

Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?*

Sumit Agarwal[†] Souphala Chomsisengphet[‡] Neale Mahoney[§] Johannes Stroebel[¶]

Abstract

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 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.

Keywords: Pass-through, Monetary Policy, Credit Cards, Asymmetric Information

*This version: May 12, 2017. This paper 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.

[†]Georgetown University. Email: ushakri@yahoo.com

[‡]Office of the Comptroller of the Currency. Email: souphala.chomsisengphet@occ.treas.gov

[§]University of Chicago, Booth School of Business and NBER. Email: neale.mahoney@gmail.com

[¶]New York University, Stern School of Business, NBER, and CEPR. Email: johannes.stroebel@nyu.edu

During the Great Recession, policymakers 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 paper, 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 – e.g., 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 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:

$$-\frac{dq}{dc} = \int_i \underbrace{-\frac{dCL_i}{dc}}_{\text{MPL}} \times \underbrace{\frac{dq_i}{dCL_i}}_{\text{MPB}} \quad (1)$$

We operationalize our framework by estimating heterogeneous MPBs and MPLs using panel data on all

¹For 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 non-financial 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 Appendix A for more information.


credit cards issued by the 8 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 ($\text{FICO} \leq 660$), 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 tradeoff. 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.

Figure 1: Pass-Through of Reduction in Cost of Funds into Credit Limits



Note: 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 ($CL^* \rightarrow CL^{**}$). 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 5 for more details.

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 1, 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 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

²This 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.

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).

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 1. 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 1 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 paper as making three contributions. First, our paper 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 (2016), 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, Acharya et al., 2015; Dell'Ariscia, Laeven and Suarez, 2016; Jiménez et al., 2012, 2014; Kashyap and Stein, 1994; Kishan and Opiela, 2000).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 paper 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 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

³Also see Agarwal, Marvell and McGranahan (Forthcoming), Agarwal and Qian (2017), Agarwal et al. (2015a), Agarwal, Liu and Souleles (2007), Agarwal and Qian (2014), Baker (2013), Bhutta and Keys (2016), Blundell, Pistaferri and Preston (2008), Dobbie and Skiba (2013), Gelman et al. (2015), Hsieh (2003), Johnson, Parker and Souleles (2006), Parker (2015), Sahm, Shapiro and Slemrod (2015), Souleles (1999), Stephens (2003, 2008), and Zeldes (1989). Jappelli and Pistaferri (2010) and Zinman (2015) review this literature. See Carroll (1997, 2001) for theoretical foundations.

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

⁵It also relates to recent research by Scharfstein and Sunderam (2013), who show that the pass-through of credit expansion is also affected by regional variation in the competitive environment.

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

important channels for monetary policy transmission.⁷ However, we think that our finding of lower pass-through to less creditworthy borrowers – e.g., 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 paper 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 paper proceeds as follows: Section 1 presents background information on the determinants of credit limits and describes our credit card data. Section 2 discusses our regression discontinuity research design. Section 3 verifies the validity of this research design. Section 4 presents our estimates of the marginal propensity to borrow. Section 5 provides a model of credit limits. Section 6 presents our estimates of the marginal propensity to lend. Section 7 concludes.

1 Background and Data

Our research design exploits quasi-random variation in the credit limits set by credit card lenders (see Section 2). 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.

1.1 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 uti-

⁷According 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.

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

lization, 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 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 foregone 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.

1.2 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 8 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,

⁹The 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.

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.

1.3 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. 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 (pre-approved, 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 post-origination 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 2 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 blue lines; the number of accounts originated are shown with grey 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.

While 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 2 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.

1.4 Summary Statistics

Table 1 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: ≤ 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 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 Appendix Figure A1).

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.

2 Research Design

Our identification strategy exploits the credit limit quasi-experiments identified in Section 1 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, \bar{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:

$$\tau = \frac{\lim_{x \downarrow \bar{x}} E[y|x] - \lim_{x \uparrow \bar{x}} E[y|x]}{\lim_{x \downarrow \bar{x}} E[cl|x] - \lim_{x \uparrow \bar{x}} E[cl|x]}. \quad (2)$$

The denominator is always non-zero because of the known discontinuity in the credit supply function at \bar{x} . The parameter τ identifies the local average treatment effect (LATE) of extending more credit to people with FICO scores in the vicinity of \bar{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 \bar{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, $\tilde{y} \in \{cl, y\}$:

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

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; 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:

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

2.1 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 $FICO_k$, $k = 1, \dots, 4$ be indicator variables that take on a value of 1 when the FICO score of the discontinuity for quasi-experiment j falls into one of our FICO score groups (≤ 660 , 661-700, 701-740, > 740). We recover

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

heterogeneity in treatment effects by regressing τ_j on the FICO score group dummies and controls:

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

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 Appendix Section E, we examine the relationship between our LATEs and FICO scores using non-parametric 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 measuring the precision of our sample average treatment effects for the population averages.

3 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.

3.1 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 3 shows average credit limits at origination within 50 FICO score points of the quasi-experiments

¹²Following 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.

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 3). Panel A of Figure 4 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 4 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 4 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 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 5.4.

3.2 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 3, we 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

¹³For all RD graphs we control for additional discontinuous jumps in credit limits as discussed in Section 2.

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 3 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) and manage exposure to risk (set credit limits).¹⁵ Table 3 shows that, for the average (median) experiment, the APR increases by 1.7 basis points (declines by 0.5 basis points) at the FICO score cutoff; these changes are economically tiny relative to an average APR of 15.4%. Panel E of Figure 3 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 5 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 3, show that there are no discontinuous changes around the cutoff in any of these borrower characteristics.¹⁷

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

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

¹⁵We 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.

¹⁶A 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 3 shows average contract characteristics at all placebo experiments. All characteristics trend smoothly through the FICO score cutoff at banks with no quasi-experiments.

¹⁷Appendix Figure A2 shows similar graphs for six additional borrower characteristics, all of which trend smoothly through the FICO score cutoff.

Panel F of Figure 5 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 3, 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.

4 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.

4.1 Average Borrowing and Spending

We start by presenting basic summary statistics on credit card utilization. The left column of Table 2 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 five 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 (≤ 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 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

¹⁸In 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

Table 2 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.

4.2 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.

Average daily balances. We first examine the effects on ADB on the treated credit card. Panel A of Figure 6 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 7 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 5 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 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 7 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 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.

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.

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 6 and 7 plots the effects on interest bearing debt. Panel B of Table 5 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 paper, 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 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

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 across all credit cards is the primary object of interest for policymakers wanting to stimulate household borrowing and spending. Panel C of Figures 6 and 7 plots the effects on total balances across all cards. Panel C of Table 5 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

¹⁹Positive 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 downpayment).

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 2). Instead, the high FICO score group has on average \$44,813 in credit limits across all of their credit cards (see Table 1), indicating that these households are not credit constrained on the margin.²⁰

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 6 and 7 shows the effect of higher credit limits on cumulative purchase volume on the treated card. Panel D of Table 5 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.

Robustness and additional heterogeneity. In Appendix Section E, we show that the patterns documented above are robust to non-parametric 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 ac-

²⁰The fact that we observe total credit card balances and not total ADB in the credit bureau data (see footnote 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.

counts 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.

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 (≤ 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 policymakers 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.

5 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 the degree of asymmetric information, create heterogeneity in this effect. In Section 6, 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. Appendix Figure A3 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 re-

sponds 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.

5.1 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 Appendix Figure A4). 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; Parlour and Rajan, 2001; Stango, 2000; Stavins, 1996). In contrast to interest rates, credit limits vary significantly over time. Appendix Figure A5 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.

5.2 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 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 to 660, 661 to 700, etc.) 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 non-interest 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 maximize profits.²²

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

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

²¹Importantly, 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 5.1.

²²The 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.

marginal cost:

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

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 4, 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.

5.3 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 this, consider Figure 1 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:

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

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

²³These 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.

holders borrow more on the margin. This is also the reason why the vertical axis in Figure 1 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 disproportionately raise borrowing among consumers with a greater probability of default. This increases the 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,

²⁴See Einav, Finkelstein and Cullen (2010) and Mahoney and Weyl (2016) for a more in-depth discussion of how the slope of marginal costs parameterizes the degree of selection in a market.

²⁵In principle, selection could also be advantageous, with higher credit limits disproportionately 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.

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 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 6, 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.

5.4 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 6 that realized marginal profits at prevailing credit limits were indeed very close to zero, suggesting that

²⁶If greater debt levels reduce the rate of default – e.g., 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.

²⁷In 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.

²⁸While initial credit limits are highly persistent (see Section 3.1), 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.

banks' expectations during our time period were approximately correct. We think this is not 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., Cabral, Geruso and Mahoney, 2014; Einav, Finkelstein and Cullen, 2010; 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 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 $AP(CL)/CL$ has a larger value. Since we find in our data that $AP(CL)/CL$ is larger for lower FICO score borrowers (see Table 2), our finding that the slope of marginal profits is steeper for lower FICO score borrowers is qualitatively robust.

6 Marginal Propensity to Lend

In Section 5, 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 4 to estimate heterogeneity in the MPL.

6.1 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 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 Appendix Section G for additional discussion of how we measure profitability components at the account level.

The middle section of Table 2 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 (≤ 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.

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

The right section of Table 2 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.

6.2 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 infra-marginal borrowing, further increasing chargeoffs and the slope of marginal costs.³¹

Figure 8 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

³⁰When 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.

³¹Mathematically, 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.

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 6 shows the corresponding estimates. Appendix Figure A6 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.

