	26.68	20.94
Credit Card Limits	20.03	4.835	4.839	4.806
Number of Auto Trades	1.923	1.613	1.618	1.584
N	148666	24985	21452	3533

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 made any loan payments from 2010 onward. 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 three years years prior and two years following the phone call, hold no private or Direct loans, are not recorded as non-english speakers, and were assigned to agents with at least 40 observed phone calls. IDR is and 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]. Initial zip income is the median 2010 income for the borrower's 5-digit zip code as of her first payment. Days delinquent is the maximum number of days the borrower was ever past due on payments as of the end of the month, and ever delinquent is an indicator for whether days delinquent is greater than 10. Number of calls is the total number of outgoing calls recorded for the borrower. Credit scores, bankruptcies, derogatory ratings, credit card, mortgage, and auto loan information are taken from TransUnion credit bureau data.

Table 2: Balance Test

	(1) Agent Score∗100	(2) IDR∗100
Female	0.0192267	1.9057462^{***}
	(0.0364808)	(0.4107607)
Amount Borrowed	−0.0002625	-0.0753578^{***}
	(0.0019822)	(0.0215861)
Age	0.0000204	−0.0349901
	(0.0014395)	(0.0235802)
Lag Zip Median Income	-0.0020918^{***}	-0.0265073^{***}
	(0.0007458)	(0.0094746)
Lag Minimum Payment	0.2533364	6.3175518^{***}
	(0.1757279)	(1.9366660)
Lag Remaining Balance	0.0006257	0.0869712^{***}
	(0.0011584)	(0.0160278)
Lag Credit Score	-0.0005425^{**}	0.0356230^{***}
	(0.0002723)	(0.0040647)
Lag Credit Card Balances	−0.0069284	−0.0396615
	(0.0054446)	(0.0724199)
Lag Number of Auto Trades	−0.0025185	-0.2650774^{**}
	(0.0078358)	(0.1166223)
Lag Any Mortgage	0.0469042	−0.9608241
	(0.0554542)	(0.8276821)
Lag Any Credit Card	−0.0058517	0.2461155
	(0.0362830)	(0.5559533)
Lag Mortgage Balances	−0.0005516	-0.0108382^{**}
	(0.0004285)	(0.0047757)
Lag Number of Credit Cards	−0.0079420	0.0485748
	(0.0050360)	(0.0868937)
Lag Credit Card Limits	0.0035912^{*}	−0.0407309
	(0.0021653)	(0.0392379)
Mean Dep.	-0.096	10.864
F-stat	1.17	15.86
P-value	0.3029	0.0000
R-squared	0.014	0.016
N	31113	31113

Note: This table reports balance test results. The regressions are estimated on the call sample described in the notes to Table 1. Agent score is estimated using data from other phone calls placed by the same agent following the procedure described in Section 5.2. Column 1 reports the estimated coefficients from an OLS regression of agent score multiplied by 100 against the variables listed and call year, month, and hour fixed effects. 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 in the rows. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 3: First Stage

	(1) IDR	(2) IDR
Agent Score	0.6179307***	0.6042973***
	(0.1320810)	(0.1309608)
Female		0.0189413***
		(0.0040865)
Amount Borrowed		−0.0007520***
		(0.0002164)
Age		−0.0003500
		(0.0002361)
Lag Zip Median Income		−0.0002524***
		(0.0000948)
Lag Minimum Payment		0.0616446***
		(0.0194076)
Lag Remaining Balance		0.0008659***
		(0.0001603)
Lag Credit Score		0.0003595***
		(0.0000403)
Lag Credit Card Balances		−0.0003547
		(0.0007229)
Lag Number of Auto Trades		−0.0026356**
		(0.0011642)
Lag Any Mortgage		−0.0098917
		(0.0082987)
Lag Any Credit Card		0.0024965
		(0.0055521)
Lag Mortgage Balances		−0.0001050**
		(0.0000480)
Lag Number of Credit Cards		0.0005337
		(0.0008721)
Lag Credit Card Limits		−0.0004290
		(0.0003948)
Mean Dep.	0.109	0.109
F-stat	21.89	21.29
P-value	0.0000	0.0000
R-squared	0.010	0.018
N	31113	31113

Note: This table reports first stage results. The regressions are estimated on the call sample described in the notes to Table 1. Agent score is estimated using data from other phone calls placed by the same agent following the procedure described in Section 5.2. 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 4: Difference-in-Differences and Instrumental Variables Estimates of the Effect of IDR on Loan Repayment Outcomes

		Difference-in-Differences			Instrumental Variables		
Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mean $t = -1$	Mos. 4 to 7	Mos. 8 to 11	Mos. 12 to 15	Mos. 4 to 7	Mos. 8 to 11	Mos. 12 to 15
Minimum Payment	0.22	−0.14***	−0.14***	−0.05***	−0.20***	−0.22***	−0.14**
		(0.00)	(0.00)	(0.00)	(0.02)	(0.03)	(0.06)
Remaining Balance	27.83	−0.04	−0.38***	−0.29***	−0.27	−1.35	−1.32
		(0.03)	(0.04)	(0.05)	(0.79)	(1.21)	(1.37)
Δ Remaining Balance	0.13	−0.09***	−0.08***	0.07***	−0.15	−0.11	0.03
		(0.01)	(0.01)	(0.01)	(0.16)	(0.15)	(0.13)
10+ Days Delinquent	0.22	−0.21***	−0.16***	−0.04***	−0.37**	−0.29***	−0.09
		(0.00)	(0.00)	(0.00)	(0.15)	(0.11)	(0.13)
90+ Days Delinquent	0.02	−0.08***	−0.08***	−0.07***	−0.13	−0.07	0.06
		(0.00)	(0.00)	(0.00)	(0.10)	(0.09)	(0.09)
Call Time FE		Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes
N	31113	124447	124447	124447	124447	124447	124447

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 2-4 report coefficients on the effect of IDR in consecutive three-month periods, beginning with four months following the delinquency call from the pooled OLS regression specified in Equation 18. Each of Columns 5-7 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 with the agent score calculated using data from other calls made by the agent following the procedure described in Section 5.2. All specifications include controls from 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 Estimates of the Effect of IDR on Financial Outcomes

		Diff-in-Diff			Diff-in-Diff w Linear Trend		
Dependent Variable	(1) Mean $t = -1$	(2) Year 0	(3) Year 1	(4) Year 2	(5) Year 0	(6) Year 1	(7) Year 2
Credit Score	593.26	5.01***	7.53***	5.93***	4.95***	7.48***	5.87***
		(1.06)	(1.40)	(1.63)	(1.53)	(1.82)	(1.98)
Bankruptcy	0.19	−0.00	−0.00	−0.00	−0.00	−0.01	−0.01
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Any Mortgage	0.20	0.00	0.02**	0.02***	0.00	0.02**	0.02***
		(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Any Auto Trade	0.71	−0.00	−0.00	0.01	−0.00	−0.00	0.01
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)

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 2-4 report coefficients on the effect of IDR in consecutive years, beginning with the year of the delinquency call from the pooled OLS regression specified in Equation 18. Columns 5-7 report coefficents 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 124,447 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 6: Difference-in-Differences Estimates of the Effect of IDR on Credit Cards

		Diff-in-Diff			Diff-in-Diff w Linear Trend		
Dependent Variable	(1) Mean $t = -1$	(2) Year 0	(3) Year 1	(4) Year 2	(5) Year 0	(6) Year 1	(7) Year 2
Credit Card Balances	1.42	0.05	0.24***	0.40***	0.02	0.21**	0.37***
		(0.05)	(0.08)	(0.11)	(0.07)	(0.09)	(0.12)
Log Credit Card Balances	-2.51	0.07	0.33***	0.38***	0.06	0.32***	0.37***
		(0.05)	(0.07)	(0.07)	(0.06)	(0.08)	(0.09)
Number of Credit Cards	3.00	−0.01	0.07	0.14**	−0.03	0.06	0.12**
		(0.03)	(0.04)	(0.06)	(0.03)	(0.05)	(0.06)
Credit Card Limits	2.92	−0.13	0.14	0.34	−0.23	0.04	0.24
		(0.12)	(0.16)	(0.21)	(0.14)	(0.18)	(0.23)

Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on yearly credit cards. Column 1 reports the dependent variable mean in the year prior to receiving a delinquency call. Columns 2-4 report coefficients on the effect of IDR in consecutive years, beginning with the year of the delinquency call from the pooled OLS regression specified in Equation 18. Columns 5-7 report coefficents 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 124,447 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

