Decomposing Present Value Effects on Defaults: Evidence from a large-scale restructuring experiment

November 18, 2018

1 Introduction

During the U.S. Great Recession, many influential voices called for policies that would reduce the debt obligations of financially distressed consumers as a way of reducing economy-wide defaults and increasing aggregate demand.[1] The debate focused on the relative merits of different forms of debt relief. In particular, was it preferable to reduce principal or to provide temporary payment relief? If liquidity constraints play a dominant role, then the latter have the benefit of front-loading the spending response (Eberly and Krishnamurthy [2014]). Indeed, studying the policy variation in the US Home Affordable Mortgage Program (HAMP), Ganong and Noel [2017] convincingly established that principal reductions had very limited effects over and beyond payment reductions. But a better understanding of consumer responses, and financial institution incentives to offer the precise form of payment reductions, is needed in order to guide the design of future large-scale debt relief programs, as well as the day-to-day restructurings of distressed debt. This is the question we address in this paper.

More precisely, our analysis seeks to disentangle the effects of the three most common forms of debt relief—reductions in loan contract rates, maturity extensions, and temporary payment deferrals -while holding consumers' outstanding debt balance constant. The latter is often a practical constraint on loan restructurings, since banks tend to keep loans at face value on their balance sheet, and are very averse to realizing face-value losses. In theory, each of these three types of debt relief each affect consumers' decisions to spend and default in different ways, since they do not have the same effect on the intertemporal path of

[1] See Mian and Sufi [2009] and Dynan [2012] for the effect of debt on weak recovery.

payments that they face. As we explicate, reductions in the loan contract rates, similar to debt write-downs, have large effects on the present value of the loan calculated at the consumers' own discount rate, and comparatively small effects on their monthly payments. By contrast, temporary relief has large effects on monthly payments, and a negligable effect on the present value. Understanding how consumers respond in practice to these two types of debt relief therefore allows us to decompose present-value effects, in ways that the literature has not yet addressed.

In order to shed light on these questions, we design and implement a large scale field experiment at a large European retail bank, where we deliberately create variation in restructured contract terms. Our subject pool consists of about 25,000 pre-existing consumers whom are in delinquent status on their personal loans, and are headed for default if not offered debt relief. Our design randomizes the contract term independently along three dimensions: the interest rate, the amount of temporary payment reductions, and the maturity of restructured loans. Consumers in the rate relief group are pushed to receive a rate reduction relative to treatment, and consumers in the relief group are offered a 3 month interest-on-principal only payment reduction.

The rate reduction is quantitatively large, 480bps APR, a discernible change to the old contract rate sufficiently large to wake the consumers up. Importantly, both the payment and present-value shocks are unexpected -they offered by the financial instution for the first time as a part of the debt relief program designed for this trial, and are not offered by other institutions. To the best of our knowledge, this is the first paper to create truly exogenous, unexpected variation *seperately* in debt service flow and present-value stock, in order to better understand their relative contributions.

We then confront our empirical results to the predictions of a structural model of consumption and default. It allows us to contrast the aggregate effects of debt restructuring policies to the effects of more standard stabilization policies such as fiscal and monetary policy, as well as to study their interaction. Finally, it allows us to quantify the effect of debt restructuring on consumer welfare, consumption and wealth inequality.

2 Data and Environment

For the purposes of the field experiment, we collaborate with a large retail bank. The financial institution is one of the largest retail banking platforms in Europe, with more than 10 million customers, and a 20% market share in the local retail market, with a customer

base that is representative of the local banked population. It offers a multitude of financial products, including credit cards, deposit and investment accounts, consumer loans, payment services, and insurance.

We merge four datasets. First, we have detailed information on the loan contract before and after the restructuring, including contract terms, such as the rate, maturity, outstanding balance and the payment schedule, as well as the performance of the restructured loan, such as the date at which it reached 30 plus or 90 plus days late non-performing status. Second, we have a rich set of demographic data, including age, gender, marital status, education, city of residence and profession. Third, credit card data includes card limits, within month expenditures in 18 spending categories, end-of-month balances, payments made toward the balances.