6.3 Slope of Marginal Profits and Components

The top row of Figure 9 considers the effects of increasing credit limits on marginal costs and marginal chargeoffs. For each FICO score group, the grey 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 4-year horizon, although we will show robustness to different time horizons. Columns 1 to 4 of Table 7 show the corresponding estimates, and Panels A and B of Appendix Figure A7 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 (≤ 660), a \$1 increase in credit limits raises cumulative costs over 4 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 5, 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 9 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 7 show the corresponding estimates, and Panels C and D of Appendix Figure A7 present the standard RD plots for the pooled sample of all quasi-experiments. Marginal revenue at prevailing credit limits, shown by the grey bars, is decreasing in FICO score. For the lowest FICO score group, a \$1 increase 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 9 brings these results together into an analysis of cumulative marginal profits at 48 months since account origination.³³ Columns 9 and 10 of Table 7 show the corresponding estimates and Panel E of Appendix Figure A7 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

³²Marginal 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.

³³We 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.

banks from exploiting profitable consumer lending opportunities.³⁴

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 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 $AP(CL)/CL$ has a larger value. Using the values in Table 2 at a 48 month horizon, we calculate that $AP(CL)/CL$ 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.

6.4 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 5).

Figure 10 shows the effect on credit limits of a permanent 1 percentage point decrease in the cost of funds by FICO score group.³⁶ For each FICO score group, we show estimates using 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

³⁴This 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 is greater, and they are being prudent in their approach to lending." (Wall Street Journal, 2009).

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

³⁶While we consider the effect of a uniform 1 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 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.

³⁷In 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 1 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

6.5 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 11 shows the effect of a 1 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 8 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 8 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 1.4). 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 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 Appendix Figure A20,

³⁸Using cumulative flows over different time horizons involves a tradeoff. 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.

³⁹This 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.

we find that banks' cost of funds declined by 0.96 percentage points 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 8, 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 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, 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 of 2008. Parker et al. (2013) estimate a marginal propensity to consume out of these payments of 50 to 90 percent for combined non-durable 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

⁴⁰The 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 non-durable 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.

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.

7 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 8 largest U.S. banks together with 743 credit limit regression discontinuities to estimate the heterogeneity in banks' marginal propensity to lend (MPL) to different households, and heterogeneity in these households' marginal propensity to borrow (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 month after account origination by 59 cents for households with the lowest FICO scores (≤ 660) while having no effect on households with the highest FICO scores (> 740). Banks' MPLs are negatively correlated with these MPBs, with a 1 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.

References

- Acharya, Viral V, Björn Imbierowicz, Sascha Steffen, and Daniel Teichmann.** 2015. "Does the Lack of Financial Stability Impair the Transmission of Monetary Policy?" Mimeo.
- Adams, William, Liran Einav, and Jonathan Levin.** 2009. "Liquidity Constraints and Imperfect Information in Subprime Lending." *American Economic Review*, 99(1): 49–84.
- Agarwal, Sumit, and Wenlan Qian.** 2014. "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore." *American Economic Review*, 104(12): 4205–30.
- Agarwal, Sumit, and Wenlan Qian.** 2017. "Access to Home Equity and Consumption: Evidence from a Policy Experiment." *Review of Economics and Statistics*, 99(1): 40–52.
- Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles.** 2007. "The Reaction of Consumer Spending and Debt to Tax Rebates—Evidence from Consumer Credit Data." *Journal of Political Economy*, 115(6): 986–1019.
- Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tomasz Piskorski, Amit Seru, and Vincent Yao.** 2015a. "Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program." *NBER Working Paper*.
- Agarwal, Sumit, Nathan Marvell, and Leslie McGranahan.** Forthcoming. "Consumption Response to Temporary Tax Incentives: Evidence from State Sales Tax Holidays." *American Economic Journal: Economic Policy*.
- Agarwal, Sumit, Souphala Chomsisengphet, and Chunlin Liu.** 2010. "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(4): 743–754.
- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel.** 2015b. "Regulating Consumer Financial Products: Evidence from Credit Cards." *Quarterly Journal of Economics*, 130(1): 111–164.
- Auclert, Adrien.** 2014. "Monetary Policy and the Redistribution Channel." MIT mimeo.
- Ausubel, Lawrence M.** 1991. "The Failure of Competition in the Credit Card Market." *American Economic Review*, 81(1): 50–81.
- Aydin, Deniz.** 2016. "The Marginal Propensity to Consume out of Liquidity." *Working Paper*.
- Baker, Scott R.** 2013. "Debt and the Consumption Response to Household Income Shocks." *Working Paper*.
- Bhutta, Neil, and Benjamin J. Keys.** 2016. "Interest Rates and Equity Extraction During the Housing Boom." *American Economic Review*, 106(7): 1742–74.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston.** 2008. "Consumption Inequality and Partial Insurance." *American Economic Review*, 98(5): 1887–1921.
- Cabral, Marika, Michael Geruso, and Neale Mahoney.** 2014. "Does Privatized Health Insurance Benefit Patients or Producers? Evidence from Medicare Advantage." *NBER Working Paper*.
- Calem, Paul S., and Loretta J. Mester.** 1995. "Consumer Behavior and the Stickiness of Credit-Card Interest Rates." *American Economic Review*, 85(5): 1327–1336.

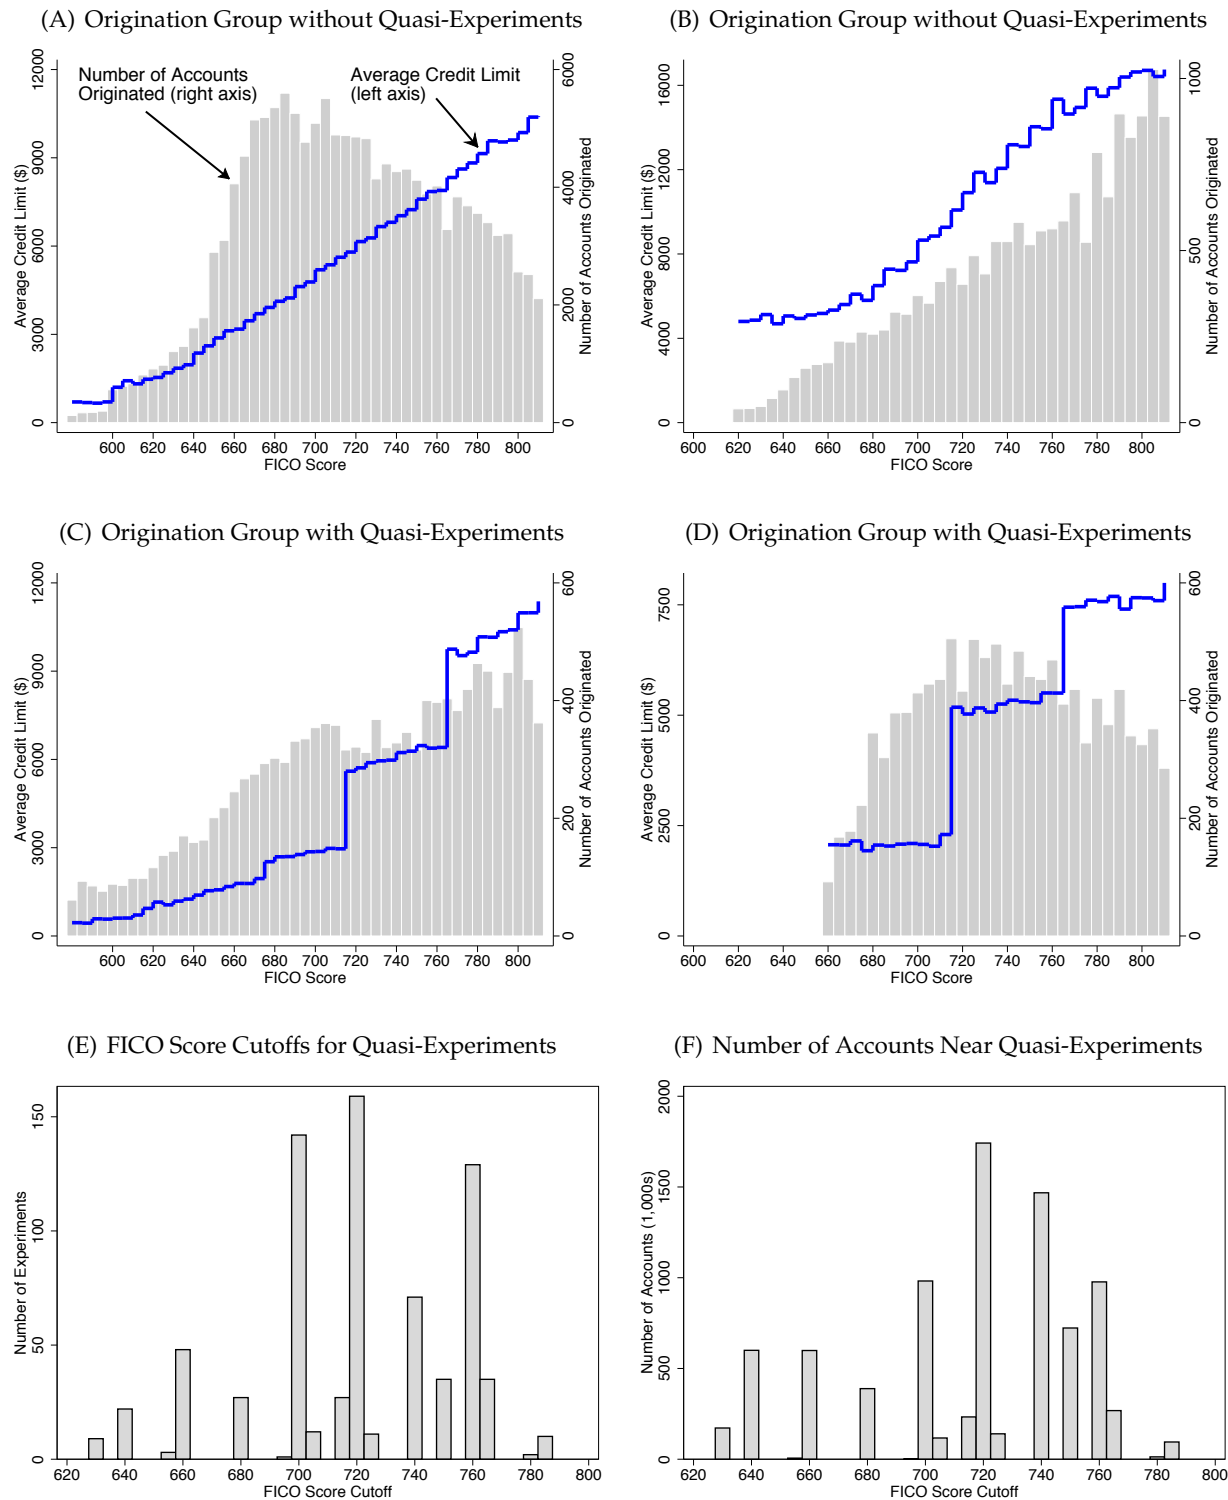
- Carroll, Christopher D.** 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *The Quarterly Journal of Economics*, 112(1): 1–55.
- Carroll, Christopher D.** 2001. "A Theory of the Consumption Function, with and without Liquidity Constraints." *The Journal of Economic Perspectives*, 15(3): 23–45.
- Chakraborty, Indraneel, Itay Goldstein, and Andrew MacKinlay.** 2016. "Monetary Stimulus and Bank Lending." *Working Paper*.
- Chetty, Raj.** 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics*, 1(1): 451–488.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff.** 2014. "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates." *American Economic Review*, 104(9): 2593–2632.
- Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia.** 2012. "Innocent Bystanders? Monetary Policy and Inequality in the US." *NBER Working Paper*.
- Dell’Ariccia, Giovanni, Luc Laeven, and Gustavo Suarez.** 2016. "Bank Leverage and Monetary Policy’s Risk-Taking Channel: Evidence from the United States." *ECB Working Paper*.
- Di Maggio, Marco, Amir Kermani, and Rodney Ramcharan.** 2014. "Monetary Policy Pass-Through: Household Consumption and Voluntary Deleveraging." *Columbia Business School Research Paper*.
- Dobbie, Will, and Paige Marta Skiba.** 2013. "Information Asymmetries in Consumer Credit Markets: Evidence from Payday Lending." *American Economic Journal: Applied Economics*, 5(4): 256–282.
- Doepke, Matthias, and Martin Schneider.** 2006. "Inflation and the Redistribution of Nominal Wealth." *Journal of Political Economy*, 114(6): 1069–1097.
- Drechsler, Itamar, Alexi Savov, and Philipp Schnabl.** 2016. "The Deposits Channel of Monetary Policy." *Working Paper*.
- Eggertsson, Gauti B., and Paul Krugman.** 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *Quarterly Journal of Economics*, 127(3): 1469–1513.
- Einav, Liran, Amy Finkelstein, and Mark R. Cullen.** 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics*, 125(3): 877–921.
- Einav, Liran, Mark Jenkins, and Jonathan Levin.** 2012. "Contract Pricing in Consumer Credit Markets." *Econometrica*, 80(4): 1387–1432.
- Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan.** 2015. "Assessing Sale Strategies in Online Markets Using Matched Listings." *American Economic Journal: Microeconomics*, 7(2): 215–47.
- Fay, Scott, Erik Hurst, and Michelle J. White.** 2002. "The Household Bankruptcy Decision." *American Economic Review*, 92(3): 706–718.
- FDIC.** 2007. "Scoring and Modeling." Federal Deposit Insurance Corporation (FDIC) Credit Card Activities Manual.
- Fuster, Andreas, and Paul Willen.** 2010. "\$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(4).
- GAO.** 2009. "Government Accountability Office: Rising Interchange Fees Have Increased Costs for Merchants, but Options for Reducing Fees Pose Challenges." *GAO-10-45*.