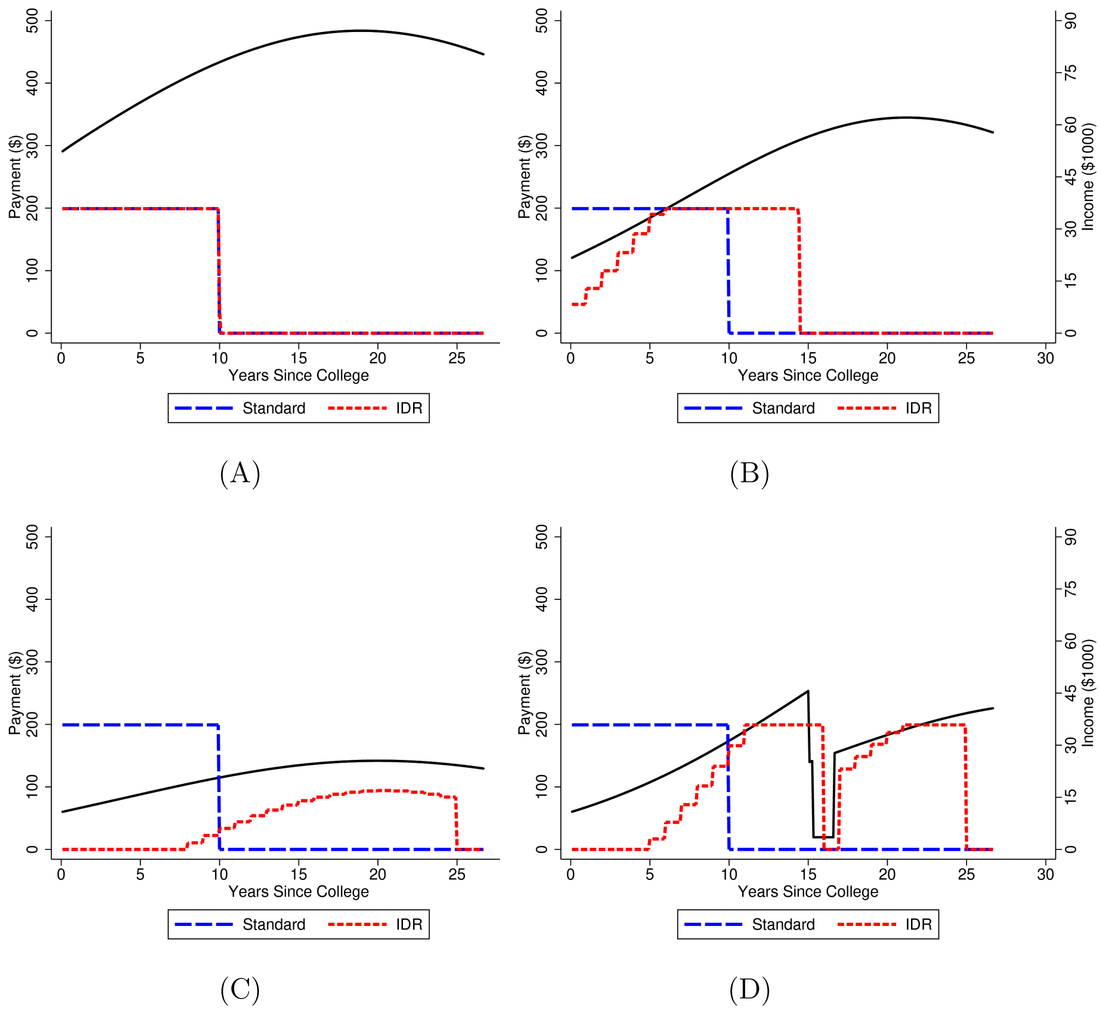

Note: This figure plots hypothetical repayment paths for standard and IDR plans under various income scenarios. Each panel represents an alternative income/repayment scenario for a borrower holding $18,000 of student debt at the time they leave college. The solid black line, plotted against the right 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 Repayment Outcomes

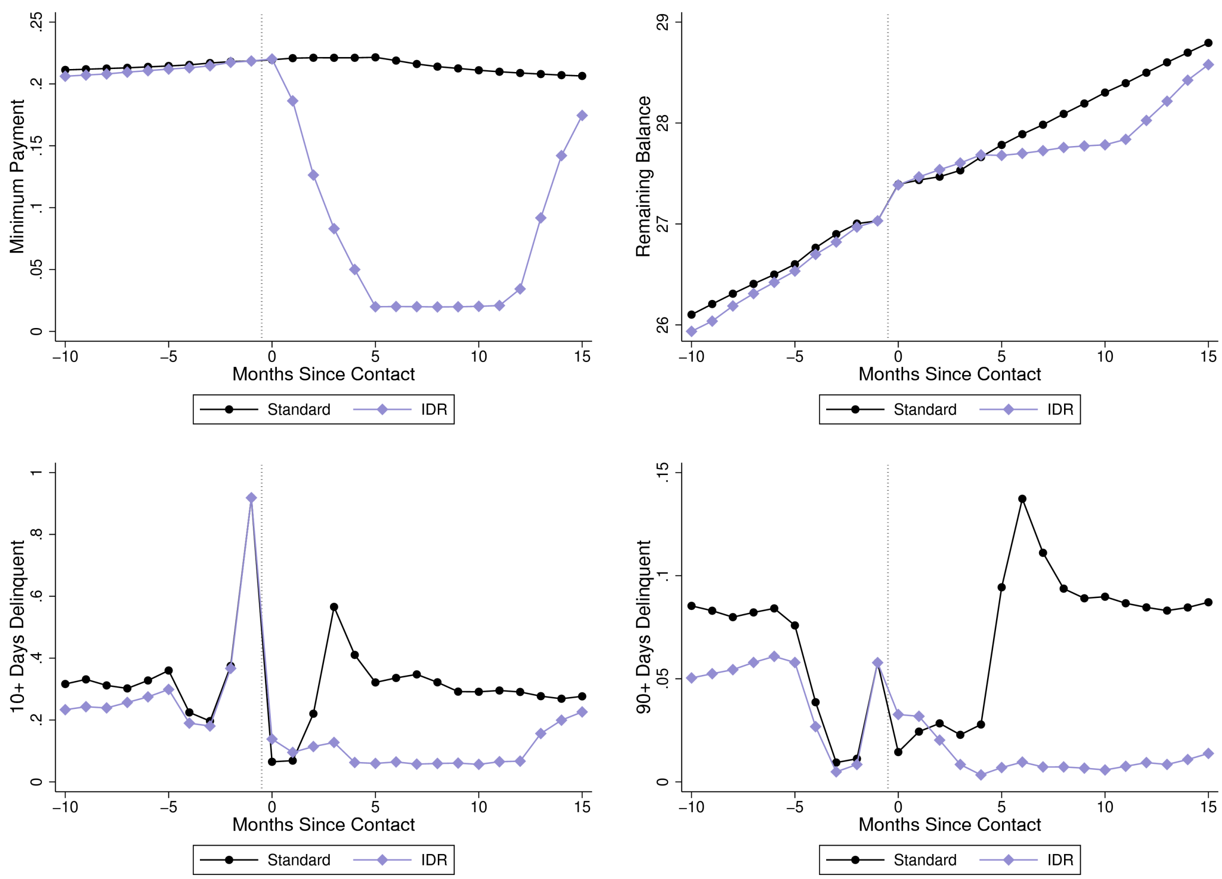

Note: This figure plots the average loan repayment outcomes for IDR enrollees and non-enrollees in the analysis sample. The horizontal axis denotes time, in months, relative to the month of the loan servicing call. Outcomes are normalized to the average value of the outcome for non-enrollees in the month prior to the call. See Table 1 notes for additional details on the outcome measures and sample.

Figure 3: Pre/Post-Call Trends in Credit Scores and Bankruptcies

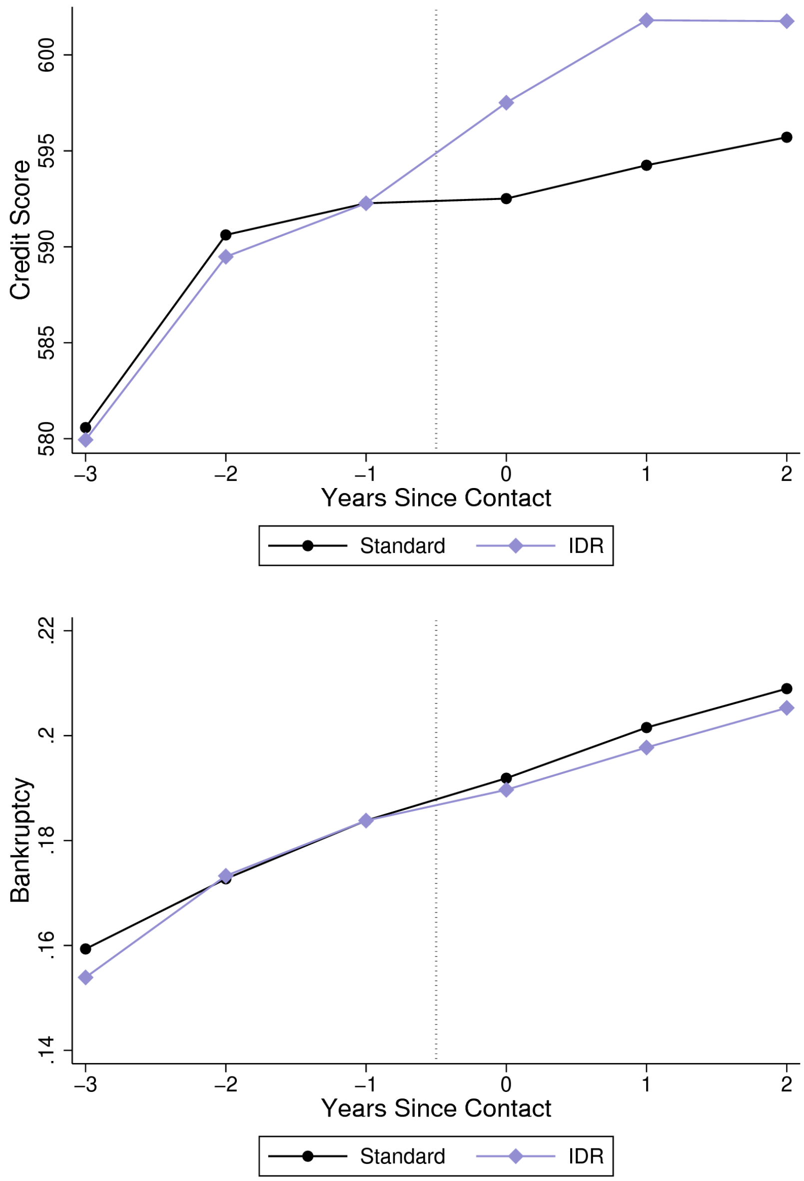

Note: This figure plots the average credit scores and bankruptcies for IDR enrollees and non-enrollees in the analysis sample. The horizontal axis denotes time, in years, relative to the year of the loan servicing call. Outcomes are normalized to the average value of the outcome for non-enrollees in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

