DO BANKS PASS THROUGH CREDIT EXPANSIONS TO CONSUMERS WHO WANT TO BORROW?*

SUMIT AGARWAL
SOUPHALA CHOMSISENGPHET
NEALE MAHONEY
JOHANNES STROEBEL

We propose a new approach to studying the pass-through of credit expansion policies that focuses on frictions, such as asymmetric information, that arise in the interaction between banks and borrowers. We decompose the effect of changes in banks' cost of funds on aggregate borrowing into the product of banks' marginal propensity to lend (MPL) to borrowers and those borrowers' marginal propensity to borrow (MPB), aggregated over all borrowers in the economy. We apply our framework by estimating heterogeneous MPBs and MPLs in the U.S. credit card market. Using panel data on 8.5 million credit cards and 743 credit limit regression discontinuities, we find that the MPB is declining in credit score, falling from 59% for consumers with FICO scores below 660 to essentially zero for consumers with FICO scores above 740. We use a simple model of optimal credit limits to show that a bank's MPL depends on a small number of parameters that can be estimated using our credit limit discontinuities. For the lowest FICO score consumers, higher credit limits sharply reduce profits from lending, limiting banks' optimal MPL to these consumers. The negative correlation between MPB and MPL

*This article was previously circulated as "Do Banks Pass Through Credit Expansions? The Marginal Profitability of Consumer Lending During the Great Recession." For helpful comments, we are grateful to Viral Acharya, Scott Baker, Eric Budish, Charles Calomiris, Chris Carroll, Liran Einav, Alex Frankel, Erik Hurst, Anil Kashyap, Theresa Kuchler, Randall Kroszner, Marco di Maggio, Matteo Maggiori, Atif Mian, Rick Mishkin, Christopher Palmer, Jonathan Parker, Thomas Philippon, Amit Seru, Andrei Shleifer, Amir Sufi, Alessandra Voena, and Arlene Wong, as well as seminar and conference participants at the 2017 AEA Meetings, Bank of England, Banque de France, Bank for International Settlements, Bank of Italy, Baruch, Berkeley Econ, Berkeley Haas, Brown University, Boston University, Chicago Booth, Columbia University, Columbia GSB, Federal Reserve Bank of Philadelphia, Federal Reserve Bank of St. Louis, Financial Conduct Authority, Goethe University Frankfurt, HEC Paris, Ifo Institute, ITAM, LMU Munich, Mannheim University, MIT, NBER Summer Institute, NYU Stern, Northwestern University, SAIF, SED 2015, Stanford University, University of Chicago, University of Minnesota, UT Austin, Yale University, and Yale SOM. We thank Regina Villasmil, Mariel Schwartz, Yin Wei Soon, Andreas Weber, and Hanbin Yang for truly outstanding and dedicated research assistance. The views expressed are those of the authors alone and do not necessarily reflect those of the Office of the Comptroller of the Currency.

The Quarterly Journal of Economics (2018), 129–190. doi:10.1093/qje/qjx027.
Advance Access publication on July 10, 2017.

reduces the impact of changes in banks' cost of funds on aggregate household borrowing, and highlights the importance of frictions in bank-borrower interactions for understanding the pass-through of credit expansions. *JEL Codes*: D14, E51, G21.

I. Introduction

During the Great Recession, policy makers sought to stimulate the economy by providing banks with lower-cost capital and liquidity. One goal was to encourage banks to expand credit to households and firms that would, in turn, increase their borrowing, spending, and investment.[1] Yet empirically analyzing the strength of this "bank lending channel" is challenging. For example, there was a large drop in U.S. banks' cost of funds in the fall of 2008, when the Federal Funds Rate was cut to zero in response to the financial crisis. However, this was exactly the period when lenders and borrowers were updating their expectations about the economy, making it practically impossible to use time-series analysis to isolate the effect of the change in monetary policy on borrowing volumes.

In this article, we propose a new empirical approach to studying the bank lending channel that focuses on frictions, such as asymmetric information, that arise in bank-borrower interactions. Our approach is based on the observation that the effect on aggregate borrowing of a change in banks' (shadow) cost of funds—for example, due to an easing of monetary policy, a reduction in capital requirements, or a market intervention that reduces financial frictions—can be expressed as a function of the supply and demand for credit by different agents in the economy. This approach is empirically useful because it allows us to quantify the pass-through of credit expansion policies by decomposing the overall effect into objects that can be estimated using micro-data on lending and quasi-exogenous variation in contract terms. This approach is also conceptually useful because understanding the

1. For example, when introducing the Financial Stability Plan, Geithner (2009) argued that "the capital will come with conditions to help ensure that every dollar of assistance is used to generate a level of lending greater than what would have been possible in the absence of government support." In Europe, similar schemes were put in place in order to reduce the cost of capital for banks that expand lending to nonfinancial firms and households (e.g., the "Funding for Lending Scheme" of the Bank of England, and the "Targeted Longer-Term Refinancing Operation" of the ECB). See Online Appendix A for more information.

relative importance of these supply and demand factors is independently important for designing effective policies.

We apply our framework to the U.S. credit card market. As we discuss below, in this market credit limits are a key determinant of credit supply and the primary margin of adjustment to changes in the cost of funds. Let c denote the banks' cost of funds, CL_i the credit limit of consumer i, and q_i the borrowing of that consumer. The effect of a change in c on total borrowing q can be expressed as the product of banks' marginal propensity to lend (MPL) to consumer i and that consumer's marginal propensity to borrow (MPB), aggregated across all the consumers in the economy:

$$-\frac{dq}{dc} = \int_i \underbrace{-\frac{dCL_i}{dc}}_{\text{MPL}} \times \underbrace{\frac{dq_i}{dCL_i}}_{\text{MPB}}. \tag{1}$$

We operationalize our framework by estimating heterogeneous MPBs and MPLs using panel data on all credit cards issued by the eight largest U.S. banks. These data, assembled by the Office of the Comptroller of the Currency (OCC), provide us with monthly account-level information on contract terms, utilization, payments, and costs for more than 400 million credit card accounts between January 2008 and December 2014. The data are merged with credit bureau information, allowing us to track balances across consumers' broader unsecured credit portfolios.

Our research design exploits the fact that banks sometimes set credit limits as discontinuous functions of consumers' FICO credit scores. For example, a bank might grant a \$2,000 credit limit to consumers with a FICO score below 720 and a \$5,000 credit limit to consumers with a FICO score of 720 or above. We show that other borrower and contract characteristics trend smoothly through these cutoffs, allowing us to use a regression discontinuity strategy to identify the causal impact of providing extra credit at prevailing interest rates. We identify a total of 743 credit limit discontinuities in our data, which are distributed across the range of the FICO score distribution. We observe 8.5 million new credit cards issued to borrowers within 50 FICO score points of a cutoff.

Using this regression discontinuity design, we estimate substantial heterogeneity in MPBs across the FICO score distribution. For the least credit-worthy consumers (FICO $\leqslant$ 660), a \$1 increase in credit limits raises borrowing volumes on the treated credit card by 58 cents at 12 months after origination. This effect is due to increased spending and is not explained by a shifting of borrowing across credit cards. For the highest FICO score consumers ($>$740), we estimate a 23% effect on the treated card that is entirely explained by a shifting of borrowing across credit cards, with an increase in credit limits having no effect on total borrowing.

We next analyze how banks pass through credit expansions to different consumers. As discussed above, estimating the MPL directly using observed changes in the cost of funds is challenging, because such changes are typically correlated with shifts in the economic environment that also affect borrowing and lending decisions. We use economic theory and our quasi-exogenous variation in credit limits to address this identification problem. In particular, we write down a simple model of optimal credit limits to show that a bank's MPL depends on a small number of "sufficient statistics" that can be estimated directly using our regression discontinuities. Our approach involves a trade-off. To avoid the standard identification problem, we need to assume that banks respond optimally to changes in the cost of funds and that we can measure the incentives faced by banks. We think both assumptions are reasonable: credit card lending is highly sophisticated and our estimates of bank incentives are fairly precise. Indeed, we show that observed credit limits are close to the optimal credit limits implied by the model.

In our model, banks set credit limits at the level where the marginal profit from a further increase in credit limits is zero. A decrease in banks' cost of funds reduces the cost of extending a given unit of credit and corresponds to an outward shift in the marginal profit curve. As shown in Figure I, a reduction in the cost of funds has a larger effect on optimal credit limits when the marginal profit curve is relatively flat (Panel A) than when it is relatively steep (Panel B).

What are the economic forces that determine the slope of marginal profits? One important factor is the degree of adverse selection. With adverse selection, higher credit limits are disproportionately taken up by consumers with higher probabilities of default. These higher default rates lower the marginal profit of lending, thereby generating more steeply downward-sloping marginal profits. Higher credit limits can also lower marginal profits holding the distribution of marginal borrowers fixed. For example, if higher debt levels have a *causal* effect on the

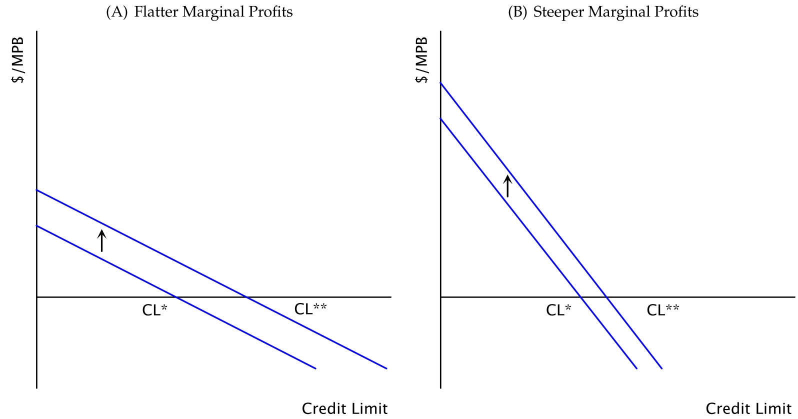

FIGURE I

Pass-Through of Reduction in Cost of Funds into Credit Limits

Figure shows marginal profits for lending to observationally identical borrowers. A reduction in the cost of funds shifts the marginal profit curve outward, and raises equilibrium credit limits (CL* → CL**). Panel A considers a case with a relatively flat marginal profit curve; Panel B considers a case with a steeper marginal profit curve. The vertical axis is divided by the MPB because a given decrease in the cost of funds induces a larger shift in marginal profits when credit card holders borrow more on the margin. See Section VI for more details.

probability of default—as they do, for example, in the strategic bankruptcy model of Fay, Hurst, and White (2002)—then higher credit limits, which increase debt levels, will also raise default rates. As before, this lowers the marginal profit of lending, generating more steeply downward-sloping marginal profits.[2]

The effect of these (and other) frictions in the bank-borrower relationship on the pass-through of credit expansions is fully captured by the slope of the marginal profit curve. Indeed, by estimating this slope, we can quantify the pass-through of credit expansion policies without requiring strong assumptions on the underlying micro-foundations of consumer behavior. This approach of estimating sufficient statistics rather than model-dependent structural parameters builds on approaches that are increasingly popular in the public finance literature (see Chetty 2009).

2. This mechanism also arises in models of myopic behavior, in which consumers, faced with a higher credit limit, borrow more than they can repay because they do not fully internalize having to repay their debt in the future.

How do we estimate the slope of marginal profits? Conceptually, each quasi-experiment provides us with two moments at the prevailing credit limit: marginal profits, which can be estimated using our regression discontinuities, and average profits, which can be directly observed in our data and which correspond to the area under the marginal profit curve in Figure I. With these two moments, we can identify any two-parameter curve for marginal profits. Intuitively, for a given credit limit, larger average profits correspond to a steeper slope of the marginal profit curve.

To obtain quantitative estimates of the MPL, we parameterize the marginal profit curve using a linear functional form. We find that marginal profits are most steeply downward-sloping for consumers with the lowest FICO scores, consistent with significant asymmetric information in this segment of the population. Consequently, a one percentage point reduction in the cost of funds increases optimal credit limits by $253 for borrowers with FICO scores below 660, compared with $1,224 for borrowers with FICO scores above 740. While these precise estimates rely on our linear functional form assumption, we prove that, given the moments in our data, our finding of larger pass-through to higher FICO score borrowers is qualitatively robust to any functional form that satisfies an appropriately defined single-crossing condition.

Taken together, our estimates imply that MPBs and MPLs are negatively correlated across consumers. This negative correlation is economically significant. Suppose one incorrectly calculated the impact of a decrease in the shadow cost of funds as the product of the average MPL and the average MPB in the population. This would generate an estimate of the effect on total borrowing that is approximately twice as large as an estimate that accounts for this correlation.

We view our article as making three contributions. First, our article builds on a literature that has estimated marginal propensities to consume (MPCs) and MPBs using shocks to income and liquidity. Our finding of substantial heterogeneity in MPBs by FICO score complements recent papers by Parker et al. (2013) and Jappelli and Pistaferri (2014) that have shown substantial heterogeneity in MPCs out of income shocks, and recent work by Mian and Sufi (2011) and Mian, Rao, and Sufi (2013), who have shown substantial heterogeneity in MPCs out of shocks to housing prices and wealth. Most closely related are Gross and Souleles (2002), who estimate MPBs using time-series variation in credit limits but do not have the power to identify heterogeneous effects, and Aydin (2017), who estimates MPBs using a credit limit experiment in Turkey.[3] We advance this literature by providing the first joint estimates of consumers' MPBs and banks' MPLs. Estimating both objects together is important because it allows for an evaluation of credit expansion policies that are intermediated by banks. We show that the interaction between MBPs and MPLs across different types of consumers is key to understanding the aggregate impact of these policies.[4]

Second, our approach to estimating banks' MPLs highlights the importance of frictions in bank-borrower interactions—such as asymmetric information—in determining the strength of the bank lending channel. This complements research on how variation in capital and liquidity levels or risk across banks mediates the strength of the bank lending channel (see, among others, Kashyap and Stein 1994; Kishan and Opiela 2000; Jiménez et al. 2012, 2014; Acharya et al. 2015; Dell'Ariccia, Laeven, and Suarez 2016).[5] In our model, forces like liquidity levels affect banks' shadow cost of funds, c, and are therefore conceptually separable from the bank-borrower interactions that we focus on.

Third, our article contributes to a literature that has identified declining household borrowing volumes as a proximate cause of the Great Recession.[6] Within this literature, there is considerable debate over the relative importance of supply versus demand factors in explaining the reduction in aggregate borrowing. Our estimates suggest that both explanations have merit, with credit

3. Also see Zeldes (1989); Souleles (1999); Hsieh (2003); Stephens (2003, 2008); Johnson, Parker, and Souleles (2006); Agarwal, Liu, and Souleles (2007); Blundell, Pistaferri, and Preston (2008); Dobbie and Skiba (2013); Agarwal and Qian (2014); Agarwal et al. (2015a); Baker (2015); Gelman et al. (2015); Parker (2015); Sahm, Shapiro, and Slemrod (2015); and Bhutta and Keys (2016). Jappelli and Pistaferri (2010) and Zinman (2015) review this literature. See Carroll (1997, 2001) for theoretical foundations.

4. A related literature has analyzed heterogeneity in the transmission of monetary policy through other channels. See Doepke and Schneider (2006); Coibion et al. (2012); Di Maggio, Kermani, and Ramcharan (2014); Keys et al. (2014); Auclert (2016); Hurst et al. (2016); Chakraborty, Goldstein, and MacKinlay (2017); and Drechsler, Savov, and Schnabl (2017).

5. It also relates to recent research by Scharfstein and Sunderam (2016), who show that the pass-through of credit expansion is also affected by regional variation in the competitive environment.

6. See, for example, Mian and Sufi (2010, 2014); Guerrieri and Lorenzoni (2011); Hall (2011); Philippon and Midrigan (2011); Eggertsson and Krugman (2012); Mian, Rao, and Sufi (2013); Korinek and Simsek (2016).

supply being the limiting factor at the bottom of the FICO score distribution and credit demand being the limiting factor at higher FICO scores.

There are a number of caveats for using our estimates to obtain a complete picture of the effectiveness of monetary policy during the Great Recession. First, we only study one market. While the credit card market is of stand-alone interest because credit cards are the marginal source of credit for many U.S. households, other markets, such as mortgage lending and small-business lending, are probably more important channels for monetary policy transmission.[7] However, we think that our finding of lower pass-through to less creditworthy borrowers—for example, because of asymmetric information—is likely to apply across this broader set of markets, all of which feature significant potential for adverse selection and moral hazard.[8] A second caveat is that our article does not assess the desirability of stimulating household borrowing from a macroeconomic stability or welfare perspective. For example, while extending credit to low FICO score households might lead to more borrowing and consumption in the short run, we do not evaluate the consequences of the resulting increase in household leverage. Our results also do not capture general equilibrium effects that might arise from the increased spending of low-FICO-score households.

The rest of the article proceeds as follows. Section II presents background information on the determinants of credit limits and describes our credit card data. Section III discusses our regression discontinuity research design. Section IV verifies the validity of this research design. Section V presents our estimates of the marginal propensity to borrow. Section VI provides a model of credit limits. Section VII presents our estimates of the marginal propensity to lend. Section VIII concludes.

7. According to the 2010 Survey of Consumer Finances, 68% of households had a credit card versus 10.3% for a home equity line of credit and 4.1% for other lines of credit. Moreover, credit cards were particularly important during the Great Recession when many homeowners were underwater and unable to borrow against home equity. In our sample, credit cards issued to consumers with FICO scores above 740 had, on average, $1,294 of interest-bearing debt at one year after origination, indicating that credit cards were a key source of credit even in the upper range of the FICO distribution.

8. See, for example, Petersen and Rajan (1994); Adams, Einav, and Levin (2009); Karlan and Zinman (2009); Keys et al. (2010); Kurlat and Stroebel (2015); Stroebel (2015); Hertzberg, Liberman, and Paravisini (2016).

II. Background and Data

Our research design exploits quasi-random variation in the credit limits set by credit card lenders (see Section III). In this section, we describe the process by which banks determine these credit limits and introduce the data we use in our empirical analysis. We then describe our process for identifying credit limit discontinuities and present summary statistics on our sample of quasi-experiments.