- Geithner, Timothy.** 2009. "Secretary Geithner Introduces Financial Stability Plan."
- Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis.** 2015. "How Individuals Smooth Spending: Evidence from the 2013 Government Shutdown Using Account Data." NBER Working Paper.
- Gross, David B, and Nicholas S. Souleles.** 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics*, 117(1): 149–185.
- Guerrieri, Veronica, and Guido Lorenzoni.** 2011. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *NBER Working Paper*.
- Hackmann, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski.** 2015. "Adverse Selection and an Individual Mandate: When Theory Meets Practice." *American Economic Review*, 105(3): 1030–1066.
- Hall, Robert E.** 2011. "The Long Slump." *American Economic Review*, 101(2): 431–69.
- Hertzberg, Andrew, Andres Liberman, and Daniel Paravisini.** 2016. "Adverse Selection and Maturity Choice in Consumer Credit Markets: Evidence from an Online Lender." *Columbia Business School Research Paper*.
- Hsieh, Chang-Tai.** 2003. "Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund." *American Economic Review*, 93(1): 397–405.
- Hurst, Erik, Benjamin J. Keys, Amit Seru, and Joseph Vavra.** 2016. "Regional Redistribution Through the US Mortgage Market." *American Economic Review*, 106(10): 2982–3028.
- Imbens, Guido, and Karthik Kalyanaraman.** 2011. "Optimal Bandwidth Choice for the Regression Discontinuity Estimator." *Review of Economic Studies*, 79(3): 933–959.
- Jappelli, Tullio, and Luigi Pistaferri.** 2010. "The Consumption Response to Income Changes." *Annual Review of Economics*, 2(1): 479–506.
- Jappelli, Tullio, and Luigi Pistaferri.** 2014. "Fiscal Policy and MPC Heterogeneity." *American Economic Journal: Macroeconomics*, 6(4): 107–136.
- Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina.** 2012. "Credit Supply and Monetary Policy: Identifying the Bank Balance-Sheet Channel with Loan Applications." *American Economic Review*, 102(5): 2301–26.
- Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina.** 2014. "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(2): 463–505.
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles.** 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review*, 96(5): 1589–1610.
- Karlan, Dean, and Jonathan Zinman.** 2009. "Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment." *Econometrica*, 77(6): 1993–2008.
- Kashyap, Anil K, and Jeremy C Stein.** 1994. "Monetary Policy and Bank Lending." In *Monetary policy*. 221–261. The University of Chicago Press.
- Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig.** 2010. "Did Securitization Lead to Lax Screening? Evidence from Subprime Loans." *Quarterly Journal of Economics*, 125(1): 307–362.

- Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vincent Yao.** 2014. "Mortgage Rates, Household Balance Sheets, and the Real Economy." *NBER Working Paper*.
- Kishan, Ruby P, and Timothy P Opiela.** 2000. "Bank Size, Bank Capital, and the Bank Lending Channel." *Journal of Money, Credit and Banking*, 32(1): 121–141.
- Korinek, Anton, and Alp Simsek.** 2016. "Liquidity Trap and Excessive Leverage." *American Economic Review*, 106(3): 699–738.
- Kurlat, Pablo, and Johannes Stroebel.** 2015. "Testing for Information Asymmetries in Real Estate Markets." *Review of Financial Studies*, 28(8): 2429–2461.
- Lee, David S., and Thomas Lemieux.** 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature*, 48(2): 281–355.
- Livshits, Igor, James C. Macgee, and Michèle Tertilt.** 2016. "The Democratization of Credit and the Rise in Consumer Bankruptcies." *The Review of Economic Studies*, 83(4): 1673–1710.
- Mahoney, Neale, and E. Glen Weyl.** 2016. "Imperfect Competition in Selection Markets." *Review of Economics and Statistics*, Forthcoming.
- Mian, Atif, and Amir Sufi.** 2010. "Household Leverage and the Recession of 2007–09." *IMF Economic Review*, 58(1): 74–117.
- Mian, Atif, and Amir Sufi.** 2011. "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis." *American Economic Review*, 101(5): 2132–56.
- Mian, Atif, and Amir Sufi.** 2012. "What Explains High Unemployment? The Aggregate Demand Channel." *NBER Working Paper*.
- Mian, Atif, Kamalesh Rao, and Amir Sufi.** 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics*, 128(4): 1687–1726.
- Office of the Comptroller of the Currency.** 2015. *Comptroller's Handbook: Credit Card Lending*.
- Parker, Jonathan.** 2015. "Why Don't Households Smooth Consumption? Evidence from a 25 Million Dollar Experiment." *NBER Working Paper*.
- Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland.** 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review*, 103(6): 2530–2553.
- Parlour, Christine A., and Uday Rajan.** 2001. "Competition in Loan Contracts." *American Economic Review*, 91(5): pp. 1311–1328.
- Petersen, Mitchell A., and Raghuram G. Rajan.** 1994. "The Benefits of Lending Relationships: Evidence from Small Business Data." *Journal of Finance*, 49(1): 3–37.
- Philippon, Thomas, and Virgiliu Midrigan.** 2011. "Household Leverage and the Recession." *NBER Working Paper*.
- Sahm, Claudia, Matthew D. Shapiro, and Joel Slemrod.** 2015. "Balance-Sheet Households and Stimulus: Lessons from the Payroll Tax Cut and Its Expiration." *NBER Working Paper*.
- Scharfstein, David S., and Adi Sunderam.** 2013. "Concentration in Mortgage Lending, Refinancing Activity and Mortgage Rates." *NBER Working Paper*.
- Souleles, Nicholas S.** 1999. "The Response of Household Consumption to Income Tax Refunds." *American Economic Review*, 89(4): 947–958.