Figure 4: Pre/Post-Call Trends in Mortgages and Auto Loans

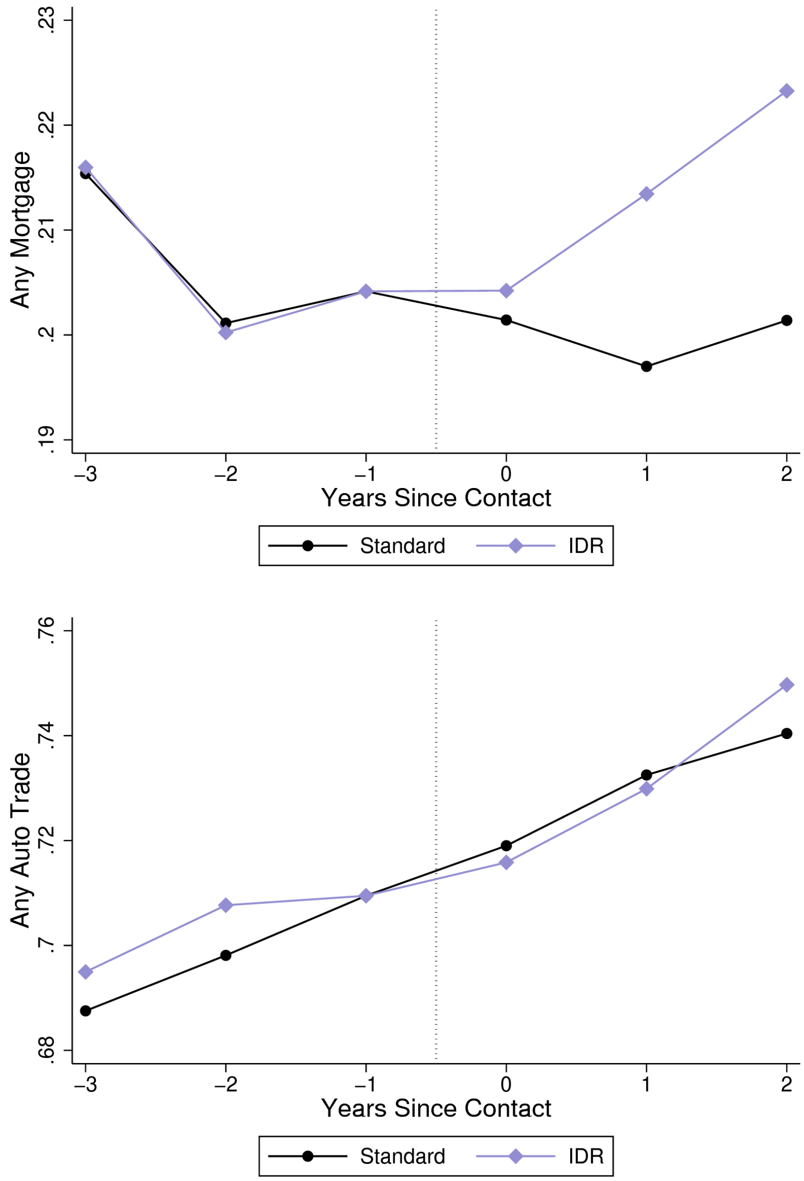

Note: This figure plots the share of IDR enrollees and non-enrollees holding mortgages or auto loans in the analysis sample. The horizontal axis denotes time, in years, relative to the year of the loan servicing call. Outcomes are normalized to the average value of the outcome for non-enrollees in the year 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 Cards

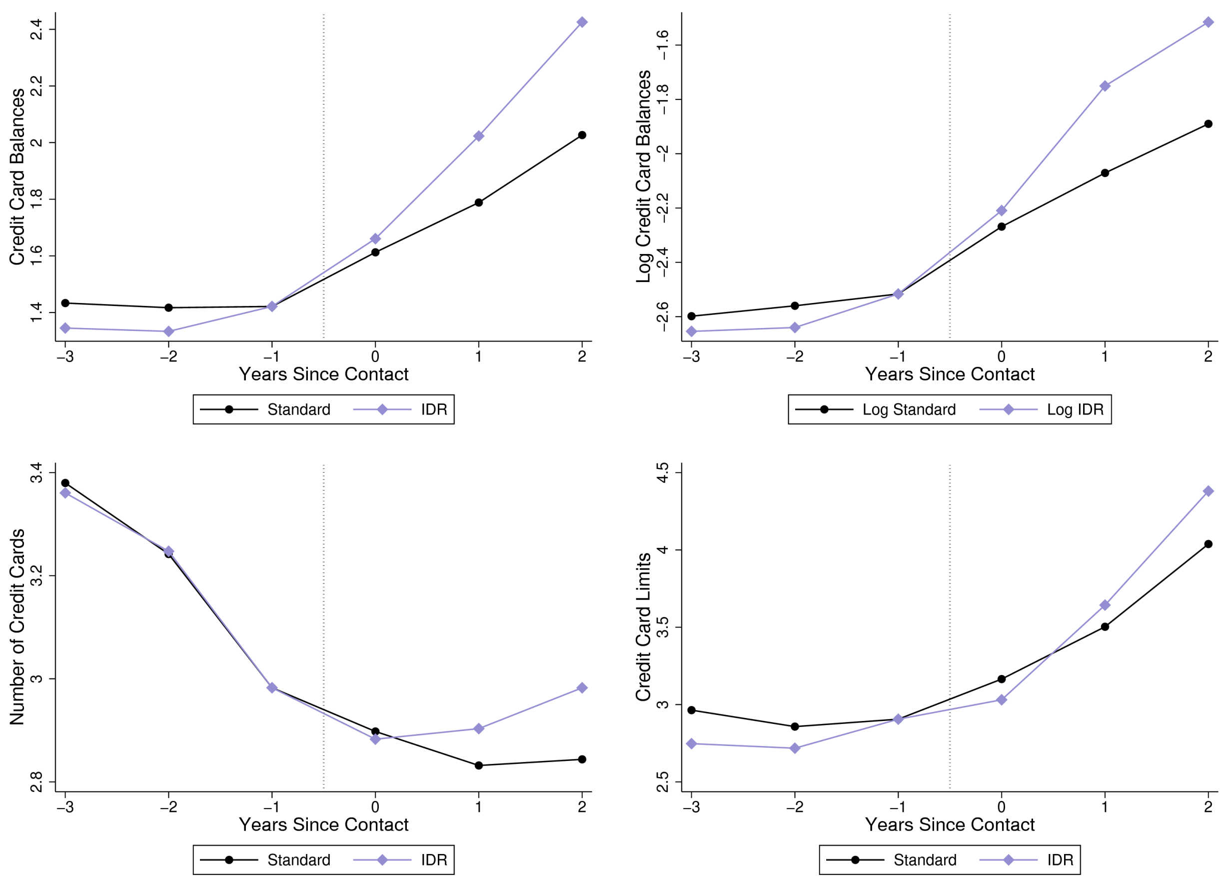

Note: This figure plots the average credit card outcomes for IDR enrollees and non-enrollees in the analysis sample. The horizontal axis denotes time, in years, relative to the year of the loan servicing call. Outcomes are normalized to the average value of the outcome for non-enrollees in the year prior to the call. See Table 1 notes for additional details on the outcome measures and sample.

Figure 6: First Stage

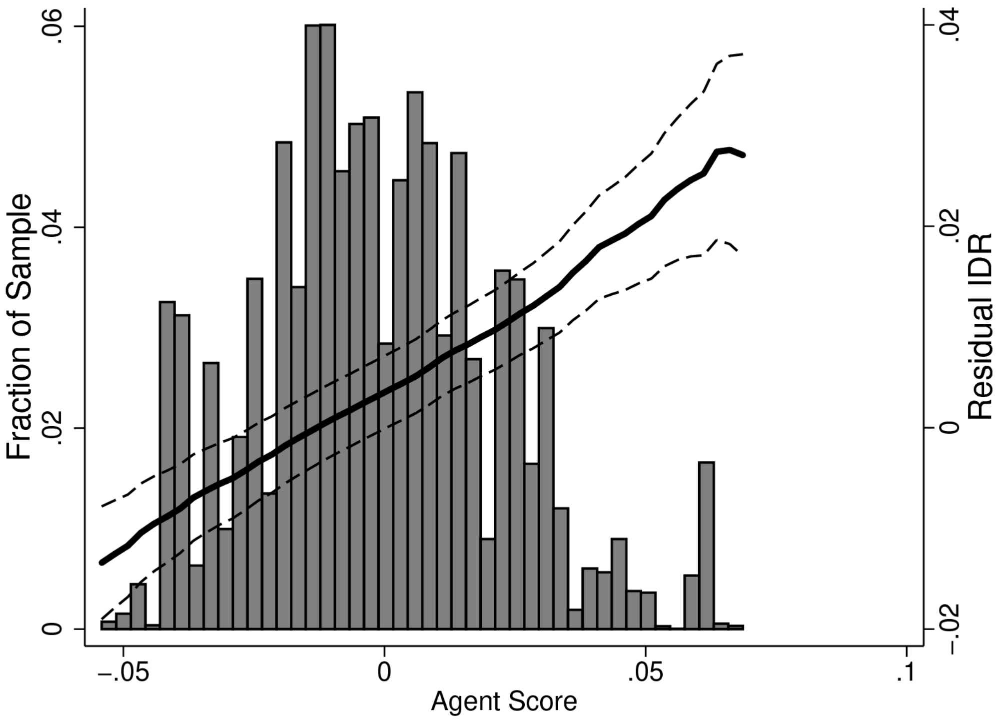

Note: This figure reports first-stage effects and distribution of agent scores across delinquency calls, where agent score is the leave-out mean IDR take-up calculated using data from other calls made by the agent following the procedure described in Section 5.2. The solid and dashed lines, plotted against 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 7: Difference-in-Differences: Repayment Outcomes