The data is of monthly frequency and the unit of analysis is an individual. If an individual has multiple accounts or had multiple loans consolidated, then accounts are matched and variables are aggregated using a unique citizenship number, and verified using a customer identification number, ensuring perfect match quality. Information regarding balance sheet variables are end-of-month calculations; however, credit card variables are end-of-billing-cycle calculations. Due to the administrative nature of our data, it has far fewer problems with attrition, non-response, and measurement error than survey data sources.

3 Conceptual Framework

The consumers we consider all have a non-collateralized personal loans, which is a fixed-rate loan with fixed nominal payments. Write D for the initial outstanding principal on this loan, T for its maturity, and R for its contractual interest rate. The initial monthly amortizing payments $\text{Pay} = \text{Pay}(D, T, R)$ are given by the standard mortgage formula,

$$\text{Pay} = DR \left[1 - \frac{1}{(1+R)^T} \right]^{-1} \tag{1}$$

Consider an agent (the consumer or the bank) discounting loans at discount rate R^*. Given payments Pay and maturity T, we write the present value of the stream of payments at the inception of the loan, $\text{PV} = \text{PV}(\text{Pay}, T, R^*)$ which is equal to

$$PV = \frac{\text{Pay}}{R^*} \left[1 - \frac{1}{(1+R^*)^T} \right] \tag{2}$$

From the perspective of the lender, R^* is the risk-adjusted cost of funding. For lenders, the risk adjusted discount rate reflects the opportunity cost of funds and significantly exceeds the risk-free cost-of-funding, as the loan is highly likely to default at restructuring. From the perspective of the borrower, R^* is the subjective discount rate. In some models, this corresponds to his own cost of marginal funding at each date. In others, such as models with liquidity constraints, R^* incorporates the shadow cost of that constraint. For a borrowing constrained consumer who is on the verge of defaulting, this number could also potentially exceed the contract rate.

Our experiments create separate variation in Pay and PV. For any debt restructuring that we consider, we can calculate the change in payments Pay resulting from the restructuring using formula (1). Then, given various assumptions about the discount rate R^*, we can calculate PV using formula (2). One way to understand this is to consider the approximation to (1) and (2)

$$\text{Pay} \simeq D\left(\frac{R}{2}+\frac{1}{T}\right)$$
$$PV \simeq D\left(1+\frac{R-R^*}{2}T\right)$$

In our sample, a typical annual interest rate is $R = 16\%$ and a typical initial maturity is $T = 2$ years. Hence an interest contributes $\frac{8\%}{8\%+50\%} \simeq \frac{1}{6}$ to Pay and amortizing principal payments contribute the rest.

Although rate reductions keep the face-value of the loan constant, the present-value of the stream of payments is significantly altered. This present-value effect lowers the market value of the loan from the lender's perspective, and the present-value of payments from the borrowers perspective. Therefore a reduction in payment stream achieved through a face value reduction could also be replicated via a change in the contract rate. Borrowers see effect in current and future payments in both scenarios. Of course in an amortizing framework, the composition of principal and interest rates differ,

In many benchmark models of consumption, such as the permanent-income hypothesis, the consumption response of households is entirely dictated by PV, and there is no effect on Pay conditional on PV. Another alternative is models where consumers only default if current cash-on-hand less than debt-service amount. For many low-income, high-risk individuals already in delinquent status with little access to credit markets, such non-forward-looking behavior. In such an environment only periodic payment flows might matter, and

longer-run obligations captured by present-value may have no effect on household borrowing and saving decisions. These could be models with severe liquidity constraints, or models with myopic behavior, say if consumers focus only on period flow payment profiles. However in general, even when liquidity is scarce, the entire profile of payment obligations may matter for both the consumption and the default decision.