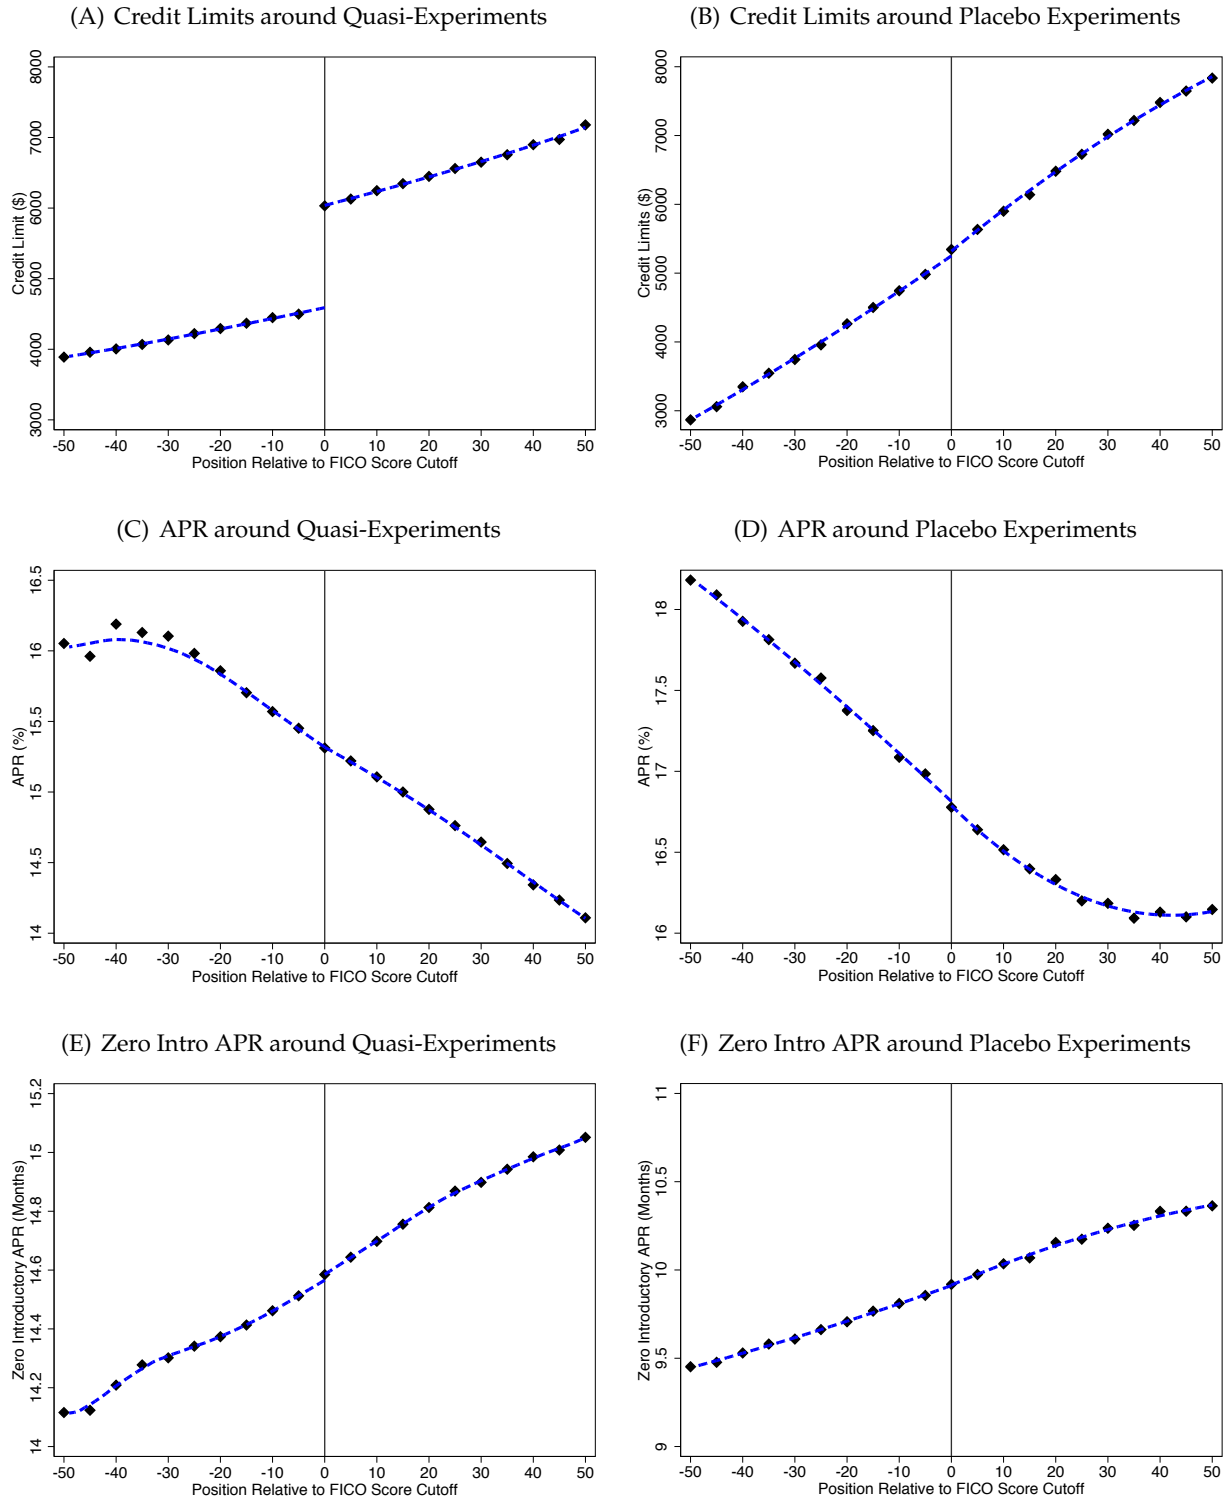
- Stango, Victor.** 2000. "Competition and Pricing in the Credit Card Market." *Review of Economics and Statistics*, 82(3): 499–508.
- Stavins, Joanna.** 1996. "Can Demand Elasticities Explain Sticky Credit Card Rates?" *New England Economic Review*, Jul: 43–54.
- Stephens, Melvin.** 2003. "'3rd of the Month': Do Social Security Recipients Smooth Consumption Between Checks?" *American Economic Review*, 93(1): 406–422.
- Stephens, Melvin.** 2008. "The Consumption Response to Predictable Changes in Discretionary Income: Evidence from the Repayment of Vehicle Loans." *Review of Economics and Statistics*, 90(2): 241–252.
- Stroebel, Johannes.** 2015. "Asymmetric Information About Collateral Values." *Journal of Finance*, 71(3): 1071–1112.
- Wall Street Journal.** 2009. "Lending Declines as Bank Jitters Persist."
- Wooldridge, Jeffrey M.** 2003. "Cluster-Sample Methods in Applied Econometrics." *The American Economic Review*, 93(2): 133–138.
- Zeldes, Stephen P.** 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy*, 97(2): 305–46.
- Zinman, Jonathan.** 2015. "Household Debt: Facts, Puzzles, Theories, and Policies." *Annual Review of Economics*, 7(1): 251–276.

Figure 2: Credit Limit Quasi-Experiments: Examples and Summary Statistics



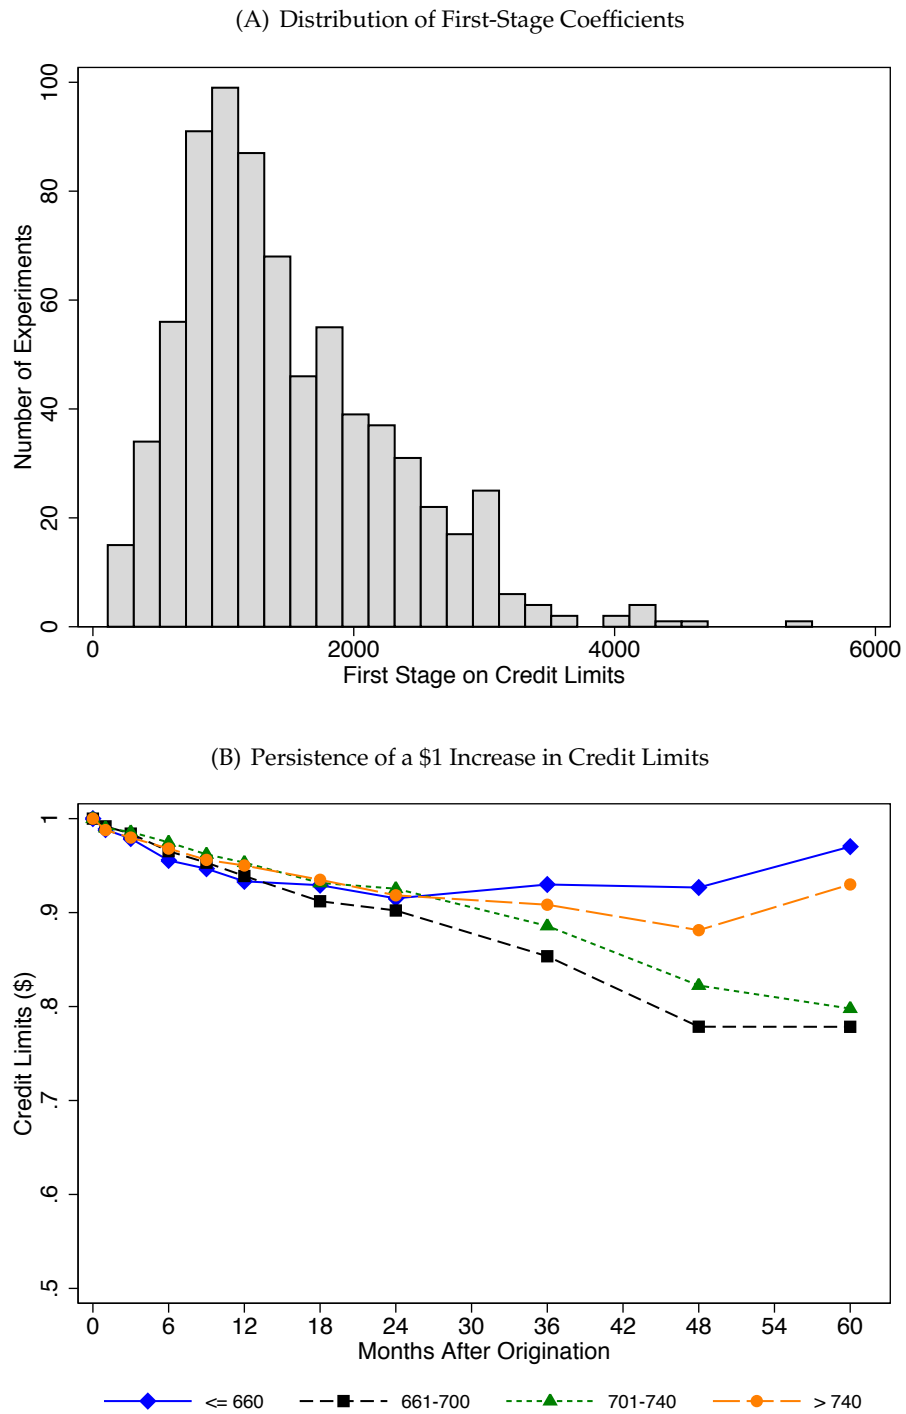
Note: Panels A to D show examples of average credit limits by FICO score for accounts in “origination groups” with and without credit limit quasi-experiments. Origination groups are defined as all credit cards of the same product-type originated by the same bank in the same month through the same loan channel. The horizontal axis shows FICO score at origination. The blue line plots the average credit limit for accounts in FICO score buckets of 5 (left axis); grey 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.

