American Economic Review 2020, 110(4): 984–1018
https://doi.org/10.1257/aer.20171541

Targeted Debt Relief and the Origins of Financial Distress: Experimental Evidence from Distressed Credit Card Borrowers[†]

By Will Dobbie and Jae Song*

We study the drivers of financial distress using a large-scale field experiment that offered randomly selected borrowers a combination of (i) immediate payment reductions to target short-run liquidity constraints and (ii) delayed interest write-downs to target long-run debt constraints. We identify the separate effects of the payment reductions and interest write-downs using both the experiment and cross-sectional variation in treatment intensity. We find that the interest write-downs significantly improved both financial and labor market outcomes, despite not taking effect for three to five years. In sharp contrast, there were no positive effects of the more immediate payment reductions. These results run counter to the widespread view that financial distress is largely the result of short-run constraints. (*JEL* G56, K35)

Financial distress is extraordinarily common in the United States. Over one-third of Americans have a debt in collections, and more than one in ten will file for bankruptcy at some point during their lives. Americans are also severely liquidity constrained, with approximately one-quarter of households unable to come up with \$2,000 to cope with an unexpected need (Lusardi, Schneider, and Tufano 2011). As a result, there is a widespread view that liquidity constraints are the most important driver of financial distress and that debt relief will be most effective when it targets

* Dobbie: Harvard Kennedy School and NBER (email: will_dobbie@hks.harvard.edu); Song: Social Security Administration (email: jae.song@ssa.gov). A previous version of this paper was circulated under the title "Debt Relief or Debt Restructuring? Evidence from an Experiment with Distressed Credit Card Borrowers." Esther Duflo was the coeditor for this article. We gratefully acknowledge four anonymous referees for many valuable insights and suggestions. We are also grateful to Ann Woods and Robert Kaplan at Money Management International, David Jones at the Association of Independent Consumer Credit Counseling Agencies, Ed Falco at Auriemma Consulting Group, Jennifer Werkley at TransUnion, and Gerald Ray and David Foster at the Social Security Administration for their help and support. We thank Tal Gross, Matthew Notowidigdo, and Jialan Wang for providing the bankruptcy data used in this analysis. We also thank Josh Angrist, Leah Platt Boustan, Hank Farber, James Feigenbaum, Paul Goldsmith-Pinkham, Tal Gross, Peter Hull, Larry Katz, Ben Keys, Patrick Kline, Ilyana Kuziemko, Michal Kolesár, Alex Mas, Jesse Shapiro, Andrei Shleifer, Crystal Yang, Jonathan Zinman, Eric Zwick, and numerous seminar participants for helpful comments and suggestions. Emily Battaglia, Kevin DeLuca, Nicole Gandre, Daniel Herbst, Disa Hynsjo, Samsun Knight, Ashley Litwin, James Reeves, Kevin Tang, Daniel Van Deusen, Amy Wickett, Alice Wu, and Yining Zhu provided excellent research assistance. Financial support from the Washington Center for Equitable Growth is gratefully acknowledged. The project received IRB approval at Princeton University (6657). The experiment was registered in the AEA RCT Registry (AEARCTR-0005045). Correspondence can be addressed to the authors by email. Any opinions expressed herein are those of the authors and not those of the Social Security Administration.

[†] Go to https://doi.org/10.1257/aer.20171541 to visit the article page for additional materials and author disclosure statements.

these short-run constraints. This view has important implications for understanding both the growing level of financial distress in the United States and the optimal design of debt relief programs such as consumer bankruptcy. In this paper, however, we show that this view significantly overstates the benefits of debt relief targeting short-run liquidity constraints, while significantly understating the benefits of debt relief targeting longer-run financial constraints, such as the distortionary effects of excessive debt (so-called "debt overhang").

Estimating the effects of targeted debt relief is challenging because most debt relief programs are designed to address both short- and long-run financial constraints at the same time. For example, consumer bankruptcy protection offers both lower minimum payments (to address short-run liquidity constraints) and generous debt write-downs (to address longer-run debt overhang). As a result, standard "black box" estimates of consumer bankruptcy cannot be used to predict the effects of specific types of targeted debt relief or to understand the relative importance of addressing short- or long-run financial constraints alone. An added complication is that most debt relief recipients are negatively selected, biasing cross-sectional comparisons, and many of the most proximate causes of seeking debt relief, such as job loss and expense shocks, also impact later outcomes, biasing within-individual comparisons.

In this paper, we overcome these challenges using information from a randomized field experiment matched to administrative tax, bankruptcy, and credit records (Dobbie and Woods 2019). The experiment was designed and implemented by a large nonprofit credit counseling organization in the context of an important but understudied debt relief program called the Debt Management Plan (DMP). The DMP is a structured repayment program that allows distressed borrowers to simultaneously repay all of their outstanding credit card debt over a three-to-five-year period. In exchange for enrolling in a DMP, credit card issuers will lower the minimum payment amount at the beginning of the repayment program (to address short-run liquidity constraints) and provide a partial write-down of interest payments and late fees at the end of the repayment program (to address longer-run debt overhang). Each year, more than 600,000 individuals repay between \$1.5 and \$2.5 billion of credit card debt through these repayment programs (Wilshusen 2011).

During the experiment, borrowers in both the treatment and control groups were offered a repayment program. While control borrowers were offered the status quo repayment program that had been offered to all borrowers prior to the randomized trial, treated borrowers were offered a much more generous repayment program that included a combination of two different types of targeted debt relief: (i) immediate minimum payment reductions meant to address short-run liquidity constraints and (ii) delayed interest write-downs meant to address longer-run debt overhang. The additional debt relief provided by the experiment was substantial: the maximum payment reductions in the treatment group were \$92 (21.0 percent) per month larger than those in the status quo program, while the maximum interest write-downs in the treatment group were \$4,302 (100.0 percent) larger than those in the status quo program.

We estimate the separate impact of the interest write-downs and minimum payment reductions using variation in potential treatment intensity across individuals.

The variation in treatment intensity comes from two sources. First, each of the credit card issuers participating in the randomized trial offered a different combination of interest write-downs and minimum payment reductions to treated borrowers. Second, individual borrowers made different decisions about how much to borrow from each of these credit card issuers before the experiment began. These decisions translated into economically significant differences in the interest write-downs and minimum payment reductions offered to the treatment group. For each individual in the sample, how much they borrowed from each participating creditor determined what specific treatment intensity they would receive in these two dimensions, if treated. For example, treated borrowers at the seventy-fifth percentile of the interest write-down distribution received write-downs that were \$1,521 larger than treated borrowers at the twenty-fifth percentile of the distribution. Similarly, treated borrowers at the seventy-fifth percentile of the minimum payment distribution received payment reductions that were \$33 per month larger than treated borrowers at the twenty-fifth percentile of the distribution. We compare the effects of the randomized treatment eligibility across individuals with these higher and potential lower treatment intensities to identify the separate effects of the interest write-downs and minimum payment reductions.

The critical identification assumption underlying our research design is that the causal effects being offered treatment are uncorrelated with the specific treatment bundle offered (treatment intensity). This is a strong assumption that could plausibly be violated for several reasons. For example, individuals who borrow from credit card issuers offering more generous interest write-downs and payment reductions in the treatment group may be more responsive to debt relief than individuals who borrow from issuers offering less generous interest write-downs in the treatment group. To account for such selection concerns, our preferred specifications compare individuals with the same set of credit cards, but different proportions of debt on each credit card. These specifications weaken the identifying assumption by only requiring that the proportion of debt with each card issuer be as-good-as-randomly assigned with respect to treatment effect heterogeneity, and not the initial choice of which credit cards to hold. Consistent with this weaker version of the identifying assumption, we find that treatment intensity is uncorrelated with observable borrower characteristics within these "creditor risk sets" and that our estimates are robust to a set of sharp overidentification tests once we include an exhaustive set of creditor risk set fixed effects.[1]

We begin our empirical analysis by estimating the effect of treatment eligibility on repayment, bankruptcy, collections debt, credit scores, and labor market outcomes. We find that treatment eligibility increased the probability of starting and finishing the repayment program, and decreased the probability of filing for bankruptcy. There were no detectable effects of treatment eligibility on collections debt, credit scores, or labor market outcomes, although large standard errors mean that we cannot rule out economically significant effects in both directions.

[1] Our approach builds on recent work identifying the causal effect of a single treatment relative to multiple fallback options (e.g., Kline and Walters 2016; Kirkebøen, Leuven, and Mogstad 2016; Hull 2018) and potential mediators in multi-site experiments (e.g., Reardon and Raudenbush 2013).

We then estimate the separate impact of the minimum payment reductions and interest write-downs using the interaction of treatment eligibility and treatment intensity. We find that the interest write-downs significantly improved both financial and labor market outcomes, particularly for the highest-debt borrowers, despite not taking effect for three to five years. For these high-debt borrowers, we find that the maximum interest write-down in the treatment group increased the probability of finishing a repayment program by 4.4 percentage points (30.8 percent) and decreased the probability of filing for bankruptcy by 3.5 percentage points (33.3 percent). The probability of having collections debt also decreased by 1.2 percentage points (3.1 percent) for these high-debt borrowers, while the probability of being employed increased by 4.2 percentage points (5.1 percent). The estimated effects of the interest write-downs on credit scores and earnings are smaller and not statistically significant for all borrowers. Taken together, however, our results indicate that there are significant benefits of debt relief targeting long-run debt overhang.

In sharp contrast, we find no positive effects of the more immediate payment reductions targeting short-run liquidity constraints. The maximum payment reduction in the treatment group only increased the probability of finishing a repayment program by a statistically insignificant 0.3 percentage points (2.1 percent) in the full sample, with similar null effects among the highest-debt borrowers. The maximum payment reduction in the treatment group also *increased* the probability of filing for bankruptcy in this sample by a statistically significant 2.3 percentage points (21.9 percent) and *decreased* credit scores by a statistically insignificant 1.9 points in our preferred specifications. There are also no detectable positive effects of the payment reductions on collections debt, employment, or earnings. In sum, there is no evidence that borrowers in our sample benefited from the payment reductions, and even some evidence that borrowers seem to have been hurt by these reductions.

We show that these null results can be explained by the unintended negative effect of increasing the number of months a borrower remains in the repayment program. The payment reductions increased the length of the repayment program in the treatment group by an average of four months and, as a result, increased the number of months where a treated borrower could be hit by an adverse shock that causes default (e.g., job loss). We find that the positive effects of increased liquidity in the treatment group were nearly exactly offset by the negative effects of this increased exposure to default risk. These results help to reconcile our findings with the vast literature documenting liquidity constraints in a variety of settings (e.g., Gross and Souleles 2002; Johnson, Parker, and Souleles 2006; Agarwal, Liu, and Souleles 2007; Parker et al. 2013; Agarwal et al. 2015; Gross, Notowidigdo, and Wang forthcoming), while indicating that the potential benefits of targeting these short-run constraints may have been significantly overstated, at least in our setting.

Our results contribute to an emerging literature estimating the "black box" effects of consumer bankruptcy protection, which, as mentioned above, addresses both short- and long-run financial constraints at the same time. Consistent with our findings, bankruptcy protection increases post-filing earnings and decreases both post-filing mortality and financial distress (Dobbie and Song 2015; Dobbie, Keys, and Mahoney 2017). There is also evidence that the availability of consumer bankruptcy as an outside option provides implicit health (Gross and Notowidigdo 2011, Mahoney 2015), consumption (Auclert, Dobbie, and Goldsmith-Pinkham 2019), and mortgage insurance (Li, White, and Zhu 2011). However, none of these papers are able to identify the effects of targeting either short-run liquidity constraints or long-run debt overhang alone.

This paper is also related to recent work estimating the effects of debt relief in the mortgage market. Mortgage modifications made through the Home Affordable Modification Program (HAMP) modestly decreased both mortgage and nonmortgage defaults, although it is unclear whether the effects were driven by lower minimum payments or lower debt burdens (Agarwal et al. 2017). More recent work suggests that the principal write-downs made through HAMP had no impact on underwater borrowers (Ganong and Noel 2018), while both cross-sectional regressions and theoretical work suggest that principal forgiveness may be effective for borrowers who are not underwater (Eberly and Krishnamurthy 2014; Haughwout, Okah, and Tracy 2016). While our results are broadly consistent with this literature, we caution against generalizing our results to the mortgage market. It is possible, for example, that liquidity constraints may be more important in the mortgage market, where delinquent borrowers often have fewer outside options than otherwise similar credit card borrowers.