II.A. How Do Banks Set Credit Limits?

Most credit card lenders use credit-scoring models (also called "scorecards") to make their pricing and lending decisions. These models are developed by analyzing the correlation between cardholder characteristics, contract terms, and outcomes such as default and profitability. Banks use both internally developed and externally purchased credit-scoring models. The most commonly used external credit scores are called FICO scores, which are developed by the Fair Isaac Corporation. FICO scores are used by the majority of financial institutions and take into account a consumer's payment history, credit utilization, length of credit history, and the opening of new accounts. Scores range between 300 and 850, with higher scores indicating a lower probability of default. The vast majority of the population has scores between 550 and 800.

Each bank develops its own policies and risk tolerance for credit card lending, with lower credit limits generally assigned to consumers with lower credit scores. Setting cutoff scores is one way that banks assign credit limits. For example, banks might split their customers into groups based on their FICO scores and assign each group a different credit limit (FDIC 2007). In Online Appendix B, we show how such a contract-setting process can be optimal in the presence of fixed costs for determining optimal contract terms for a set of observationally similar individuals. We also show that the magnitude of profits forgone by suboptimally pricing individuals close to credit limit discontinuities is small relative to industry estimates of the fixed cost of determining optimal contract terms for similar individuals.

II.B. Data

Our main data source is the Credit Card Metrics (CCM) data set assembled by the U.S. Office of the Comptroller of the

Currency (OCC).[9] The CCM data set has two components. The main data set contains account-level panel information on credit card utilization (e.g., purchase volume, measures of borrowing volume such as ADB), contract characteristics (e.g., credit limits, interest rates), charges (e.g., interest, assessed fees), performance (e.g., chargeoffs,[10] days overdue), and borrower characteristics (e.g., FICO scores) for all credit card accounts at the eight largest U.S. banks. The second data set contains portfolio-level information for each bank on items such as operational costs and fraud expenses across all credit cards managed by these banks. Both data sets are submitted monthly; reporting started in January 2008 and continues through the present. We use data from January 2008 to December 2014 for our analysis. In the average month, we observe account-level information on over 400 million credit cards. See Agarwal et al. (2015b) for more details on these data and summary statistics on the full sample.

To track changes in borrowing across the consumers' broader credit portfolios, we merge the CCM data to quarterly credit bureau data using a unique identifier. The credit bureau data we observe were collected to study credit card borrowing and contain rich information on individuals' unsecured-borrowing behavior across all lenders (e.g., the total number of credit cards, total credit limits, total balances, length of credit history, and credit performance measures such as whether the borrower was ever more than 90 days past due on an account). We do not observe borrowing on secured credit products such as mortgages or auto loans.

II.C. Identifying Credit Limit Discontinuities

In our empirical analysis, we focus on credit cards that were originated during our sample period, which started in January 2008. Our data do not contain information on the credit supply functions of banks when the credit cards were originated.

9. The OCC supervises and regulates nationally chartered banks and federal savings associations. In 2008, the OCC initiated a request to the largest banks that issue credit cards to submit data on general purpose, private label, and small-business credit cards. The purpose of the data collection was to have more timely information for bank supervision.

10. "Chargeoffs" refer to an expense incurred on the lender's income statement when a debt is deemed uncollectible for being sufficiently long past due. For an open-ended account such as a credit card, regulatory rules usually require a lender to charge off balances after 180 days of delinquency.

Therefore, the first step involves backing out these credit supply functions from the observed credit limits offered to individuals with different FICO scores. To do this, we jointly consider all credit cards of the same type (co-branded, oil and gas, affinity, student, or other), issued by the same bank, in the same month, and through the same loan channel (preapproved, invitation to apply, branch application, magazine and internet application, or other). It is plausible that the same credit supply function was applied to each card within such an "origination group." Since our data end in December 2014, we only consider credit cards originated until November 2013 to ensure that we observe at least 12 months of postorigination data for each account. For each of the more than 10,000 resulting origination groups between January 2008 and November 2013, we plot the average credit limit as a function of the FICO score.

Panels A to D of Figure II show examples of such plots. Since banks generally adjust credit limits at FICO score cutoffs that are multiples of 5 (e.g., 650, 655, 660), we pool accounts into such buckets. Average credit limits are shown with dark lines; the number of accounts originated are shown with gray bars. Panels A and B show examples where there are no discontinuous jumps in the credit supply function. Panels C and D show examples of clear discontinuities. For instance, in Panel C, a borrower with a FICO score of 714 is offered an average credit limit of approximately \$2,900, while a borrower with a FICO score of 715 is offered an average credit limit of approximately \$5,600.

Although continuous credit supply functions are significantly more common, we detect a total of 743 credit limit discontinuities between January 2008 and November 2013. We refer to these cutoffs as "credit limit quasi-experiments" and define them by the combination of origination group and FICO score. Panel E of Figure II shows the distribution of FICO scores at which we observe these quasi-experiments. They range from 630 to 785, with 660, 700, 720, 740, and 760 being the most common cutoffs. Panel F shows the distribution of quasi-experiments weighted by the number of accounts originated within 50 FICO score points of the cutoffs, which is the sample we use for our regression discontinuity analysis. We observe more than 1 million accounts within 50 FICO score points of the most prominent cutoffs. Our experimental sample has 8.5 million total accounts, or about 11,400 per quasi-experiment.

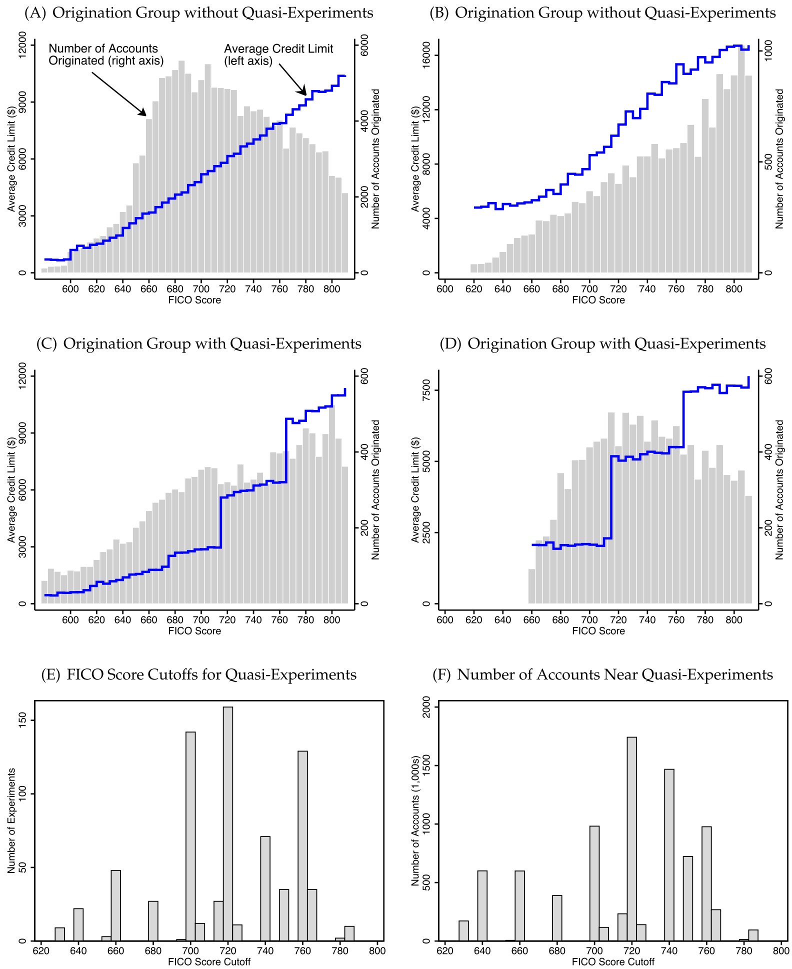

FIGURE II

Credit Limit Quasi-Experiments: Examples and Summary Statistics

Panels A to D show examples of average credit limits by FICO score for accounts in "origination groups" with and without credit limit quasi-experiments. Origination groups are defined as all credit cards of the same product-type originated by the same bank in the same month through the same loan channel. The horizontal axis shows FICO score at origination. The dark line plots the average credit limit for accounts in FICO score buckets of 5 (left axis); gray bars show the total number of accounts originated in those buckets (right axis). Panels E and F show summary statistics for the quasi-experiments. Panel E plots the number of quasi-experiments at each FICO score cutoff. Panel F plots the number of accounts within 50 FICO score points of these quasi-experiments for each FICO score cutoff.

TABLE I
QUASI-EXPERIMENT-LEVEL SUMMARY STATISTICS, AT ORIGINATION

	Average	Std. dev.		Average	Std. dev.
Credit limit on treated card ($)			Total balances across all credit card accounts ($)		
Pooled	5,265	2,045	Pooled	9,551	3,469
⩽660	2,561	674	⩽660	5,524	2,324
661–700	4,324	1,090	661–700	9,956	2,680
701–740	4,830	1,615	701–740	10,890	3,328
>740	6,941	1,623	>740	9,710	3,326
APR on treated card (%)			Credit limit across all credit card accounts ($)		
Pooled	15.38	3.70	Pooled	33,533	14,627
⩽660	19.63	5.43	⩽660	12,856	5,365
661–700	14.50	3.65	661–700	26,781	7,524
701–740	15.35	3.11	701–740	32,457	8,815
>740	14.70	2.52	>740	44,813	12,828
Number of credit card accounts			Number times 90+ DPD in last 24 months		
Pooled	11.00	2.93	Pooled	0.17	0.30
⩽660	7.13	1.18	⩽660	0.51	0.31
661–700	10.22	1.68	661–700	0.21	0.16
701–740	11.12	2.34	701–740	0.14	0.10
>740	12.63	2.92	>740	0.05	0.08
Age oldest account (months)			Number accounts currently 90+ DPD		
Pooled	190.1	29.1	Pooled	0.03	0.03
⩽660	162.0	26.3	⩽660	0.10	0.05
661–700	180.1	19.9	661–700	0.02	0.02
701–740	184.7	24.0	701–740	0.02	0.02
>740	208.6	25.7	>740	0.01	0.01

Notes. Table shows quasi-experiment-level summary statistics at the time of account origination, both pooled across our 743 quasi-experiments and split by FICO score groups. For each quasi-experiment, we first calculate the mean value for a given variable across all of the accounts within five FICO score points of the cutoff. We then show the means and standard deviations of these values across our 743 quasi-experiments. We follow the same procedure to obtain the means and standard deviations by FICO score group.

II.D. Summary Statistics

Table I presents summary statistics for the accounts in our sample of quasi-experiments at the time the accounts were originated. In particular, to characterize the accounts that are close to the discontinuities, we calculate the mean value for a given variable across all accounts within 5 FICO score points of the cutoff for each quasi-experiment. We then show the means and standard deviations of these values across the 743 quasi-experiments in our data. We also show summary statistics separately for each of the four FICO score groups that we use to explore heterogeneity in the data: $\leqslant$660, 661–700, 701–740, and >740. These ranges were chosen to split our quasi-experiments into roughly equal-sized groups, but we show in Online Appendix E that our conclusions are not sensitive to the exact grouping of experiments. In the entire sample, 28% of credit cards were issued to borrowers with FICO scores up to 660; 16% and 19% were issued to borrowers with FICO score ranges of 661–700 and 701–740, respectively; and 37% of credit cards were issued to borrowers with FICO scores above 740 (see Online Appendix Figure A.I).

At origination, accounts at the average quasi-experiment have a credit limit of $5,265 and an annual percentage rate (APR) of 15.4%. Average credit limits increase from $2,561 to $6,941 across FICO score groups, while average APRs decline from 19.6% to 14.7%. In the merged credit bureau data, we observe utilization on all credit cards held by the borrower. At the average quasi-experiment, account holders have 11 credit cards, with the oldest account being more than 15 years old. Across these credit cards, account holders have $9,551 in total balances and $33,533 in credit limits. Total balances are hump-shaped in FICO score, while total credit limits are monotonically increasing. In the credit bureau data, we also observe historical delinquencies and default. At the average quasi-experiment, account holders have been more than 90 days past due (90+ DPD) 0.17 times in the previous 24 months. This number declines from 0.51 to 0.05 across the FICO score groups.

III. Research Design

Our identification strategy exploits the credit limit quasi-experiments identified in Section II using a fuzzy regression discontinuity (RD) research design (see Lee and Lemieux 2010). In our setting, the "running variable" is the FICO score. The treatment effect of a $1 change in credit limit is determined by the jump in the outcome variable divided by the jump in the credit limit at the discontinuity.

We first describe how we recover the treatment effect for each quasi-experiment and then discuss how we aggregate across the 743 quasi-experiments in the data. For a given quasi-experiment, let x denote the FICO score, $\overline{x}$ the cutoff FICO level, cl the credit limit, and y the outcome variable of interest (e.g., borrowing volume). The fuzzy RD estimator, a local Wald estimator, is given by:

$$\tau = \frac{\lim_{x \downarrow \overline{x}} E[y|x] - \lim_{x \uparrow \overline{x}} E[y|x]}{\lim_{x \downarrow \overline{x}} E[cl|x] - \lim_{x \uparrow \overline{x}} E[cl|x]}. \tag{2}$$

The denominator is always nonzero because of the known discontinuity in the credit supply function at $\overline{x}$. The parameter τ identifies the local average treatment effect (LATE) of extending more credit to people with FICO scores in the vicinity of $\overline{x}$. We estimate the limits in equation (2) using locally linear regressions. Specifically, let i denote a credit card account and $\mathbb{I}$ the set of accounts within 50 FICO score points on either side of $\overline{x}$. For each quasi-experiment, we fit a locally linear regression that solves the following objective function separately for observations i on either side of the cutoff, $d \in \{l, h\}$, for the variables, $\tilde{y} \in \{cl, y\}$:

$$\min_{\alpha_{\tilde{y},d}, \beta_{\tilde{y},d}} \sum_{i \in \mathbb{I}} [\tilde{y}_i - \alpha_{\tilde{y},d} - \beta_{\tilde{y},d}(x_i - \overline{x})]^2 \mathbf{1}_{(|x_i - \overline{x}| < b)} \quad \text{for } d \in \{l, h\}. \tag{3}$$

In our baseline results we use the optimal bandwidth b from Imbens and Kalyanaraman (2011).[11] For those quasi-experiments where we identify an additional jump in credit limits within our 50-FICO-score-point window, we include an indicator variable in equation (3) that is equal to 1 for all FICO scores above this second cutoff; Online Appendix C shows that this approach allows us to recover unbiased estimates of the actual treatment effect. Given these estimates, the LATE is given by:

$$\tau = \frac{\hat{\alpha}_{y,h} - \hat{\alpha}_{y,l}}{\hat{\alpha}_{cl,h} - \hat{\alpha}_{cl,l}}. \tag{4}$$

III.A. Heterogeneity by FICO Score

Our objective is to estimate the heterogeneity in treatment effects by FICO score (see Einav et al. 2015, for a discussion of estimating treatment effect heterogeneity across experiments). Let j indicate quasi-experiments, let τ_j be the LATE for quasi-experiment j estimated using equation (4), and let $FICO_k$, $k = 1, \ldots, 4$ be indicator variables that take on a value of 1

11. Our results are robust to using different specifications. For example, we obtain similar estimates when we run second-order local polynomial regressions with a triangular kernel.

when the FICO score of the discontinuity for quasi-experiment j falls into one of our FICO score groups ($\leqslant$660, 661–700, 701–740, $>$740). We recover heterogeneity in treatment effects by regressing τ_j on the FICO score group dummies and controls:

$$\tau_j = \left(\sum_{k=1}^{4} \beta_k FICO_{j,k} \right) + X_j' \delta_X + \epsilon_j. \tag{5}$$

In our baseline specification, X_j includes fully interacted controls for origination quarter, bank, and a "zero initial APR" dummy that captures whether the account has a promotional period during which no interest is charged, and additively separable fully interacted loan channel by "zero initial APR" fixed effects.[12] The β_k are the coefficients of interest and capture the mean effect for accounts in FICO score group k, conditional on the other covariates. In Online Appendix Section E, we examine the relationship between our LATEs and FICO scores using nonparametric binned scatter plots, and show our results are robust to the choice of FICO score groups in the baseline analysis.

We construct confidence intervals by bootstrapping over the 743 quasi-experiments. In particular, we draw 500 samples of local average treatment effects with replacement, and estimate the coefficients of interest, β_k, in each sample. Our reported 95% confidence intervals give the range from the 2.5th percentile of estimates to the 97.5th percentile of estimates. Conceptually, we think of the local average treatment effects τ_j as "data" that are drawn from a population distribution of treatment effects. We are interested in the average treatment effect in the population for a given FICO score group. Our confidence intervals can be interpreted as

12. Following Wooldridge (2003), we give each of the underlying quasi-experiments equal weight in the regression specification. As he describes, in a two-step estimation procedure it is efficient to weight the second stage observations differently if there is a small number of observations in each of the underlying groups (quasi-experiments in our context). The reason is that the small number of observations will create measurement error in these estimates that should be accounted for by the efficient estimator. However, if the number of underlying observations is large, then this estimation error is likely to be second order, and it is efficient to weight the observations equally. Since we have a very large number of observations in each of our quasi-experiments (on average 11,400 per quasi-experiment), we follow Wooldridge's suggestion and weight each observation equally. To deal with outliers in the estimated treatment effects from equation (4), we Winsorize the values of τ_j at the 2.5% level.

measuring the precision of our sample average treatment effects for the population averages.

IV. Validity of Research Design

The validity of our research design rests on two assumptions: First, we require a discontinuous change in credit limits at the FICO score cutoffs. Second, other factors that could affect outcomes must trend smoothly through these thresholds. Below we present evidence in support of these assumptions.