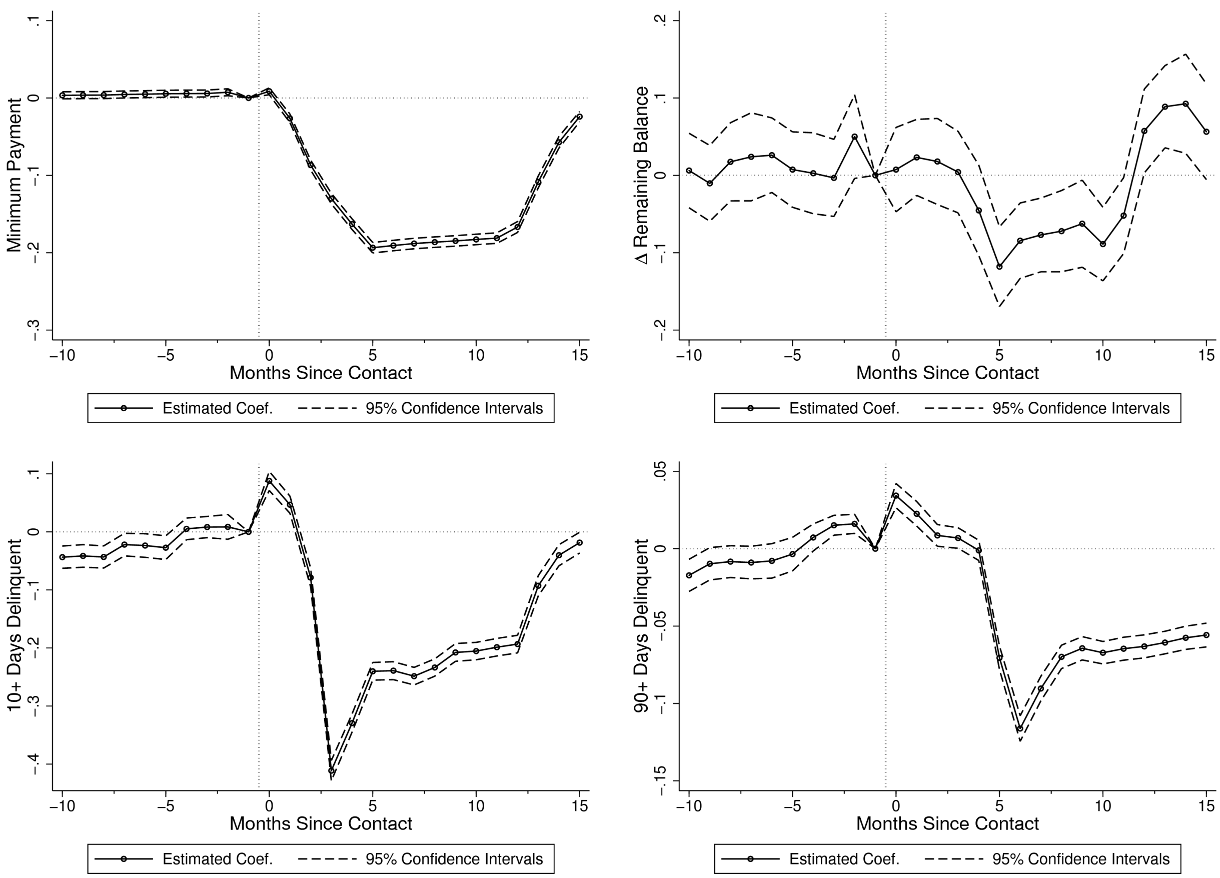

Note: This figure reports estimated coefficients from a series of two-stage least squares regressions. Each point represents the estimated effect of post-call IDR status on the outcome variable at a given time period relative to the date of delinquency call. Relative months are plotted along the x-axis. Dashed lines represent 95% confidence intervals. All regressions include the full set of call date and time fixed effects. Robust standard errors are two-way clustered at the borrower and agent level.

Figure 8: Two-Stage Least Squares: Repayment Outcomes

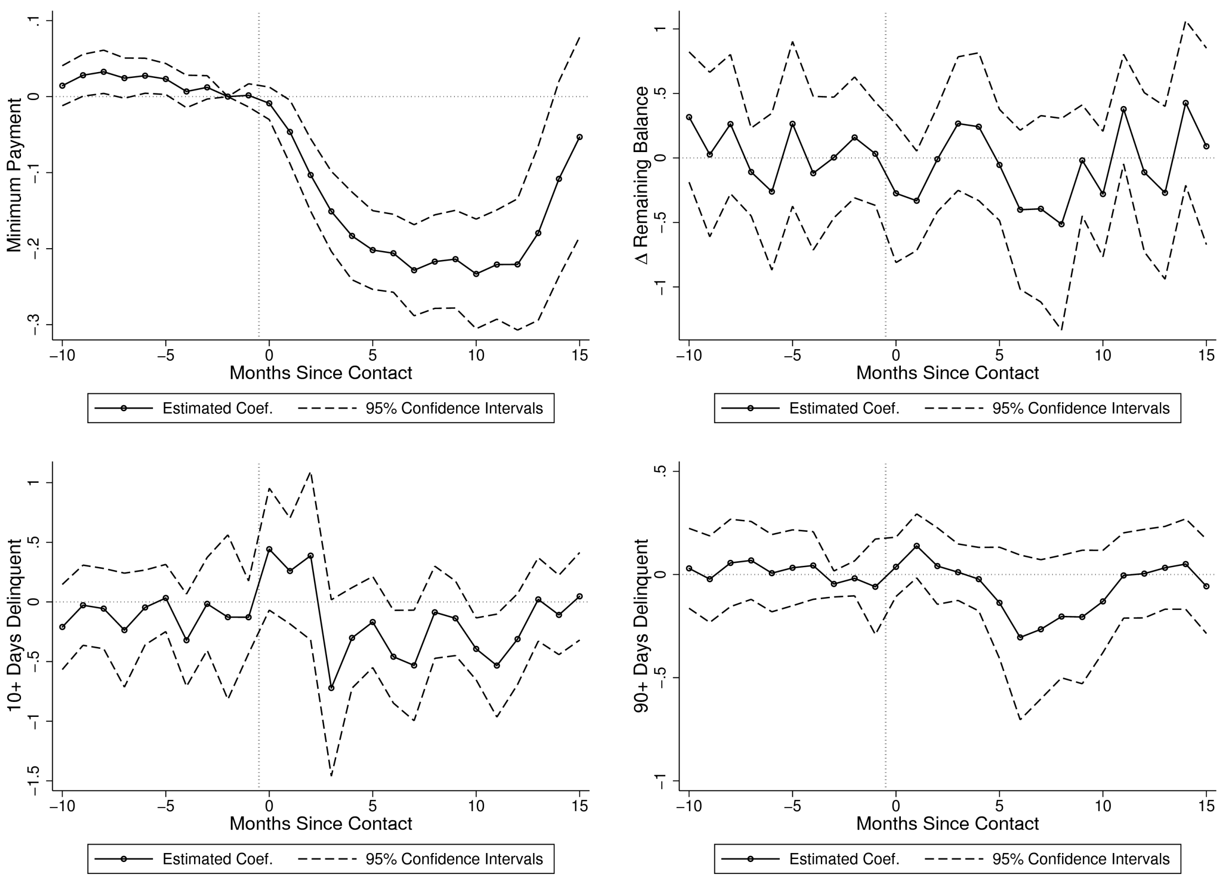

Note: This figure reports estimated coefficients from a series of two-stage least squares regressions. Each point represents the estimated effect of post-call IDR status on the outcome variable at a given time period relative to the date of delinquency call. Relative months are plotted along the x-axis. Dashed lines represent 95% confidence intervals. All regressions include the full set of call date and time fixed effects. Robust standard errors are two-way clustered at the borrower and agent level.