The remainder of this paper is structured as follows. Section I describes the institutional setting, data, and experimental design. Section II describes our empirical design. Section III presents our main results, Section IV explores potential mechanisms, and Section V concludes. The online Appendix provides additional results and econometric proofs.

I. Background and Experimental Design

A. *Background*

The randomized experiment described in this paper was designed and implemented by Money Management International (MMI), the largest nonprofit credit counseling agency in the United States. In the early 1950s, the first nonprofit credit counseling organizations were established to increase credit card repayment rates and decrease the number of new bankruptcy filings. Today, nonprofit credit counseling organizations such as MMI provide a wide range of services to their clients via phone and in-person sessions, including credit counseling, bankruptcy counseling, and foreclosure counseling.

One of the most important products offered by nonprofit credit counselors is the debt management plan (DMP), a structured repayment program that simultaneously repays all of a borrower's outstanding credit card debt over three to five years.[2] Under the DMP, the credit counseling agency negotiates directly with each of the borrower's credit card issuers to lower the minimum payment amount

[2] Under current regulatory guidelines, the term length for a DMP cannot exceed five years. If borrowers cannot fully repay their credit card debts within this five-year limit, they cannot participate in a DMP unless the creditor is willing to write off a portion of the original balance and recognize the loan as impaired. To date, however, creditors have typically been unwilling to do this (Wilshusen 2011).

to address short-run liquidity constraints and partially write-down interest payments and late fees to address longer-run debt overhang. In most cases, credit card issuers will also agree to stop recording the debt as delinquent on the borrower's credit report. Compared to making only the minimum payment on a credit card, enrolling in a DMP will reduce the average borrower's monthly payments by about 10 to 15 percent and reduce the total cost of repayment by about 20 to 40 percent. Following the negotiations with the credit card issuers, the borrower makes one monthly payment to the credit counseling agency that is disbursed to his or her creditors according to the terms of the restructured agreements. The minimum monthly payment for each credit card account is typically about two to three percent of the original balance, although borrowers can make additional payments to reduce the length of the repayment program. In our sample, the average minimum monthly payment for the control group is 2.38 percent of the original balance, or about $440, and the average length of the repayment program is 52 months.

Creditors will usually allow borrowers to resume the repayment program if they miss just one or two payments. However, if a borrower misses multiple payments or withdraws from the program, the remaining credit card debt is usually sent to collections and borrowers' credit scores are likely to be adversely affected, at least in the short run.

To help ensure that creditors benefit from their participation in the repayment program, the counseling agency screens potential clients to assess whether the borrower has a sufficient cash flow to repay his or her debts over the three-to-five-year period of the repayment program, but not enough to reasonably repay his or her debts without the repayment program. In our data, nearly all potential clients are approved during the screening process, with approved clients having only slightly better labor market and financial outcomes than rejected clients (see online Appendix Table A1). Our sample of potential clients is broadly similar to other financially distressed populations in the United States, suggesting that the screening process does not substantially impact our experimental population.

Credit card issuers' participation in a DMP is voluntary, and card issuers may choose to participate in only a subset of the DMPs proposed by the credit counseling agencies. In principle, a credit card issuer will only participate in a repayment program if doing so increases the expected repayment rate, presumably because the borrower is less likely to default or file for bankruptcy (Wilshusen 2011). Consistent with this view, individuals enrolled in a DMP are less likely to file for bankruptcy (Staten and Barron 2006) and less likely to report financial distress (O'Neill et al. 2006) than observably similar individuals who are not enrolled in a DMP. Credit card issuers can also directly refer borrowers to a credit counseling agency if the risk of default or bankruptcy is particularly high. In our sample, approximately 33.7 percent of individuals report that they learned about MMI from an internet search, 19.8 percent from a family member or friend, 20.0 percent from a paid advertisement, and 15.5 percent from a card issuer.

Each year, MMI administers over 75,000 DMPs that repay nearly $600 million in unsecured debt. Nationwide, it is estimated that nonprofit credit counselors administer approximately 600,000 DMPs that repay credit card issuers between $1.5 and $2.5 billion each year (Hunt 2005, Wilshusen 2011).

B. *A Nationwide Randomized Debt Restructuring Experiment*

Overview.—In 2003, MMI and 11 large credit card issuers agreed to offer more generous minimum payment reductions and interest write-downs to a subset of borrowers interested in a structured repayment program. The purpose of the experiment was to evaluate the effect of more generous debt relief on repayment rates, particularly for the most financially distressed borrowers.

The resulting nationwide experiment was conducted between January 2005 and August 2006. The experimental population consisted of the near universe of prospective clients that contacted MMI during this time period. There were two main restrictions to the experimental sample. First, the experiment was restricted to individuals contacting MMI for the first time during this time period; individuals who had already enrolled in a DMP before January 2005 were excluded from the randomized trial. Second, the experiment was restricted to individuals assigned to counselors with more than six months of experience. We also drop the small number of individuals who were rejected during the screening process described above. In total, the estimation sample includes 78,438 borrowers assigned to 1,099 different counselors. The sample includes individuals from all 50 states and the District of Columbia, with a similar proportion of individuals across most states (see online Appendix Figure A1).

MMI worked with the participating credit card issuers to design the experiment, but none of the card issuers were directly involved with the day-to-day management of the experiment. The participating credit card issuers were, however, given regular updates during the experiment, including the number of individuals enrolling in and completing a repayment program in both the treatment and control groups. The participating credit card issuers were also given a final report following the completion of the experiment, although those results were not shared publicly until the publication of this paper. Following the experiment, many of the participating credit card issuers began offering a combination of more generous minimum payment reductions and interest write-downs to the highest-debt borrowers interested in a DMP. However, it is not clear if this new policy was due to the experiment itself or the onset of the financial crisis.

The research team was not involved in the design, implementation, or original analysis of the experiment. Information on the experimental design comes from MMI records documenting the experimental procedures and results and conversations with both the card issuers and the MMI administrator tasked with implementing the experiment. We were also provided with administrative data from MMI (described in detail below) that allowed us to verify the most important details of the experimental design, including the random assignment of clients to credit counselors, the rotation of credit counselors between the treatment and control groups at two-week intervals, the accuracy of the algorithm used to calculate the individual-specific terms of the repayment program, and the effective random assignment of clients to the treatment and control groups during the experiment.

Sequence of the Experiment.—First, each prospective client was randomly assigned to a credit counselor conditional on the contact date, state, and reference channel (i.e., web versus phone). For each counselor, the MMI computer system would automatically switch the counselor from the control group repayment program to the treatment group repayment program every two weeks. This automated rotation procedure was meant to ensure that experimental protocols were followed by the counselors and that any counselor-specific effects would not bias the experiment. The rotation procedure was staggered across counselors so that, on any given day, approximately 50 percent of prospective clients were assigned to the treatment group with the remaining 50 percent assigned to the control group. Counselors were strictly instructed not to inform prospective clients of the experiment, and a senior credit counselor conducted frequent audits of the counselors to ensure that the experimental protocols were followed. In practice, the automated rotation procedure also meant that the counselors were largely unaware of when they were in the control group and when they were in the treatment group. There was also no publicly available information available regarding the experiment or the credit card issuers' concessions during or immediately after the experiment. These institutional features mean that the experiment was effectively double-blind, making it unlikely that our results are driven by Hawthorne or "observer" effects, where treated individuals change their behavior only because they are aware of being studied. We test for Hawthorne effects directly by estimating the effect of treatment eligibility on borrowers who were offered the "control" repayment program even when they were assigned to the treatment group because they did not have credit card debt with an issuer participating in the experiment. We find statistically and economically insignificant effects of treatment eligibility on these borrowers (see panel B of online Appendix Table A2).

Following the assignment of an individual to a credit counselor, the assigned counselor collected information on the prospective client's unsecured debts, assets, liabilities, monthly income, monthly expenses, homeownership status, number of dependents, and so on. Identical information was collected from both the treatment and control groups, and there was no indication of treatment status communicated to individuals. Using the information collected by the counselor, the MMI computer system would then calculate the individual-specific terms of the repayment program, including the minimum payment amount, the length of the program, and the total financing fees. These terms depended on the amount of debt with each credit card issuer and whether the individual was assigned to the treatment or control group.

Next, the credit counselor would explain the individual's options for repaying his or her debts. The details of this process closely followed MMI's usual procedures and were identical for the treatment and control groups. In most cases, the repayment options were explained in the following way. First, individuals were told that they could liquidate their assets and repay their debts immediately, although relatively few individuals in our sample had enough assets to make this a viable option. Second, individuals were told that they could file for bankruptcy protection, which would allow them to discharge their unsecured debts and avoid debt collection in exchange for any non-exempt assets and the required court fees. Third, individuals were told what would happen if they continued paying only the minimum payment on their credit cards. In a representative call provided to the research team, the MMI counselor explained that "if you continue making the minimum payment of $350, it will take you 348 months to repay your credit cards and you will have to spend about $21,300 in financing charges." Finally, individuals were told about the benefits of enrolling in a structured repayment program. In the same representative call, the MMI counselor explained that "if you enroll in a debt management plan, your payments will drop to $301, you will repay all of your credit cards in 56 months, and you will only have $3,800 in financing charges. That is a savings of about $17,500."

Finally, the individual would indicate whether he or she wished to enroll in the offered repayment program following the counselor's explanation of the repayment options. Individuals could also call back at a later date to enroll in the repayment program under the same terms.

Treatment Intensity.—Table 1 illustrates how the experiment impacted the typical borrower's repayment program. Each row presents DMP terms for a hypothetical borrower as if he or she was assigned to the control group, as if he or she was assigned to the treatment group and only eligible for an interest write-down, and as if he or she was assigned to the treatment group and only eligible for a minimum payment reduction. We use the control mean for credit card debt acquired before the experiment ($18,470) throughout. We first calculate the DMP terms for this hypothetical borrower as if he or she had been assigned to the control group using the control mean for the minimum payment requirement (2.38 percent of initial debt) and an interest rate of 9.9 percent. For this hypothetical borrower, the control repayment program requires making minimum payments of $439.59 for 51.80 months, with $4,302 in financing fees.

In the second row of Table 1, we recalculate the DMP terms for this hypothetical borrower using the maximum possible interest write-down in the treatment group (9.9 percentage points), holding the minimum payment constant. The maximum interest write-down would decrease the financing fees for this hypothetical borrower by $4,302, or 100.0 percent, by dropping the last nine to ten payments of the repayment program. However, the interest write-down does not affect the borrower's minimum payment amount. Thus, the interest write-down will only increase enrollment in the repayment program if borrowers value debt forgiveness at the end of the repayment program, about three to five years in the future.

In the last row of Table 1, we recalculate the DMP terms using the maximum payment reduction in the treatment group (0.5 percentage points), holding the interest write-down constant. The maximum minimum payment reduction in the treatment group would decrease the hypothetical borrower's minimum payment by $92.35, or 21.0 percent, by adding an additional 18 months to the repayment program. The longer repayment period would also increase the financing fees for this borrower by $1,645, or 38.2 percent. Thus, the payment reductions may decrease liquidity-based defaults at the beginning of the repayment program by lowering the minimum payment amount, but may increase defaults at the end of the repayment program by increasing the exposure to default risk.