IV.A. First-Stage Effect on Credit Limits

We first verify that there is a discontinuous change in credit limits at our quasi-experiments. Panel A of Figure III shows average credit limits at origination within 50 FICO score points of the quasi-experiments together with a local linear regression line estimated separately on each side of the cutoff. Initial credit limits are smoothly increasing except at the FICO score cutoff, where they jump discontinuously by \$1,472. The magnitude of this increase is significant relative to an average credit limit of \$5,265 around the cutoff (see Table II). Panel A of Figure IV shows the distribution of first-stage effects from RD specifications estimated separately for each of the 743 quasi-experiments in our data. These correspond to the denominator of equation (4). The first-stage estimates are fairly similar in size, with an interquartile range of \$677 to \$1,755 and a standard deviation of \$796.[13]

Panel B of Figure IV examines the persistence of the jump in the initial credit limit. It shows the RD estimate of the effect of a \$1 increase in initial credit limits on credit limits at different time horizons following account origination. The initial effect is highly persistent and very similar across FICO score groups, with a \$1 higher initial credit limit raising subsequent credit limits by \$0.85 to \$0.93 at 36 months after origination. Table III shows the corresponding regression estimates.

In the analysis that follows, we estimate the effect of a change in *initial* credit limits on outcomes at different time horizons. A natural question is whether it would be preferable to scale our estimates by the change in contemporaneous credit limits instead of the initial increase. We think the initial increase in credit limits

13. For all RD graphs we control for additional discontinuous jumps in credit limits as discussed in Section III.

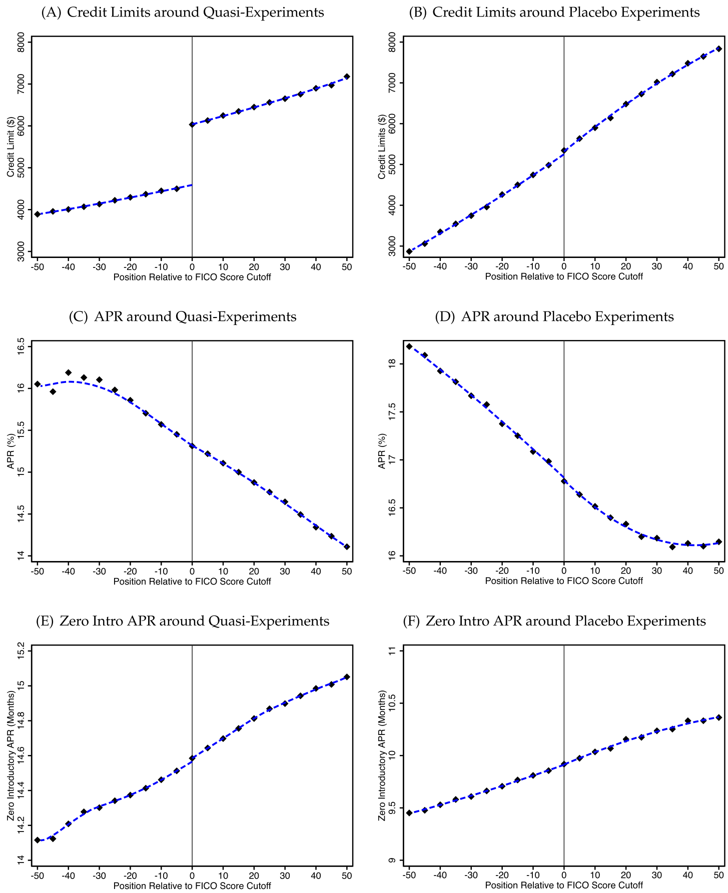

FIGURE III

Credit Limits and Cost of Credit Around Credit Limit Quasi-Experiments and Placebo Experiments

Figure plots average credit limits (Panels A and B), average APR (Panels C and D), and average number of months with zero introductory APR (Panels E and F; limited to originations with zero introductory APR). The left column plots these outcomes around our 743 pooled quasi-experiments. We also control for other quasi-experiments within 50 FICO score points in the same origination group. The right column plots the same outcomes around the same FICO score cutoffs but for "placebo experiments" originated in the same month as the quasi-experiments in the left column but for origination groups with no quasi-experiments at that FICO score. The horizontal axis shows FICO score at origination, centered at the FICO score cutoff. Scatter plots show means of outcomes for 5-point FICO score buckets. Dashed lines show predicted values from locally linear regressions estimated separately on either side of the cutoff using the Imbens and Kalyanaraman (2011) optimal bandwidth.

TABLE II
VALIDITY OF RESEARCH DESIGN: DISCONTINUOUS INCREASE AT FICO CUTOFF

	Distribution of jump across quasi-experiments			
	Average	Median	Standard devation	Baseline
Credit limit	1,472	1,282	796	5,265
APR (%)	0.017	−0.005	0.388	15.38
Months to rate change	0.027	0.016	0.800	13.37
Number of credit card accounts	0.060	0.031	0.713	11.00
Total credit limit—all accounts	151	28	2,791	33,533
Age oldest account (months)	1.034	0.378	11.072	190.11
Number times 90+ DPD—last 24 months	0.010	0.002	0.111	0.169
Number accounts 90+ DPD—at origination	0.001	0.001	0.017	0.026
Number accounts 90+ DPD—ever	0.004	0.003	0.095	0.245
Number of accounts originated	10.21	4.38	47.61	580.12

Notes. Table shows the reduced-form discontinuous increase ("jump") in credit limits and outcome variables at the FICO score cutoff (see equation (4)). All variables are measured at account origination, allowing us to inspect the validity of the research design. We present the average, median, and standard deviation of this jump across our 743 quasi-experiments. We also present the average value of the variable at the cutoff ("baseline"), allowing us to judge the economic significance of any differences.

is the appropriate denominator because subsequent credit limits are endogenously determined by household responses to the initial increase. We discuss this issue further in Section VI.D.

IV.B. Other Characteristics Trend Smoothly through Cutoffs

For our research design to be valid, the second requirement is that all other factors that could affect the outcomes of interest trend smoothly through the FICO score cutoff. These include contract terms, such as the interest rate (Assumption 1), characteristics of borrowers (Assumption 2), and the density of new account originations (Assumption 3). Because we have 743 quasi-experiments, graphically assessing the validity of our identifying assumptions for each experiment is not practical. Therefore, we show results graphically that pool across all of the quasi-experiments in the data, estimating a single pooled treatment effect and pooled locally linear regression line. In Table II, we

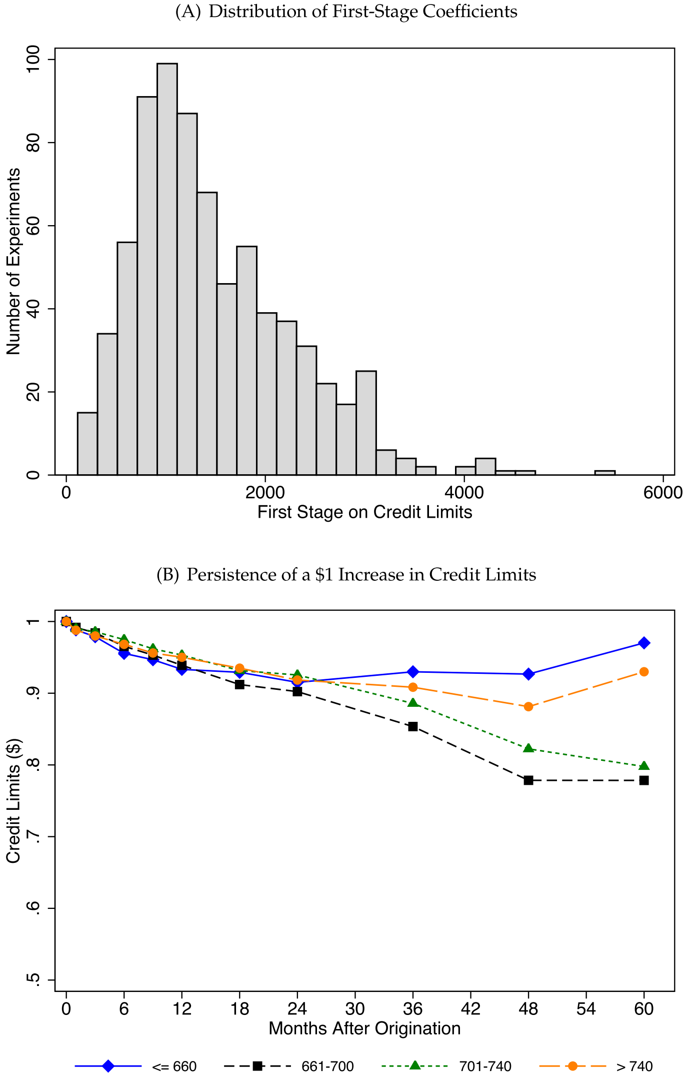

FIGURE IV

Effect of FICO Score Cutoff on Credit Limits

Panel A shows the distribution of credit limit increases at the FICO score cutoffs across our 743 credit limit quasi-experiments. Panel B shows regression discontinuity estimates of the effect of a $1 increase in initial credit limits on credit limits at different time horizons after account origination. Estimates are shown for FICO score groups, defined at account origination. The corresponding estimates are shown in Table III.

TABLE III
PERSISTENCE OF CREDIT LIMIT EFFECT

	Months after account origination				
	12	24	36	48	60
FICO					
⩽660	0.93	0.92	0.93	0.93	0.97
	[0.91, 0.96]	[0.87, 0.96]	[0.86, 0.99]	[0.83, 1.04]	[0.79, 1.19]
661–700	0.94	0.90	0.85	0.78	0.78
	[0.90, 0.94]	[0.88, 0.93]	[0.86, 0.90]	[0.71, 0.85]	[0.67, 0.90]
701–740	0.95	0.93	0.89	0.82	0.80
	[0.94, 0.97]	[0.90, 0.94]	[0.86, 0.90]	[0.77, 0.87]	[0.67, 0.90]
>740	0.95	0.92	0.91	0.88	0.93
	[0.94, 0.96]	[0.90, 0.94]	[0.87, 0.93]	[0.81, 0.94]	[0.83, 1.11]

Notes. Table shows regression discontinuity estimates of the effect of a $1 increase in initial credit limits on credit limits at different time horizons after account origination and by FICO score group, defined at account origination. 95% confidence intervals are constructed by bootstrapping over quasi-experiments, and are presented in square brackets.

present summary statistics on the distribution of these treatment effects across the 743 individual quasi-experiments.

ASSUMPTION 1. Credit limits are the only contract characteristic that changes at the cutoff.

The interpretation of our results requires that credit limits are the only contract characteristic that changes discontinuously at the FICO score cutoffs. For example, if the cost of credit also changed at our credit limit quasi-experiments, an increase in borrowing around the cutoff might not only result from additional access to credit, but could also be explained by lower borrowing costs.

Panel C of Figure III shows the average APR around our quasi-experiments. APR is defined as the initial interest rate for accounts with a positive interest rate at origination, and the "go-to" rate for accounts which have a zero introductory APR.[14] As one would expect, the APR is declining in the FICO score. Importantly, there is no discontinuous change in the APR around our credit limit quasi-experiments. This is consistent with the standard practice of using different models to price credit (set APRs)

14. The results look identical when we remove quasi-experiments for accounts with an initial APR of zero.

and manage exposure to risk (set credit limits).[15] Table II shows that, for the average (median) experiment, the APR increases by 1.7 basis points (declines by 0.5 basis point) at the FICO score cutoff; these changes are economically tiny relative to an average APR of 15.4%. Panel E of Figure III shows the length of the zero introductory APR period for the 248 quasi-experiments with a zero introductory APR. The length of the introductory period is increasing in FICO score, but there is no jump at the credit limit cutoff.[16]

ASSUMPTION 2. All other borrower characteristics trend smoothly through the cutoff.

We next examine whether borrowers on either side of the FICO score cutoff looked similar on observable characteristics in the credit bureau data when the credit card was originated. Panels A and B of Figure V show the total number of credit cards and the total credit limit on those credit cards, respectively. Both are increasing in the FICO score, and there is no discontinuity around the cutoff. Panel C shows the age of the oldest credit card account for consumers, capturing the length of the observed credit history. We also plot the number of payments for each consumer that were 90 or more days past due (90+ DPD), both over the entire credit history of the borrower (Panel D), as well as in the 24 months prior to origination (Panel E). These figures, and the information in Table II, show that there are no discontinuous changes around the cutoff in any of these borrower characteristics.[17]

15. We initially identified a few instances in which the APR also changed discontinuously at the same cutoff where we detected a discontinuous change in credit limits. These quasi-experiments were dropped in our process of arriving at the sample of 743 quasi-experiments that are the focus of our empirical analysis.

16. A related concern is that while contract characteristics other than credit limits are not changing at the cutoff for the bank with the credit limit quasi-experiment, they might be changing at other banks. If this were the case, the same borrower might also be experiencing discontinuous changes in contract terms on his other credit cards, which would complicate the interpretation of our estimates. To test whether this is the case, for every FICO score where we observe at least one bank discontinuously changing the credit limit for one card, we define a "placebo experiment" as all other cards that are originated around the same FICO score at banks without an identified credit limit quasi-experiment. The right column of Figure III shows average contract characteristics at all placebo experiments. All characteristics trend smoothly through the FICO score cutoff at banks with no quasi-experiments.

17. Online Appendix Figure A.II shows similar graphs for six additional borrower characteristics, all of which trend smoothly through the FICO score cutoff.

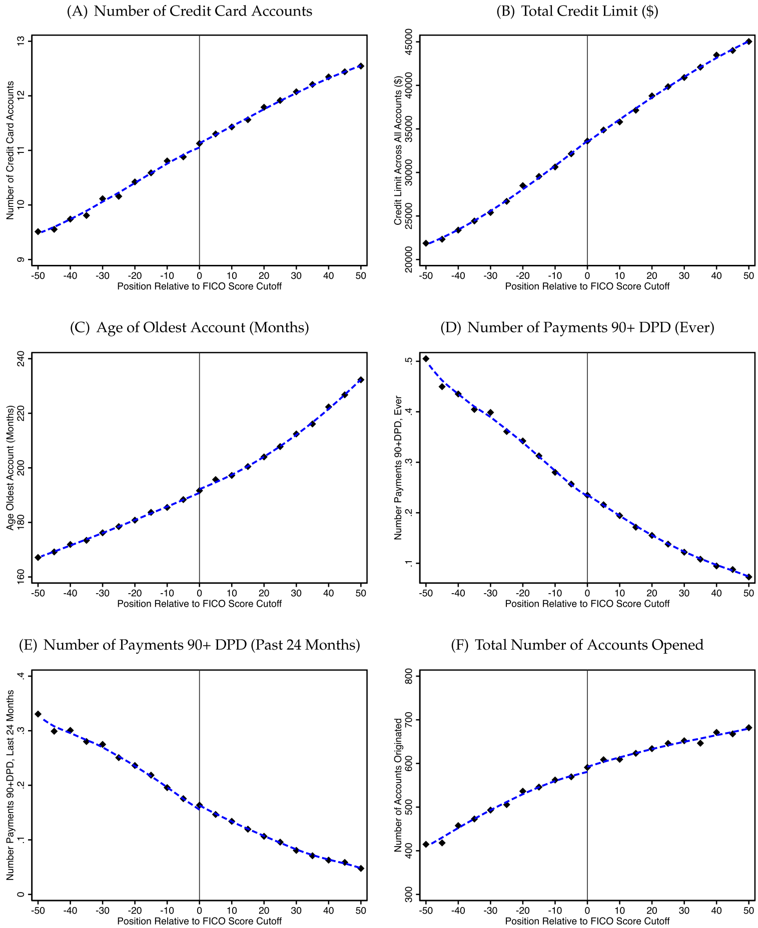

FIGURE V

Initial Borrower Characteristics around Credit Limit Quasi-Experiments

Figure plots average borrower characteristics around our 743 pooled credit limit quasi-experiments. The horizontal axis shows FICO score at origination, centered at the FICO score cutoff. The vertical axis shows the number of credit card accounts (Panel A), total credit limit across all credit card accounts (Panel B), age of the oldest account (Panel C), number of payments ever 90+ days past due (Panel D), number of payments 90+ days past due in last 24 months (Panel E), and the total number of accounts opened in the origination group where we observe the credit limit quasi-experiment (Panel F). All borrower characteristics are as reported to the credit bureau at account origination. Scatter plots show means of outcomes for 5-point FICO score buckets. Dashed lines show predicted values from locally linear regressions estimated separately on either side of the cutoff using the Imbens and Kalyanaraman (2011) optimal bandwidth.

ASSUMPTION 3. The number of originated accounts trends smoothly through the cutoff.

Panel F of Figure V shows that the number of originated accounts trends smoothly through the credit score cutoffs. This addresses a number of potential concerns with the validity of our research design.

First, regression discontinuity designs are invalid if individuals are able to precisely manipulate the forcing variable. In our setting, the lack of strategic manipulation is unsurprising. Since the banks' credit supply functions are unknown, individuals with FICO scores just below a threshold are unaware that marginally increasing their FICO scores would lead to a significant increase in their credit limits. Moreover, even if consumers knew of the location of these thresholds, since the FICO score function is proprietary, it would be very difficult for consumers to manipulate their FICO scores in a precise manner.

A second concern in our setting is that banks might use the FICO score cutoff to make extensive margin lending decisions. For example, if banks relaxed some other constraint once individuals crossed a FICO score threshold, more accounts would be originated for households with higher FICO scores, but households on either side of the FICO score cutoff would differ along that other dimension. In Figure III, we already documented that there are no changes in observable characteristics around the FICO score cutoffs. The smooth trend in the number of accounts further indicates that banks do not select borrowers on an unobservable dimension as well.

Finally, we would observe fewer accounts to the left of the threshold if there was a "demand response," whereby consumers were more likely to turn down credit card offers with lower credit limits. However, in this market, consumers do not know their exact credit limits when they apply for a credit card and only learn of their credit limits when they have been approved and receive a credit card in the mail. Since consumers have already paid the sunk cost of applying, it is not surprising that consumers with lower credit limits do not immediately cancel their cards, which would generate a discontinuity in the number of accounts.

V. BORROWING AND SPENDING

Having established the validity of our research design, we turn to estimating the causal impact of an increase in credit limits on borrowing and spending, focusing on how these effects vary across the FICO score distribution.