Figure 3: Credit Limits and Cost of Credit Around Credit Limit Quasi-Experiments and Placebo Experiments



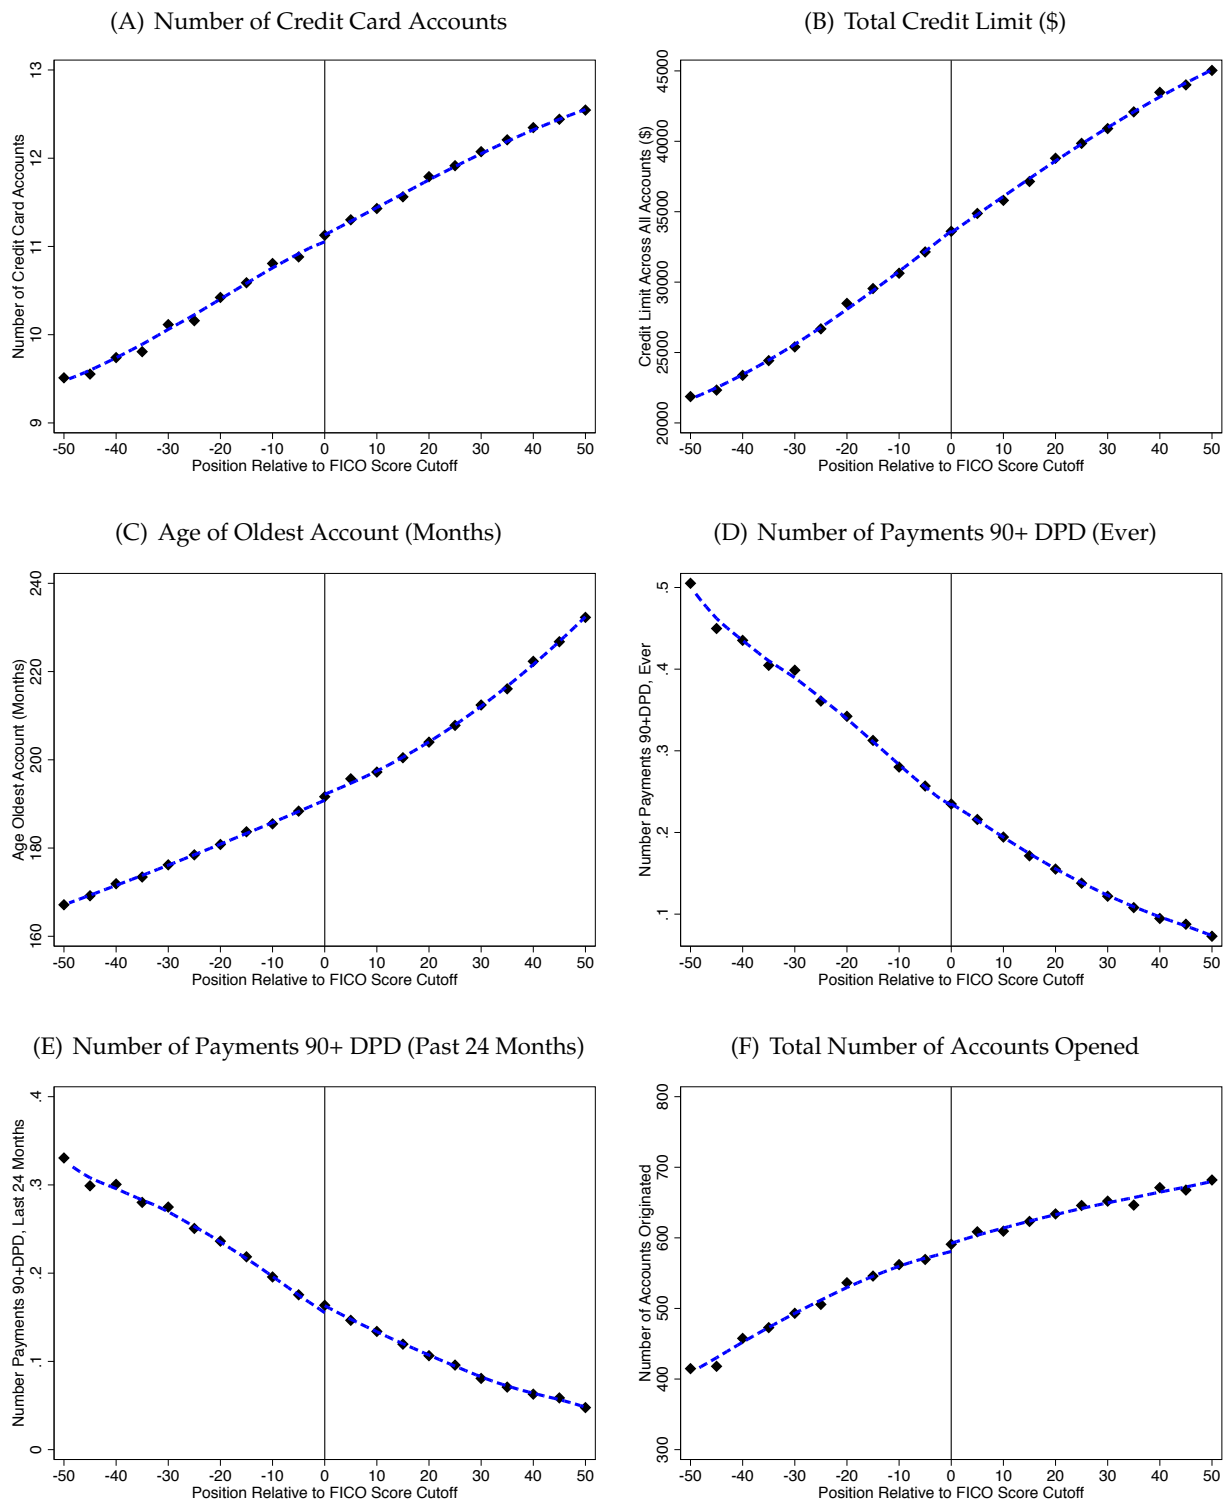
Note: Figure plots average credit limits (Panels A and B), average APR (Panels C and D), and average number of months with zero introductory APR (Panels E and F; limited to originations with zero introductory APR). The left column plots these outcomes around our 743 pooled quasi-experiments. We also 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 cutoff. 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. Blue lines show predicted values from locally linear regressions estimated separately on either side of the cutoff using the Imbens and Kalyanaraman (2011) optimal bandwidth.

Figure 4: Effect of FICO Score Cutoff on Credit Limits




Note: 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 4.