Treatment Costs.—To understand the relative magnitudes of these interest write-downs and minimum payment reductions in the treatment group, columns 6–8 of Table 1 and online Appendix Figure A2 present estimates of the net present costs of providing each treatment for a range of discount rates. For each discount rate, we calculate the net present difference between the total payments made by our hypothetical borrower if assigned to the control group and if assigned to the treatment

Table 1—Examples of the Randomized Treatments

Treatments		Program characteristics			Discounted cost to lender		
Interest write-down (1)	Payment reduction (2)	Minimum payment (3)	Financing fees (4)	Total months (5)	0% discount rate (6)	9.9% discount rate (7)	20% discount rate (8)
—	—	$439.59	$4,302	51.80	—	—	—
9.90%	—	$439.59	$0	42.00	$3,621	$2,175	$1,255
—	0.50%	$347.24	$5,947	70.32	$461	$1,257	$1,636

Notes: This table illustrates how treatment eligibility impacts repayment program characteristics and lender costs. The first row reports program characteristics for the baseline case in the control group. The second row reports program characteristics applying the maximum interest rate write-down in the treatment group. The third row reports program characteristics applying the maximum minimum payment reduction in the treatment group. Columns 1 and 2 describe the interest write-down and payment reduction received if in the treatment group. Column 3 presents the minimum required payment in dollars. Column 4 presents the total cost of all interest rate payments and late fees assuming no early payments. Column 5 presents the total number of months before the program is complete assuming no early payments. Columns 6–8 present the net present cost of providing each treatment relative to the baseline case, using the control mean for the monthly default rate during the repayment program (1.12 percent). All program characteristics and lender costs are calculated using the control means for debt ($18,470) and minimum payment (2.38 percent of debt), and a baseline interest rate of 9.90 percent.

group. For the control group calculation, we assume an interest rate of 9.9 percent and the control means for debt ($18,470), minimum payment (2.38 percent of debt), and monthly default rate (1.12 percent), which we include to account for the mechanical change in default associated with a shorter or longer repayment program.

The net present cost of providing the interest write-downs in the treatment group largely comes from the discounted value of the financing fees that are written down at the end of the repayment program, a clear economic cost that is decreasing in the discount rate. In contrast, the net present cost of providing the payment reductions in the treatment group largely comes from shifting debt payments from the present to the future at the prevailing interest rate, an unclear economic cost that is increasing in the discount rate. In practice, we find that the net present costs of providing the interest write-downs and payment reductions are identical when the discount rate is equal to 16.4 percent (see online Appendix Figure A2). The costs of the interest write-downs in the treatment group are $918 higher than the costs of the payment reductions in the treatment group when the discount rate is around 10 percent ($2,175 versus $1,257), but $381 lower when the discount rate is 20 percent ($1,255 versus $1,636). We interpret these calculations as suggesting that the experiment provides a reasonably similar comparison of these two different types of debt relief.

C. *Data and Experiment Validity*

To estimate the impact of the randomized experiment, we match counseling data from MMI to administrative bankruptcy, credit, and tax records. This section describes the construction and matching of each dataset and provides tests of experimental validity.

Data Sources and Sample Construction.—The counseling data provided by MMI include information on all prospective clients eligible for the randomized trial

(Money Management International 2014). The data include information on each individual's total unsecured debts, assets, liabilities, monthly income, monthly expenses, homeownership status, number of dependents, treatment status, enrollment in a repayment program, the repayment program characteristics, and the amount of debt repaid through the repayment program. The data also include information on the date of first contact, state of residence, who referred the individual to MMI, and the assigned counselor. Finally, the data include detailed information on the amount of unsecured debt with the 19 largest creditors in the sample, including all 11 of the credit card issuers participating in the experiment. We top-code all continuous variables at the ninety-ninth percentile of our estimation sample.

Information on bankruptcy filings comes from individual-level PACER bankruptcy records. The bankruptcy records are available from 2000 to 2011 for the 81 (out of 94) federal bankruptcy courts that allow full electronic access to their dockets. These data represent approximately 87 percent of all bankruptcy filings during our sample period.[3] We match the credit counseling data to Chapter 7 and Chapter 13 filings in the PACER data using name and the last four digits of the social security number. We assume that unmatched individuals did not file for bankruptcy protection during the sample period, and control for state fixed effects in all specifications to account for the fact that we do not observe filings in all states. We therefore explicitly control for any potential selection bias due to the incomplete nature of the bankruptcy data. In addition, we allow individuals to be matched to multiple bankruptcy filings to account for the fact that many individuals file multiple times during our sample period. We find nearly identical results if we limit the sample to borrowers living in states with PACER data coverage or if we only match to the first bankruptcy filing observed for each individual in the PACER data.

Information on collections debt and credit scores come from individual-level credit reports from TransUnion (TU). The TU data are derived from public records, collections agencies, and trade lines data from lending institutions (TransUnion 2015). The collections data contain information on any unpaid bills that have been sent to collections agencies, including the date of collections and the current amount owed. The credit score we use is calculated by TU to predict the probability that a consumer will become delinquent on a new loan within the next 24 months. Since credit scores are used in the vast majority of lending decisions, improvements in credit scores should directly translate into increased credit availability, lower interest rates, or both (e.g., Dobbie et al. 2016). TransUnion was able to successfully match 86.7 percent of the credit counseling data to the credit bureau data, with a small number of observations matched but without credit scores. The probability of being matched to the credit report data is not significantly related to treatment status (see Table 2). No personally identifiable information (PII) was provided to us by TransUnion.

Information on labor market outcomes and 401(k) contributions comes from administrative tax records from the SSA (Social Security Administration 2014). The SSA data are available from 1978 to 2013 for every individual who has ever acquired an SSN, including those who are institutionalized. The SSA data include

[3] See Gross, Notowidigdo, and Wang (2014) for additional details on the bankruptcy data used in our analysis. We thank these authors for providing the bankruptcy data used in our analysis.

Table 2—Descriptive Statistics and Balance Tests

	Control mean (1)	Treatment versus control (2)
Baseline characteristics		
Age	40.827	−0.0004
		(0.0002)
Male	0.365	−0.0026
		(0.0045)
White	0.640	−0.0066
		(0.0076)
Black	0.168	−0.0089
		(0.0089)
Hispanic	0.090	−0.0203
		(0.0105)
Number of dependents	2.179	−0.0016
		(0.0017)
Homeowner	0.419	0.0056
		(0.0068)
Renter	0.435	0.0017
		(0.0062)
Monthly income (1,000s)	2.498	−0.0001
		(0.0018)
Debt in repayment (1,000s)	18.470	0.0003
		(0.0001)
Percent with exp. creditors	0.445	0.0010
		(0.0076)
Baseline outcomes		
Bankruptcy	0.003	−0.0045
		(0.0359)
Nonzero collections debt	0.248	0.0007
		(0.0057)
Credit score	586.665	−0.0000
		(0.0000)
Employment	0.848	0.0034
		(0.0074)
Earnings (1,000s)	23.702	−0.0001
		(0.0001)
Data quality		
Matched to SSA data	0.952	−0.0787
		(0.1832)
Matched to TU data	0.869	−0.0052
		(0.0225)
p-value from joint F-test	—	[0.4215]
Observations	39,855	78,438

Notes: This table reports descriptive statistics and balance tests for the estimation sample. Information on age, gender, race, earnings, employment, and 401(k) contributions is only available for individuals matched to the SSA data and information on collections debt and credit score is only available for individuals matched to the TU data. Each baseline outcome is for the year before the experiment. Column 1 reports the mean for the control group. Column 2 reports the difference between the treatment and control groups controlling for randomization strata fixed effects and clustering standard errors at the counselor level. The p-value is from an F-test of the joint significance of the variables listed.

information on all formal sector earnings and 401(k) contributions from annual W-2s and self-employment earnings from annual 1040s at the IRS. Individuals with no W-2 or self-employment earnings in any particular year are assumed to have had no formal sector earnings in that year. Individuals with no W-2 are also assumed to have had no 401(k) contributions in that year. The 401(k) variable includes all conventional, pre-tax contributions, but does not include contributions to Roth accounts. Individuals with zero earnings or zero 401(k) contributions are included in all regressions throughout the paper.[4] We match the credit counseling data to the tax data using the full social security number. We were able to successfully match 95.3 percent of the counseling data to the SSA data. The probability of being matched to the SSA data is also not significantly related to treatment status (see Table 2).

We make three main sample restrictions to the estimation sample. First, we drop individuals who are not randomly assigned to counselors because they need specialized services such as bankruptcy counseling or housing assistance. Second, we drop individuals with less than \$1,000 in unsecured debt or more than \$100,000 in unsecured debt to minimize the influence of outliers. These cutoffs correspond to the first and ninety-ninth percentiles of the control group, respectively. Third, we drop the small number of individuals who were rejected during the MMI screening process, as well as a small number of observations where the counseling data provided by MMI were corrupted. The resulting estimation sample consists of 39,855 individuals in the control group and 38,583 individuals in the treatment group. Our sample for the labor market and 401(k) outcomes is further restricted to 74,738 individuals matched to the SSA data, while our sample for the collections debt and credit score outcomes is restricted to the 68,000 individuals matched to the TU data.

Descriptive Statistics.—Table 2 presents descriptive statistics. Column 1 presents means for the control group. The average individual in the control group is just over 40 years old with 2.18 dependents. Thirty-seven percent of individuals in the control group are men, 64.0 percent are white, 16.8 percent are black, and 9.0 percent are Hispanic. Forty-two percent are homeowners, 43.5 percent are renters, and the remainder live with either a family member or friend. Individuals in our sample are highly indebted before contacting MMI, with the typical individual in the control group holding \$18,470 in unsecured debt, with 44.5 percent of that debt being held by a credit card issuer participating in the randomized experiment. Not surprisingly, individuals in our sample are also severely financially distressed before contacting MMI. Baseline credit scores in the control group are about 587 points, with 24.8 percent of individuals in the control group having nonzero collections debt. Eighty-five percent of individuals in the control group have nonzero earnings in the SSA data, with average annual earnings of approximately \$23,700 (including 0s). Baseline bankruptcy rates are very low in the control group, however, at 0.3 percent, likely because individuals are unlikely to contact a credit counselor if they have already filed for bankruptcy protection.

[4] The SSA data also include information on mortality and Disability Insurance receipt. Very few individuals in our data die or receive Disability Insurance during our sample period and estimates for these outcomes are small and not statistically different from zero.

II. Empirical Strategy

Overview.—Consider a model that relates outcomes such as debt repayment to interest write-downs ($WriteDown_i$) and minimum payment reductions ($Payment_i$):

$$y_{it} \;=\; \beta_0 + \beta_1\, WriteDown_i + \beta_2 Payment_i + \beta_3 \mathbf{X}_i + \varepsilon_{it}, \tag{1}$$

where y_{it} is the outcome of interest for individual i in year t, $\mathbf{X}_i$ is a vector of individual-level controls, and ε_{it} is an error term. The key problem for inference is that ordinary least squares (OLS) estimates of equation (1) in non-experimental populations may be biased if the interest write-downs and minimum payment reductions are correlated with the unobservable determinants of later outcomes. For example, individuals borrowing from the credit card issuers offering more generous interest write-downs and payment reductions may be unobservably different than individuals borrowing from the credit card issuers offering less generous repayment terms.

We are also unable to separately identify the causal effects of the interest write-downs and minimum payment reductions using purely experimental variation (comparing the treatment and control groups), as treated borrowers were offered a repayment program that included a combination of both the more generous interest write-downs and the more generous minimum payment reductions. As a result, intent-to-treat estimates will measure the combined effect of both forms of debt relief, not the separate impact of the interest write-downs and minimum payment reductions. A second complicating factor is that over 25 percent of borrowers in our sample also had no credit card debt with the 11 credit card issuers participating in the experiment and, as a result, were offered the "control" repayment program even when they were assigned to the treatment group. In total, nearly 90 percent of borrowers received a less intensive treatment than originally intended because they had at least some credit card debt with a nonparticipating issuer. The fact that these borrowers were only partially treated by the experiment means that intent-to-treat estimates will understate the true impact of targeted debt relief.

To address these issues, we estimate the separate impact of the interest write-downs and minimum payment reductions combining variation from the randomized experiment and cross-sectional differences in the composition of credit card debt (how much each participant borrowed from each creditor), which determined the potential package that each participant was eligible for (and received, if treated). The critical identification assumption underlying our approach is that the baseline debt composition (and hence the specific treatment intensity) is uncorrelated with the effect of receiving an offer to participate in the restructuring program. This is a strong assumption that could plausibly be violated for several reasons. For example, individuals who borrow from credit card issuers offering more generous interest write-downs in the treatment group may be more responsive to debt relief than individuals who borrow from issuers offering less generous interest write-downs in the treatment group. To account for such selection concerns, our preferred specifications compare individuals with the same set of credit cards, but different proportions of debt on each credit card. These specifications weaken the identifying assumption by only requiring that the proportion of debt with each card issuer be as-good-as-randomly assigned with respect to treatment effect heterogeneity, not the initial choice of which credit cards to hold.