V.A. Average Borrowing and Spending

We start by presenting basic summary statistics on credit card utilization. The left column of Table IV shows average borrowing by FICO score group at different time horizons after account origination. To characterize the credit cards that identify the causal estimates, we again restrict the sample to accounts within 5 FICO score points of a credit limit quasi-experiment.

Average daily balances (ADB) are the industry standard measure of borrowing, and are defined as the arithmetic mean of end-of-day balances over the billing cycle. If interest charges are assessed, they are calculated as a percentage of ADB. We find that ADB are hump-shaped in FICO score. At 12 months after origination, ADB increase from \$1,260 for the lowest FICO score group ($\leqslant$660), to more than \$2,150 for the middle FICO score groups, before falling to \$2,101 for the highest FICO score group ($>$740). ADB are fairly flat over time for the lowest FICO score group but drop more sharply for accounts with higher FICO scores.

Accounts can have positive ADB even though no interest charges are incurred, for example during periods with zero introductory interest rates. To measure borrowing for which interest charges are assessed, we construct a variable called *interest-bearing debt*. This measure is equal to the ADB if the account holder is assessed positive interest charges in that billing period and zero if no interest charges are assessed. At 12 months after origination, interest-bearing debt is approximately half as large as ADB, mainly due to zero introductory rate periods, and is relatively smaller for higher FICO score groups. At longer time horizons, ADB and interest-bearing debt are very similar, with interest-bearing debt approximately 8% smaller than ADB across FICO score groups and years.

One interesting question is whether the relatively high average measures of interest-bearing debt, in particular for the high FICO score groups, are the result of a few accounts with large balances, or whether these balances are more evenly distributed across the sample. To address this question, we measure the fraction of accounts that had positive interest-bearing debt at least once over a given period. The summary statistics on the cumulative probability of interest-bearing debt show that, at 24 months

TABLE IV

QUASI-EXPERIMENT-LEVEL SUMMARY STATISTICS, POSTORIGINATION

	FICO score group			
	⩽660	661–700	701–740	>740
Credit limit ($)				
After 12 months	2,617	4,370	4,964	6,980
After 24 months	2,414	4,306	4,946	7,071
After 36 months	2,301	4,622	5,047	7,005
After 48 months	2,252	4,525	4,985	6,944
After 60 months	2,290	4,449	4,601	6,839
ADB ($)				
After 12 months	1,260	2,160	2,197	2,101
After 24 months	1,065	1,794	1,719	1,524
After 36 months	1,164	1,734	1,481	1,343
After 48 months	1,079	1,501	1,260	1,064
After 60 months	1,050	1,465	1,097	1,084
Average interest bearing debt ($)				
After 12 months	864	903	811	672
After 24 months	1,040	1,676	1,557	1,294
After 36 months	1,068	1,615	1,344	1,135
After 48 months	1,044	1,416	1,144	924
After 60 months	1,020	1,388	1,001	941

	FICO score group			
	⩽660	661–700	701–740	>740
Cumulative purchase volume ($)				
After 12 months	2,212	2,579	2,514	2,943
After 24 months	2,447	3,956	3,791	4,374
After 36 months	3,240	5,023	4,253	4,521
After 48 months	3,741	5,154	4,919	4,845
After 60 months	4,524	5,598	5,121	5,626
Cumulative total costs ($)				
After 12 months	122	172	169	147
After 24 months	281	451	433	304
After 36 months	459	710	644	395
After 48 months	588	845	808	488
After 60 months	712	962	901	583
Cumulative chargeoffs ($)				
After 12 months	47	67	61	35
After 24 months	178	259	245	124
After 36 months	306	443	403	190
After 48 months	403	552	524	261
After 60 months	483	634	602	322

	FICO score group			
	⩽660	661–700	701–740	>740
Cumulative cost of funds ($)				
After 12 months	14	16	16	15
After 24 months	23	29	28	25
After 36 months	28	38	36	31
After 48 months	31	43	41	34
After 60 months	33	46	44	36
Cumulative total revenue ($)				
After 12 months	233	192	181	175
After 24 months	474	503	439	347
After 36 months	740	793	663	449
After 48 months	953	971	863	563
After 60 months	1,148	1,126	965	669
Cumulative interest charge revenue ($)				
After 12 months	106	61	52	42
After 24 months	297	295	243	159
After 36 months	484	520	420	243
After 48 months	625	669	578	340
After 60 months	760	794	657	429

TABLE IV

(CONTINUED)

	FICO score group					FICO score group					FICO score group			
	⩽660	661–700	701–740	>740		⩽660	661–700	701–740	>740		⩽660	661–700	701–740	>740
Cumulative prob positive interest-bearing debt (%)					Cumulative prob 60+ DPD (%)					Cumulative fee revenue ($)				
After 12 months	58.4	36.1	31.6	26.9	After 12 months	6.4	4.1	3.6	1.6	After 12 months	73	79	79	74
After 24 months	75.4	73.0	64.9	50.3	After 24 months	12.0	9.3	8.2	3.8	After 24 months	129	129	121	101
After 36 months	84.0	79.4	72.3	61.6	After 36 months	15.1	12.2	10.9	5.2	After 36 months	192	173	157	116
After 48 months	87.4	84.0	78.1	69.8	After 48 months	16.5	13.6	12.2	5.9	After 48 months	254	199	187	126
After 60 months	90.1	86.3	81.3	75.2	After 60 months	17.2	14.4	12.9	6.2	After 60 months	364	310	211	101
Total balances across all cards ($)					Cumulative prob 90+ DPD (%)					Cumulative profits ($)				
After 12 months	6,155	10,546	11,411	10,528	After 12 months	4.8	3.3	2.9	1.3	After 12 months	111	21	12	30
After 24 months	5,919	10,521	11,307	10,703	After 24 months	10.2	8.1	7.2	3.2	After 24 months	194	56	9	46
After 36 months	6,387	10,716	11,702	11,267	After 36 months	13.2	10.9	9.7	4.5	After 36 months	281	91	23	59
After 48 months	6,698	10,437	11,665	11,137	After 48 months	14.5	12.2	10.9	5.1	After 48 months	365	126	55	75
After 60 months	7,566	10,591	11,972	12,490	After 60 months	15.2	12.9	11.5	5.4	After 60 months	436	164	63	87

Notes. Table shows quasi-experiment-level summary statistics at different horizons after account origination by FICO score group. For each quasi-experiment, we calculate the mean value for a given variable across all of the accounts within 5 FICO score points of the cutoff. We then show the means and standard deviations of these values across the available quasi-experiments. Since later quasi-experiments are observed for shorter periods of time only, the set of experiments contributing to the averages across different horizons is not constant. FICO score groups are defined at account origination.

after origination, approximately three-quarters of accounts have had positive interest-bearing debt in at least one billing cycle. Even in the highest FICO score group, more than half of accounts were charged interest at least once. This suggests that our analysis considers a sample of credit card holders that regularly use their cards to borrow, and might therefore be responsive in their borrowing behavior to expansions in their credit limits.

Total balances across all credit cards are between $10,400 and $12,500 for borrowers with FICO scores above 660, and do not vary substantially with the time since the treated card was originated; for accounts with FICO scores below 660, total balances are about $6,500.[18] The top panel of the middle column of Table IV shows summary statistics on cumulative purchase volume. Despite large differences in credit limits by FICO score, purchase volumes over the first 12 months since origination are fairly similar, ranging from $2,212 to $2,943 across FICO score groups. Higher FICO score borrowers spend somewhat more on their cards over longer time horizons, but even at 60 months after origination, cumulative purchase volumes range between $4,524 and $5,626 across FICO score groups.

V.B. Marginal Propensity to Borrow (MPB)

We next exploit our credit limit quasi-experiments to estimate the marginal propensity to borrow out of an increase in credit limits. We examine effects on four outcome variables: (i) ADB on the treated credit card, (ii) interest-bearing debt on the treated card, (iii) total balances across all cards, and (iv) cumulative purchase volume on the treated card. Each of these outcome variables highlights different aspects of consumer borrowing and spending. While, in principle, our findings could differ across these outcomes, the effects we estimate are actually very similar.

18. In the CCM data, we can construct clean measures of interest-bearing debt. In the credit bureau data, we observe the account balances at the point the banks report them to the credit bureau. These account balances will include interest-bearing debt, but can also include balances incurred during the credit card cycle, but repaid at the end of the cycle, and therefore not considered debt. This explains why the level of credit bureau account balances is higher than the amount of total credit card borrowing that households report, for example, in the Survey of Consumer Finances. We discuss below why this does not affect our interpretation of a marginal increase in total balances as a marginal increase in total credit card borrowing.

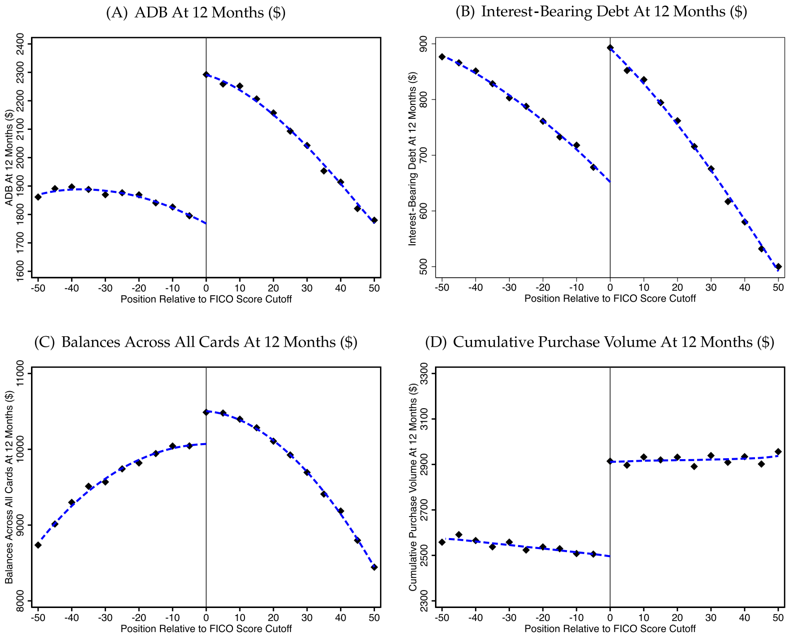

FIGURE VI

Borrowing and Spending Around Credit Limit Quasi-Experiments

Figure shows changes in borrowing quantities after 12 months around our 743 pooled credit limit quasi-experiments; these plots are constructed as described in Figure III. Panel A shows average daily balances on the treated credit card. Panel B shows interest-bearing debt on the treated credit card. Panel C shows total balances aggregated across all credit cards held by the account holder. Panel D shows cumulative purchase volume on the treated credit card.

1. Average Daily Balances. We first examine the effects on ADB on the treated credit card. Panel A of Figure VI shows the effect on ADB at 12 months after account origination in the pooled sample of all quasi-experiments. ADB increase sharply at the credit limit discontinuity but otherwise trend smoothly in FICO score. Panel A of Figure VII decomposes this average effect, showing the impact of a $1 increase in credit limits on ADB at different time horizons after account origination and for different FICO score groups. Panel A of Table V shows the corresponding RD estimates and confidence intervals. Higher credit limits generate a sharp increase in ADB on the treated credit card for all FICO score groups. Within 12 months, the lowest FICO score group

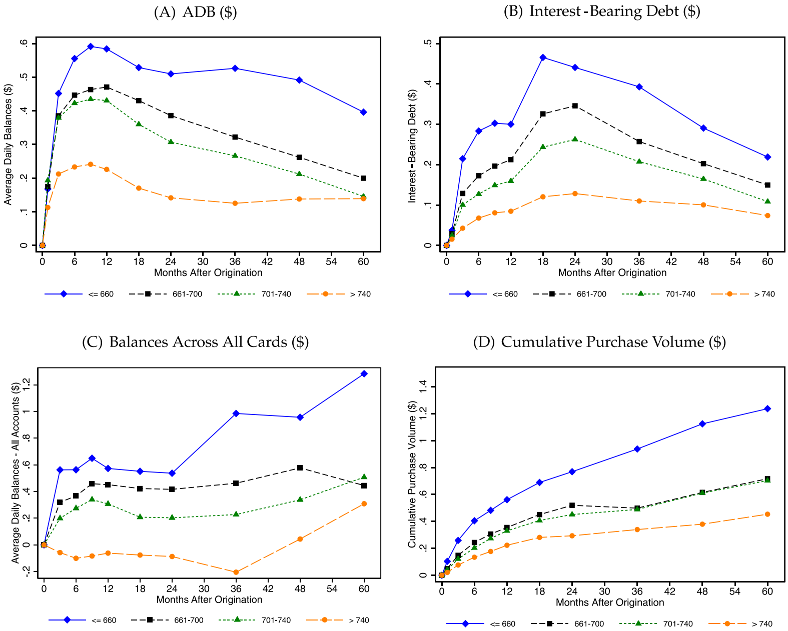

FIGURE VII

Marginal Propensity to Borrow

Figure shows the effects of credit limits on borrowing and spending. We show regression discontinuity estimates of the effect of a $1 increase in credit limits for different FICO score groups and different time horizons after account origination. FICO score groups are determined by FICO score at account origination. Panel A shows effects on average daily balances on the treated credit card. Panel B shows effects on interest-bearing debt on the treated card. Panel C shows effects on total balances aggregated across all credit cards held by the account holder. Panel D shows effects on cumulative purchase volume on the treated card. The corresponding estimates are shown in Table V.

raises ADB by 58 cents for each additional dollar in credit limits. The effect is decreasing in FICO score, but even borrowers in the highest FICO score group increase their ADB by 23 cents for each additional dollar in credit limits. Panel A of Figure VII also reveals interesting patterns in borrowing effects over time. For the lowest FICO score group, the initial increase in ADB is quite persistent, declining by less than 20% between the first and fourth year following account origination. This is consistent with these low FICO score borrowers using the increase in credit to fund immediate spending and then "revolving" their debt in future periods. For the higher FICO score groups, the MPB drops

TABLE V

MARGINAL PROPENSITY TO BORROW

	Months after account origination				
	12	24	36	48	60
Panel A: Average daily balance					
FICO					
⩽660	0.58	0.51	0.53	0.49	0.40
	[0.54, 0.63]	[0.46, 0.57]	[0.46, 0.59]	[0.39, 0.58]	[0.32, 0.48]
661–700	0.47	0.39	0.32	0.26	0.20
	[0.44, 0.49]	[0.35, 0.41]	[0.28, 0.35]	[0.22, 0.30]	[0.15, 0.25]
701–740	0.43	0.31	0.27	0.21	0.15
	[0.40, 0.45]	[0.28, 0.33]	[0.23, 0.29]	[0.18, 0.25]	[0.10, 0.20]
>740	0.23	0.14	0.13	0.14	0.14
	[0.20, 0.25]	[0.11, 0.17]	[0.09, 0.16]	[0.09, 0.18]	[0.08, 0.20]
Panel B: Interest-bearing debt					
FICO					
⩽660	0.30	0.46	0.40	0.35	0.33
	[0.26, 0.35]	[0.42, 0.51]	[0.36, 0.45]	[0.31, 0.40]	[0.28, 0.38]
661–700	0.21	0.34	0.30	0.28	0.27
	[0.19, 0.23]	[0.32, 0.37]	[0.27, 0.32]	[0.25, 0.31]	[0.24, 0.30]
701–740	0.16	0.27	0.23	0.21	0.19
	[0.14, 0.18]	[0.24, 0.29]	[0.20, 0.26]	[0.18, 0.23]	[0.17, 0.22]
>740	0.08	0.13	0.12	0.12	0.11
	[0.07, 0.10]	[0.10, 0.15]	[0.09, 0.14]	[0.09, 0.14]	[0.09, 0.13]
Panel C: Total balance across all cards					
FICO					
⩽660	0.59	0.54	1.00	0.96	1.27
	[0.34, 0.84]	[0.18, 0.94]	[0.51, 1.48]	[0.12, 1.97]	[−0.16, 2.51]
661–700	0.46	0.42	0.48	0.59	0.43
	[0.31, 0.59]	[0.26, 0.58]	[0.26, 0.70]	[0.13, 0.97]	[−0.42, 1.11]
701–740	0.32	0.21	0.24	0.35	0.49
	[0.16, 0.47]	[0.03, 0.37]	[0.03, 0.44]	[0.00, 0.65]	[−0.52, 1.30]
>740	−0.05	−0.08	−0.19	0.05	0.29
	[−0.15, 0.08]	[−0.26, 0.10]	[−0.47, 0.08]	[−0.42, 0.39]	[−0.47, 1.03]
Panel D: Cumulative purchase volume					
FICO					
⩽660	0.56	0.77	0.94	1.12	1.24
	[0.48, 0.67]	[0.61, 0.94]	[0.68, 1.20]	[0.70, 1.51]	[0.75, 1.69]
661–700	0.35	0.52	0.50	0.62	0.72
	[0.31, 0.40]	[0.44, 0.59]	[0.38, 0.58]	[0.47, 0.75]	[0.52, 0.91]
701–740	0.33	0.45	0.49	0.61	0.70
	[0.29, 0.38]	[0.38, 0.52]	[0.38, 0.59]	[0.45, 0.78]	[0.50, 0.95]
>740	0.22	0.29	0.34	0.38	0.45
	[0.19, 0.26]	[0.23, 0.36]	[0.26, 0.42]	[0.25, 0.51]	[0.21, 0.71]

Notes. Table shows regression discontinuity estimates of the effect of a $1 increase in credit limits on borrowing and spending. Panel A shows effects on average daily balances on the treated credit card. Panel B shows effects on total interest-bearing debt on the treated credit card. Panel C shows effects on total balances across all credit cards held by the account holder. Panel D shows effects on cumulative purchase volume on the treated credit card. Columns show effects at different time horizons after account origination. Within each panel, rows show effects for different FICO score groups, defined at account origination. 95% confidence intervals are constructed by bootstrapping over quasi-experiments, and are presented in square brackets.

more rapidly over time. This is consistent with these high FICO score borrowers making large purchases during zero introductory rate periods and then repaying this debt relatively quickly as the introductory rate period expires.