A Additional Tables and Figures

Table A1: Difference-in-Differences Estimates of the Effect of IDR on Loan Repayment Outcomes: Predicted Non-Zero Payments

		Diff-in-Diff			Diff-in-Diff w Linear Trend		
Dependent Variable	(1) Mean $t = -1$	(2) Mos. 4 to 7	(3) Mos. 8 to 11	(4) Mos. 12 to 15	(5) Mos. 4 to 7	(6) Mos. 8 to 11	(7) Mos. 12 to 15
Minimum Payment	0.55	-0.25^{***}	-0.25^{***}	-0.11^{***}	-0.27^{***}	-0.27^{***}	-0.13^{***}
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Remaining Balance	60.55	-0.13	-0.84^{***}	-0.79^{***}	-0.29^{**}	-1.00^{***}	-0.95^{***}
		(0.17)	(0.18)	(0.21)	(0.15)	(0.16)	(0.20)
Δ Remaining Balance	0.26	-0.07	-0.11^{**}	0.16^{**}	-0.12^{***}	-0.16^{***}	0.11^{**}
		(0.06)	(0.06)	(0.07)	(0.04)	(0.02)	(0.04)
10+ Days Delinquent	0.21	-0.17^{***}	-0.12^{***}	-0.03^{**}	-0.18^{***}	-0.13^{***}	-0.05^{***}
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
90+ Days Delinquent	0.01	-0.05^{***}	-0.04^{***}	-0.03^{***}	-0.06^{***}	-0.05^{***}	-0.05^{***}
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)

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 2-4 report coefficients on the effect of IDR in consecutive three-month periods, beginning with four months following the delinquency call from the pooled OLS regression specified in Equation 18. Columns 5-7 report coefficents 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 3,338 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 A2: Difference-in-Differences Estimates of the Effect of IDR on Loan Repayment Outcomes: High Balance Borrowers

		Diff-in-Diff			Diff-in-Diff w Linear Trend		
Dependent Variable	(1) Mean $t = -1$	(2) Mos. 4 to 7	(3) Mos. 8 to 11	(4) Mos. 12 to 15	(5) Mos. 4 to 7	(6) Mos. 8 to 11	(7) Mos. 12 to 15
Minimum Payment	0.33	-0.21^{***}	-0.20^{***}	-0.06^{***}	-0.21^{***}	-0.21^{***}	-0.07^{***}
		(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)
Remaining Balance	46.46	-0.13^{***}	-0.70^{***}	-0.59^{***}	-0.26^{***}	-0.83^{***}	-0.72^{***}
		(0.05)	(0.06)	(0.08)	(0.04)	(0.06)	(0.07)
Δ Remaining Balance	0.21	-0.15^{***}	-0.14^{***}	0.10^{***}	-0.16^{***}	-0.15^{***}	0.09^{***}
		(0.02)	(0.02)	(0.02)	(0.01)	(0.01)	(0.02)
10+ Days Delinquent	0.22	-0.21^{***}	-0.15^{***}	-0.04^{***}	-0.21^{***}	-0.16^{***}	-0.05^{***}
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
90+ Days Delinquent	0.03	-0.07^{***}	-0.07^{***}	-0.06^{***}	-0.08^{***}	-0.08^{***}	-0.07^{***}
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on monthly loan repayment outcomes for borrowers with balances of $20,000 or more. Column 1 reports the dependent variable mean in the month prior to receiving a delinquency call. Columns 2-4 report coefficients on the effect of IDR in consecutive three-month periods, beginning with four months following the delinquency call from the pooled OLS regression specified in Equation 18. Columns 5-7 report coefficents 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 15,002 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: Difference-in-Differences Estimates of the Effect of IDR on Loan Repayment Outcomes: Low Balance Borrowers

		Diff-in-Diff			Diff-in-Diff w Linear Trend		
Dependent Variable	(1) Mean $t = -1$	(2) Mos. 4 to 7	(3) Mos. 8 to 11	(4) Mos. 12 to 15	(5) Mos. 4 to 7	(6) Mos. 8 to 11	(7) Mos. 12 to 15
Minimum Payment	0.11	-0.07^{***}	-0.07^{***}	-0.03^{***}	-0.08^{***}	-0.08^{***}	-0.04^{***}
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Remaining Balance	10.48	-0.06^{***}	-0.20^{***}	-0.19^{***}	-0.10^{***}	-0.23^{***}	-0.23^{***}
		(0.02)	(0.03)	(0.03)	(0.02)	(0.03)	(0.03)
Δ Remaining Balance	0.05	-0.04^{***}	-0.02^{***}	0.02^{**}	-0.04^{***}	-0.02^{***}	0.02^{***}
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
10+ Days Delinquent	0.21	-0.22^{***}	-0.17^{***}	-0.03^{***}	-0.23^{***}	-0.18^{***}	-0.04^{***}
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
90+ Days Delinquent	0.02	-0.08^{***}	-0.08^{***}	-0.08^{***}	-0.09^{***}	-0.09^{***}	-0.08^{***}
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on monthly loan repayment outcomes for borrowers with balances less than \$20,000. Column 1 reports the dependent variable mean in the month prior to receiving a delinquency call. Columns 2-4 report coefficients on the effect of IDR in consecutive three-month periods, beginning with four months following the delinquency call from the pooled OLS regression specified in Equation 18. Columns 5-7 report coeffecients 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 16,111 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: Difference-in-Differences Estimates of the Effect of IDR on Financial Outcomes: High Balance Borrowers

		Diff-in-Diff			Diff-in-Diff w Linear Trend		
Dependent Variable	(1) Mean $t = -1$	(2) Year 0	(3) Year 1	(4) Year 2	(5) Year 0	(6) Year 1	(7) Year 2
Credit Score	597.46	5.26***	8.05***	4.58**	5.17**	7.97***	4.49*
		(1.41)	(1.87)	(2.22)	(2.04)	(2.42)	(2.71)
Bankruptcy	0.20	0.00	−0.01	−0.01	0.00	−0.00	−0.01
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Any Mortgage	0.25	0.00	0.01	0.02	0.01	0.02	0.02*
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Any Auto Trade	0.75	0.01	−0.00	0.01	0.01	−0.00	0.01
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)

Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on yearly financial outcomes for borrowers with balances of $20,000 or more. Column 1 reports the dependent variable mean in the year prior to receiving a delinquency call. Columns 2-4 report coefficients on the effect of IDR in consecutive years, beginning with the year of the delinquency call from the pooled OLS regression specified in Equation 18. Columns 5-7 report coefficents 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 60,006 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: Difference-in-Differences Estimates of the Effect of IDR on Financial Outcomes: Low Balance Borrowers

		Diff-in-Diff			Diff-in-Diff w Linear Trend		
Dependent Variable	(1) Mean $t = -1$	(2) Year 0	(3) Year 1	(4) Year 2	(5) Year 0	(6) Year 1	(7) Year 2
Credit Score	588.95	4.53***	6.83***	7.35***	4.48*	6.78**	7.30**
		(1.61)	(2.13)	(2.38)	(2.31)	(2.77)	(2.89)
Bankruptcy	0.19	−0.01	−0.01	−0.01	−0.01	−0.01	−0.02
		(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
Any Mortgage	0.15	0.00	0.03***	0.03***	−0.00	0.02**	0.03**
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Any Auto Trade	0.66	−0.02*	−0.00	0.00	−0.01	−0.00	0.01
		(0.01)	(0.01)	(0.02)	(0.01)	(0.01)	(0.02)

Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on yearly financial outcomes for borrowers with balances less than \$20,000. Column 1 reports the dependent variable mean in the year prior to receiving a delinquency call. Columns 2-4 report coefficients on the effect of IDR in consecutive years, beginning with the year of the delinquency call from the pooled OLS regression specified in Equation 18. Columns 5-7 report coefficents 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 64,441 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 A6: Difference-in-Differences Estimates of the Effect of IDR on Credit Cards: High Balance Borrowers

		Diff-in-Diff			Diff-in-Diff w Linear Trend		
Dependent Variable	(1) Mean $t = -1$	(2) Year 0	(3) Year 1	(4) Year 2	(5) Year 0	(6) Year 1	(7) Year 2
Credit Card Balances	1.77	0.00	0.26^{**}	0.48^{***}	0.03	0.29^{**}	0.51^{***}
		(0.08)	(0.12)	(0.16)	(0.10)	(0.13)	(0.17)
Log Credit Card Balances	-2.20	0.07	0.33^{***}	0.38^{***}	0.10	0.37^{***}	0.42^{***}
		(0.07)	(0.09)	(0.10)	(0.09)	(0.11)	(0.12)
Number of Credit Cards	3.42	-0.02	0.07	0.17^{**}	-0.02	0.07	0.16^{*}
		(0.04)	(0.06)	(0.08)	(0.05)	(0.07)	(0.08)
Credit Card Limits	3.73	-0.26	0.04	0.23	-0.30	-0.01	0.19
		(0.19)	(0.25)	(0.32)	(0.23)	(0.28)	(0.35)

Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on yearly credit cards for borrowers with balances of $20,000 or more. Column 1 reports the dependent variable mean in the year prior to receiving a delinquency call. Columns 2-4 report coefficients on the effect of IDR in consecutive years, beginning with the year of the delinquency call from the pooled OLS regression specified in Equation 18. Columns 5-7 report coefficents 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 60,006 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 A7: Difference-in-Differences Estimates of the Effect of IDR on Credit Cards: Low Balance Borrowers

		Diff-in-Diff			Diff-in-Diff w Linear Trend		
Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mean $t = -1$	Year 0	Year 1	Year 2	Year 0	Year 1	Year 2
Credit Card Balances	1.06	0.10	0.18^{*}	0.25^{*}	-0.01	0.08	0.14
		(0.08)	(0.11)	(0.13)	(0.10)	(0.13)	(0.15)
Log Credit Card Balances	-2.83	0.07	0.32^{***}	0.36^{***}	0.00	0.26^{**}	0.30^{**}
		(0.07)	(0.10)	(0.11)	(0.09)	(0.11)	(0.13)
Number of Credit Cards	2.57	0.00	0.10	0.13^{*}	-0.04	0.05	0.09
		(0.04)	(0.06)	(0.08)	(0.04)	(0.06)	(0.08)
Credit Card Limits	2.09	0.01	0.24	0.42^{*}	-0.15	0.08	0.25
		(0.11)	(0.17)	(0.23)	(0.14)	(0.19)	(0.25)