In the remainder of this section, we discuss our treatment intensity measures, our identifying assumptions, and our regression and matching estimators that compare individuals with the same set of credit cards but different proportions of debt on each card. Additional details and econometric proofs are in the online Appendix.

Treatment Intensity Calculation.—We construct our potential interest write-down and payment reduction measures using the difference between hypothetical treatment and hypothetical control repayment program offers for each individual in our sample. That is, we first calculate the hypothetical interest write-downs and minimum payments for all individuals in our sample as if they had been assigned to the control group, using exactly the same calculation and information that MMI uses to calculate the terms of the structured repayment program in the control group. We then calculate the hypothetical write-downs and minimum payments for those same individuals as if they had been assigned to the treatment group, now using the exact same calculation and information that MMI uses in the treatment group. Following MMI's internal guidelines, we assume an interest rate of 6.7 percent and a minimum payment of 2.25 percent for the 16.7 percent of the debt in our sample that is held by small creditors for which we do not have information on interest rates or minimum payments. Finally, we calculate the difference between these hypothetical control and treatment write-downs and hypothetical control and treatment payment reductions for each individual in our sample, dividing each measure by the maximum possible change in the treatment group (9.9 percentage points for the interest write-downs and 0.5 percentage points for the minimum payment reductions). Following this renormalization, we can interpret our estimates as the effect of the maximum possible interest write-down in the treatment group and the maximum possible payment reduction in the treatment group.[5]

Treatment Intensity Variation.—Online Appendix Figure A4 plots the distribution of potential interest write-downs and payment reductions in our estimation sample. As discussed above, the variation in potential treatment intensity comes from the fact that each of the credit card issuers participating in the experiment offered a different combination of interest write-downs and minimum payment reductions to treated borrowers, and that individual borrowers made different decisions about how much to borrow from each of these card issuers before the experiment. These decisions translated into approximately 50,000 different unique combinations of potential interest write-downs and payment reductions in our sample, with considerable support over the entire range of possible treatment intensities. For example, one of the credit card issuers offered the largest interest write-down (9.9 percentage points) and no minimum payment reduction to treated borrowers, while another offered the largest minimum payment reduction (0.5 percentage points) and the smallest

[5] To confirm the accuracy of our calculations, online Appendix Figure A3 plots predicted DMP characteristics against actual DMP characteristics for the control and treatment groups. There is nearly a one-to-one relationship between predicted and actual payments in both groups. There is a similarly tight relationship between predicted and actual plan length for shorter programs, with a weaker relationship for the longer programs where borrowers are more likely to make extra payments.

interest write-down (4.0 percentage points). Online Appendix Table A3 lists the treatment and control group offers for each of the 11 credit card issuers participating in the experiment.

The variation in potential treatment intensity is also large in economic terms. The difference between the twenty-fifth percentile and seventy-fifth percentile interest write-downs within the treatment group, for example, is roughly equivalent to the difference between the median control group write-down and the median treatment group write-down ($1,521 versus $1,712). Similarly, the difference between the twenty-fifth percentile and seventy-fifth percentile minimum payment reductions within the treatment group is slightly larger than the difference between the median control group reduction and the median treatment group reduction ($33 per month versus $26 per month).

As discussed above, the critical identification assumption underlying our research design is that the causal effects of the interest write-downs and payment reductions are uncorrelated with this variation in potential treatment intensity. To better understand what this identifying assumption entails, Table 3 reports results from an OLS regression of potential treatment intensity on all baseline characteristics and outcomes. We begin by controlling only for date-by-state-by-reference group "randomization strata" fixed effects that account for the level at which individuals are randomly assigned to counselors. Borrowers with larger potential interest write-downs have less debt in repayment, have higher baseline earnings, and are somewhat more likely to be matched to the TU credit report data (column 2). Borrowers with larger potential minimum payment reductions are also older and less likely to own a home (column 4). The proportion of debt with credit card issuers participating in the experiment is also (mechanically) correlated with both the potential interest write-downs and potential minimum payment reductions. The p-value from an F-test of the joint significance of all of the variables listed is 0.000 (column 2) and 0.002 (column 4) for the interest write-downs and payment reductions, respectively, even when omitting the proportion of debt with participating card issuers. These results suggest that our identifying assumption is likely to be violated when we only control for the randomization strata fixed effects.[6]

Estimation and Identifying Assumptions.—We therefore develop two complementary estimators that compare individuals with the exact same set of credit cards, but different proportions of debt on each credit card. Both estimators weaken the key identifying assumption by only requiring that the proportion of debt with each card issuer be as-good-as-randomly with respect to treatment effect heterogeneity, not the initial choice of which credit cards to hold. We do this by controlling for the number and identity of each individual's credit card issuers, or what we call the "creditor risk set." For each individual, we define the creditor risk set as the list of all credit cards that an individual holds from participating card issuers, as well as an indicator for holding at least one credit card from a nonparticipating card issuer. With 11 participating card issuers and 1 aggregate nonparticipating card issuer,

[6] Online Appendix Table A4 reports results from OLS regressions of an indicator for having any debt with each card issuer on all baseline characteristics and outcomes, i.e., the extensive margin version of Table 3. The results largely follow those reported in Table 3.

Table 3—Correlates of Potential Treatment Intensity

	Control mean	Max interest write-down × 100		Max payment reduction × 100	
	(1)	(2)	(3)	(4)	(5)
Baseline characteristics					
Age	40.827	−0.0155	−0.0092	0.0441	0.0012
		(0.0177)	(0.0140)	(0.0210)	(0.0150)
Male	0.365	−0.2777	−0.0026	0.6668	−0.0338
		(0.4948)	(0.3832)	(0.4889)	(0.3502)
White	0.640	0.1346	0.1671	0.5885	0.0425
		(0.8282)	(0.6002)	(0.9558)	(0.6777)
Black	0.168	0.1433	−0.2747	0.4299	0.1025
		(0.9857)	(0.7208)	(0.2591)	(0.7089)
Hispanic	0.090	0.1414	0.2718	−0.3346	0.0623
		(1.0137)	(0.7663)	(1.0847)	(0.8168)
Number of dependents	2.179	−0.1782	−0.1062	−0.0452	−0.0981
		(0.1693)	(0.1223)	(0.1830)	(0.1253)
Homeowner	0.419	−0.5708	−0.2717	−1.2060	0.0762
		(0.7538)	(0.5546)	(0.7574)	(0.5527)
Renter	0.435	0.1596	−0.0286	0.7964	0.6429
		(0.6417)	(0.4702)	(0.6885)	(0.4819)
Monthly income (1,000s)	2.498	0.1538	−0.0950	0.0762	0.0669
		(0.1833)	(0.1429)	(0.1956)	(0.1486)
Debt in repayment (1,000s)	18.470	−0.0675	−0.0171	−0.0313	0.0021
		(0.0153)	(0.0124)	(0.0151)	(0.0152)
Percent with exp. creditors	0.445	61.2246	64.1221	41.5915	39.9184
		(0.9304)	(1.0954)	(1.0743)	(1.1760)
Baseline outcomes					
Bankruptcy	0.003	0.8269	2.4481	0.4035	−0.2327
		(2.9804)	(2.2521)	(2.4277)	(2.0392)
Nonzero collections debt	0.248	−0.2509	−0.0113	0.0047	0.5707
		(0.5425)	(0.4062)	(0.5635)	(0.4192)
Credit score	586.665	−0.0118	−0.0007	0.0059	0.0044
		(0.0042)	(0.0029)	(0.0037)	(0.0028)
Employment	0.848	−0.1695	−0.3153	−0.4177	−0.3814
		(0.7514)	(0.5952)	(0.9123)	(0.6257)
Earnings (1,000s)	23.702	0.0110	0.0063	−0.0137	−0.0002
		(0.0124)	(0.0090)	(0.0135)	(0.0097)
Data quality					
Matched to SSA data	0.952	−10.7645	16.3510	−16.3291	−11.7238
		(24.7010)	(19.1893)	(19.7504)	(15.0040)
Matched to TU data	0.869	7.2316	0.6201	−2.3265	−2.2481
		(2.5278)	(1.8077)	(2.2688)	(1.7471)
p-value from joint F-test	–	[0.0000]	[0.8802]	[0.0022]	[0.8947]
Creditor risk set fixed effects	–	No	Yes	No	Yes
Observations	39,855	78,438	78,438	78,438	78,438

Notes: This table describes correlates of potential treatment intensity with and without controls for the creditor risk set. The dependent variable for columns 2 and 3 is the maximum potential change in interest rates × 100. The dependent variable for columns 4 and 5 is the maximum potential change in minimum payments × 100. All regressions control for randomization strata fixed effects and cluster standard errors at the counselor level. Columns 3 and 5 also control for creditor risk set fixed effects. The p-value is from an F-test of the joint significance of all the variables listed except the percent of debt with experimental creditors.

there are $2^{12} = 4{,}096$ possible creditor risk sets, although only 436 creditor risk sets include at least 1 treatment observation and at least 1 control observation in our estimation sample.

Our first estimator uses the standard regression framework to identify the causal effects of the interest write-downs and minimum payment reductions by directly controlling for the creditor risk set. Our regression estimator allows us to estimate a weighted average of the risk set-specific treatment effects by simply adding creditor risk set fixed effects to equation (1), where the weights are proportional to the variation in $WriteDown_i$ and $Payment_i$ in each risk set. The regression estimator is simple to implement, the standard errors can be calculated using conventional statistical packages, and it is straightforward to examine treatment intensities are correlated with baseline covariates and outcomes after conditioning on the creditor risk set fixed effects. But, the weighting scheme underlying the regression estimator may not be economically relevant, complicating the interpretation of these estimates. In addition, the weighting scheme used for the interest write-downs and minimum payment reductions estimates may not be identical, as the relative variation in $WriteDown_i$ and $Payment_i$ may differ across the creditor risk sets.

In contrast, our second estimator builds on the matching framework developed by Angrist (1998) and Abadie and Imbens (2002), among many others, by estimating equation (1) separately within each creditor risk set and then imposing our own weighting scheme. In practice, we use the number of treated borrowers in each risk set as weights, yielding estimates with a clear economic interpretation and identical weights for the interest write-downs and payment reductions estimates. We calculate standard errors using a Bayesian bootstrap procedure that adjusts for first-step error in the estimation of the risk-set-specific estimates (Rubin 1981).[7] While the matching estimator allows us to impose identical and economically relevant weights, it may be infeasible when there are many small creditor risk sets. Testing whether the treatment intensities are correlated with baseline covariates and outcomes within each creditor risk set is particularly challenging in our setting, for example, as these tests require a relatively large number of observations in each risk set. We therefore view the regression and matching estimators as complementary and present estimates from both throughout much of the paper. See the online Appendix for additional details on these estimators and econometric proofs.

Our regression and matching estimators identify the causal effects of interest write-downs and minimum payment reductions if the following conditions hold within the creditor risk sets: (i) treatment eligibility only impacts outcomes through the change in interest write-downs and minimum payment reductions, (ii) the causal effects of the write-downs and payment reductions are linear and additively separable, and (iii) the causal effects of the write-downs and payment reductions are uncorrelated with treatment intensity. Further, the regression estimator relies on the functional form assumption that all baseline controls enter linearly, and both

[7] The Bayesian bootstrap smooths bootstrap samples by reweighting rather than resampling observations, preventing the omission of small randomization strata that would occasionally be dropped in a standard nonparametric bootstrap. The Bayesian bootstrap used here is implemented by drawing vectors of Dirichlet $(1, \ldots, 1)$ weights, reestimating equation (1) for all risk sets using the Dirichlet weights, and then aggregating the risk-set-specific estimates using the number of treated borrowers as weights. We repeat this procedure 500 times and report the standard deviation of the bootstrap estimates.

estimators require that the correlation between potential outcomes and the potential treatment intensity measures is linear and additively separable (see the online Appendix for additional details). We now consider whether each of the three main conditions holds in our data.