The model that we write down is an extension of the models of liquidity constraints of Deaton [1991] and Carroll [1992] that incorporate a default option (such as in Chatterjee et al. [2007] and Livshits et al. [2007]). These models nest the permanent-income hypothesis as a special case when liquidity is large. Depending on the importance of liquidity and endogenous default constraints, Pay or PV may matter more, reflecting the relative importance of cash flow versus present-value constraints.

In order to discipline our model, we therefore first explore in reduced form the relationship between the consumption response with Pay and PV. In a second step, we will ensure that this relationship is accurately captured by our structural model. Estimating equation (**??**) using our experiment to create random variation in Pay_{it} and PV_{it} allows us to disentangle the key factors driving consumer behavior, and in particular the relative role of liquidity constraints in affecting consumption behavior at various horizons. While the estimated coefficients $\widehat{\phi}_j$ and $\widehat{\psi}_j$ are interesting on their own, we will make sure that these same coefficients are generated by our model. Having validated our model in this way, we will be able to use the model for the counterfactual analyses described above.

4 Experimental Design

Our design varies contract terms along three dimensions in a 2-by-2-by-2 design. We take consumer's outstanding debt balance D_0 as given, and exogenously vary interest rates, maturity and payment schedules.

The assignment of consumers to treatment legs is done in advance to the consumer becoming delinquent. We assign all of the customers at the bank (1m+) to 8 bins, a {0,1} bin with respect each contract feature: rate, term, relief. We do this by drawing three random numbers for each contract feature, and assigning a customer to group 1 if the customer is above a specific threshold. This threshold equals 0.5 for rate and maturity, 0.66 for relief. This leads to 8 groups of participants.

When randomizing, we stratify the participants into non-overlapping and exhaustive bins with respect to variables S. The variable set S includes the number of late pay days; out-

standing balance D and city of residence, latter included in order to capture local economic shocks. Call these bins $b \in B$ and let N_b indicate the number of cardholders in each bin. Finally we check for balance between H and L treatments, for each treatment leg, by comparing the difference in means, and the regression coefficients of assignment to treatment on observables. We do K draws and pick the draw with the minimum maximum t-statistic.

The participants in the experiment are 25,000 consumers that fall into delinquent status as of or after June 2017. These customers have failed to make a payment on at least one debt contract and are being monitored by the bank. They have been previously nudged via text messages, and are contacted via phone calls to remind them of their payments.

At the point when the customer is late, we anticipate a restructuring and assign a new rate, and maturity for the potentially restructured contract.

- *Interest rates.* The delinquent contract features an interest R and the restructuring features an interest rate reductionto $R' < R$. Consumers in the four high rate treatment legs all get a 60bps APR rate reduction, and the typical contract in the low rate treatment group is intended to get a 540bps APR rate reduction. However, the new minimum rate is bounded below by a minimum set by the financial institution.

- *Maturity.* The restructured contract features a maturity extension offer to $T' > T$. As we describe later, this number need not be the term of the restructured contract, but is a recommendation made by the bank representative while the final term is bilaterally negotiated. When assigning customers to maturity offer bins, we first group customers to grids of width 12 with respect to the remaining maturity of the old contract. Name these bins with the largest element in each bin $\bar{T}_k$ -for example, those with 1 to 12 and 13 to 24 months remaining are in the 12 and 24 month bin respectively. We then offer T^{offer}, a maturity extension to $\bar{T}_k$ times 150% to those in the low maturity treatment, and a $\bar{T}_k$ times 200% to those in the high maturity treatment. The offer is bounded above by the maximum maturity set by the regulatory authority, 48 or 72 months depending on when the old contract is originated.

- *Payment schedules.* We select a subset of customers for temporary payment relief offers. For these customers, we offer to postpone the payment of the principal in the first quarter, and the total payments for this group in the first quarter then equals interest payment on the principal, DR'. The amortizing payments for these consumers start in month 4.