Note: This table reports difference-in-differences estimates of the effect of IDR enrollment on yearly credit cards for borrowers with balances less than $20,000. Column 1 reports the dependent variable mean in the year prior to receiving a delinquency call. Columns 2-4 report coefficients on the effect of IDR in consecutive years, beginning with the year of the delinquency call from the pooled OLS regression specified in Equation 18. Columns 5-7 report coefficents 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 64,441 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 A8: Difference-in-Differences and Instrumental Variables Estimates of the Effect of IDR on Financial Outcomes

Dependent Variable	(1) Mean $t = -1$	Difference-in-Differences (2) Year 0	(3) Year 1	(4) Year 2	Instrumental Variables (5) Year 0	(6) Year 1	(7) Year 2
Credit Score	593.26	5.01***	7.53***	5.93***	−8.83	−20.07	−28.65
		(1.06)	(1.40)	(1.63)	(18.29)	(18.40)	(28.09)
Bankruptcy	0.19	−0.00	−0.00	−0.00	0.10	0.14	0.17
		(0.00)	(0.00)	(0.00)	(0.15)	(0.17)	(0.19)
Any Mortgage	0.20	0.00	0.02**	0.02***	0.03	−0.02	−0.07
		(0.00)	(0.01)	(0.01)	(0.09)	(0.12)	(0.13)
Any Auto Trade	0.71	−0.00	−0.00	0.01	0.03	−0.12	−0.16
		(0.01)	(0.01)	(0.01)	(0.16)	(0.16)	(0.15)
Call Time FE		Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes
N	15953	15953	15953	15953	15953	15953	15953

Note: This table reports difference-in-differences and two-stage least squares 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 2-4 report coefficients on the effect of IDR in consecutive years, beginning with the year of the delinquency call from the pooled OLS regression specified in Equation 18. Each of Columns 5-7 report estimates from separate two-stage least squares regressions on outcomes in the same years. Regressions are estimated on the analysis sample as described in the notes to Table 1. Two-stage least squares models instrument for IDR with the agent score calculated using data from other calls made by the agent following the procedure described in Section 5.2. All specifications include controls from 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 A9: Difference-in-Differences and Instrumental Variables Estimates of the Effect of IDR on Credit Cards

		Difference-in-Differences			Instrumental Variables		
Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mean $t = -1$	Year 0	Year 1	Year 2	Year 0	Year 1	Year 2
Credit Card Balances	1.42	0.05	0.24***	0.40***	−1.55	−1.52	−0.44
		(0.05)	(0.08)	(0.11)	(1.13)	(1.65)	(1.86)
Log Credit Card Balances	-2.51	0.07	0.33***	0.38***	−0.53	−0.03	0.47
		(0.05)	(0.07)	(0.07)	(1.05)	(1.32)	(1.42)
Number of Credit Cards	3.00	−0.01	0.07	0.14**	−0.39	−0.36	0.15
		(0.03)	(0.04)	(0.06)	(0.49)	(0.79)	(1.02)
Credit Card Limits	2.92	−0.13	0.14	0.34	−3.72**	−2.58	−1.06
		(0.12)	(0.16)	(0.21)	(1.88)	(2.65)	(3.49)
Call Time FE		Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes	Yes	Yes	Yes	Yes	Yes
N	15953	15953	15953	15953	15953	15953	15953

Note: This table reports difference-in-differences and two-stage least squares estimates of the effect of IDR enrollment on yearly credit cards. Column 1 reports the dependent variable mean in the year prior to receiving a delinquency call. Columns 2-4 report coefficients on the effect of IDR in consecutive years, beginning with the year of the delinquency call from the pooled OLS regression specified in Equation 18. Each of Columns 5-7 report estimates from separate two-stage least squares regressions on outcomes in the same years. Regressions are estimated on the analysis sample as described in the notes to Table 1. Two-stage least squares models instrument for IDR with the agent score calculated using data from other calls made by the agent following the procedure described in Section 5.2. All specifications include controls from 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.

Figure A1: Standard versus IDR Payments among IDR-Enrollees

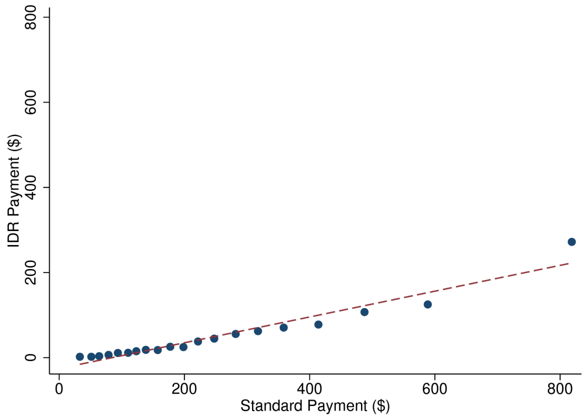

(A) Relative Payment Size

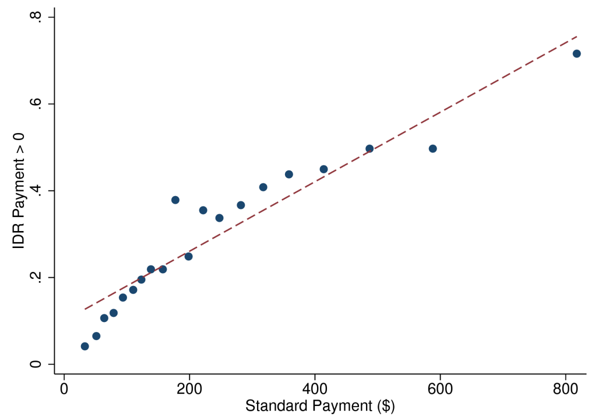

(B) Share with Positive IDR Payments

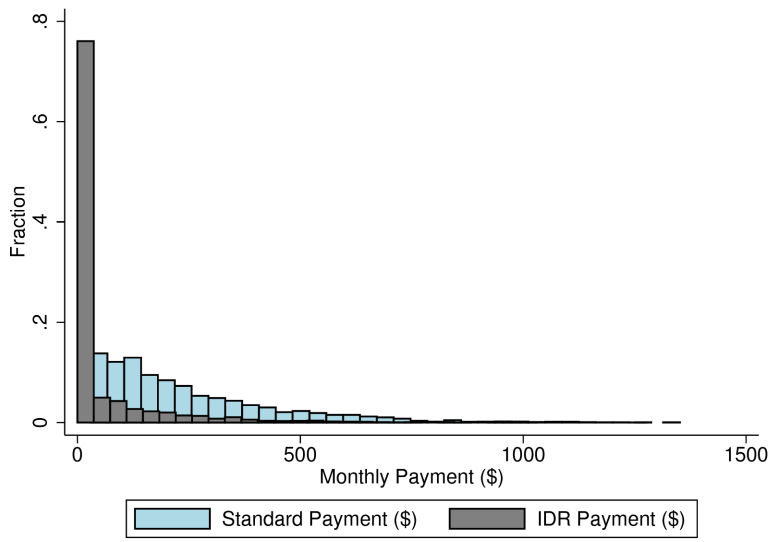

(C) Distribution of Payments

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. Panel A plots average standard payment size against average IDR payment size. Panel B plots average standard payment size against the share of individuals with IDR payments greater than zero. Panel C plots histograms for standard and IDR payments. See Table 1 notes for additional details on the sample.