Exclusion Restriction.—The first condition needed for our research design is an exclusion restriction that treatment eligibility only impacts outcomes through the change in interest write-downs and minimum payment reductions. Table 2 partially tests this assumption by verifying that treatment eligibility is randomly assigned in the full sample after we condition on date-by-state-by-reference group fixed effects that account for the level at which individuals are randomly assigned to counselors. Column 2 of Table 2 reports results from an OLS regression of treatment eligibility on all baseline characteristics and these randomization strata fixed effects. Standard errors are clustered at the counselor level. The means of all of the baseline variables are similar in the treatment and control groups and the p-value from an F-test of the joint significance of all of the variables listed is 0.422, suggesting that the randomization was successful. We find similar results if we add the creditor risk set fixed effects, while online Appendix Table A5 verifies that the randomization was also successful within narrowly defined treatment intensity bins.

Linear and Additively Separable Treatment Effects.—The second condition needed for our research design is that the causal effects of the write-downs and payment reductions are linear and additively separable. To partially test this assumption, online Appendix Table A6 presents nonparametric estimates of the interest write-downs and minimum payment reductions in our experiment. We estimate these nonparametric treatment effects by grouping our treatment intensity measures into equally sized bins for both the interest write-downs and minimum payment reductions. We report the interaction of treatment eligibility and each treatment intensity bin, controlling for the treatment intensity bins and the randomization strata fixed effects described above. The results are consistent with linear and additively separable treatment effects, although large standard errors mean that we cannot rule out modest nonlinearities or interaction effects. We present additional evidence in support of linear treatment effects in our robustness checks.

Conditional Ignorability.—The final condition needed to interpret our estimates as the causal effects of the interest write-downs and payment reductions is that potential treatment intensity is uncorrelated with the treatment effects within each creditor risk set. A sufficient condition would be that potential treatment intensity is as-good-as-randomly assigned within each creditor risk set. Columns 3 and 5 of Table 3 provide a partial test of this assumption, reporting results from an OLS regression of potential treatment intensity on the baseline characteristics and outcomes from Table 2, the randomization strata fixed effects, and the creditor risk set fixed effects. We only use our regression estimator to implement these baseline tests as using our matching estimator would require a relatively large number of observations in each risk set. After controlling for these creditor risk sets fixed effects, there is no discernible relationship between potential treatment intensity and the baseline controls and the p-value from an F-test of the joint significance of all of the variables listed (omitting the proportion of debt with participating issuers) is 0.880 for the interest write-downs (column 3) and 0.895 for the payment reductions (column 5). These results suggest that our identifying assumption is more likely to hold after we account for the creditor risk sets. In robustness checks, we show that our results are also robust to a set of overidentification tests once we condition on the creditor risk set fixed effects.

III. Results

In this section, we examine the effects of targeted debt relief using the empirical strategy described above. We first analyze the effects of targeted debt relief on debt repayment, before turning to its effects on bankruptcy, financial outcomes, and labor market outcomes.

A. *Debt Repayment*

Table 4 presents estimates of the impact of being offered more generous interest write-downs and minimum payment reductions on starting and completing a structured repayment program over about the next five years. Columns 1 and 5 report intent-to-treat estimates of the impact of treatment eligibility. Columns 2 and 6 report our baseline estimates of treatment eligibility interacted with the potential interest write-down and treatment eligibility interacted with the potential minimum payment reduction. The potential interest write-down and payment reduction variables are scaled such that our estimates can be interpreted as the causal effect of being offered the maximum interest write-down in the treatment group and maximum payment reduction in the treatment group. Columns 3 and 7 report estimates from our preferred regression specification that adds the creditor risk set fixed effects described above. Columns 4 and 8 report estimates from our matching specification that uses the number of treated individuals in each creditor risk set cell as weights. All specifications control for potential treatment intensity, the baseline controls in Table 2, and the date-by-state-by-reference group randomization strata fixed effects that account for the level at which individuals are randomly assigned to counselors. The ITT, baseline, and preferred regression specifications report standard errors clustered at the counselor level, and the matching specifications report standard errors from the Bayesian bootstrap procedure described above.

The intent-to-treat estimates show that there is an economically and statistically significant effect of treatment eligibility on starting and completing the repayment program. Treatment eligibility increased the probability of starting a repayment program by 1.9 percentage points (column 1), a 5.8 percent increase from the control mean of 32.8 percent. The probability of finishing a repayment program also increased by 1.0 percentage points (column 4), a 7.0 percent increase from the control mean of 14.3 percent. These results confirm that more generous debt relief increases debt repayment at the margin but do not distinguish between the effects of the interest write-downs and minimum payment reductions.

The reduced-form estimates in columns 2–4 and 6–8 report our estimates of the separate impact of the interest write-downs and minimum payment reductions. We find that the interest write-downs significantly increased debt repayment in both

Table 4—Targeted Debt Relief and Repayment

	Start repayment				Finish repayment			
	ITT (1)	Baseline (2)	Regression (3)	Matching (4)	ITT (5)	Baseline (6)	Regression (7)	Matching (8)
Treatment eligibility	0.019				0.010			
	(0.004)				(0.003)			
Treatment × max interest write-down		0.044	0.039	0.051		0.029	0.027	0.029
		(0.017)	(0.014)	(0.027)		(0.012)	(0.011)	(0.017)
Treatment × max payment reduction		0.022	0.013	−0.009		0.005	0.003	−0.008
		(0.020)	(0.018)	(0.029)		(0.015)	(0.014)	(0.023)
Control group mean	0.328	0.328	0.328	0.328	0.143	0.143	0.143	0.143
Observations	78,438	78,438	78,438	78,438	78,438	78,438	78,438	78,438

Notes: This table reports reduced-form estimates of the impact of targeted debt relief on repayment. Information on repayment comes from administrative records at the credit counseling organization. Columns 1 and 5 report intent-to-treat estimates. Columns 2 and 6 report our baseline regression estimates. Columns 3 and 7 report regression estimates with creditor risk set fixed effects. Columns 4 and 8 report matching estimates run at the creditor-risk-set-level as described in the text. All specifications control for potential treatment intensity, the baseline controls in Table 2, and randomization strata fixed effects. ITT, baseline, and regression specifications report standard errors clustered at the counselor level and matching specifications report standard errors from the Bayesian bootstrap procedure described in the text.

the short and long run despite not taking effect until three to five years after the experiment. In our preferred regression specification controlling for creditor risk set fixed effects, for example, we find that the maximum interest write-down in the treatment group (9.90 percentage points) increased the probability of starting a structured repayment program by 3.9 percentage points (column 3), an 11.9 percent increase from the control mean. The probability of finishing the program also increased by 2.7 percentage points (column 6), an 18.9 percent increase from the control mean. Our matching specification yields similar results, implying that the maximum interest write-down in the treatment group increased the probability of starting a structured repayment program by 5.1 percentage points (column 4), and of finishing the repayment program by 2.9 percentage points (column 8). Taken together, the results from Table 4 suggest that there may be significant benefits of debt relief targeting longer-run financial constraints such as debt overhang.

In sharp contrast, we find no positive effects of the minimum payment reductions targeting short-run liquidity constraints. In our preferred regression specification with creditor risk set fixed effects, we find that the maximum monthly payment reduction in the treatment group (0.5 percentage points) increased the probability of completing a structured repayment program by only 0.3 percentage points (column 6). The effect on starting a repayment program is slightly larger at 1.3 percentage points (column 3), but still statistically insignificant. While the 95 percent confidence intervals include modest effects (e.g., a 3.1 percent increase in program completion), all of the minimum payment estimates are statistically differentiable from the interest write-down estimates at the 1 percent level. Our matching specification again yields similar results, implying that the maximum payment reduction actually decreased the probability of starting a structured repayment program by a statistically insignificant 0.9 percentage points (column 4), and of finishing the repayment program by a statistically insignificant 0.8 percentage points (column 8). As discussed above, the null effect of the minimum payment reductions is surprising given a large and influential literature documenting liquidity constraints and

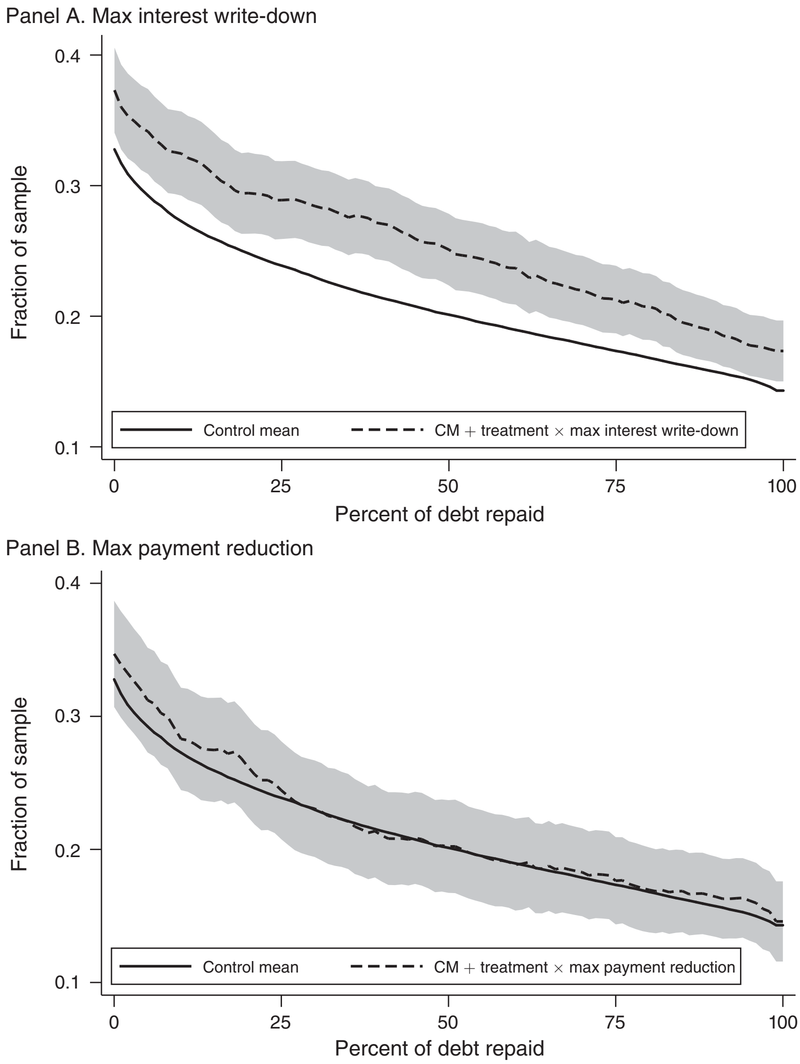

Figure 1. Debt Relief and Repayment Rates

Notes: These panels report control means and the implied treatment group means from our regression estimates. We calculate each treatment group mean using the control mean and the regression estimates described in Table 4. The shaded regions indicate the 95 percent confidence intervals. All specifications control for potential treatment intensity, the baseline controls in Table 2, randomization strata fixed effects, and creditor risk set fixed effects. Standard errors are clustered at the counselor level.

present-biased preferences in a number of otherwise similar settings. Our reduced-form results suggest that either liquidity constraints are not an important driver of borrower behavior in our data, or that a lower minimum payment is an ineffective way to alleviate these issues, at least in our setting.

To better understand these effects, Figure 1 plots the control mean and the treatment group means implied by estimated treatment effects at each percentile of debt repayment. In other words, we estimate treatment effects for each percentile of debt repayment, adding the estimated effect of the interest write-downs

Table 5—Targeted Debt Relief and Bankruptcy Filing

	Bankruptcy filing			
	ITT (1)	Baseline (2)	Regression (3)	Matching (4)
Treatment eligibility	−0.006			
	(0.002)			
Treatment × max interest write-down		−0.031	−0.030	−0.013
		(0.009)	(0.009)	(0.016)
Treatment × max payment reduction		0.023	0.023	0.013
		(0.011)	(0.011)	(0.019)
Control group mean	0.105	0.105	0.105	0.105
Observations	78,438	78,438	78,438	78,438

Notes: This table reports reduced-form estimates of the impact of targeted debt relief on bankruptcy. Information on bankruptcy comes from court records. Column 1 reports intent-to-treat estimates. Column 2 reports our baseline regression estimates. Column 3 reports regression estimates with creditor risk set fixed effects. Column 4 reports matching estimates run at the creditor-risk-set-level as described in the text. All specifications control for potential treatment intensity, the baseline controls in Table 2, and randomization strata fixed effects. ITT, baseline, and regression specifications report standard errors clustered at the counselor level and matching specifications report standard errors from the Bayesian bootstrap procedure described in the text.

and minimum payment reductions to the control mean at each percentile. We focus on our preferred regression specification with creditor risk set fixed effects. Consistent with the results from Table 4, the effect of the interest write-downs are economically and statistically significant at each point in the distribution. The immediate increase in debt repayment estimated in Table 4 remains roughly constant throughout the repayment program, with only a modest fade out of the effects in the last 25 percent of the distribution. In contrast, the small (but not statistically significant) effect of the maximum payment reduction in the short run fades out relatively quickly, with no discernible difference between the treatment and control groups after about the twenty-fifth percentile. It is also worth noting that both treatment and control borrowers exit the repayment program at high rates, with only 14.3 percent of the control group completing the repayment program. In Section IV, we will discuss what mechanisms are most consistent with these patterns.