2. Interest-Bearing Debt. We next examine the effect on interest-bearing debt on the treated credit card, which excludes borrowing during zero introductory rate periods. Panel B of Figures VI and VII plots the effects on interest-bearing debt. Panel B of Table V shows the corresponding RD estimates and confidence intervals. The response of interest-bearing debt over the first few months is smaller than the response of ADB. At 12 to 18 months after origination, we observe a sharp increase in the marginal effect on interest-bearing debt, as balances previously held under a zero introductory rate now shift into interest-bearing debt. At time horizons of 24 months and greater, the effects on ADB and interest-bearing debt are very similar. For the remainder of the article, we use the term "marginal propensity to borrow" (MPB) on the treated card to refer to the effect of a $1 increase in credit limits on ADB. The choice of ADB rather than interest-bearing debt is largely inconsequential, since at most time horizons the estimated effects on these outcomes are economically identical.

In Online Appendix D, we decompose the effect of higher credit limits on interest-bearing debt into an extensive-margin effect (encouraging credit card holders who did not previously borrow to start borrowing) and an intensive-margin effect (encouraging credit card holders that already borrow to borrow more). While there is a small extensive-margin effect, the vast majority of the effect occurs on the intensive margin.[19]

3. Balances Across All Cards. We next examine the effects on account balances across all credit cards held by the consumer, using the merged credit bureau data. The reason to look at this broader measure of borrowing is to account for balance shifting across credit cards. For example, a consumer who receives a higher credit limit on a new credit card might shift borrowing to this card to take advantage of a low introductory interest rate. This would result in an increase in borrowing on the treated credit card but no increase in overall balances. The response of total borrowing

19. Positive extensive-margin effects are consistent with a model of lumpy expenditure, in which some consumers borrow only if they have a high enough credit limit to fund the entire purchase amount (e.g., for a television or automobile down payment).

across all credit cards is the primary object of interest for policy makers wanting to stimulate household borrowing and spending. Panel C of Figures VI and VII plots the effects on total balances across all cards. Panel C of Table V shows the corresponding RD estimates and confidence intervals.

For all but the highest FICO score group, the marginal increase in ADB on the treated credit card corresponds to an increase in overall borrowing. Indeed, we cannot reject the null hypothesis that the increase in ADB translates one-for-one into an increase in total balances. The one exception is the highest FICO score group for which we find evidence of significant balance shifting. At one year after origination, these consumers exhibit a 23% MPB on the treated card but essentially zero MPB across all their accounts (the statistically insignificant point estimate is −5%). This is not because high FICO score consumers do not borrow. Indeed, consumers with high FICO scores have sizable average interest-bearing debt on the treated credit card (see Table IV). Instead, the high FICO score group has on average $44,813 in credit limits across all of their credit cards (see Table I), indicating that these households are not credit constrained on the margin.[20]

4. Purchase Volume. The increase in borrowing on both the treated credit card and across all credit cards suggests that higher credit limits raise overall spending. However, at least in the short run, consumers could increase their borrowing volumes by paying off their debt at a slower rate without spending more. To examine whether the increase in borrowing is indeed due to higher spending rather than slower debt repayment, Panel D of Figures VI and

20. The fact that we observe total credit card balances and not total ADB in the credit bureau data (see note 18) does not affect our interpretation of the marginal increase in balances as a marginal increase in borrowing. In particular, one might worry that the response of balances in the credit bureau data picks up an increase in credit card spending, without an increase in total credit card borrowing. Such a response, which would not generate a stimulative effect on the economy, could result if people switched their method of payment from cash to credit cards. However, in our setting this is unlikely to be a concern. Among high FICO score borrowers, we observe no treatment effect on balances across all cards, suggesting that neither spending nor borrowing was affected by the increase in credit limits. For lower FICO score borrowers, the increase in balances across all credit cards maps one-for-one into the observed increase in ADB and interest-bearing debt on the treated credit card, again showing that we are not just picking up a shift of payment methods from cash to credit cards. This confirms that the change in total balances across all cards picks up the change in total borrowing across these cards.

VII shows the effect of higher credit limits on cumulative purchase volume on the treated card. Panel D of Table V shows the corresponding RD estimates.

Over the first year, the higher borrowing levels on the treated card are almost perfectly explained by an increase in purchase volume. For the lowest FICO score group, a $1 increase in credit limits raises cumulative purchase volume over the first year by 56 cents, ADB on the treated card by 58 cents, and balances across all cards by 59 cents. For the highest FICO score group, the increase in cumulative purchase volume is 22 cents, which is almost identical to the 23 cent increase in ADB on the treated card. Over longer time horizons, the cumulative increase in purchase volume outstrips the rise in ADB. This is consistent with larger effects on overall spending than borrowing. Since we do not have information on purchase volume across all credit cards or cash spending, we cannot rule out that the additional purchase volume over longer time horizons results from shifts in the payment method.

5. Robustness and Additional Heterogeneity. In Online Appendix Section E, we show that the patterns documented above are robust to nonparametric specifications of the relationship between MPB and FICO score, and we show that the main estimates do not differ by the size of the credit limit jump at the quasi-experiment, or by whether the credit card origination was consumer-initiated or bank-initiated. We also explore heterogeneity in the MPBs by borrower income and borrower credit card utilization, instead of by borrower FICO score, and we document that the estimated MPBs are relatively constant across accounts originated at different points during our sample period. Finally, we show that our results are robust to the distribution of FICO scores at which we observe the credit limit quasi-experiments.

6. MPB Take-Away. The quasi-experimental variation in credit limits provides evidence of a large average MPB and substantial heterogeneity in the MPB across FICO score groups. For the lowest FICO score group ($\leqslant$660), we find that a $1 increase in credit limits raises total borrowing by 59 cents at 12 months after origination. This effect is explained by more spending rather than less pay-down of debt. For the highest FICO score group ($>$740), we estimate a 23% effect on the treated credit card that is entirely explained by balance shifting, with a $1 increase in credit limits having no effect on total borrowing. Of course, these estimates are for the set of new credit card applicants, and are not the appropriate estimates for a representative population.

However, among credit card holders, this is the group that is likely to be more responsive to credit expansions, and is thus of particular relevance to policy makers hoping to stimulate borrowing and spending through the banking sector.

Our findings thus suggest that the effects of bank-mediated stimulus on borrowing and spending will depend on whether credit expansions reach those low FICO score borrowers with large MPBs. On the other hand, extending extra credit to low FICO score households who are more likely to default might well conflict with other policy objectives, such as reducing the riskiness of bank balance sheets.

VI. A Model of Optimal Credit Limits

We next present a model of optimal credit limits. We use this model to examine (i) the effect of a change in the cost of funds on credit limits and (ii) how primitives, such as the degree of asymmetric information, create heterogeneity in this effect. In Section VII, we estimate the parameters of this model, allowing us to characterize banks' marginal propensity to lend (MPL) to borrowers with different FICO scores.

To see the value of our approach, consider the alternative of estimating pass-through of declines in the cost of funds using time-series data. Online Appendix Figure A.III shows average credit limits for different FICO score groups over time as well as the cost of funds as reported by banks to the OCC. The plots show that at the onset of the financial crisis, there was a sharp drop in the cost of funds and a sharp drop in credit limits. Of course, the drop in credit limits was due, at least in part, to banks anticipating worse future loan performance. However, a bivariate time-series analysis of these data would generate negative estimates of pass-through. Even with controls, a time-series analysis that is unable to perfectly control for changes in expectations about future loan performance would generate biased estimates.

Naturally, our approach requires us to make alternative assumptions: namely, that bank lending responds optimally to changes in the cost of funds and that we can measure the incentives faced by banks. We think both assumptions are reasonable in our setting: credit card lending is highly sophisticated and our estimates of bank incentives are fairly precise. Indeed, we show that realized marginal profits at prevailing credit limits were close to zero, indicating that the observed credit limits were close to the optimal credit limits implied by our model.

VI.A. Credit Limits as the Primary Margin of Adjustment

In principle, banks could respond to a decline in the cost of funds by adjusting any number of contract terms, including credit limits, interest rates, rewards, and fees. Because of well-known issues of equilibrium existence and uniqueness, the empirical literature on contract pricing in credit markets typically restricts attention to a single margin of adjustment. For example, recent research on the auto market focuses on the determination of down-payment requirements for subprime auto loans (Adams, Einav, and Levin 2009; Einav, Jenkins, and Levin 2012).

An attractive feature of studying the credit card market is that, according to a large body of evidence, interest rates are relatively sticky and credit limits are the primary margin of adjustment. This research on interest rate stickiness builds on the seminal work of Ausubel (1991), which shows that credit card interest rates do not vary with changes in the cost of funds (also see Online Appendix Figure A.IV). The literature has proposed a number of explanations for this interest rate stickiness, including limited interest rate sensitivity by borrowers, collusion by credit card lenders, default externalities across credit card lenders, and an adverse selection story whereby lower interest rates disproportionately attract borrowers with higher default rates (Ausubel 1991; Calem and Mester 1995; Stavins 1996; Stango 2000; Parlour and Rajan 2001). In contrast to interest rates, credit limits vary significantly over time. Online Appendix Figure A.V shows credit limits and interest rates between 2000 and 2015, where for comparability the contract terms in year 2000 are normalized to 100%. Credit limits vary substantially, with a peak-to-trough range of 86% of the initial value. Interest rates vary much less, with a peak-to-trough range of 15% of the initial value.

For the analysis that follows, we therefore focus on credit limits as the single dimension of adjustment. We emphasize, however, that our empirical framework can be applied to other markets, including those where there are other primary margins of adjustment (e.g., the mortgage market). For instance, Fuster and Willen (2010) show that most of the mortgage refinancing in response to the Federal Reserve's quantitative easing programs was done by households with higher FICO scores, with limited refinancing by lower FICO score households. Our framework could be used to determine the extent to which adverse selection in the lower FICO score segment of the market can provide an explanation for this result.

VI.B. Model Setup

Consider a one-period lending problem in which a bank chooses a single credit limit CL for an exogenously defined group of observationally similar borrowers, such as all consumers with the same FICO score, to maximize profits. In Online Appendix B, we show that this optimization problem corresponds to the second stage of a two-stage model of credit card lending, along the lines of the model proposed by Livshits, MacGee, and Tertilt (2016). In this model, banks need to pay a fixed cost to develop a scorecard for lending to a given group of borrowers. Because of this fixed cost, in the first stage banks group borrowers based on FICO score ranges (e.g., 621–660, 661–700) and in the second stage banks set an optimal credit limit for each group. In this section, we only model the second stage of setting credit limits for a group of borrowers, since this is the decision that is most directly affected by changes in the cost of funds.

Let $q(CL)$ describe how the quantity of borrowing depends on the credit limit, and let $MPB = q'(CL)$ indicate the consumers' marginal propensity to borrow out of a credit expansion. Let r denote the interest rate, which, as discussed above, is fixed and determined outside of the model.[21] Let $\tilde{R}(CL) \equiv \bar{R} + R(CL)$ denote noninterest revenue. This includes revenue components such as interchange income and fee revenue, which vary with credit limits, as well as fixed revenue components such as the benefit from cross-selling other products to credit card users. Let c denote the bank's cost of funds, which can be thought of as a refinancing cost, but more broadly captures anything that affects the bank's cost of lending, including capital requirements and financial frictions. Let $\tilde{C}(CL) \equiv \bar{C} + C(CL)$ denote all other costs. These include components such as chargeoffs, which vary with credit limits, as well as potentially fixed costs, such as the cost of originating credit cards. The objective for the bank is to choose a credit limit to

21. Importantly, this does not mean that interest rates have to be the same across the FICO score distribution. Instead, it implies that interest rates for a given FICO score do not change in response to a change in the cost of funds, consistent with the evidence discussed in Section VI.A.

maximize profits:[22]

$$\max_{CL} q(CL)(r - c) + \tilde{R}(CL) - \tilde{C}(CL). \tag{6}$$

The optimal credit limit sets marginal profits to zero, or, equivalently, sets marginal revenue equal to marginal cost:

$$\underbrace{q'(CL)r + R'(CL)}_{=MR(CL)} = \underbrace{q'(CL)c + C'(CL)}_{=MC(CL)}. \tag{7}$$

We assume that profits are weakly positive, and that marginal revenue crosses marginal cost "from above" (i.e., $MR(0) > MC(0)$ and $MR'(CL) < MC'(CL)$) so we are guaranteed to have an interior optimal credit limit. We note that fixed components of revenue and costs (i.e., $\bar{R}$ and $\bar{C}$) drop out of the first-order condition and will therefore have no impact on the MPL.[23]

We are interested in the effect on borrowing of a decrease in the cost of funds, which is given by the total derivative $-\frac{dq}{dc}$. As described in the introduction, and shown in equation (1), this can be decomposed into the product of the marginal propensity to lend (MPL) and the marginal propensity to borrow (MPB). In Section V, we estimated the MPB directly using the quasi-experimental variation in credit limits. We next discuss how we use our variation to estimate the MPL.

VI.C. Pass-Through of a Decrease in the Cost of Funds

A decrease in the cost of funds reduces the marginal cost of extending each unit of credit, and can be thought of as a downward shift in the marginal cost curve and an upward shift in the marginal profit curve. Since equilibrium credit limits are set where marginal profits are equal to zero (see equation (7)), the slope of marginal profits determines the increase in equilibrium credit limits in response to a decline in the cost of funds. To see

22. The model abstracts from the extensive-margin decision of whether or not to offer a credit card. To capture this margin, the model could be extended to include a fixed cost of originating and maintaining a credit card account. In such a model, borrowers would only receive a credit card if expected profits exceeded this fixed cost.

23. These components do affect the overall profitability of credit card lending, and therefore the bank's decision of whether to originate a card in the first place. But, conditional on a card being originated, they will have no effect on the pass-through of changes in the cost of funds.

this, consider Figure I from the introduction. In Panel A, marginal profits are relatively flat, and a given upward shift in the marginal profit curve leads to a large increase in equilibrium credit limits. In Panel B, where marginal profits are relatively steep, the same upward shift in the marginal profit curve leads to a smaller increase in optimal credit limits.

Mathematically, the effect on credit limits of a decrease in the cost of funds can be derived by applying the implicit function theorem to the first-order condition shown in equation (7):

$$(8) \qquad MPL = -\frac{dCL}{dc} = -\frac{q'(CL)}{MR'(CL) - MC'(CL)} = -\frac{q'(CL)}{MP'(CL)}.$$

The numerator is the marginal propensity to borrow ($q'(CL) \equiv MPB$) and scales the size of the effect because a given decrease in the cost of funds induces a larger shift in marginal costs when credit card holders borrow more on the margin. This is also the reason why the vertical axis in Figure I is divided by the MPB. The denominator is the slope of marginal profits: $MP'(CL) = MR'(CL) - MC'(CL)$. The existence assumption ($MR'(CL) < MC'(CL)$) guarantees the denominator is negative and thus implies positive pass-through ($MPL > 0$). The MPL is decreasing as the downward-sloping marginal profit curve becomes steeper. Economically, we view the MPB and the slope of marginal profits as "sufficient statistics" that capture the effect on pass-through of a number of underlying features of the credit card market without requiring strong assumptions on the underlying model of consumer behavior (see Chetty 2009 for more on this approach).

Perhaps the most important of these features is asymmetric information, which includes both selection and moral hazard.[24] Since banks can adjust credit limits based on observable borrower characteristics, they determine the optimal credit limit separately for each group of observably identical borrowers. By selection, we therefore mean selection on information that the lender does not observe or is prohibited from using by law. With adverse selection, higher credit limits disproportionally raise borrowing among consumers with a greater probability of default. This increases the

24. See Einav, Finkelstein, and Cullen (2010) and Mahoney and Weyl (2017) for a more in-depth discussion of how the slope of marginal costs parameterizes the degree of selection in a market.

marginal cost and thus reduces the marginal profit of extending more credit. This could occur because forward-looking consumers, who anticipate defaulting in the future, strategically increase their borrowing. Alternatively, it could occur because there are some consumers that are always more credit constrained, and these consumers borrow more today and have a higher probability of default in the future. Regardless of the channel, adverse selection translates into a more positively sloped marginal cost curve, a more negatively sloped marginal profit curve, and less pass-through.[25]

Higher credit limits could also affect marginal costs, and thus marginal profits, holding the composition of marginal borrowers fixed. For instance, in Fay, Hurst, and White's (2002) model of consumer bankruptcy, the benefits of filing for bankruptcy are increasing in the amount of debt while the costs of filing are fixed. The implication is that higher credit limits, which raise debt levels, lead to higher default probabilities, a more positively sloped marginal cost curve, and a lower rate of pass-through. This mechanism is sometimes called moral hazard because borrowers do not fully internalize the cost of their decisions when choosing how much to borrow and whether to default. However, a positive effect of credit limits on borrowing does not require strategic behavior on the part of the borrower. For example, myopic consumers might borrow heavily out of an increase in credit limits, not because they anticipate defaulting next period, but because they down-weight the future.[26]

The slope of marginal revenue is equally significant in determining the MPL, and revenue from fees (e.g., annual fees, late fees) is a key determinant of the slope of marginal revenue. In particular, fee revenue does not scale one-for-one with credit card utilization. On the margin, an increase in credit limits might increase fee revenue (e.g., by raising the probability a consumer renews her card and pays next year's annual fee) but not by a

25. In principle, selection could also be advantageous, with higher credit limits disproportionally raising borrowing among households with a lower default probability. In this case, more advantageous selection would translate into a less negatively sloped marginal profit curve, and more pass-through.

26. If greater debt levels reduce the rate of default—for example, because increased credit access allows households to "ride out" temporary negative shocks without needing to default—an increase in credit limits would result in lower default probabilities, a less negatively sloped marginal profit curve, and more pass-through.

large amount. A decline in marginal fee revenue at higher credit limits would translate into a more negatively sloped marginal revenue curve, a more negatively sloped marginal profit curve, and less pass-through.