Figure A2: Pre/Post-Call Trends in Repayment Outcomes: High Balance Borrowers

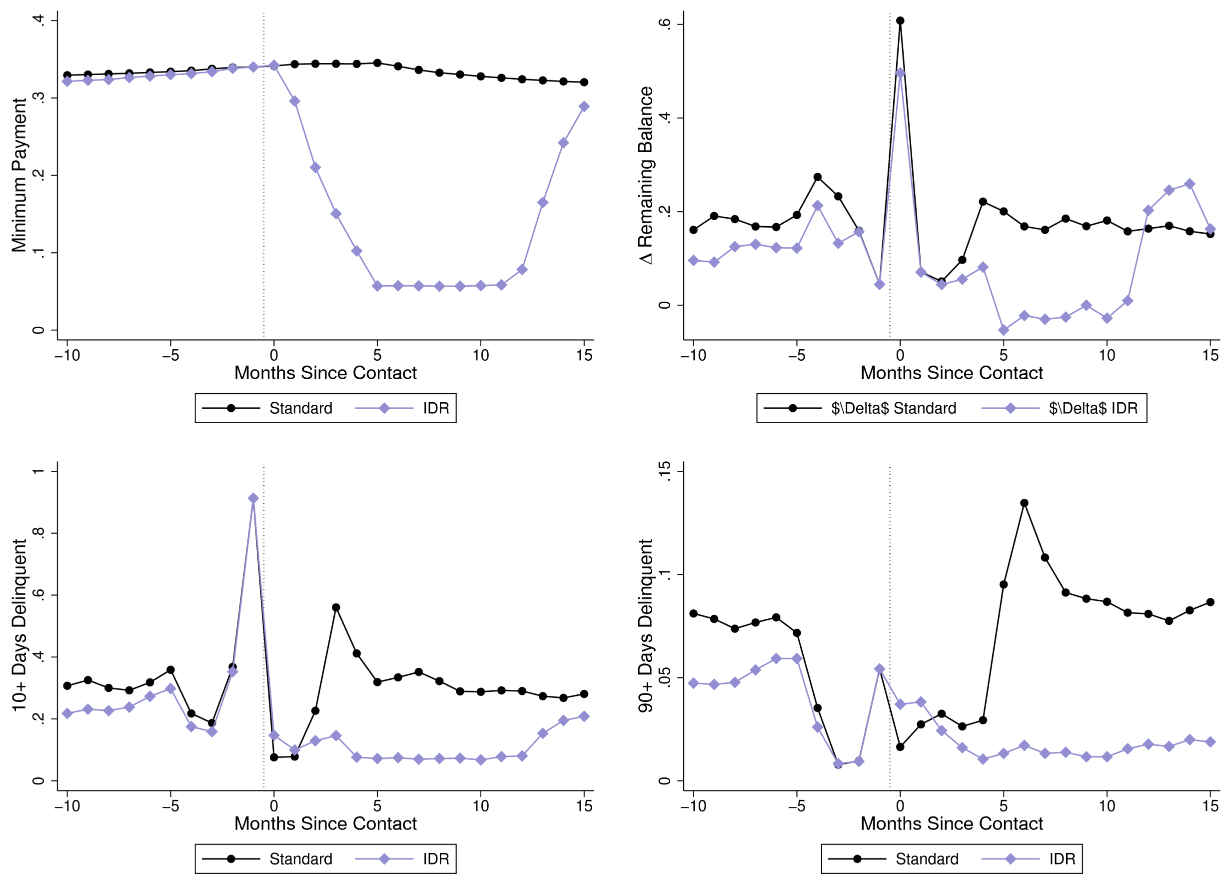

Note: This figure plots the average loan repayment outcomes for IDR enrollees and non-enrollees with $20,000 or more in debt balance in the analysis sample. The horizontal axis denotes time, in months, relative to the month of the loan servicing call. Outcomes are normalized to the average value of the 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 A3: Pre/Post-Call Trends in Repayment Outcomes: Low Balance Borrowers

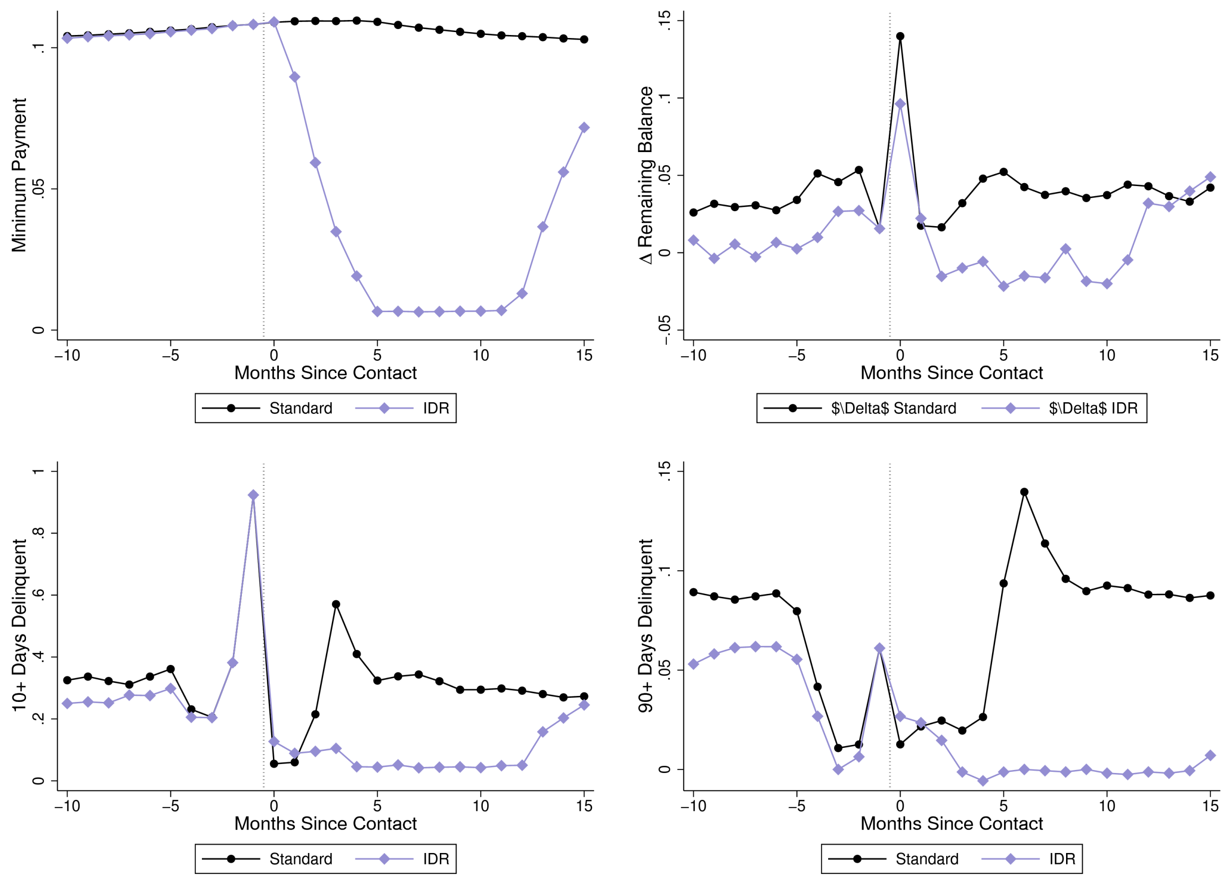

Note: This figure plots the average loan repayment outcomes for IDR enrollees and non-enrollees with less than $20,000 in debt balance in the analysis sample. The horizontal axis denotes time, in months, relative to the month of the loan servicing call. Outcomes are normalized to the average value of 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 A4: Difference-in-Differences: Financial Outcomes

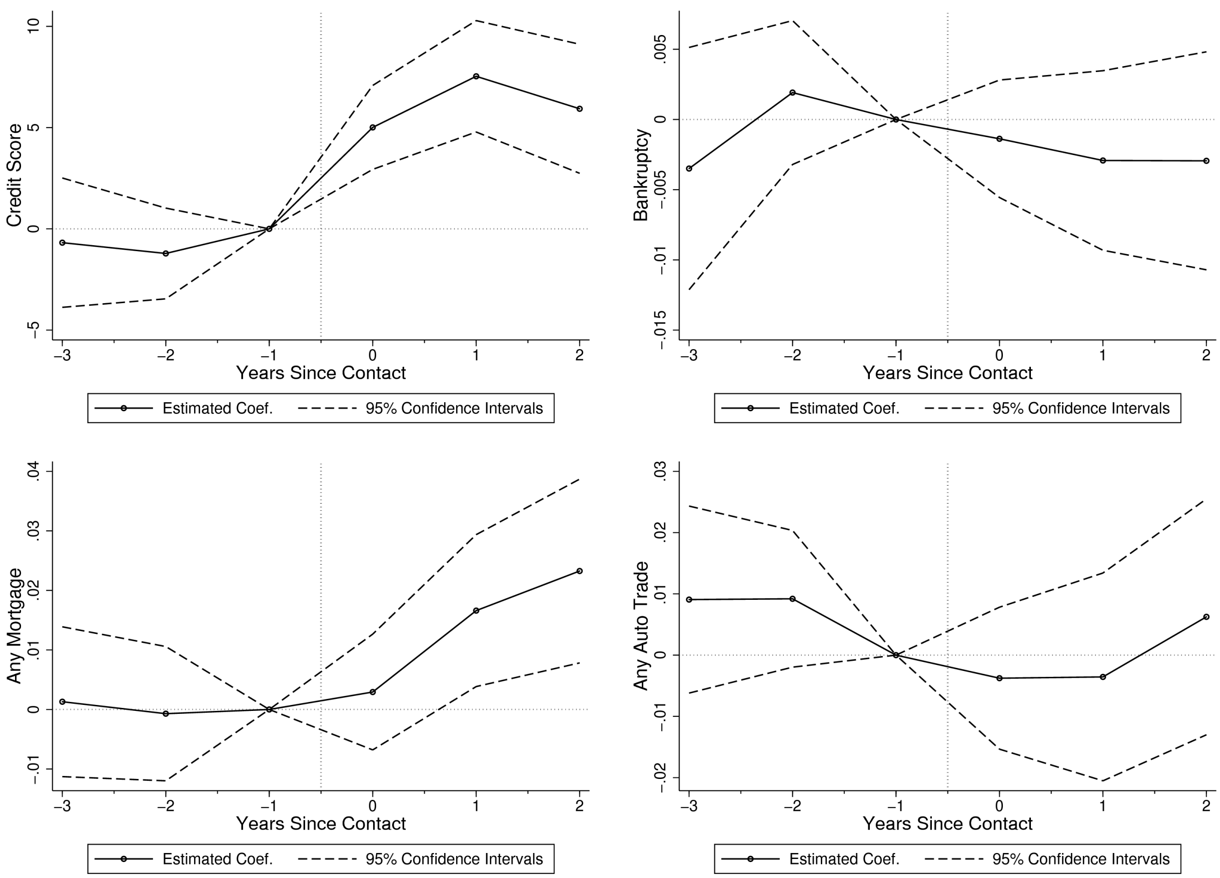

Note: This figure reports estimated difference-in-differences coefficients for the effect of IDR on financial outcomes. Each point represents the estimated effect of post-call IDR status on the outcome variable at a given time period relative to the date of delinquency call. Relative years are plotted along the x-axis. Dashed lines represent 95% confidence intervals. All regressions include the full set of call date and time fixed effects. Robust standard errors are two-way clustered at the borrower and agent level.