B. *Bankruptcy*

Table 5 presents results for bankruptcy filing in the first five years following the experiment, an important outside option for borrowers in our sample. MMI discusses both the costs and benefits of bankruptcy with prospective clients, and 10.5 percent of the control group files for bankruptcy in the first five years following the experiment. In our setting, we interpret bankruptcy as an alternative and potentially more costly form of debt forgiveness and debt restructuring.

The intent-to-treat estimates in Table 5 show that treatment eligibility decreased the probability of filing for bankruptcy protection by 0.60 percentage points over the first five years following the experiment in the pooled sample (column 1), a 5.7 percent decrease from the control mean of 10.5 percent. The effects of treatment

TABLE 6—TARGETED DEBT RELIEF AND FINANCIAL OUTCOMES

	Nonzero collections debt				Credit score			
	ITT (1)	Baseline (2)	Regression (3)	Matching (4)	ITT (5)	Baseline (6)	Regression (7)	Matching (8)
Treatment eligibility	0.000				−0.324			
	(0.003)				(0.530)			
Treatment		−0.006	−0.006	−0.027		2.339	1.898	3.550
× max interest write-down		(0.012)	(0.012)	(0.020)		(2.274)	(2.272)	(3.341)
Treatment		0.002	0.002	0.019		−2.092	−1.913	−4.124
× max payment reduction		(0.016)	(0.016)	(0.025)		(2.716)	(2.766)	(4.261)
Control group mean	0.389	0.389	0.389	0.389	604.099	604.099	604.099	604.099
Observations	68,000	68,000	68,000	68,000	67,705	67,705	67,705	67,705

Notes: This table reports reduced-form estimates of the impact of targeted debt relief on collections debt and credit score. Information on outcomes comes from credit records at TransUnion. Columns 1 and 5 report intent-to-treat estimates. Columns 2 and 6 report our baseline regression estimates. Columns 3 and 7 report regression estimates with creditor risk set fixed effects. Columns 4 and 8 report matching estimates run at the creditor-risk-set-level as described in the text. All specifications control for potential treatment intensity, the baseline controls in Table 2, and randomization strata fixed effects. ITT, baseline, and regression specifications report standard errors clustered at the counselor level and matching specifications report standard errors from the Bayesian bootstrap procedure described in the text.

eligibility on bankruptcy filing are again driven by the interest write-downs. Over the first five years following the experiment, we find that the maximum interest write-down in the treatment group decreased the probability of filing for bankruptcy by 1.3 to 3.0 percentage points in the pooled sample (columns 3 and 4), a 12.4 to 28.6 percent decrease from the control mean. Conversely, the maximum payment reduction in the treatment group actually *increased* the probability of filing for bankruptcy over the first five years following the experiment by 1.3 to 2.3 percentage points (columns 3 and 4), a 12.4 to 21.9 percent increase.

In online Appendix Table A7, we find that the effects of the interest write-downs on bankruptcy filing are largest in the first two to three years following the experiment, likely because this is the time period when the control group is most likely to file for bankruptcy. Online Appendix Table A8 also reveals that there are modestly larger effects of the interest write-downs for individuals contacting MMI prior to the 2005 Bankruptcy Reform that increased the financial and administrative costs of filing for bankruptcy protection (column 3), likely because it became more difficult for the control group to file for bankruptcy following the reform. In contrast, the estimated effects of the minimum payment reductions are statistically similar both across years and before and after the 2005 Bankruptcy Reform, in part because we lack the statistical power to detect modest differences in the estimated effects.

C. *Collections Debt and Credit Score*

Table 6 presents results for average collections debt and credit scores over the first five years following the experiment, both important proxies for financial distress and access to credit. In theory, the experiment could either improve borrowers' financial health by increasing debt repayment and decreasing collections activity, or have no impact if the experiment crowds out other debt payments.

Table 7—Targeted Debt Relief and Labor Market Outcomes

	Employment				Earnings (1,000s)			
	ITT	Baseline	Regression	Matching	ITT	Baseline	Regression	Matching
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment eligibility	−0.002				−0.074			
	(0.002)				(0.112)			
Treatment × max interest write-down		0.005	0.006	0.016		−0.668	−0.749	−1.010
		(0.007)	(0.007)	(0.012)		(0.464)	(0.468)	(0.717)
Treatment × max payment reduction		−0.008	−0.007	−0.003		0.413	0.526	0.605
		(0.009)	(0.009)	(0.015)		(0.620)	(0.629)	(0.919)
Control group mean	0.821	0.821	0.821	0.821	27.148	27.148	27.148	27.148
Observations	74,738	74,738	74,738	74,738	74,738	74,738	74,738	74,738

Notes: This table reports reduced-form estimates of the impact of targeted debt relief on employment and earnings. Information on outcomes comes from records at the Social Security Administration. Columns 1 and 5 report intent-to-treat estimates. Columns 2 and 6 report our baseline regression estimates. Columns 3 and 7 report regression estimates with creditor risk set fixed effects. Columns 4 and 8 report matching estimates run at the creditor-risk-set-level as described in the text. All specifications control for potential treatment intensity, the baseline controls in Table 2, and randomization strata fixed effects. ITT, baseline, and regression specifications report standard errors clustered at the counselor level and matching specifications report standard errors from the Bayesian bootstrap procedure described in the text.

There are no statistically or economically significant effects of treatment eligibility on collections debt or credit scores over the first five years following the experiment (columns 1 and 5). The maximum interest write-down in the treatment group decreased the probability of having nonzero collections debt by a statistically insignificant 0.6 to 2.7 percentage points (columns 3 and 4), a 1.5 to 6.9 percent change, and increased average credit scores by an insignificant 1.9 to 3.5 points (columns 7 and 8). The maximum payment reduction had little to no impact on collections debt (columns 3 and 4) and a statistically insignificant negative effect on credit scores (columns 7 and 8). Thus, while underpowered, these results are broadly consistent with our earlier results showing that the interest write-downs modestly improved borrowers' financial health while the lower minimum payments had little positive effect on borrowers.

D. *Labor Market Outcomes*

Table 7 presents results for average employment and earnings over the first five years following the experiment. The experiment could affect labor market outcomes through a number of different channels. For example, enrollment in the repayment program could increase labor supply by decreasing the frequency of wage garnishment orders that occur when an employer is compelled by a court order to withhold a portion of an employee's earnings to repay delinquent debt. The experiment could also impact labor market outcomes through its effects on credit scores (e.g., Herkenhoff 2019; Bos, Breza, and Liberman 2018; Herkenhoff, Phillips, and Cohen-Cole 2016; Dobbie et al. 2016) or productivity (e.g., Mullainathan and Shafir 2013).

There are no statistically or economically significant effects of treatment eligibility on either employment or earnings over the first five years following the experiment (columns 1 and 5). The estimated effects of both the interest write-downs and minimum payment reductions are also small and imprecisely estimated, with opposite signs on the employment and earnings estimates in all specifications. For our preferred regression specification with creditor risk set fixed effects, the 95 percent confidence interval for the employment estimate ranges from −0.8 to 2.0 percentage points (column 3), while the 95 percent confidence interval for the earnings estimate ranges from −\$1,666 to \$168 (column 6). For the maximum payment reduction, the 95 percent confidence interval ranges from −2.5 to 1.1 percentage points for the employment estimate (column 3) and −\$707 to \$1,759 for the earnings estimate (column 6). None of the estimates suggest economically meaningful effects on labor market outcomes in the pooled sample.

To better understand these results, panel A of online Appendix Table A9 presents results separately for borrowers who were and were not employed in the year prior to the experiment. We again focus on our preferred regression specification with creditor risk set fixed effects. We find that the maximum interest write-down decreased annual earnings by \$2,250 for borrowers who were not employed in the year prior to the experiment, while having essentially no effect on borrowers employed at baseline. The employment effects are also negative for nonemployed borrowers, but the point estimate is not statistically significant. These subsample results suggest that the kind of debt forgiveness provided by the interest write-downs may decrease labor supply for borrowers most on the margin of any work.[8]

E. *Subsample Results*

Table 8 presents subsample estimates for our preferred regression specification by baseline debt-to-income and baseline credit scores, which we interpret as proxies for baseline financial distress and baseline liquidity, respectively. The split by baseline debt-to-income is motivated by the fact that only the most financially distressed borrowers were meant to be eligible for treatment in the original experimental design and, following the experiment, many credit card issuers began offering more borrower-friendly terms to financially distressed borrowers. The split by baseline credit scores is motivated by the idea that individuals who are not liquidity constrained are unlikely to benefit from the minimum payment reductions.

The effects of the interest write-downs are consistently larger for individuals with above-median debt-to-income, although the differences are not statistically significant for all outcomes. For example, the maximum write-down increases the probability of starting and completing a repayment program by 6.2 percentage points (18.9 percent) and 4.4 percentage points (30.8 percent), respectively, for individuals with above-median debt-to-income. In comparison, the maximum write-down only increases the probability of starting and completing a repayment program by 1.2 percentage points (3.7 percent) and 0.7 percentage points (4.9 percent) for those with below-median debt-to-income. For those with above-median

[8] In contrast to the relatively modest labor market effects documented here, Dobbie and Song (2015) finds that Chapter 13 bankruptcy protection increases annual earnings by \$5,562 and annual employment by 6.8 percentage points. These contrasting results are most likely due to differences in the intensity of the debt relief provided by consumer bankruptcy and our experiment. Chapter 13 bankruptcy, for example, provides a write-down of approximately 80 to 85 percent of the typical filer's unsecured debt. Conversely, the maximum write-down in the treatment group forgives about 25.84 percent of unsecured debt. In addition, Chapter 13 bankruptcy protects future wages from garnishment, while our experiment did not.

Table 8—Subsample Regression Estimates

	Start payment (1)	Finish payment (2)	Bankrupt (3)	Collections debt (4)	Credit score (5)	Employed (6)	Earnings (7)
Panel A. Baseline debt-to-income							
Treatment × max interest × high DTI	0.062	0.044	−0.035	−0.012	2.374	0.042	−0.816
	(0.017)	(0.014)	(0.012)	(0.014)	(2.740)	(0.019)	(0.521)
Treatment × max interest × low DTI	0.012	0.007	−0.023	0.003	1.321	0.008	−0.671
	(0.019)	(0.016)	(0.011)	(0.017)	(2.860)	(0.010)	(0.674)
p-value on difference	[0.024]	[0.065]	[0.385]	[0.407]	[0.747]	[0.826]	[0.845]
Treatment × max payment × high DTI	0.015	−0.002	0.022	0.014	−4.124	−0.005	0.574
	(0.020)	(0.018)	(0.014)	(0.018)	(3.353)	(0.013)	(0.710)
Treatment × max payment × low DTI	0.007	0.005	0.024	−0.012	1.062	−0.011	0.473
	(0.025)	(0.019)	(0.013)	(0.019)	(3.522)	(0.011)	(0.873)
p-value on difference	[0.778]	[0.770]	[0.902]	[0.161]	[0.205]	[0.703]	[0.915]
Panel B. Baseline credit score							
Treatment × max interest × high score	0.051	0.058	−0.027	−0.016	2.825	0.024	−0.479
	(0.020)	(0.017)	(0.013)	(0.015)	(2.882)	(0.010)	(0.639)
Treatment × max interest × low score	0.034	0.009	−0.032	0.006	0.964	0.010	−0.928
	(0.017)	(0.014)	(0.011)	(0.015)	(2.860)	(0.008)	(0.540)
p-value on difference	[0.445]	[0.018]	[0.746]	[0.189]	[0.604]	[0.531]	[0.527]
Treatment × max payment × high score	0.015	−0.015	0.018	0.031	−4.055	−0.024	0.275
	(0.021)	(0.019)	(0.014)	(0.018)	(3.477)	(0.013)	(0.822)
Treatment × max payment × low score	0.009	0.013	0.027	−0.026	0.273	0.009	0.726
	(0.023)	(0.017)	(0.014)	(0.020)	(3.323)	(0.011)	(0.705)
p-value on difference	[0.806]	[0.204]	[0.594]	[0.008]	[0.280]	[0.024]	[0.600]

Notes: This table reports subsample regression estimates. Panel A reports estimates for individuals with above- and below-median debt-to-income. Panel B reports estimates for individuals with prime and subprime credit scores. All specifications control for potential treatment intensity, the baseline controls in Table 2, randomization strata fixed effects, and creditor risk set fixed effects. Standard errors are clustered at the counselor level.

debt-to-income, the maximum write-down also decreases the probability of having any collections debt by 1.2 percentage points (3.1 percent) and increases the probability of being employed by 4.2 percentage points (5.1 percent), compared to small and statistically insignificant effects for those with below-median debt-to-income. In contrast, the effects of the minimum payment reductions are small and either of the wrong sign or statistically insignificant for both groups.