After a new contract rate and maturity has been assigned, if the consumer still has not made a payment after 60 days, the financial institution contacts the consumer through an in-house call-center for a potential restructuring, in order work-out a plan for the customer to make debt payments. It is very common to observe the ringing phone syndrome -that a delinquent borrower wouldn't pickup when banks would call. The call-center representative asks the customer about the nature of their financial distress, and asks if the customer is interested in restructuring the contract.

At the restructuring screen, the call-center employee sees the outstanding balance, the new interest rate and a maturity offer. The employee and customer discuss how much the customer could pay each month, where the employee offers maturity T^{offer}.

If the customer and the bank representative agree on the contract term, immediately before they finalize the restructuring, the customers in the relief group are offered 3 months of principle payment relief. The customer has the option to either accept or reject the relief offer. If the customer rejects the relief offer, the payment schedule determined by the mortgage formula, given D, T' and R'. If the customer takes the relief offer, payments in the first three months equal DR, interest on the principal, and payments after month $t = 3$ are given by the mortgage formula, given D, $T' - 3$ and R', as the payment relief is offered only after the customer agrees on an interest rate and maturity. After the payment schedule is determined, the contract is forwarded to processing.

Prior to our controlled trial offering various types of debt relief, the financial institution had a practice of offering maturity extensions to delinquent borrowers. However, the bank did not offer interest rate reductions and temporary relief, so the principal levers that create variation in short-term payments and long-term present-value of obligations could be considered unanticipated.

Covariate balance. First, we show that different treatment legs have statistically indistinguishable ex-ante covariates and pre-trends. Remember that the assignment of consumers to each treatment leg is done in advance of the customer falling into delinquent status and restructuring. Therefore, the take-up of the restructured contract could vary with respect to the contract offer. For example, consumers in the high interest rate bin may be less likely to restructure, or under a high contract rate, consumers could opt to shorten debt maturity. higher interest rates may lead customers to pick shorter maturity contracts. Similarly, we could observe adverse selection -borrowers interested in restructuring into higher rate contracts could be the high risk types whom are least likely to pay back. Such selection effects could be detected from the covariate balance table.

5 Event Study

In our empirical analysis, we focus on three classes of outcome variables: delinquencies, recovery and spending. Our measure of delinquencies is late payments, and defaults, i.e. 30 days and 90 days past due. The former captures, arrears, part of a debt that is overdue after missing one or more required payments. The latter captures contracts that are forwarded to collection by the financial institution. Our measure of recovery is the present value, calculated from the perspective of the lender at the time of restructuring, of the payments made by the customer through the life-cycle of the loan until it falls into collection status, defined as 90 days past due. Our primary spending measure is total spending by the consumers incurred on the bank's credit cards cards. As discussed in Section 2, we are able to disaggregate this total into conventional credit card purchases and installment expenditures; as well as durable/nondurable/services spending.

References

Christopher D. Carroll. The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence. *Brookings Papers on Economic Activity*, 1992(2):61–156, January 1992.

Satyajit Chatterjee, Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull. A Quantitative Theory of Unsecured Consumer Credit with Risk of Default. *Econometrica*, 75(6): 1525–1589, November 2007.

Angus Deaton. Saving and Liquidity Constraints. *Econometrica*, 59(5):1221–1248, 1991.

Karen Dynan. Is a Household Debt Overhang Holding Back Consumption? *Brookings Papers on Economic Activity*, pages 299–362, April 2012. ISSN 0007-2303.

Janice C. Eberly and Arvind Krishnamurthy. Efficient Credit Policies in a Housing Debt Crisis. *Brookings Papers on Economic Activity*, pages 73–118, 2014.

Peter Ganong and Pascal Noel. The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession. *Manuscript*, January 2017.

Igor Livshits, James MacGee, and Michele Tertilt. Consumer Bankruptcy: A Fresh Start. *American Economic Review*, 97(1):402–418, March 2007.

Atif Mian and Amir Sufi. The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis. *Quarterly Journal of Economics*, 124(4):1449–1496, November 2009.