Figure A5: Difference-in-Differences: Credit Cards

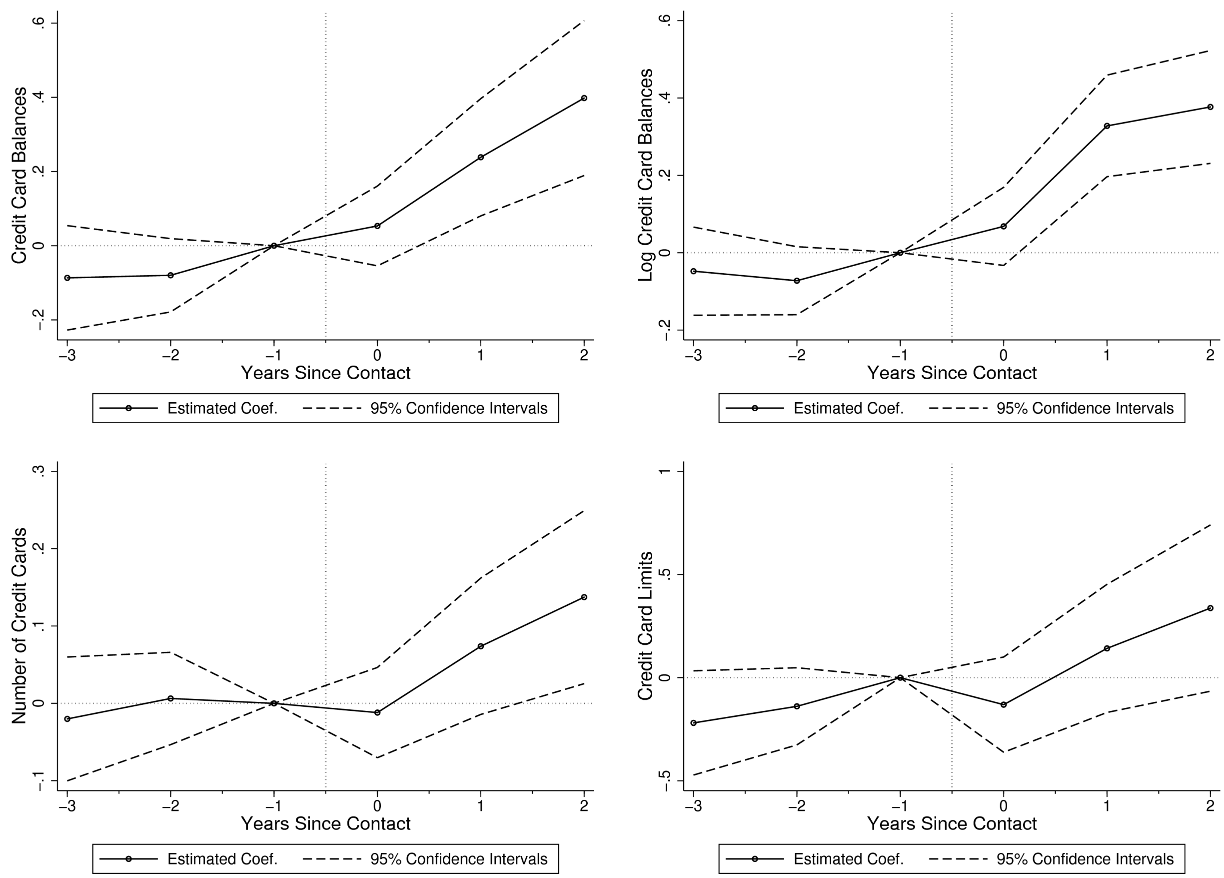

Note: This figure reports estimated difference-in-differences coefficients for the effect of IDR on credit card outcomes. Each point represents the estimated effect of post-call IDR status on the outcome variable at a given time period relative to the date of delinquency call. Relative years are plotted along the x-axis. Dashed lines represent 95% confidence intervals. All regressions include the full set of call date and time fixed effects. Robust standard errors are two-way clustered at the borrower and agent level.

B Mathematical Appendix

B.1 Model Solution

The agent's problem is given by:

$$V\left(y_t, D_t, A_t\right) = \max_{c_t, A_{t+1}} \left\{u\left(c_t\right) + E\left[V\left(y_{t+1}, D_{t+1}, A_{t+1}\right)\right]\right\} \tag{24}$$

$$\begin{aligned} A_{t+1} &\geq L & (25)\\ c_t &= A_t - A_{t+1} + y_t - x_t & (26)\\ D_{t+1} &= D_t - x_t & (27) \end{aligned}$$

The Lagrangian for this problem can be written as:

$$\begin{aligned} \mathcal{L} = {} & u\left(c_t\right) + E\left[V\left(y_{t+1}, D_{t+1}, A_{t+1}\right)\right] \\ & - \lambda_1\left(L - A_{t+1}\right) - \lambda_2\left(A_t - A_{t+1} + y_t - x_t - c_t\right) - \lambda_3\left(D_t - D_{t+1} - x_t\right) \end{aligned} \tag{28}$$

where $\lambda_1 \geq 0$, $\lambda_2 > 0$, and $\lambda_3 > 0$. The first-order condition for consumption c_t is:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial c_t}\left(y_{t+1}, D_{t+1}, A_{t+1}, A_t, c_t\right) &= 0 \\ u'(c_t) &= -\lambda_2 & (29) \end{aligned}$$

and the first-order condition for next period's assets A_{t+1} is:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial A_{t+1}}\left(y_{t+1}, D_{t+1}, A_{t+1}, A_t, c_t\right) &= 0 \\ \frac{dE\left[V\left(y_{t+1}, D_{t+1}, A_{t+1}\right)\right]}{dA_{t+1}} &= -\lambda_1 - \lambda_2 & (30) \end{aligned}$$

Applying the envelope theorem to Equation 30 and combining with Equation 29 yields:

$$u'(c_t) = E\left[\frac{\partial V}{\partial A}(y_{t+1}, A_{t+1}, D_{t+1})\right] + \lambda_1 \tag{31}$$

Since $\lambda_1 > 0$ if and only if $A_{t+1} = L$, Equation 31 implies:

$$u'(c_t) = E\left[\frac{\partial V}{\partial A}(y_{t+1}, A_{t+1}, D_{t+1})\right] \text{ if } A_{t+1} > L \tag{32}$$

$$u'(c_t) > E\left[\frac{\partial V}{\partial A}(y_{t+1}, A_{t+1}, D_{t+1})\right] \text{ if } A_{t+1} = L \tag{33}$$

which, along with conditions 25, 26, and 27, characterize the solution to the model.

B.2 Proof of Liquidity Effects

Note that the solution from the previous section implies:

$$c_t = \begin{cases} E\left[c_{t+1}\right] & \text{if } A_{t+1} > L \\ A_t - L + y_t - x_t & \text{if } A_{t+1} = L \end{cases} \tag{34}$$

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

$$A_t + y_t - x_t - c_t^* > L \tag{35}$$

$$c_t^* < A_t - L + y_t - x_t \tag{36}$$

$$E\left[c_{t+1}\right] < A_t - L + y_t - x_t \tag{37}$$

We can therefore rewrite (34) as:

$$c_t = \min\left\{A_t - L + y_t - x_t, E\left[c_{t+1}\right]\right\} \tag{38}$$

Iterating forward, we have

$$c_t = \min\left\{A_t - L + y_t - x_t, E\left[\min\left\{A_{t+1} - L + y_{t+1} - x_{t+1}, E\left[c_{t+2}\right]\right\}\right]\right\} \tag{39}$$

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:

$$E\left[x_t^I\right] = E\left[\min\left\{\theta y_t, \frac{D_0}{N}\right\}\right] \leq \frac{D_0}{N} = E\left[x_t^S\right] \qquad \forall t \leq N \tag{40}$$

$$E\left[x_t^I\right] = E\left[\min\left\{\theta y_t, \frac{D_0}{N}, D_t\right\}\right] \geq 0 = E\left[x_t^S\right] \qquad \forall t > N \tag{41}$$

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 38 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 34.

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 39 an additional period, so:

$$c_t = E\left[\min\left\{A_{t+1} - L + y_{t+1} - x_{t+1}, E\left[\min\left\{A_{t+2} - L + y_{t+2} - x_{t+2}, E\left[c_{t+3}\right]\right\}\right]\right\}\right] \tag{42}$$

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 42:

$$\frac{dc_t}{dx_q} \geq \frac{dc_t}{dx_{q'}} \quad \forall\, q' > q \tag{43}$$

Any change in payments such that $\Delta x_q < \Delta x_{q'}$ but $\sum_t^\infty \Delta x_q = 0$ 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.

B.2.1 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, $m \equiv y - x$.

$$\begin{aligned}
\text{Var}(m^I) - \text{Var}(m^S) &= \int_{-\infty}^{\infty} \left(y - x^I\right)^2 dF_y - \int_{-\infty}^{\infty} \left(y - x^S\right)^2 dF_y &(44)\\
&= \int_{-\infty}^{\infty} \left(y - \max\left\{(1-\theta)y, y - x_t^S\right\}\right)^2 dF_y - \int_{-\infty}^{\infty} \left(y - x^S\right)^2 dF_y &(45)\\
&= \int_{-\infty}^{\frac{x^S}{\theta}} \left[(y - \theta y)^2 - \left(y - x^I\right)^2\right] dF_y < 0 &(46)
\end{aligned}$$

Rewriting Equation 39 in terms of net income, we have:

$$c_t = \min\left\{A_t - L + m_t, E\left[\min\left\{A_{t+1} - L + m_{t+1}, E\left[c_{t+2}\right]\right\}\right]\right\} \tag{47}$$

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 $E\left[\min\left\{A_{t+1} - L + m_{t+1}, E\left[c_{t+2}\right]\right\}\right]$, inducing less precautionary savings and more consumption in period t.