There is a similar, if less stark, pattern by baseline credit scores. The maximum interest write-down increases the probability of completing a repayment program, for example, by 5.8 percentage points (40.6 percent) for individuals with above-median credit scores, compared to only 0.9 percentage points (6.3 percent) for those with below-median credit scores. The effects are more comparable for starting a repayment plan, however, suggesting that liquidity may be particularly important for successfully making all of the required payments. The effects on other outcomes are also larger for the high credit score group, but none of the differences are statistically significant. We also find similar effects of the payment reductions among high and low credit score borrowers, although the effects are generally more positive for low credit score borrowers, consistent with the idea that the payment reductions are most important for liquidity-constrained borrowers. However, none of the estimates suggest economically significant benefits of the payment reductions for these liquidity-constrained borrowers.

Online Appendix Table A9 presents additional subsample estimates for our preferred regression specification by gender, ethnicity, and baseline homeownership. For each of these subgroups, there are no clear theoretical predictions as to which group will benefit most from the experiment. The interest write-downs have somewhat larger effects for women compared to men, but there are no systematic patterns by either ethnicity or baseline homeownership. Moreover, all of these results should be interpreted with some caution given that we are likely to find a number of statistically significant estimates purely by chance when performing multiple hypothesis tests. We therefore interpret these results as suggesting relatively similar effects of targeted debt relief across these groups.

F. *Robustness and External Validity*

In this section, we discuss robustness checks and how the details of the experimental design may affect the external validity of our results.

Overidentification Tests.—Online Appendix Table A10 presents results from a set of overidentification tests of our main results. Panel A replicates our preferred regression estimates controlling for creditor risk set fixed effects. Panel B adds controls for treatment eligibility interacted with indicator variables for gender, race, baseline homeownership, baseline credit scores, baseline earnings, and baseline debt-to-income. Panel C instead adds controls for treatment eligibility interacted with credit card issuer fixed effects. Panel D adds both the treatment eligibility × baseline demographic variables and treatment eligibility × credit card issuer variables. Consistent with our identifying assumption, our main results are generally robust to the inclusion of treatment eligibility × baseline demographic effects, treatment eligibility × credit card issuer effects, and both treatment eligibility × baseline demographic and treatment eligibility × issuer effects. In a series of F-tests of the joint significance of the treatment eligibility × issuer and treatment eligibility × baseline demographic effects, we also find that these interactions are generally not statistically significant. Taken together, we interpret these results as indicating that our identifying assumption is likely to hold after we account for the creditor risk sets described above.

Permutation Test.—Online Appendix Table A11 presents a second set of robustness checks where the p-values from our preferred regression specification are calculated using a nonparametric permutation test that accounts for the fact that we have run regressions with a number of outcomes and subsamples. That is, we create 1,000 "placebo" samples where we randomly reassign treatment status to individuals within the randomization strata. We then calculate the fraction of treatment effects from these 1,000 placebo samples that are larger (in absolute value) than the treatment effects from the true sample. We find that our main results are robust to this alternative method of calculating p-values. If anything, we obtain smaller p-values from the nonparametric permutation procedure than implied by conventional standard errors.

Framing Effects.—As discussed above, MMI emphasized the monthly payment amount, time to repayment, and financing fees when explaining the repayment program to both the treatment and control groups during the experiment. While the internal validity of the experiment is not affected by these details of the experimental design, it is possible that the effects of the interest write-downs and minimum payment reductions are mediated by these institutional details. For example, it is possible that emphasizing the monthly payment amount increases the perceived value of a minimum payment reduction. It is also possible that emphasizing financing fees, rather than the total amount of debt repaid, either increases or decreases the perceived value of an interest write-down. Importantly, however, these experimental procedures closely followed both MMI's usual procedures and the way in which the write-downs and payment reductions would be implemented at scale through existing credit counseling organizations. Our estimates therefore measure the impact of targeted debt relief in one of the most policy-relevant contexts. Nevertheless, all of our results should be interpreted with these potential framing effects in mind.

Nonlinear Treatment Effects.—Another potential concern is that we estimate the impact of interest write-downs and minimum payment reductions at the margin of an existing debt relief program, making it impossible to estimate the impact of the first dollar of an interest write-down or the first dollar of a payment reduction using our experimental data. We also do not observe the kind of extremely large write-downs or minimum payment reductions needed to estimate, for example, a nearly complete write-down of the original balance. As a result, out-of-sample predictions based on our experimental estimates will be biased if there is a nonlinear effect of either the interest write-downs or the minimum payment reductions. In addition, we assume linear treatment effects when extrapolating the effects of maximum treatment intensities in our main results. Online Appendix Table A6 provides some evidence of linear effects, but these "binned" estimates are too imprecise to conclusively rule out nonlinear treatment effects.

To provide additional evidence on this issue, online Appendix Figure A5 presents nonparametric estimates of the interest write-downs and minimum payment reductions in our experiment. We estimate these nonparametric treatment effects by grouping our treatment intensity measures into equally sized bins for both the interest write-downs and minimum payment reductions, but, unlike online Appendix Table A6, do not allow for interactions between the interest write-downs and minimum payment reductions. We report the interaction of treatment eligibility and each treatment intensity bin, controlling for potential treatment intensity, the randomization strata fixed effects, and the creditor risk set fixed effects. We also plot the OLS best-fit line weighted by the standard error for each point estimate. The results in online Appendix Figure A5 are consistent with linear treatment effects over the range of treatment intensities observed in our data, as well as the results from online Appendix Table A6 discussed earlier. None of our results suggest the kind of nonlinear treatment effects that would bias our estimates or impact out-of-sample predictions based on our experimental estimates. Of course, we cannot test whether there are nonlinear effects for treatment intensities that we do not observe in the data and all of our results should be interpreted with this caveat in mind.

Representativeness of the Sample.—A final concern is that we estimate the impact of targeted debt relief within the sample of individuals who pass the screening procedure described above. Recall that credit counseling agencies screen potential clients to assess whether the individual has a sufficient cash flow to repay his or her debts over the three- to five-year period of the repayment program, but not enough to reasonably repay his or her debts without the repayment program. It is possible that potential clients who pass this screening process are, for example, less liquidity-constrained or more forward-looking than the broader sample of individuals seeking targeted debt relief.

To provide some evidence on the types of individuals entering our experimental sample, online Appendix Table A12 provides descriptive statistics for our experimental sample, a random sample of all credit users, a random sample of credit users with a serious delinquency occurring in the next calendar year, and a sample of credit users with a bankruptcy flag in the next calendar year. Information on all baseline outcomes comes from the TransUnion credit records described above. Online Appendix Table A12 reveals that while our experimental sample is much more financially distressed than the typical credit user, it is broadly similar to other financially distressed populations in the United States, at least on observables. In our experimental sample, for example, the average credit score in the year before contacting MMI is 586.4, compared to a credit score of 572.3 in the delinquency sample and 580.8 in the bankruptcy sample. Credit card balances are somewhat higher in the experimental sample compared to the delinquency and bankruptcy samples, while credit card utilization and delinquencies are both somewhat lower in the experimental sample. For both auto and mortgage loans, the experimental sample falls in between the delinquency and bankruptcy samples. In sum, the experimental sample appears approximately representative of the financially distressed population in the United States.

IV. Mechanisms

In this section, we investigate the potential mechanisms that can explain our interest write-down and minimum payment results.

A. *Overview*

In theory, the interest write-downs can impact debt repayment through two distinct effects. The first is a forward-looking debt overhang effect that decreases the treatment group's incentive to strategically default while both treatment and control groups are enrolled in the repayment program. The second is a mechanical exposure effect that decreases the treatment group's exposure to default risk while the control group is still enrolled in the repayment program and the treatment group is not. We can test the relative importance of these competing channels using treatment effects at the beginning and end of the repayment program. The interest write-downs do not affect the minimum payment requirements early in the repayment program, leaving forward-looking behavior as the only explanation for any interest write-down effects early in the program. We can therefore test for these forward-looking effects using interest write-down treatment effects at the end of the repayment program for the interest write-down group (but not the control group). Then, because the total interest write-down estimate includes the effects of both channels, we can estimate the exposure effect alone using the difference between the total interest write-down estimate and the forward-looking estimate.[9]

The minimum payment reductions can similarly impact debt repayment through two distinct effects. The first is a liquidity effect that, in general, decreases the treatment group's probability of nonstrategic or liquidity-based default while both the treatment and control groups are enrolled in the repayment program. The second is another mechanical exposure effect that increases the treatment group's exposure to default risk while the treatment group is still enrolled in the repayment program and control group is not. Following the same logic as above, we can test for liquidity effects using payment reduction treatment effects at the end of the repayment program for the control group (but not the payment reduction group), as the only difference between the treatment and control groups to this point is the lower minimum payment. Then, because the total payment reduction estimate includes the effects of both channels, we can estimate the exposure effect alone using the difference between the total payment reduction estimate and the liquidity estimate.[10]

B. *Estimation*

We implement these empirical tests using a five-step process. First, we calculate how long the repayment plan would have been had the individual been assigned to the treatment group and how long the repayment plan would have been had the individual been assigned to the control group. The treatment plans are shorter for individuals with relatively larger interest write-downs and longer for individuals with relatively larger minimum payment reductions. For example, individuals with the largest write-downs have treatment plans that are up to 20 percent shorter than their control plans, while individuals with the smallest write-downs and largest minimum payment reductions have treatment plans that are up to 100 percent longer than their control plans. Second, we create an indicator for staying enrolled in the repayment program until the minimum of the treatment plan length and the control plan length. This indicator variable measures payment at the treatment plan length for individuals with the shorter treatment plans (i.e., relatively larger write-downs) and payment at the control program length for individuals with the longer treatment plans (i.e., relatively larger minimum payment reductions). Third, we estimate treatment effects using this new indicator variable as the dependent variable. These reduced-form estimates measure the effect of write-downs at the treatment plan length and the effect of lower minimum payments at the control plan length. Fourth, we take the difference between the reduced-form treatment effects for full

[9] One potential concern is that dynamic selection may lead to a different composition of treated and control borrowers later in the repayment program, thereby biasing our exposure estimates. We investigate this issue in online Appendix Table A13 and find no evidence that the experiment significantly altered the composition of borrowers completing the repayment program.

[10] Our estimates of the forward-looking effect are a lower bound on the true effect because the control group can still make forward-looking default decisions after the interest write-down group finishes their repayment. Similarly, the liquidity effect is an upper bound because the payment reduction group can still make liquidity-based default decisions during their remaining time in the program. For the same reasons, our estimate of the mechanical exposure effect is an upper bound for the interest write-downs and a lower bound for the payment reductions.