In Section VII, we directly estimate heterogeneity in the slope of marginal costs, marginal revenue, and marginal profits by FICO score. This approach allows us to quantify the joint effect of a broad set of factors such as moral hazard and adverse selection on pass-through without requiring us to untangle their relative importance.

VI.D. Empirical Implementation

Taking the model to the data involves three additional steps. First, our model of optimal credit limits has one period, while our data are longitudinal with monthly outcomes for each account. To align the data with the model, we aggregate the monthly data for each outcome into discounted sums over various horizons, using a monthly discount factor of 0.996, which translates into an annual discount factor of 0.95.[27] With these aggregated data, the objective function for the bank is to set initial credit limits to maximize the discounted flow of profits, which is a one-period problem.[28]

A second issue involves the potential divergence between expected and realized profits. In our model, marginal profits can be thought of as the expectation of marginal profits when the bank sets initial credit limits. In the data, we do not observe these expected marginal profits but instead observe the marginal profits realized by the bank. The simplest way to take our model to the data is to assume that expected and realized marginal profits were the same during our time period. We show in Section VII that realized marginal profits at prevailing credit limits were indeed very close to zero, suggesting that banks' expectations during our time period were approximately correct. We think this is not

27. In 2009, the weighted average cost of capital for the banking sector was 5.86%, in 2010 it was 5.11%, and in 2011 it was 4.27% (http://pages.stern.nyu.edu/adamodar/). Our results are not sensitive to the choice of discount factor.

28. While initial credit limits are highly persistent (see Section IV.A), credit limits can be changed following origination, which affects the discounted sums. We assume that banks set initial credit limits in a dynamically optimal way, taking into account their ability to adjust credit limits in the future. The envelope theorem then allows us to consider the optimization problem faced by a bank at card origination without specifying the dynamic process of credit limit adjustment.

surprising, given the sophisticated, data-driven nature of credit card underwriting, with lenders using randomized trials to continuously learn about the degree of selection and the profitability of adjusting credit limits and other contract terms (e.g., Agarwal, Chomsisengphet, and Liu 2010).

Third, we need to estimate the slopes of outcomes, such as the discounted flow of marginal profits, with respect to a change in credit limits. Our approach to estimating these slopes closely follows the approach used in recent empirical papers on selection in health insurance markets (e.g., Einav, Finkelstein, and Cullen 2010; Cabral, Geruso, and Mahoney 2015; Hackmann, Kolstad, and Kowalski 2015). Conceptually, our approach starts with the observation that each quasi-experiment provides us with two moments. For example, we can recover marginal profits at the prevailing credit limit using our credit-limit regression discontinuities, and we can calculate average profits per dollar of credit limits by dividing total profits by the prevailing credit limit. With two moments, we can then identify any two-parameter curve for marginal profits, such as a linear specification that allows for a separate intercept and slope.

Our baseline specification is to assume that marginal profits, and other outcomes, are linear in credit limits. This specification is advantageous because it allows for internally consistent aggregation across outcomes; for instance, linear marginal costs and linear marginal revenue imply linear marginal profits. The linear specification is also particularly transparent because the slope is captured by a single parameter that can be recovered in closed form. Specifically, if marginal profits are given by $MP(CL) = \alpha + \beta CL$, then average profits per dollar of credit limits are given by $AP(CL) = \frac{\int_{X=0}^{CL} \alpha + \beta X \, dX}{CL} = \alpha + \frac{1}{2}\beta CL$, and the slope of marginal profits is therefore $\beta = \frac{2(MP(CL) - AP(CL))}{CL}$. Intuitively, if marginal profits are much smaller than average profits ($MP(CL) \ll AP(CL)$), the marginal profitability of lending must be rapidly declining in credit limits and marginal profits must be steeply downward-sloping ($MP'(CL) = \beta < 0$). Alternatively, if marginal profits are fairly similar to average profits ($MP(CL) \approx AP(CL)$), then marginal profits must be relatively flat ($MP'(CL) = \beta \approx 0$).

In Online Appendix F, we show that while our precise quantitative estimates of the MPL depend on our linear functional form assumption, our results are qualitatively robust to a wide class of functional forms. Specifically, we prove that as long as the marginal profit function satisfies an appropriately-defined single crossing condition, then the optimal marginal profit function is steeper if and only if $\frac{AP(CL)}{CL}$ has a larger value. Since we find in our data that $\frac{AP(CL)}{CL}$ is larger for lower FICO score borrowers (see Table IV), our finding that the slope of marginal profits is steeper for lower FICO score borrowers is qualitatively robust.

VII. Marginal Propensity to Lend

In Section VI, we showed how the MPL is determined by the negative ratio of the MPB and the slope of marginal profits. In this section, we use the quasi-experimental variation in credit limits to estimate how the slope of marginal profits varies across the FICO score distribution. We then combine these slopes with our estimates of the MPB from Section V to estimate heterogeneity in the MPL.

VII.A. Average Costs, Revenues, and Profits

To provide context, we first present basic facts on the profitability of the credit cards in our sample. We define profits for a credit card account as the difference between total revenue and total costs.

Total revenue is the sum of interest charge revenue, fee revenue, and interchange income. We observe interest charge revenue and fee revenue for each account in our data. Interchange fees are charged to merchants for processing credit card transactions and scale proportionally with spending. Following Agarwal et al. (2015b), we calculate interchange income for each account as 2% of purchase volume.

Total costs are the sum of chargeoffs, the cost of funds, rewards and fraud expenses, and operational costs such as costs for debt collection, marketing, and customer acquisition. We observe chargeoffs for each account in our data.[29] We observe the cost of funds at the bank-month level in the portfolio data and construct an account-level measure of the cost of funds by apportioning these costs based on each account's share of ADB. We calculate that reward and fraud expenses are 1.4% of purchase

29. We use the term "chargeoffs" to indicate gross chargeoffs minus recoveries, which are both observed in our data.

volume and operational costs are 3.5% of ADB in the portfolio data, and construct account-level values by applying these percentages to account-level purchase volume and ADB. See Online Appendix Section G for additional discussion of how we measure profitability components at the account level.

The middle section of Table IV shows cumulative total costs and its key components by FICO score group at different time horizons after account origination. As before, we restrict the sample to credit cards originated within five FICO score points of a credit limit quasi-experiment. Cumulative total costs rise fairly linearly over time and are hump-shaped in FICO score. At 48 months after origination, cumulative total costs are \$588 for the lowest FICO score group ($\leqslant$660), slightly more than \$800 for the middle groups, and \$488 for the highest FICO score group ($>$740). Cumulative chargeoffs generally account for more than half of these costs, although they are more important for lower FICO score accounts and become relatively more important at longer time horizons. The cumulative cost of funds declines from about 10% of total costs at 12 months after origination to about 5% at 60 months after origination.

The right section of Table IV shows cumulative total revenue and profits. Cumulative revenue, like cumulative costs, grows fairly linearly over time. However, while cumulative costs are hump-shaped in FICO score, cumulative revenue is decreasing. For instance, at 48 months after origination, cumulative total revenue is more than \$950 for the two lowest FICO score groups, \$863 for the second highest FICO score group, and \$563 for accounts in the highest FICO score group. Excluding the first year, interest charges make up approximately two-thirds of cumulative total revenue; fee revenue accounts for approximately one-quarter and is particularly important for the lowest FICO score group. Both interest charges and fees are somewhat less important for the highest FICO score group. For these accounts, interchange income is relatively more important, contributing approximately one-fifth of total revenue.

The data on revenue and costs combine to produce average profits that are U-shaped in FICO score. At 48 months, cumulative profits are \$365 for the lowest FICO score group, \$126 and \$55 for the middle two FICO score groups, and \$75 for accounts with the highest FICO score. Cumulative profits within a FICO score group increase fairly linearly over time.

VII.B. Marginal Probability of Default

We begin our analysis of pass-through by examining the causal effect of an increase in credit limits on the probability of delinquency and default.[30] A larger effect on default probabilities, all else equal, corresponds to more steeply upward-sloping marginal costs for two reasons: First, higher default probabilities lead to higher chargeoffs on marginal borrowing, raising marginal costs. Second, higher default probabilities lead to higher losses on inframarginal borrowing, further increasing chargeoffs and the slope of marginal costs.[31]

Figure VIII shows that an increase in credit limits has a large effect on the probability of delinquency for the lower FICO score account holders and virtually no effect for the accounts with the highest FICO scores. Panels A and B show the effect on the probability that the account is at least 60 days past due (60+ DPD) and at least 90 days past due (90+ DPD), respectively. For the lowest FICO score group, a $1,000 increase in credit limits raises the probability of moderate delinquency (60+ DPD) within 4 years by 1.21 percentage points, on a base of 16.5%, and raises the probability of a more serious delinquency (90+ DPD) within 4 years by 1.16 percentage points, on a base of 14.5%. The effect is less than two-thirds as large for accounts with an intermediate FICO score, and close to zero for accounts in the highest FICO score group. Table VI shows the corresponding estimates.

30. When a credit card borrower stops making at least the minimum monthly payment, the account is considered delinquent, or "past due." The regulator requires banks to "charge off" the account balance if an account is severely delinquent, or more than 180 days past due. This requires them to record the outstanding receivables as a loss. Although banks charge off severely delinquent accounts, the underlying debt obligations remain legally valid and consumers remain obligated to repay the debts. As discussed above, our measure of the impact of delinquency on profits is the amount of chargeoffs net any recoveries. We analyze the impact of higher credit limits both on intermediate delinquency stages (the probabilities of being more than 60 or more than 90 days past due), as well as on chargeoffs, which are a key driver of marginal profits.

31. Mathematically, if we express total chargeoffs as $C(CL) = d(CL)q(CL)$, where $d(CL)$ is a default indicator and $q(CL)$ the amount of borrowing, then the slope of marginal chargeoffs is given by $C''(CL) = 2d'(CL)MPB(CL) + d''(CL)q(CL) + d(CL)MPB'(CL)$. Since $MPB(CL) > 0$, a larger effect on the probability of default (larger $d'(CL)$) corresponds to more upward-sloping marginal chargeoffs (larger $C''(CL)$) and thus more upward-sloping marginal costs, holding the other terms constant.

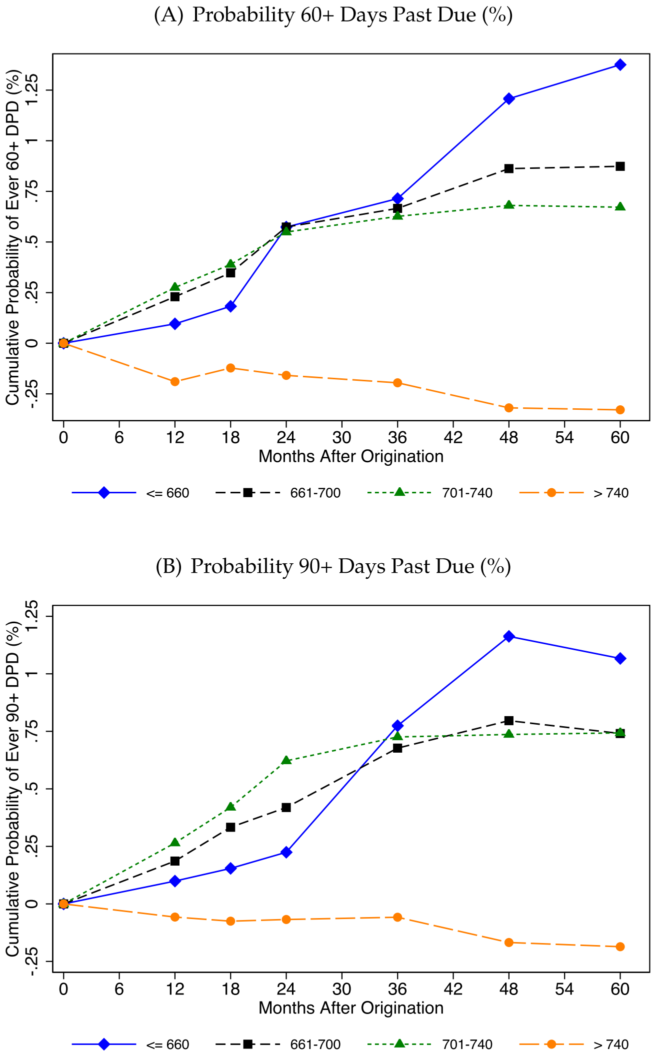

FIGURE VIII

Probability of Delinquency

Figure shows the effects of a $1,000 increase in credit limits on the probability of delinquency for different FICO score groups and different time horizons after account origination. Panel A shows effects on the probability of an account being more than 60 days past due (60+ DPD) within the time horizon, Panel B shows the probability of being more than 90 days past due (90+ DPD) within the time horizon. FICO score groups are determined by FICO score at account origination. The corresponding estimates are shown in Table VI.

TABLE VI
PROBABILITY OF DELINQUENCY

	Months after account origination				
	12	24	36	48	60
Panel A: 60+ days past due (%)					
FICO					
⩽660	0.10	0.57	0.71	1.21	1.38
	[−0.46, 0.67]	[−0.16, 1.30]	[−0.13, 1.63]	[0.25, 2.06]	[0.61, 2.19]
661–700	0.23	0.57	0.67	0.86	0.87
	[−0.05, 0.55]	[0.23, 0.93]	[0.29, 1.06]	[0.44, 1.24]	[0.46, 1.27]
701–740	0.28	0.55	0.63	0.68	0.67
	[0.04, 0.52]	[0.23, 0.86]	[0.24, 0.99]	[0.30, 1.04]	[0.30, 1.01]
>740	−0.19	−0.16	−0.20	−0.32	−0.33
	[−0.39, −0.04]	[−0.46, 0.05]	[−0.49, 0.06]	[−0.61, −0.05]	[−0.61, −0.08]
Panel B: 90+ days past due (%)					
FICO					
⩽660	0.10	0.22	0.77	1.16	1.07
	[−0.37, 0.67]	[−0.60, 0.94]	[−0.08, 1.61]	[0.34, 1.88]	[0.32, 1.79]
661–700	0.19	0.42	0.68	0.80	0.74
	[−0.02, 0.45]	[0.14, 0.78]	[0.35, 1.05]	[0.50, 1.18]	[0.45, 1.11]
701–740	0.26	0.62	0.73	0.74	0.74
	[0.08, 0.48]	[0.31, 0.91]	[0.37, 1.02]	[0.38, 1.04]	[0.41, 1.07]
>740	−0.06	−0.07	−0.06	−0.17	−0.19
	[−0.23, 0.08]	[−0.33, 0.12]	[−0.34, 0.17]	[−0.45, 0.07]	[−0.44, 0.05]

Notes. Table shows regression discontinuity estimates of the effect of a $1,000 increase in credit limits on the probability of delinquency. Panel A shows the effects on the probability that the account is at least 60 days past due (60+ DPD); Panel B shows effects on the probability that the account is at least 90 days past due (90+ DPD). Columns show effects at different time horizons after account origination. Within each panel, rows show effects for different FICO score groups, defined at account origination. 95% confidence intervals are constructed by bootstrapping over quasi-experiments, and are presented in square brackets.

Online Appendix Figure A.VI shows RD plots for the pooled sample of all quasi-experiments.

We view this evidence as complementary to our main analysis of the slopes of marginal profits. Large effects on the probability of delinquency among low FICO score borrowers indicate, holding other terms equal, that the slope of marginal chargeoffs is steeper in the bottom part of the FICO distribution. However, while the effects on delinquency are intuitive and straightforward to estimate, they are not sufficient statistics for pass-through. First, the effects need to be dollarized to capture their influence on marginal profits. Second, the estimates do not incorporate the effects of selection. For instance, if borrowers with a higher default probability increase borrowing more strongly when credit limits increase, marginal costs can be upward sloping with no effect on the probability of default. For these reasons, we next estimate the slope of marginal profits, which is directly informative for the MPL.

VII.C. Slope of Marginal Profits and Components

The top row of Figure IX considers the effects of increasing credit limits on marginal costs and marginal chargeoffs. For each FICO score group, the gray bars on the left show the marginal effects of a $1 increase in credit limits at prevailing equilibrium credit limits; the black bars on the right show the response of those marginal effects to a $1,000 increase in credit limits. The capped vertical lines show 95% confidence intervals constructed by bootstrapping over quasi-experiments. The estimates are based on cumulative outcomes over a four-year horizon, although we will show robustness to different time horizons. Columns (1) to (4) of Table VII show the corresponding estimates, and Panels A and B of Online Appendix Figure A.VII present the standard RD plots for the pooled sample of all quasi-experiments.

Marginal costs at prevailing credit limits decrease sharply by FICO score. For the lowest FICO score borrowers ($\leqslant$660), a $1 increase in credit limits raises cumulative costs over four years by 29.6 cents, mainly due to a 21.6 cents increase in chargeoffs. For the highest FICO score group ($>$740), a $1 increase in credit limits raises cumulative costs by a much smaller 6.0 cents, with a 3.7 cents increase in chargeoffs. As discussed in Section VI, what matters for pass-through, though, is not the level of marginal costs at the prevailing credit limits, but what happens to these marginal costs as credit limits are increased. For the lowest FICO score group, marginal costs are steeply upward sloping, with a $1,000 increase in credit limits raising marginal costs by 3.3 cents, or about one-ninth of the baseline marginal effect. The upward slope is driven by higher marginal chargeoffs. For the higher FICO score groups, a $1,000 increase in credit limits has virtually no effect on marginal costs and marginal chargeoffs. These results are consistent with less selection and a smaller direct effect of credit limits on default probabilities at higher FICO scores.