Figure 5: Initial Borrower Characteristics Around Credit Limit Quasi-Experiments



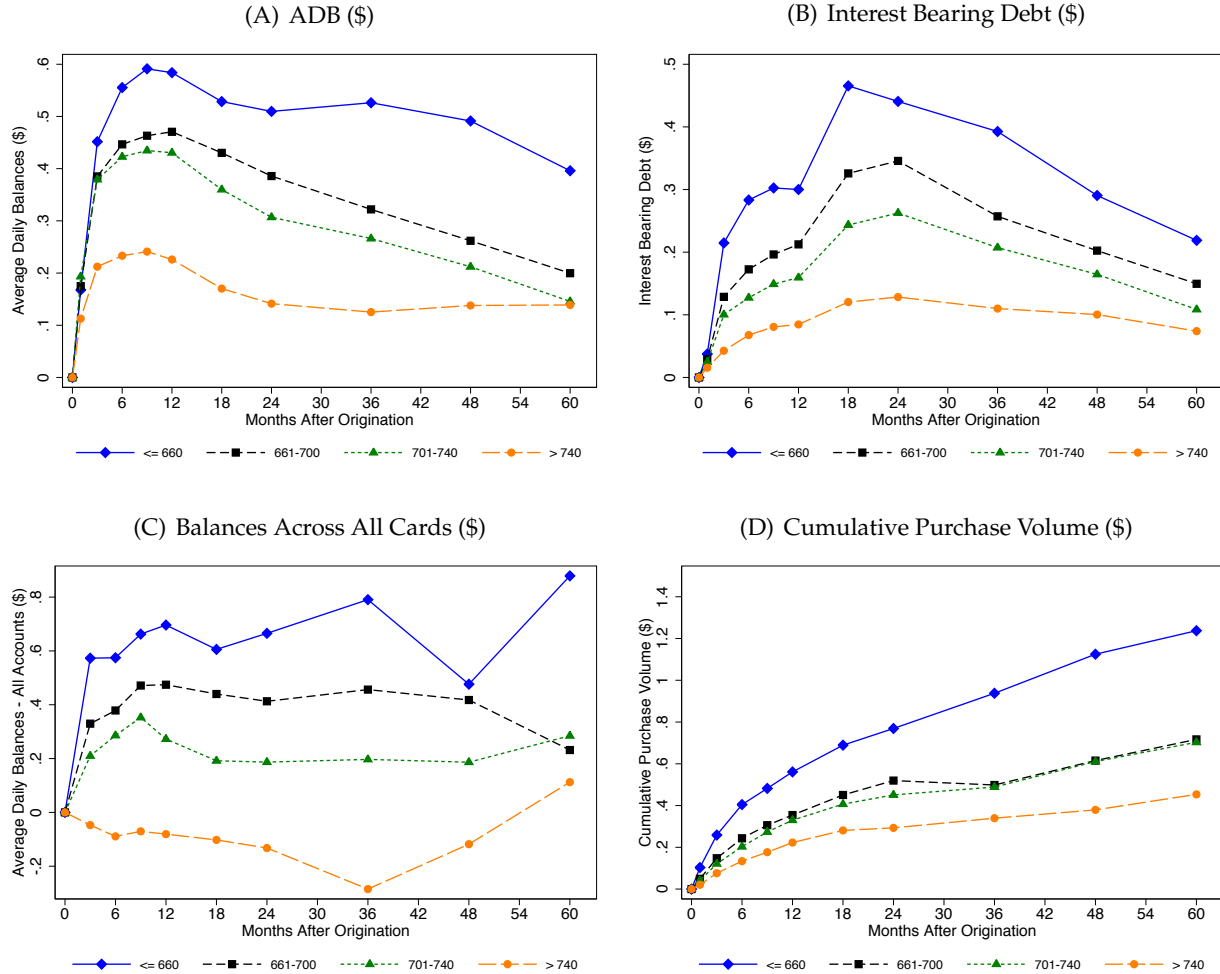
Note: Figure plots average borrower characteristics around our 743 pooled credit limit quasi-experiments. The horizontal axis shows FICO score at origination, centered at the FICO score cutoff. The vertical axis shows the number of credit card accounts (Panel A), total credit limit across all credit 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. Blue lines show predicted values from locally linear regressions estimated separately on either side of the cutoff using the Imbens and Kalyanaraman (2011) optimal bandwidth.

Figure 6: Borrowing and Spending Around Credit Limit Quasi-Experiments



Note: 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 3. 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.


Figure 7: Marginal Propensity to Borrow




Note: 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 5.

Figure 8: Probability of Delinquency

(A) Probability 60+ Days Past Due (%)

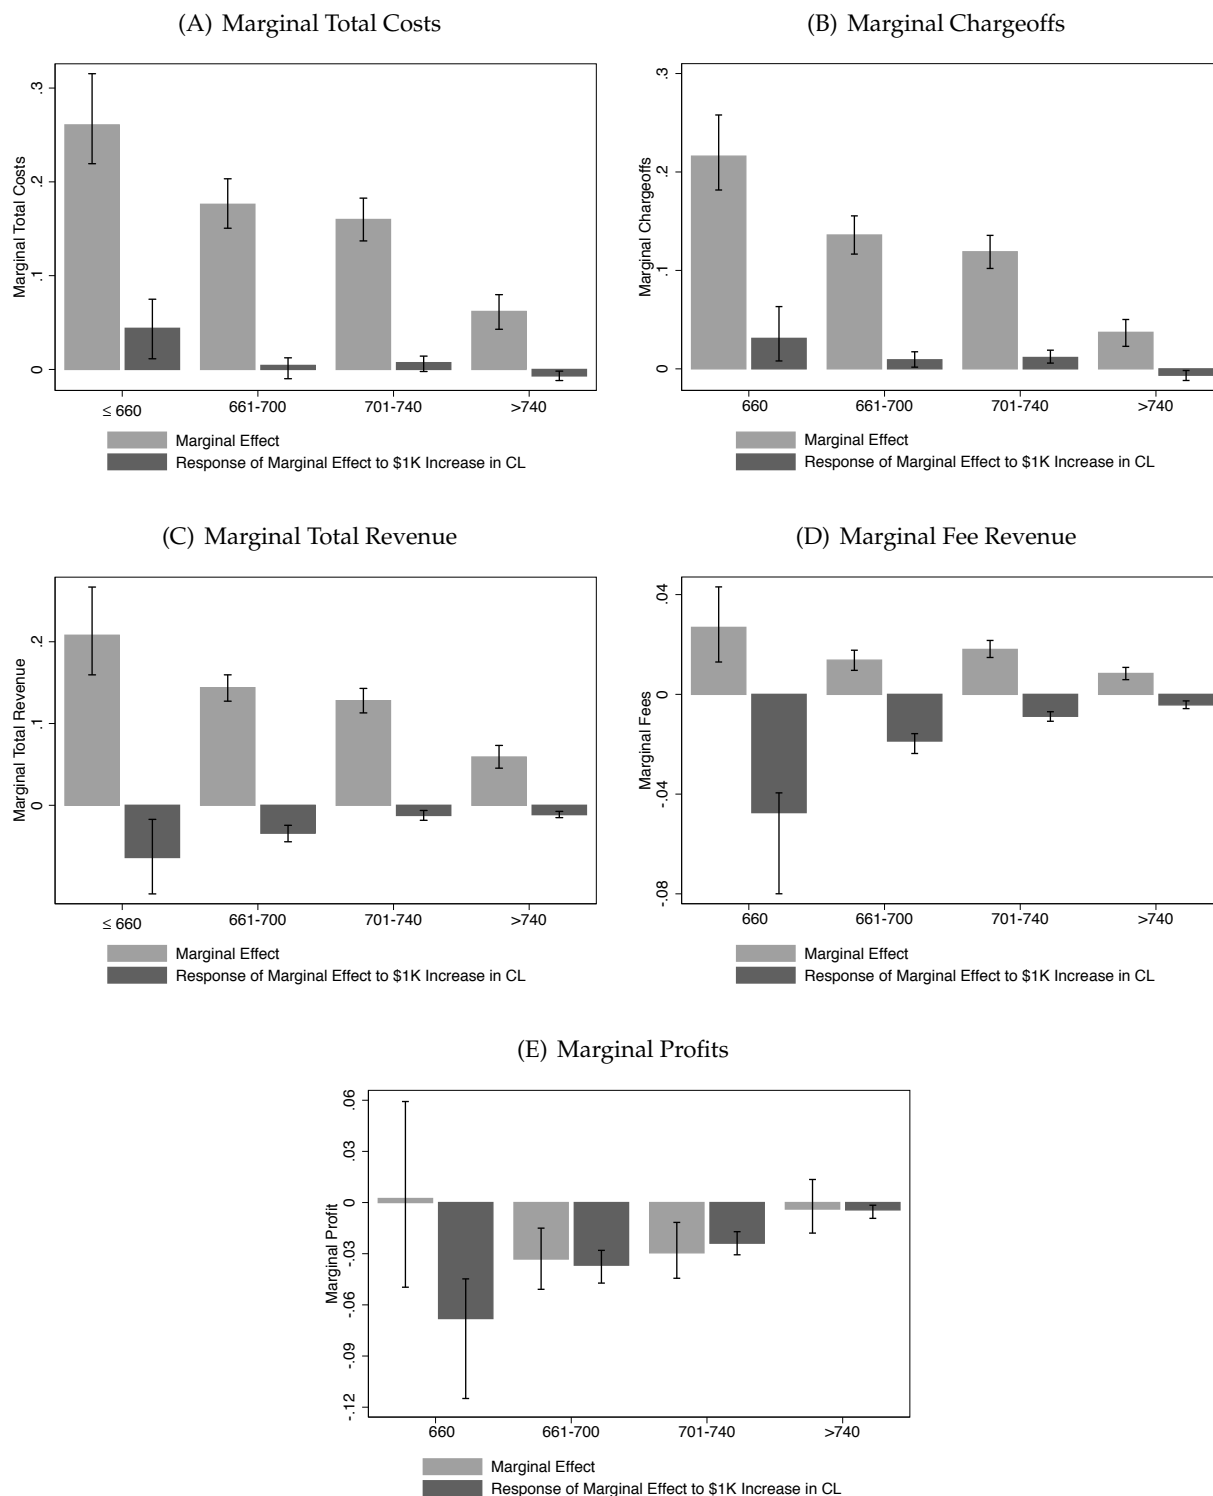


(B) Probability 90+ Days Past Due (%)




Note: 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 6.

Figure 9: Marginal Effects and Response of Marginal Effects to a \$1K Increase in Credit Limits




Note: 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 7.