Table 9—Forward-Looking, Liquidity, and Exposure Effects

	Finish repayment program			
	Total effect (1)	Forward looking (2)	Liquidity effect (3)	Exposure effect (4)
Treatment × max interest write-down	0.027	0.023		0.004
	(0.011)	(0.012)		(0.004)
Treatment × max payment reduction	0.003		0.009	−0.006
	(0.014)		(0.014)	(0.004)

Notes: This table reports the forward-looking, liquidity, and exposure effects of each treatment on finishing the repayment program. Column 1 reports the total effect of each treatment on finishing the repayment program. Columns 2 and 3 report estimates for being enrolled in the repayment program at the minimum of the treatment program length or the control program length. Column 4 reports the difference between column 1 and columns 2 and 3. All specifications control for potential treatment intensity, the baseline controls in Table 2, randomization strata fixed effects, and the creditor risk sets described in the text. Standard errors for column 4 are calculated using the bootstrap procedure described in the text. See the text for additional details.

repayment estimated in Table 4 and the new reduced-form treatment effects estimated at the shorter of the treatment and control plan lengths. Finally, we calculate the standard error of the difference by bootstrapping the entire procedure described above 500 times. We define the standard error of the treatment effect difference as the standard deviation of the resulting distribution of estimated differences. We control for the baseline controls listed in Table 2, randomization strata fixed effects, and creditor risk set fixed effects throughout.

C. *Results*

Table 9 presents estimates of the forward-looking, liquidity, and exposure effects for both the interest write-downs and minimum payment reductions. Column 1 replicates our regression estimates with creditor fixed effects from column 7 of Table 4, showing the net effect of all channels on completing the repayment program. Columns 2 and 3 report estimates for still being in the repayment program at the minimum of the treatment program length and control program length. Column 4 reports the difference between column 1 and columns 2 and 3.

We find that the positive effects of the interest write-downs can be almost entirely explained by forward-looking decisions made early in the repayment program, not the mechanical reduction in default risk from a shorter repayment program. Our estimates suggest that at least 85.2 percent of the interest write-down effect is due to the decrease in forward-looking defaults at the beginning of the repayment program (column 2). Decreased exposure to risk at the end of repayment can explain a maximum of 14.8 percent of the write-down effect (column 4), with the 95 percent confidence interval including estimates of up to 43.9 percent of the total reduced-form effect.

Figure 1 provides additional evidence in favor of forward-looking effects. First, there is an immediate impact of the interest write-downs on repayment, indicating forward-looking behavior at program sign up. Second, the effects of the interest write-downs, if anything, grow over time relative to the control mean. These results suggest additional forward-looking behavior throughout the repayment program, not just at program sign up. Taken at face value, these two findings rule out many of the most simple "behavioral" explanations for our interest write-down results, such as borrowers being "tricked" into signing up for the repayment program by some feature of the experimental design.

We also find that the null effect of the minimum payment reductions can be explained by the unintended, negative effect of increasing the number of months a borrower remains in the repayment program. Our estimates suggest that debt repayment increases by about 0.09 percentage points due to the liquidity effect (column 3), with the 95 percent confidence interval including effects as large as 3.6 percentage points. However, this positive liquidity effect is nearly exactly offset by the negative exposure effect (column 4). These estimates are also consistent with the patterns observed in Figure 1, where we see a small positive effect of the minimum payment reductions in the short run, and a precise zero effect of the payment reductions in the long run.

The results from Table 9 help to reconcile our findings with the vast literature documenting liquidity constraints in a variety of settings, while indicating that the potential benefits of targeting these liquidity constraints may have been significantly overstated, at least in our setting. Of course, the standard caveat applies that the effects of an increase in liquidity may be nonlinear or context dependent. For example, it is possible the short-run benefits from a very large increase in liquidity may outweigh the long-run costs of a much longer repayment period. It is also possible that liquidity may be more important in the mortgage or student loan markets, where borrowers usually have fewer outside options compared to the credit card borrowers that we study in this paper.

V. Conclusion

This paper uses information from a large-scale randomized experiment to estimate the effects of immediate minimum payment reductions targeting short-run liquidity constraints and delayed interest write-downs targeting longer-run debt overhang. We find that the interest write-downs significantly improved both financial and labor market outcomes, particularly for the highest-debt borrowers, despite not taking effect for three to five years. In contrast, we find no positive effects of the more immediate payment reductions on any outcome. These results stand in stark contrast to the widespread view that short-run liquidity constraints are the most important driver of borrower distress.

Our results are of particular importance in light of the ongoing debate on the relative merits of different types of debt relief. For example, current banking regulations in the United States largely prevent credit card issuers from offering more generous interest write-downs, at least in part due to the perceived unimportance of longer-run constraints such as debt overhang.[11]

[11] US banking regulations prevent credit card issuers from simultaneously reducing the original principal and lengthening the repayment period unless a debt is first classified as impaired. If the original principal is reduced without the debt being classified as impaired, borrowers are required to pay off the remaining debt in just a few months. Government regulators justify these restrictions based on concerns about when delinquent debts would be recognized on the card issuers' balance sheets.

An important open question is whether the increased repayment rates documented in our analysis are, on net, larger than the costs of the interest write-downs. While a comprehensive cost-benefit analysis is beyond the scope of this paper, we consider a partial back-of-the-envelope calculation that takes into account the ex post impact of the interest write-down treatment (2.7 percentage point increase in completed repayment) and the average completion rate in the control group (14.3 percent). To simplify the calculation, we assume that the lender is risk neutral and does not discount future payments. Based on these tentative calculations, we estimate that lenders recoup $3,195 from the typical borrower assigned to the control group, compared to only $3,072 from the typical borrower assigned to the treatment group.

There are three important caveats to our analysis. First, we are not able to estimate the impact of targeted debt relief on ex ante borrower behavior or ex ante borrowing costs or borrowing limits. Our analysis will therefore overstate the benefits of more generous interest write-downs if the ex ante availability of debt relief distorts borrower behavior in such a way that lenders must increase interest rates or decrease credit supply. Second, there may be important ex post impacts of targeted debt relief on outcomes such as post-repayment interest rates that we are unable to measure with our data. Finally, we are unable to test whether the forward-looking decisions documented in this paper are due to rational or nonrational decision making. Given these concerns, we are unable to determine the full welfare consequences of targeted debt relief using our research design.

REFERENCES

Abadie, Alberto, and Guido W. Imbens. 2002. "Simple and Bias-Corrected Matching Estimators for Average Treatment Effects." NBER Working Paper 283.

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru. 2017. "Policy Intervention in Debt Renegotiation: Evidence from the Home Affordable Modification Program." *Journal of Political Economy* 125 (3): 654–712.

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. 2015. "Regulating Consumer Financial Products: Evidence from Credit Cards." *Quarterly Journal of Economics* 130 (1): 111–64.

Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles. 2007. "The Reaction of Consumer Spending and Debt to Tax Rebates: Evidence from Consumer Credit Data." *Journal of Political Economy* 115 (6): 986–1019.

Angrist, Joshua D. 1998. "Estimating the Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants." *Econometrica* 66 (2): 249–88.

Auclert, Adrien, Will S. Dobbie, and Paul Goldsmith-Pinkham. 2019. "Macroeconomic Effects of Debt Relief: Consumer Bankruptcy Protections in the Great Recession." NBER Working Paper 25685.

Bos, Marieke, Emily Breza, and Andres Liberman. 2018. "The Labor Market Effects of Credit Market Information." *Review of Financial Studies* 31 (6): 2005–37.

Dobbie, Will, Paul Goldsmith-Pinkham, Neale Mahoney, and Jae Song. 2016. "Bad Credit, No Problem? Credit and Labor Market Consequences of Bad Credit Reports." NBER Working Paper 22711.

Dobbie, Will, Benjamin J. Keys, and Neale Mahoney. 2017. "Credit Market Consequences of Credit Flag Removals." https://ssrn.com/abstract=2991711.

Dobbie, Will, and Jae Song. 2015. "Debt Relief and Debtor Outcomes: Measuring the Effects of Consumer Bankruptcy Protection." *American Economic Review* 105 (3): 1272–1311.

Dobbie, Will, and Ann Woods. 2019. "Randomized Debt Restructuring Experiment." AEA RCT Registry. November 15. https://doi.org/10.1257/rct.5045-1.0.

Eberly, Janice, and Arvind Krishnamurthy. 2014. "Efficient Credit Policies in a Housing Debt Crisis." *Brookings Papers on Economic Activity* (Fall): 73–136.

Ganong, Peter, and Pascal Noel. 2018. "Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession." NBER Working Paper 24964.

Gross, David B., and Nicholas S. Souleles. 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics* 117 (1): 149–85.

Gross, Tal, and Matthew J. Notowidigdo. 2011. "Health Insurance and the Consumer Bankruptcy Decision: Evidence from Expansions of Medicaid." *Journal of Public Economics* 95 (7–8): 767–78.

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. 2014. "Liquidity Constraints and Consumer Bankruptcy: Evidence from Tax Rebates." *Review of Economics and Statistics* 96 (3): 431–43.

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. Forthcoming. "The Marginal Propensity to Consume over the Business Cycle." *American Economic Journal: Macroeconomics.*

Haughwout, Andrew, Ebiere Okah, and Joseph Tracy. 2016. "Second Chances: Subprime Mortgage Modification and Redefault." *Journal of Money, Credit and Banking* 48 (4): 771–93.

Herkenhoff, Kyle. 2019. "The Impact of Consumer Credit Access on Unemployment." *Review of Economic Studies* 86 (6): 2605–42.

Herkenhoff, Kyle, Gordon Phillips, and Ethan Cohen-Cole. 2016. "How Credit Constraints Impact Job Finding Rates, Sorting & Aggregate Output." NBER Working Paper 22274.

Hull, Peter. 2018. "Estimating Treatment Effects in Mover Designs." http://www.mit.edu/~hull/movers_042018.pdf.

Hunt, Robert M. 2005. "Whither Consumer Credit Counseling?" *Federal Reserve Bank of Philadelphia Business Review* Q4: 9–20.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96 (5): 1589–1610.

Kirkebøen, Lars J., Edwin Leuven, and Magne Mogstad. 2016. "Field of Study, Earnings, and Self-Selection." *Quarterly Journal of Economics* 131 (3): 1057–1111.

Kline, Patrick, and Christopher R. Walters. 2016. "Evaluating Public Programs with Close Substitutes: The Case of Head Start." *Quarterly Journal of Economics* 131 (3): 1057–1111.

Li, Wenli, Michelle J. White, and Ning Zhu. 2011. "Did Bankruptcy Reform Cause Mortgage Defaults to Rise?" *American Economic Journal: Economic Policy* 3 (4): 123–47.

Lusardi, Annamaria, Daniel Schneider, and Peter Tufano. 2011. "Financially Fragile Households: Evidence and Implications." *Brookings Papers on Economic Activity* (Spring): 83–150.

Mahoney, Neale. 2015. "Bankruptcy as Implicit Health Insurance." *American Economic Review* 105 (2): 710–46.

Money Management International. 2014. Client Counseling Database [confidential dataset]. Accessed January 2014.

Mullainathan, Sendhil, and Eldar Shafir. 2013. *Scarcity: Why Having Too Little Means So Much.* New York: Times Books.

O'Neill, Barbara, Aimee Prawitz, Benoit Sorhaindo, Jinhee Kim, and E. Thomas Garman. 2006. "Changes in Health, Negative Financial Events, and Financial Distress/Financial Well-Being for Debt Management Program Clients." *Journal of Financial Counseling and Planning* 17 (2): 46–100.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103 (6): 2530–53.

Reardon, Sean F., and Stephen W. Raudenbush. 2013. "Under What Assumptions Do Site-by-Treatment Instruments Identify Average Causal Effects?" *Sociological Methods & Research* 42 (2): 143–63.

Rubin, Donald B. 1981. "The Bayesian Bootstrap." *Annals of Statistics* 9 (1): 130–34.

Social Security Administration. 2014. Administrative Tax Records [confidential dataset]. Accessed January 2014.

Staten, Michael E., and John M. Barron. 2006. "Evaluating the Effectiveness of Credit Counseling." https://consumerfed.org/pdfs/Credit_Counseling_Report061206.pdf.

TransUnion. 2004–2007. "TransUnion Consumer Credit Report Data." https://www.povertyactionlab.org/admindatacatalog/transunion-consumercredit-report-data (accessed November 14, 2019).

Wilshusen, Stephanie M. 2011. "Meeting the Demand for Debt Relief." Federal Reserve Bank of Philadelphia Payment Cards Center Discussion Paper.