The middle row of Figure IX examines the effect of increasing credit limits on total cumulative marginal revenue and cumulative marginal fee revenue. The plots are constructed identically to the plots for costs and chargeoffs. Columns (5) to (8) of Table VII show the corresponding estimates, and Panels C and D of Online Appendix Figure A.VII present the standard RD plots for the pooled sample of all quasi-experiments. Marginal revenue at prevailing credit limits, shown by the gray bars, is decreasing in FICO score. For the lowest FICO score group, a $1 increase

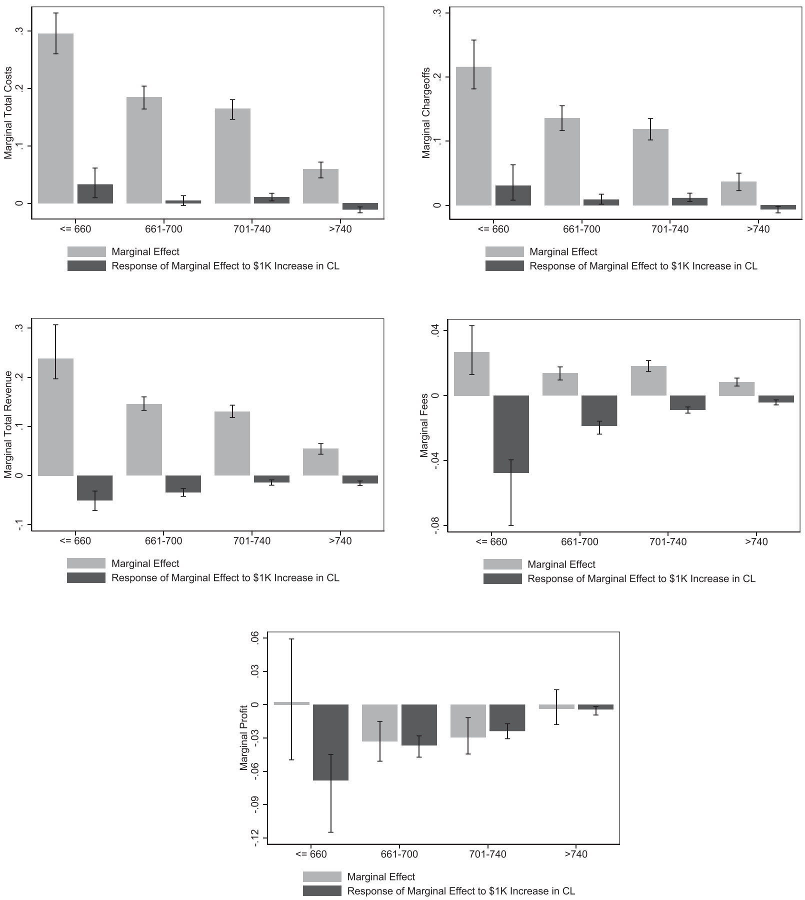

FIGURE IX

Marginal Effects and Response of Marginal Effects to a $1K Increase in Credit Limits

Figure shows marginal effects and the effect of a $1,000 increase in credit limits on marginal effects by FICO score group. We show these effects for total costs, chargeoffs (which are an important component of total costs), total revenue, fee revenue (which is an important component of total revenue), and profits (which is defined as total revenue minus total costs). We measure these variables cumulatively over a time horizon of 48 months after account origination. For each measure, the grey bars show the RD estimate of the marginal effect of a $1 increase in credit limits at the prevailing equilibrium credit limits. The black bars show the impact of a $1,000 increase in credit limits on this marginal effect. Vertical bars show 95% confidence intervals, constructed by bootstrapping across quasi-experiments. FICO score groups are determined by FICO score at account origination. The corresponding estimates are shown in Table VII.

TABLE VII

MARGINAL EFFECTS AND RESPONSE OF MARGINAL EFFECTS TO A $1,000 INCREASE IN CREDIT LIMITS

	Total costs		Chargeoffs		Total revenue		Fees		Profits	
	Marginal effect	Response of marginal effect to $1K increase	Marginal effect	Response of marginal effect to $1K increase	Marginal effect	Response of marginal effect to $1K increase	Marginal effect	Response of marginal effect to $1K increase	Marginal effect	Response of marginal effect to $1K increase
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
FICO										
⩽660	0.296	0.033	0.216	0.031	0.238	−0.051	0.027	−0.047	0.002	−0.068
	[0.261, 0.332]	[0.010, 0.061]	[0.182, 0.258]	[0.008, 0.063]	[0.197, 0.307]	[−0.071, −0.032]	[0.013, 0.043]	[−0.080, −0.040]	[−0.050, 0.059]	[−0.115, −0.045]
661–700	0.185	0.005	0.136	0.009	0.146	−0.034	0.014	−0.019	−0.033	−0.037
	[0.164, 0.204]	[−0.004, 0.013]	[0.117, 0.155]	[0.002, 0.017]	[0.133, 0.160]	[−0.042, −0.026]	[0.010, 0.018]	[−0.024, −0.016]	[−0.051, −0.015]	[−0.047, −0.028]
701–740	0.165	0.011	0.119	0.012	0.130	−0.014	0.018	−0.009	−0.029	−0.024
	[0.146, 0.181]	[0.004, 0.017]	[0.102, 0.136]	[0.006, 0.019]	[0.118, 0.143]	[−0.020, −0.009]	[0.015, 0.022]	[−0.011, −0.007]	[−0.044, −0.012]	[−0.031, −0.017]
>740	0.060	−0.011	0.037	−0.006	0.055	−0.016	0.008	−0.004	−0.004	−0.004
	[0.044, 0.072]	[−0.017, −0.006]	[0.023, 0.050]	[−0.012, −0.002]	[0.043, 0.065]	[−0.021, −0.011]	[0.006, 0.011]	[−0.006, −0.003]	[−0.018, 0.014]	[−0.009, −0.002]

Notes. Table shows marginal effects, and the response of marginal effects to a $1,000 increase in credit limits, by FICO score group. We show these effects for total costs, chargeoffs (which is an important component of total costs), total revenue, fee revenue (which is an important component of total revenue), and profits (which is defined as total revenue minus total costs). We measure these variables over a time horizon of 48 months after account origination. For each measure, the left column shows the RD estimate of the marginal effect of a $1 increase in credit limits at the prevailing equilibrium level, and the right column shows the response of that marginal effect to a $1,000 increase in credit limits. Rows show effects for different FICO score groups, defined at account origination. 95% confidence intervals are constructed by bootstrapping over quasi-experiments, and are presented in square brackets.

in credit limits raises revenue by 23.8 cents. For the highest FICO score group, a $1 increase in credit limits raises revenue by 5.5 cents.

Marginal revenue is steeply downward sloping for low FICO score borrowers and much flatter for borrowers with higher FICO scores. For the lowest FICO score group, a $1,000 increase in credit limits reduces marginal revenue by 5.1 cents, or about one-quarter of the baseline marginal effect. The majority of this decline is due to a drop in marginal fee revenue.[32] For the second lowest FICO score group, a $1,000 increase in credit limits decreases marginal revenue by only 3.4 cents, and the decrease is around 1.5 cents for the higher FICO score groups.

Panel E of Figure IX brings these results together into an analysis of cumulative marginal profits at 48 months since account origination.[33] Columns (9) and (10) of Table VII show the corresponding estimates and Panel E of Online Appendix Figure A.VII presents the standard RD plot for the pooled sample of all quasi-experiments. Marginal profits at prevailing credit limits, shown with the grey bars, are virtually zero for the lowest and highest FICO score groups (0.2 cents and −0.4 cents, respectively) and slightly negative for the middle FICO score groups (−3.3 cents and −2.9 cents, respectively), indicating that credit limits during our time period were approximately optimal ex post. While not the primary focus of our research, the implication is that banks were not forgoing profitable lending opportunities in the credit card market during our time period. This result provides support for the "no good risks" explanation for limited credit supply during the Great Recession and pushes against the argument that financial frictions prevented banks from exploiting profitable consumer lending opportunities.[34]

32. Marginal fee revenue can, in principle, be negative. For instance, a higher credit limit that reduces the frequency of over-limit fees is modeled as negative marginal fee revenue in our framework.

33. We estimate the effect on marginal profits directly rather than constructing it as the difference between marginal revenue and marginal cost. Estimating this effect directly maximizes statistical power but means that the effects do not aggregate perfectly, i.e., our point estimates for the slopes of marginal revenue and marginal cost do not combine to deliver the point estimate for the slope of marginal profit.

34. This is consistent with claims by James Chessen, the chief economist of the American Bankers Association, who explained reduced lending volumes by arguing that, "it's a very risky time for any lender because the probability of loss

The slope of marginal profits is strongly negative for the lowest FICO score borrowers and becomes less negative at higher FICO scores. For the lowest FICO score group, a \$1,000 increase in credit limits reduces cumulative marginal profits over 48 months by 6.8 cents, driven by both lower marginal revenue and higher marginal costs. In response to a \$1,000 increase in credit limits, marginal profits decline by 3.7 cents and 2.4 cents for the middle FICO score groups, and by 0.4 cents for the group with the highest FICO scores.

As we mentioned above, our qualitative finding that the slope of marginal profits is decreasing in FICO score is not dependent on our linearity assumptions. In particular, in Online Appendix Section F, we prove that as long as the marginal profit function satisfies an appropriately defined single-crossing condition, then the optimal marginal profit function is steeper if and only if $\frac{AP(CL)}{CL}$ has a larger value. Using the values in Table IV at a 48-month horizon, we calculate that $\frac{AP(CL)}{CL}$ declines monotonically from 7.2×10^{-5} for the lowest FICO group to 1.6×10^{-6} for the highest FICO group, implying that the slope of marginal profits is declining in FICO score for any marginal profit function that satisfies the single-crossing condition.[35] Thus, while our exact estimates rely on the assumed linear functional form, our basic results are qualitatively robust.

VII.D. Marginal Propensity to Lend (MPL)

The next step in our analysis is to use the estimates above to calculate the MPL in response to a decline in the cost of funds, which is given by the negative ratio of the cumulative MPB and the slope of cumulative marginal profits, measured over the same horizon: $MPL = -\frac{MPB}{MP'(CL)}$ (see Section VI).

Figure X shows the effect on credit limits of a permanent one percentage point decrease in the cost of funds by FICO score group.[36] For each FICO score group, we show estimates using

is greater, and they are being prudent in their approach to lending" (Wall Street Journal 2009).

35. Since Average Profits After 48 Months $= \frac{\text{Cumulative Profits After 48 Months}}{\text{Credit Limit After 48 Months}}$, we show values of $\frac{\text{Cumulative Profits After 48 Months}}{\text{Credit Limit After 48 Months}^2}$.

36. While we consider the effect of a uniform one percentage point decrease in the cost of funds across FICO score groups, our framework can be used to quantify the effects of reductions in the cost of funds that vary by the FICO score of the borrowers. For instance, due to higher capital charges, the cost of funds might be

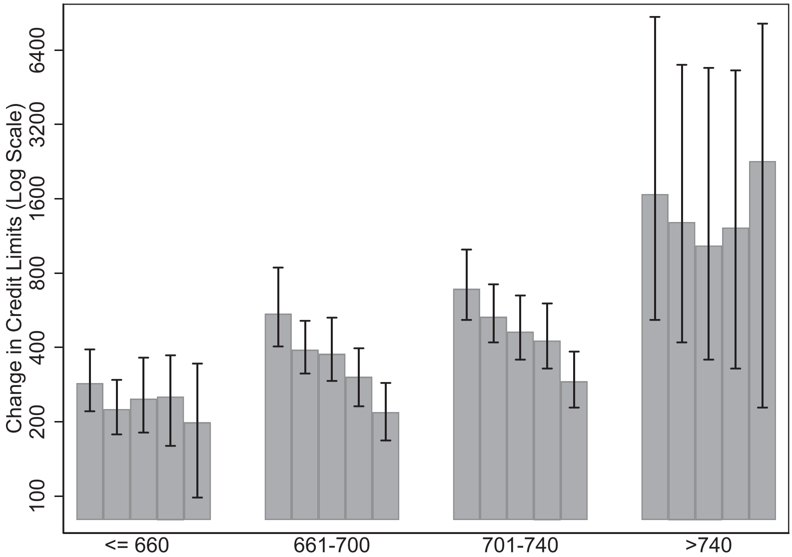

FIGURE X

Marginal Propensity to Lend (MPL)

Figure shows the implied effect of a one percentage point reduction in the cost of funds on optimal credit limits by FICO score group. Estimates are produced using equation (8), and are shown on a log scale. For each FICO score group, we show the implied increase in credit limits when measuring both the slope of cumulative marginal profits and cumulative marginal borrowing over the first 12, 24, 36, 48, and 60 months following origination (left to right). Vertical bars show 95% confidence intervals, constructed by bootstrapping across quasi-experiments. FICO score groups are determined by FICO score at account origination. The corresponding estimates are shown in Table VIII.

data on cumulative profits and ADB over time horizons of 12, 24, 36, 48, and 60 months after origination. The capped vertical lines show 95% confidence intervals constructed by bootstrapping over quasi-experiments.[37]

higher for low FICO score borrowers. More importantly, policies such as the stress tests might have differentially increased the cost of lending to the low FICO score borrowers. Our framework allows us to account for this type of heterogeneity by rescaling our estimates of the MPL by each FICO score group's specific change in the cost of funds.

37. In particular, we draw 500 sets of quasi-experiments with replacement, and calculate $MPL = -\frac{MPB}{MP'(CL)}$ using this bootstrap sample. This procedure effectively allows the standard errors of the numerator and denominator to be correlated.

The plot shows a sharp increase in the MPL by FICO score. For the lowest FICO score group, a one percentage point decrease in the cost of funds raises credit limits by $253 when we use discounted flows over 48 months to estimate the MPB and the slope of marginal profits. For consumers in the highest FICO score group, the increase is approximately five times larger at $1,224. The estimates are stable to measuring cumulative profits and ADB over different horizons. We use the 48-month values as our preferred specification.[38]

VII.E. Effect on Aggregate Borrowing

The effect of a decline in the cost of funds on aggregate borrowing is given by the product of MPL and MPB, aggregated over all FICO groups in the economy.[39] Panel A of Figure XI shows the effect of a one percentage point decrease in the cost of funds on credit limits by FICO score group. Panel B shows the MPB across all cards at 12 months after origination by FICO score group, which captures the short-term effect on borrowing. Table VIII shows the corresponding estimates.

MPL and MPB are strongly negatively correlated, with the highest MPL occurring for the accounts with the lowest MPB. The bottom panel of Table VIII quantifies the importance of this negative correlation by estimating the impact on aggregate borrowing under alternative assumptions. The first row shows this calculation when the negative correlation is not taken into account, and the effect on borrowing is given by the weighted average MPL $\times$ weighted average MPB, where we weight FICO score groups by the total number of accounts within each group in the full sample (see Section II.D). The second row accounts for this correlation by first calculating MPL $\times$ MPB for each FICO score group and then averaging across the FICO score groups. The point estimate for MPB is sometimes slightly negative for the highest

38. Using cumulative flows over different time horizons involves a trade-off. On the one hand, using longer horizons allows us to better capture potential life-cycle effects in credit card profitability. On the other hand, focusing on longer time horizons requires us to restrict the analysis to accounts that were originated in the early part of our panel, which reduces the number of quasi-experiments we can exploit. Reassuringly, our effects are robust to the choice of time horizon.

39. This approach to calculating the effect on aggregate borrowing abstracts away from the existence of spending multipliers or other general equilibrium effects, such as the possibility that additional spending from extra credit might reduce the rate of default of other borrowers.

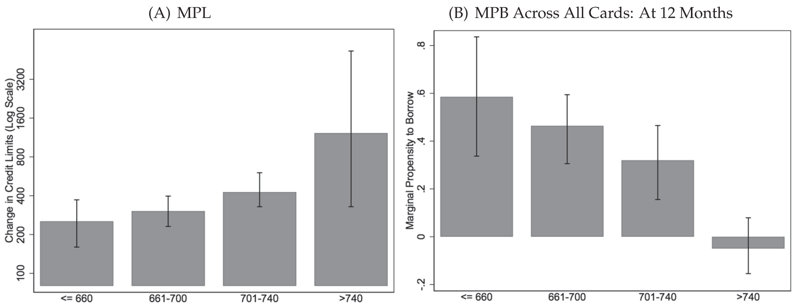

FIGURE XI

Correlation between MPL and MPB

Panel A shows the implied effect of a one percentage point reduction in the cost of funds on optimal credit limits by FICO score group. The effects are calculated using the marginal profit estimates shown in Figure IX and Table VII, and are shown on a log scale. Panel B shows the effect of a $1 increase in credit limits on borrowing across all cards by FICO score group. Vertical bars show 95% confidence intervals, constructed by bootstrapping across quasi-experiments. FICO score groups are determined by FICO score at account origination. The corresponding estimates are shown in Table VIII.

FICO score group. Therefore, the third row shows our preferred version of the calculation where we account for the correlation but bottom-code the MPB at zero. At a 12-month horizon, accounting for the correlation reduces the effect on aggregate borrowing by 51%, relative to the estimate that does not account for this correlation. This reduction is similar at longer time horizons.

We conduct two exercises to help interpret the magnitudes of our estimates. Both exercises focus on the decline in banks' cost of funds during the first few months of the 2008 financial crisis, when the federal funds rate was reduced from about 2% to 0%. As shown in Panel A of Online Appendix Figure A.XX, we find that banks' cost of funds declined by 0.96 percentage point during this time period, from an annualized rate of 3.20% in September 2008 to an annualized rate of 2.24% in January 2009.