Figure 10: Marginal Propensity to Lend (MPL)



Note: Figure shows the implied effect of a 1 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 8.

Figure 11: Correlation between MPL and MPB



Note: Panel A shows the implied effect of a 1 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 9 and Table 7, 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 8.

Table 1: Quasi-Experiment-Level Summary Statistics, At Origination

	Average	S.D.		Average	S.D.
Credit Limit on Treated Card (\$)			Total Balances Across All Credit Card Accounts (\$)		
<i>Pooled</i>	5,265	2,045	<i>Pooled</i>	9,551	3,469
<i>≤660</i>	2,561	674	<i>≤660</i>	5,524	2,324
<i>661-700</i>	4,324	1,090	<i>661-700</i>	9,956	2,680
<i>701-740</i>	4,830	1,615	<i>701-740</i>	10,890	3,328
<i>>740</i>	6,941	1,623	<i>>740</i>	9,710	3,326
APR on Treated Card (%)			Credit Limit Across All Credit Card Accounts (\$)		
<i>Pooled</i>	15.38	3.70	<i>Pooled</i>	33,533	14,627
<i>≤660</i>	19.63	5.43	<i>≤660</i>	12,856	5,365
<i>661-700</i>	14.50	3.65	<i>661-700</i>	26,781	7,524
<i>701-740</i>	15.35	3.11	<i>701-740</i>	32,457	8,815
<i>>740</i>	14.70	2.52	<i>>740</i>	44,813	12,828
Number of Credit Card Accounts			Number Times 90+ DPD In Last 24 Months		
<i>Pooled</i>	11.00	2.93	<i>Pooled</i>	0.17	0.30
<i>≤660</i>	7.13	1.18	<i>≤660</i>	0.51	0.31
<i>661-700</i>	10.22	1.68	<i>661-700</i>	0.21	0.16
<i>701-740</i>	11.12	2.34	<i>701-740</i>	0.14	0.10
<i>>740</i>	12.63	2.92	<i>>740</i>	0.05	0.08
Age Oldest Account (Months)			Number Accounts Currently 90+DPD		
<i>Pooled</i>	190.1	29.1	<i>Pooled</i>	0.03	0.03
<i>≤660</i>	162.0	26.3	<i>≤660</i>	0.10	0.05
<i>661-700</i>	180.1	19.9	<i>661-700</i>	0.02	0.02
<i>701-740</i>	184.7	24.0	<i>701-740</i>	0.02	0.02
<i>>740</i>	208.6	25.7	<i>>740</i>	0.01	0.01

Note: 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.

Table 2: Quasi-Experiment-Level Summary Statistics, Post Origination

	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
Credit Limit (\$)					Cumulative Purchase Volume (\$)					Cumulative Cost of Funds (\$)				
<i>After 12 Months</i>	2,617	4,370	4,964	6,980	<i>After 12 Months</i>	2,212	2,579	2,514	2,943	<i>After 12 Months</i>	14	16	16	15
<i>After 24 Months</i>	2,414	4,306	4,946	7,071	<i>After 24 Months</i>	2,447	3,956	3,791	4,374	<i>After 24 Months</i>	23	29	28	25
<i>After 36 Months</i>	2,301	4,622	5,047	7,005	<i>After 36 Months</i>	3,240	5,023	4,253	4,521	<i>After 36 Months</i>	28	38	36	31
<i>After 48 Months</i>	2,252	4,525	4,985	6,944	<i>After 48 Months</i>	3,741	5,154	4,919	4,845	<i>After 48 Months</i>	31	43	41	34
<i>After 60 Months</i>	2,290	4,449	4,601	6,839	<i>After 60 Months</i>	4,524	5,598	5,121	5,626	<i>After 60 Months</i>	33	46	44	36
ADB (\$)					Cumulative Total Costs (\$)					Cumulative Total Revenue (\$)				
<i>After 12 Months</i>	1,260	2,160	2,197	2,101	<i>After 12 Months</i>	122	172	169	147	<i>After 12 Months</i>	233	192	181	175
<i>After 24 Months</i>	1,065	1,794	1,719	1,524	<i>After 24 Months</i>	281	451	433	304	<i>After 24 Months</i>	474	503	439	347
<i>After 36 Months</i>	1,164	1,734	1,481	1,343	<i>After 36 Months</i>	459	710	644	395	<i>After 36 Months</i>	740	793	663	449
<i>After 48 Months</i>	1,079	1,501	1,260	1,064	<i>After 48 Months</i>	588	845	808	488	<i>After 48 Months</i>	953	971	863	563
<i>After 60 Months</i>	1,050	1,465	1,097	1,084	<i>After 60 Months</i>	712	962	901	583	<i>After 60 Months</i>	1,148	1,126	965	669
Average Interest Bearing Debt (\$)					Cumulative Chargeoffs (\$)					Cumulative Interest Charge Revenue (\$)				
<i>After 12 Months</i>	864	903	811	672	<i>After 12 Months</i>	47	67	61	35	<i>After 12 Months</i>	106	61	52	42
<i>After 24 Months</i>	1,040	1,676	1,557	1,294	<i>After 24 Months</i>	178	259	245	124	<i>After 24 Months</i>	297	295	243	159
<i>After 36 Months</i>	1,068	1,615	1,344	1,135	<i>After 36 Months</i>	306	443	403	190	<i>After 36 Months</i>	484	520	420	243
<i>After 48 Months</i>	1,044	1,416	1,144	924	<i>After 48 Months</i>	403	552	524	261	<i>After 48 Months</i>	625	669	578	340
<i>After 60 Months</i>	1,020	1,388	1,001	941	<i>After 60 Months</i>	483	634	602	322	<i>After 60 Months</i>	760	794	657	429
Cumulative Prob Positive Interest Bearing Debt (%)					Cumulative Prob 60+ DPD (%)					Cumulative Fee Revenue (\$)				
<i>After 12 Months</i>	58.4%	36.1%	31.6%	26.9%	<i>After 12 Months</i>	6.4%	4.1%	3.6%	1.6%	<i>After 12 Months</i>	73	79	79	74
<i>After 24 Months</i>	75.4%	73.0%	64.9%	50.3%	<i>After 24 Months</i>	12.0%	9.3%	8.2%	3.8%	<i>After 24 Months</i>	129	129	121	101
<i>After 36 Months</i>	84.0%	79.4%	72.3%	61.6%	<i>After 36 Months</i>	15.1%	12.2%	10.9%	5.2%	<i>After 36 Months</i>	192	173	157	116
<i>After 48 Months</i>	87.4%	84.0%	78.1%	69.8%	<i>After 48 Months</i>	16.5%	13.6%	12.2%	5.9%	<i>After 48 Months</i>	254	199	187	126
<i>After 60 Months</i>	90.1%	86.3%	81.3%	75.2%	<i>After 60 Months</i>	17.2%	14.4%	12.9%	6.2%	<i>After 60 Months</i>	364	310	211	101
Total Balances Across All Cards (\$)					Cumulative Prob 90+ DPD (%)					Cumulative Profits (\$)				
<i>After 12 Months</i>	6,155	10,546	11,411	10,528	<i>After 12 Months</i>	4.8%	3.3%	2.9%	1.3%	<i>After 12 Months</i>	111	21	12	30
<i>After 24 Months</i>	5,919	10,521	11,307	10,703	<i>After 24 Months</i>	10.2%	8.1%	7.2%	3.2%	<i>After 24 Months</i>	194	56	9	46
<i>After 36 Months</i>	6,387	10,716	11,702	11,267	<i>After 36 Months</i>	13.2%	10.9%	9.7%	4.5%	<i>After 36 Months</i>	281	91	23	59
<i>After 48 Months</i>	6,698	10,437	11,665	11,137	<i>After 48 Months</i>	14.5%	12.2%	10.9%	5.1%	<i>After 48 Months</i>	365	126	55	75
<i>After 60 Months</i>	7,566	10,591	11,972	12,490	<i>After 60 Months</i>	15.2%	12.9%	11.5%	5.4%	<i>After 60 Months</i>	436	164	63	87

Note: 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.

Table 3: Validity of Research Design: Discontinuous Increase at FICO Cutoff

	Distribution of Jump Across Quasi-Experiments			Baseline
	Average	Median	Standard Deviation	
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

Note: 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.

Table 4: Persistence of Credit Limit Effect

	Months After Account Origination				
	12	24	36	48	60
<i>FICO</i>					
<i>≤660</i>	0.93 [0.91, 0.96]	0.92 [0.87, 0.96]	0.93 [0.86, 0.99]	0.93 [0.83, 1.04]	0.97 [0.79, 1.19]
<i>661-700</i>	0.94 [0.92, 0.95]	0.90 [0.88, 0.93]	0.85 [0.81, 0.88]	0.78 [0.71, 0.85]	0.78 [0.67, 0.94]
<i>701-740</i>	0.95 [0.94, 0.97]	0.93 [0.9, 0.94]	0.89 [0.86, 0.9]	0.82 [0.77, 0.87]	0.80 [0.67, 0.9]
<i>>740</i>	0.95 [0.94, 0.96]	0.92 [0.9, 0.94]	0.91 [0.87, 0.93]	0.88 [0.81, 0.94]	0.93 [0.83, 1.11]

Note: 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.