The first exercise is to compare our estimates of the extra spending on each new credit card account to established evidence on the spending effects of fiscal stimulus payments such as tax rebates. Based on the estimates in Table VIII, our results indicate that the 0.96 percentage point decline in the cost of funds generated a $630 average increase in credit limits, and a $65 average increase in borrowing and consumption for new

TABLE VIII

MARGINAL PROPENSITY TO LEND × MARGINAL PROPENSITY TO BORROW

	MPL	MPB across all cards				
		12 months	24 months	36 months	48 months	60 months
FICO						
⩽660	253	0.59	0.54	1.00	0.96	1.27
	[160, 372]	[0.34, 0.84]	[0.18, 0.94]	[0.51, 1.48]	[0.12, 1.97]	[−0.16, 2.51]
661–700	304	0.46	0.42	0.48	0.59	0.43
	[231, 397]	[0.31, 0.59]	[0.26, 0.58]	[0.26, 0.70]	[0.13, 0.97]	[−0.42, 1.11]
701–740	427	0.32	0.21	0.24	0.35	0.49
	[329, 603]	[0.16, 0.47]	[0.03, 0.37]	[0.03, 0.44]	[0.00, 0.65]	[−0.52, 1.30]
>740	1,224	−0.05	−0.08	−0.19	0.05	0.29
	[329, 5,300]	[−0.15, 0.08]	[−0.26, 0.10]	[−0.47, 0.08]	[−0.42, 0.39]	[−0.47, 1.03]
Weighted average	655	0.28	0.23	0.33	0.45	0.62

	MPL × MPB				
	12 months	24 months	36 months	48 months	60 months
Without accounting for correlation	183.53	149.87	218.87	293.75	408.10
Accounting for correlation	67.40	38.80	28.88	148.36	282.23
Accounting for correlation + lower bound	89.76	75.84	114.14	148.36	282.23

Notes. Table shows the effects of a reduction in the cost of funds on lending and borrowing. The first column of the top panel shows the effect of a permanent one percentage point reduction in the cost of funds on optimal credit limits (MPL), constructed using cumulative marginal profits and cumulative borrowing over the first 48 months after account origination. The remaining columns reproduce the MPB estimates from Table V at different time horizons after account origination. Both estimates are shown by FICO score group, defined at account origination. 95% confidence intervals are constructed by bootstrapping over quasi-experiments, and are presented in square brackets. The bottom panel shows the implied stimulative effect at these same time horizons. The estimates that do not account for correlation are calculated as weighted average MPL × weighted average MPB. The estimates that account for this correlation are constructed by first calculating MPL × MPB for each FICO score group and then taking the weighted average. In the last row we set the (statistically insignificant) negative coefficient for MPB for high FICO score borrowers to zero. Weighted averages are produced by weighting each group by the share of credit card holders with that FICO score in our data (see Section II.D and Online Appendix Figure A.I).

cardholders over a 12-month time horizon.[40] To evaluate the size of this effect, we calculate the fiscal stimulus payment that would have been needed to generate an equivalent increase in spending. Fiscal stimulus provides an interesting comparison,

40. The size of this effect is relatively small compared to the effects of monetary policy on consumption through the mortgage market. For example, Di Maggio, Kermani, and Ramcharan (2014) find that due to reductions in the federal funds rate, borrowers with adjustable rate mortgages originated between 2005 and 2007 experienced an average drop of $900 in monthly mortgage payments upon mortgage reset. This increased monthly spending on car purchases by $140. The authors cannot measure nondurable spending. Keys et al. (2014) study a different sample of adjustable rate mortgages, and show that the reset of 5/1 ARMs lowered monthly mortgage payments by $150. They find that, two years after the reset, car loan balances are $324 higher, suggesting substantial durable goods purchases as a result of the decline in interest rates.

because in contrast to our setting—which features a mismatch between MPL and MPB—stimulus payments can be more uniformly distributed across households. We focus on stimulus payments under the 2008 Economic Stimulus Act, which provided rebate checks of $300 to $600 to individuals and $600 to $1,200 to families between May and July 2008. Parker et al. (2013) estimate a marginal propensity to consume out of these payments of 50% to 90% for combined nondurable and durable consumption. If we take the mid-point of this range, achieving a $65 increase in consumption would require a stimulus payment of $93.

The second exercise we perform is to quantify the impact of the change in the cost of funds on aggregate credit card borrowing. We can calculate the aggregate effects for new cardholders with comparatively weak assumptions. Using a representative sample of credit bureau data, we calculate that there were 57.2 million new credit card accounts opened in the 12-month period starting in October 2008. As discussed above, we find that the reduction in banks' cost of funds raised borrowing on new credit cards by $65 on average over a 12-month time horizon. If we assume that the effects for the new account holders in our sample are equal to the effects for all new credit card accounts, then this reduction in banks' cost of funds raised aggregate borrowing by new account holders by $3.7 billion over a 12-month period. If this increased borrowing translated one-to-one into an increase in consumption, as our estimates suggest, the credit expansion would have raised Personal Consumer Expenditure (PCE) over this 12-month period by 0.04%. This calculation requires us to extrapolate from our local average treatment effects, and our methodology cannot estimate the general equilibrium effects of the policy, such as multiplier or price effects.

Extrapolating from the effect estimated on the sample of new borrowers to the effect on all credit card accounts is challenging, because we need to make assumptions on how pass-through for existing accounts compares to our estimates for new account holders. Conceptually, it seems likely that existing account holders would have a lower MPB than new account holders because they are not actively applying for additional credit. Gross and Souleles (2002) find an average MPB among existing credit cards of between 10% and 14%, relative to our average MPB of 28%. Based on data from the New York Fed CCP and the credit bureau data cited above, we estimate that there were 373 million existing credit card accounts over the 12-month period starting in October 2008. If we assume that the MPB for existing accounts is one-third as large as it is for newly originated accounts, then our estimates imply an aggregate increase in borrowing of $8.1 billion. Combining the effects for new and existing accounts yields an average increase in borrowing of $11.8 billion, which would translate into a 0.12% increase in PCE over this period.

VIII. Conclusion

We propose a new empirical approach to studying the bank lending channel that focuses on frictions, such as asymmetric information, that arise in bank-borrower interactions. Our approach highlights that the effectiveness of bank-mediated stimulus in raising household borrowing depends on whether banks pass through credit expansions to households that want to borrow. We use panel data on all credit cards issued by the eight largest U.S. banks together with 743 credit limit regression discontinuities to estimate the heterogeneity in banks' MPL to different households, and heterogeneity in these households' MPB.

We find large differences in the MPB across the FICO score distribution, with a $1 increase in credit limits raising total borrowing at 12 months after account origination by 59 cents for households with the lowest FICO scores ($\leqslant$660) while having no effect on households with the highest FICO scores (>740). Banks' MPLs are negatively correlated with these MPBs, with a one percentage point reduction in the cost of funds raising optimal credit limits by $253 for households with FICO scores below 660 versus $1,224 for households with FICO scores above 740. We conclude that banks pass through credit expansions least to households that want to borrow the most, reducing the effectiveness of bank-mediated stimulus.

Georgetown University
Office of the Comptroller of the Currency
University of Chicago, Booth School of Business, and NBER
New York University, Stern School of Business, NBER, and CEPR

Supplementary Material

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online. Code replicating the tables and figures in this article can be found in Agarwal et al. (2017), in the Harvard Dataverse, doi:10.7910/DVN/LD67JZ.

References

Acharya, Viral V., Björn Imbierowicz, Sascha Steffen, and Daniel Teichmann, "Does the Lack of Financial Stability Impair the Transmission of Monetary Policy?" Working paper, 2015.

Adams, William, Liran Einav, and Jonathan Levin, "Liquidity Constraints and Imperfect Information in Subprime Lending," *American Economic Review*, 99 (2009), 49–84.

Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tomasz Piskorski, Amit Seru, and Vincent Yao, "Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program," NBER Working Paper 21512, 2015a.

Agarwal, Sumit, Souphala Chomsisengphet, and Chunlin Liu, "The Importance of Adverse Selection in the Credit Card Market: Evidence from Randomized Trials of Credit Card Solicitations," *Journal of Money, Credit and Banking*, 42 (2010), 743–754.

———, "Regulating Consumer Financial Products: Evidence from Credit Cards," *Quarterly Journal of Economics*, 130 (2015b), 111–164.

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

Agarwal, Sumit, and Wenlan Qian, "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore," *American Economic Review*, 104 (2014), 4205–4230.

Auclert, Adrien, "Monetary Policy and the Redistribution Channel," Technical report, Working paper, 2016.

Ausubel, Lawrence M., "The Failure of Competition in the Credit Card Market," *American Economic Review*, 81 (1991), 50–81.

Aydin, Deniz, "The Marginal Propensity to Consume out of Liquidity: Evidence from Random Assignment of 54,522 Credit Lines," Working paper, 2017.

Baker, Scott R., "Debt and the Consumption Response to Household Income Shocks," 2015; doi: 10.2139/ssrn.2541142.

Bhutta, Neil, and Benjamin J. Keys, "Interest Rates and Equity Extraction During the Housing Boom," *American Economic Review*, 106 (2016), 1742–1774.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, 98 (2008), 1887–1921.

Cabral, Marika, Michael Geruso, and Neale Mahoney, "Does Privatized Health Insurance Benefit Patients or Producers? Evidence from Medicare Advantage," NBER Working Paper 20470, 2015.

Calem, Paul S., and Loretta J. Mester, "Consumer Behavior and the Stickiness of Credit-Card Interest Rates," *American Economic Review*, 85 (1995), 1327–1336.

Carroll, Christopher D., "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112 (1997), 1–55.

——, "A Theory of the Consumption Function, with and without Liquidity Constraints," *Journal of Economic Perspectives*, 15 (2001), 23–45.

Chakraborty, Indraneel, Itay Goldstein, and Andrew MacKinlay. "Monetary Stimulus and Bank Lending," Working paper, 2017.

Chetty, Raj, “Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods,” *Annual Review of Economics*, 1 (2009), 451–488.
Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia, “Innocent Bystanders? Monetary Policy and Inequality in the US,” NBER Working Paper 18170, 2012.
Dell’Ariccia, Giovanni, Luc Laeven, and Gustavo Suarez, “Bank Leverage and Monetary Policy’s Risk-Taking Channel: Evidence from the United States,” ECB Working Paper, 2016.
Di Maggio, Marco, Amir Kermani, and Rodney Ramcharan, “Monetary Policy Pass-Through: Household Consumption and Voluntary Deleveraging,” Columbia Business School Research Paper, 2014.
Dobbie, Will, and Paige Marta Skiba, “Information Asymmetries in Consumer Credit Markets: Evidence from Payday Lending,” *American Economic Journal: Applied Economics*, 5 (2013), 256–282.
Doepke, Matthias, and Martin Schneider, “Inflation and the Redistribution of Nominal Wealth,” *Journal of Political Economy*, 114 (2006), 1069–1097.
Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, “The Deposits Channel of Monetary Policy,” *Quarterly Journal of Economics*, 132 (2017).
Eggertsson, Gauti B., and Paul Krugman, “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach,” *Quarterly Journal of Economics*, 127 (2012), 1469–1513.
Einav, Liran, Amy Finkelstein, and Mark R. Cullen, “Estimating Welfare in Insurance Markets Using Variation in Prices,” *Quarterly Journal of Economics*, 125 (2010), 877–921.
Einav, Liran, Mark Jenkins, and Jonathan Levin, “Contract Pricing in Consumer Credit Markets,” *Econometrica*, 80 (2012), 1387–1432.
Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan, “Assessing Sale Strategies in Online Markets Using Matched Listings,” *American Economic Journal: Microeconomics*, 7 (2015), 215–247.
Fay, Scott, Erik Hurst, and Michelle J. White, “The Household Bankruptcy Decision,” *American Economic Review*, 92 (2002), 706–718.
FDIC, “Scoring and Modeling,” Technical report, Federal Deposit Insurance Corporation (FDIC) Credit Card Activities Manual, 2007.
Fuster, Andreas, and Paul Willen, “$1.25 Trillion Is Still Real Money: Some Facts about the Effects of the Federal Reserve’s Mortgage Market Investments,” FRB of Boston Public Policy Discussion Paper 10 (2010).
Geithner, Timothy, “Secretary Geithner Introduces Financial Stability Plan,” 2009, https://www.treasury.gov/press-center/press-releases/Pages/tg18.aspx
Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis, “How Individuals Smooth Spending: Evidence from the 2013 Government Shutdown Using Account Data,” Technical report, NBER Working Paper 21025, 2015.
Gross, David B., and Nicholas S. Souleles, “Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data,” *Quarterly Journal of Economics*, 117 (2002), 149–185.
Guerrieri, Veronica, and Guido Lorenzoni, “Credit Crises, Precautionary Savings, and the Liquidity Trap,” *Quarterly Journal of Economics*, 132 (2017), 1427–1467.
Hackmann, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski, “Adverse Selection and an Individual Mandate: When Theory Meets Practice,” *American Economic Review*, 105 (2015), 1030–1066.
Hall, Robert E., “The Long Slump,” *American Economic Review*, 101 (2011), 431–469.
Hertzberg, Andrew, Andres Liberman, and Daniel Paravisini, “Adverse Selection and Maturity Choice in Consumer Credit Markets: Evidence from an Online Lender,” Columbia Business School Research Paper, 2016.

Hsieh, Chang-Tai, “Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund,” *American Economic Review*, 93 (2003), 397–405.

Hurst, Erik, Benjamin J. Keys, Amit Seru, and Joseph Vavra, “Regional Redistribution through the US Mortgage Market,” *American Economic Review*, 106 (2016), 2982–3028.

Imbens, Guido, and Karthik Kalyanaraman, “Optimal Bandwidth Choice for the Regression Discontinuity Estimator,” *Review of Economic Studies*, 79 (2011), 933–959.

Jappelli, Tullio, and Luigi Pistaferri, “The Consumption Response to Income Changes,” *Annual Review of Economics*, 2 (2010), 479–506.

——, “Fiscal Policy and MPC Heterogeneity,” *American Economic Journal: Macroeconomics*, 6 (2014), 107–136.

Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina, “Credit Supply and Monetary Policy: Identifying the Bank Balance-Sheet Channel with Loan Applications,” *American Economic Review*, 102 (2012), 2301–2326.

———, “Hazardous Times for Monetary Policy: What Do Twenty-Three Million Bank Loans Say about the Effects of Monetary Policy on Credit Risk-Taking?,” *Econometrica*, 82 (2014), 463–505.

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

Karlan, Dean, and Jonathan Zinman, “Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment,” *Econometrica*, 77 (2009), 1993–2008.

Kashyap, Anil K., and Jeremy C. Stein, “Monetary Policy and Bank Lending,” in *Monetary Policy* (Chicago: University of Chicago Press, 1994), 221–261.

Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, “Did Securitization Lead to Lax Screening? Evidence from Subprime Loans,” *Quarterly Journal of Economics*, 125 (2010), 307–362.

Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vincent Yao, “Mortgage Rates, Household Balance Sheets, and the Real Economy,” NBER Working Paper 20561, 2014.

Kishan, Ruby P., and Timothy P. Opiela, “Bank Size, Bank Capital, and the Bank Lending Channel,” *Journal of Money, Credit and Banking*, 32 (2000), 121–141.

Korinek, Anton, and Alp Simsek, “Liquidity Trap and Excessive Leverage,” *American Economic Review*, 106 (2016), 699–738.

Kurlat, Pablo, and Johannes Stroebel, “Testing for Information Asymmetries in Real Estate Markets,” *Review of Financial Studies*, 28 (2015), 2429–2461.

Lee, David S., and Thomas Lemieux, “Regression Discontinuity Designs in Economics,” *Journal of Economic Literature*, 48 (2010), 281–355.

Livshits, Igor, James C. MacGee, and Michèle Tertilt, “The Democratization of Credit and the Rise in Consumer Bankruptcies,” *Review of Economic Studies*, 83 (2016), 1673–1710.

Mahoney, Neale, and E. Glen Weyl, “Imperfect Competition in Selection Markets,” *Review of Economics and Statistics*, forthcoming, 2017.

Mian, Atif, Kamalesh Rao, and Amir Sufi, “Household Balance Sheets, Consumption, and the Economic Slump,” *Quarterly Journal of Economics*, 128 (2013), 1687–1726.

Mian, Atif, and Amir Sufi, “Household Leverage and the Recession of 2007–09,” *IMF Economic Review*, 58 (2010), 74–117.

———, “House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis,” *American Economic Review*, 101 (2011), 2132–56.

———, “What Explains the 2007–2009 Drop in Employment?” *Econometrica*, 82 (2014), 2197–2223.

Parker, Jonathan, “Why Don’t Households Smooth Consumption? Evidence from a 25 Million Dollar Experiment,” *American Economic Journal: Macroeconomics*, forthcoming, 2017.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland, "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 103 (2013), 2530–2553.
Parlour, Christine A., and Uday Rajan, "Competition in Loan Contracts," *American Economic Review*, 91 (2001), 1311–1328.
Petersen, Mitchell A., and Raghuram G. Rajan, "The Benefits of Lending Relationships: Evidence from Small Business Data," *Journal of Finance*, 49 (1994), 3–37.
Philippon, Thomas, and Virgiliu Midrigan, "Household Leverage and the Recession," NBER Working Paper, 2011.
Sahm, Claudia, Matthew D. Shapiro, and Joel Slemrod, "Balance-Sheet Households and Stimulus: Lessons from the Payroll Tax Cut and Its Expiration," NBER Working Paper 21220, 2015.
Scharfstein, David S., and Adi Sunderam, "Market Power in Mortgage Lending, Refinancing Activity and Mortgage Rates," Working Paper, 2016.
Souleles, Nicholas S., "The Response of Household Consumption to Income Tax Refunds," *American Economic Review*, 89 (1999), 947–958.
Stango, Victor, "Competition and Pricing in the Credit Card Market," *Review of Economics and Statistics*, 82 (2000), 499–508.
Stavins, Joanna, "Can Demand Elasticities Explain Sticky Credit Card Rates?," *New England Economic Review*, July (1996), 43–54.
Stephens, Melvin, "'3rd of the Month': Do Social Security Recipients Smooth Consumption between Checks?," *American Economic Review*, 93 (2003), 406–422.
———, "The Consumption Response to Predictable Changes in Discretionary Income: Evidence from the Repayment of Vehicle Loans," *Review of Economics and Statistics*, 90 (2008), 241–252.
Stroebel, Johannes, "Asymmetric Information about Collateral Values," *Journal of Finance*, 71 (2015), 1071–1112.
Wall Street Journal. "Lending Declines as Bank Jitters Persist," November 25, 2009.
Wooldridge, Jeffrey M., "Cluster-Sample Methods in Applied Econometrics," *American Economic Review*, 93 (2003), 133–138.
Zeldes, Stephen P. "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 97 (1989), 305–346.
Zinman, Jonathan, "Household Debt: Facts, Puzzles, Theories, and Policies." *Annual Review of Economics*, 7 (2015), 251–276.