Table 5: Marginal Propensity to Borrow

	Months After Account Origination				
	12	24	36	48	60
Panel A: Average Daily Balance					
<i>FICO</i>					
≤ 660	0.58 [0.54, 0.63]	0.51 [0.46, 0.57]	0.53 [0.46, 0.59]	0.49 [0.39, 0.58]	0.40 [0.32, 0.48]
661-700	0.47 [0.44, 0.49]	0.39 [0.35, 0.41]	0.32 [0.28, 0.35]	0.26 [0.22, 0.3]	0.20 [0.15, 0.25]
701-740	0.43 [0.4, 0.45]	0.31 [0.28, 0.33]	0.27 [0.23, 0.29]	0.21 [0.18, 0.25]	0.15 [0.1, 0.2]
>740	0.23 [0.2, 0.25]	0.14 [0.11, 0.17]	0.13 [0.09, 0.16]	0.14 [0.09, 0.18]	0.14 [0.08, 0.2]
Panel B: Interest Bearing Debt					
<i>FICO</i>					
≤ 660	0.30 [0.26, 0.35]	0.46 [0.42, 0.51]	0.40 [0.36, 0.45]	0.35 [0.31, 0.4]	0.33 [0.28, 0.38]
661-700	0.21 [0.19, 0.23]	0.34 [0.32, 0.37]	0.30 [0.27, 0.32]	0.28 [0.25, 0.31]	0.27 [0.24, 0.3]
701-740	0.16 [0.14, 0.18]	0.27 [0.24, 0.29]	0.23 [0.2, 0.26]	0.21 [0.18, 0.23]	0.19 [0.17, 0.22]
>740	0.08 [0.07, 0.1]	0.13 [0.1, 0.15]	0.12 [0.09, 0.14]	0.12 [0.09, 0.14]	0.11 [0.09, 0.13]
Panel C: Total Balance Across All Cards					
<i>FICO</i>					
≤ 660	0.59 [0.34, 0.84]	0.54 [0.18, 0.94]	1.00 [0.51, 1.48]	0.96 [0.12, 1.97]	1.27 [-0.16, 2.51]
661-700	0.46 [0.31, 0.59]	0.42 [0.26, 0.58]	0.48 [0.26, 0.7]	0.59 [0.13, 0.97]	0.43 [-0.42, 1.11]
701-740	0.32 [0.16, 0.47]	0.21 [0.03, 0.37]	0.24 [0.03, 0.44]	0.35 [0, 0.65]	0.49 [-0.52, 1.3]
>740	-0.05 [-0.15, 0.08]	-0.08 [-0.26, 0.1]	-0.19 [-0.47, 0.08]	0.05 [-0.42, 0.39]	0.29 [-0.47, 1.03]
Panel D: Cumulative Purchase Volume					
<i>FICO</i>					
≤ 660	0.56 [0.48, 0.67]	0.77 [0.61, 0.94]	0.94 [0.68, 1.2]	1.12 [0.7, 1.51]	1.24 [0.75, 1.69]
661-700	0.35 [0.31, 0.4]	0.52 [0.44, 0.59]	0.50 [0.38, 0.58]	0.62 [0.47, 0.75]	0.72 [0.52, 0.91]
701-740	0.33 [0.29, 0.38]	0.45 [0.38, 0.52]	0.49 [0.38, 0.59]	0.61 [0.45, 0.78]	0.70 [0.5, 0.95]
>740	0.22 [0.19, 0.26]	0.29 [0.23, 0.36]	0.34 [0.26, 0.42]	0.38 [0.25, 0.51]	0.45 [0.21, 0.71]

Note: 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.

Table 6: Probability of Delinquency

	Months After Account Origination				
	12	24	36	48	60
Panel A: 60+ Days Past Due (%)					
<i>FICO</i>					
<i>≤660</i>	0.10 [-0.46, 0.67]	0.57 [-0.16, 1.3]	0.71 [-0.13, 1.63]	1.21 [0.25, 2.06]	1.38 [0.61, 2.19]
<i>661-700</i>	0.23 [-0.05, 0.55]	0.57 [0.23, 0.93]	0.67 [0.29, 1.06]	0.86 [0.44, 1.24]	0.87 [0.46, 1.27]
<i>701-740</i>	0.28 [0.04, 0.52]	0.55 [0.23, 0.86]	0.63 [0.24, 0.99]	0.68 [0.3, 1.04]	0.67 [0.3, 1.01]
<i>>740</i>	-0.19 [-0.39, -0.04]	-0.16 [-0.46, 0.05]	-0.20 [-0.49, 0.06]	-0.32 [-0.61, -0.05]	-0.33 [-0.61, -0.08]
Panel B: 90+ Days Past Due (%)					
<i>FICO</i>					
<i>≤660</i>	0.10 [-0.37, 0.67]	0.22 [-0.6, 0.94]	0.77 [-0.08, 1.61]	1.16 [0.34, 1.88]	1.07 [0.32, 1.79]
<i>661-700</i>	0.19 [-0.02, 0.45]	0.42 [0.14, 0.78]	0.68 [0.35, 1.05]	0.80 [0.5, 1.18]	0.74 [0.45, 1.11]
<i>701-740</i>	0.26 [0.08, 0.48]	0.62 [0.31, 0.91]	0.73 [0.37, 1.02]	0.74 [0.38, 1.04]	0.74 [0.41, 1.07]
<i>>740</i>	-0.06 [-0.23, 0.08]	-0.07 [-0.33, 0.12]	-0.06 [-0.34, 0.17]	-0.17 [-0.45, 0.07]	-0.19 [-0.44, 0.05]

Note: 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.

Table 7: 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 [0.261, 0.332]	0.033 [0.01, 0.061]	0.216 [0.182, 0.258]	0.031 [0.008, 0.063]	0.238 [0.197, 0.307]	-0.051 [-0.071, -0.032]	0.027 [0.013, 0.043]	-0.047 [-0.08, -0.04]	0.002 [-0.05, 0.059]	-0.068 [-0.115, -0.045]
661-700	0.185 [0.164, 0.204]	0.005 [-0.004, 0.013]	0.136 [0.117, 0.155]	0.009 [0.002, 0.017]	0.146 [0.133, 0.16]	-0.034 [-0.042, -0.026]	0.014 [0.01, 0.018]	-0.019 [-0.024, -0.016]	-0.033 [-0.051, -0.015]	-0.037 [-0.047, -0.028]
701-740	0.165 [0.146, 0.181]	0.011 [0.004, 0.017]	0.119 [0.102, 0.136]	0.012 [0.006, 0.019]	0.130 [0.118, 0.143]	-0.014 [-0.02, -0.009]	0.018 [0.015, 0.022]	-0.009 [-0.011, -0.007]	-0.029 [-0.044, -0.012]	-0.024 [-0.031, -0.017]
>740	0.060 [0.044, 0.072]	-0.011 [-0.017, -0.006]	0.037 [0.023, 0.05]	-0.006 [-0.012, -0.002]	0.055 [0.043, 0.065]	-0.016 [-0.021, -0.011]	0.008 [0.006, 0.011]	-0.004 [-0.006, -0.003]	-0.004 [-0.018, 0.014]	-0.004 [-0.009, -0.002]

Note: 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.

Table 8: Marginal Propensity to Lend \times Marginal Propensity to Borrow

	MPL	MPB Across All Cards				
		12 Months	24 Months	36 Months	48 Months	60 Months
FICO						
≤ 660	253 [160, 372]	0.59 [0.34, 0.84]	0.54 [0.18, 0.94]	1.00 [0.51, 1.48]	0.96 [0.12, 1.97]	1.27 [-0.16, 2.51]
661-700	304 [231, 397]	0.46 [0.31, 0.59]	0.42 [0.26, 0.58]	0.48 [0.26, 0.7]	0.59 [0.13, 0.97]	0.43 [-0.42, 1.11]
701-740	427 [329, 603]	0.32 [0.16, 0.47]	0.21 [0.03, 0.37]	0.24 [0.03, 0.44]	0.35 [0, 0.65]	0.49 [-0.52, 1.3]
> 740	1224 [329, 5300]	-0.05 [-0.15, 0.08]	-0.08 [-0.26, 0.1]	-0.19 [-0.47, 0.08]	0.05 [-0.42, 0.39]	0.29 [-0.47, 1.03]
Weighted Average	655	0.28	0.23	0.33	0.45	0.62
		MPL X 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

Note: 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 1 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 5 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 \times weighted average MPB. The estimates that account for this correlation are constructed by first calculating MPL \times 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 1.4 and Appendix Figure A1).

Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?

Online Appendix

Sumit Agarwal Souphala Chomsisengphet Neale Mahoney Johannes Stroebel

A Review of Policy Interventions Partially Aimed at Stimulating Lending

This appendix describes policy interventions during the Great Recession that were at least partially aimed at encouraging more consumer lending. We analyze the objectives of policies in both the U.S. and in Europe. For the U.S., we consider programs aimed at improving banks' ability to cheaply refinance themselves in short-term funding markets, such as the Term Asset-Backed Securities Loan Facility (TALF) program and the Term Auction Facility (TAF) program (Section A.1). We also discuss programs created to increase the availability of affordable capital for U.S. banks (Section A.2), such as the Capital Purchase Program (CCP) and the Capital Assistance Program (CAP). We document that these programs had at least the partial objective of increasing credit availability for U.S. households. We also discuss the "Funding for Lending Scheme" at the Bank of England (Section A.3) and the Targeted Longer-Term Refinancing Operations (TLTRO) at the European Central Bank (Section A.4).

A.1 U.S. programs focused on short-term funding markets

In the U.S., a number of programs were set up with the explicit aim of increasing credit availability for households and firms by reducing the costs at which financial institutions could refinance themselves in short-term funding markets. These programs can be viewed within the framework in Section 5 as attempts to reduce the cost of funds, *c*.

The **Term Asset-Backed Securities Loan Facility (TALF)** was announced on November 25, 2008, and was aimed at supporting the issuance of asset-backed securities (ABS) collateralized by student loans, auto loans, credit card loans, and loans guaranteed by the Small Business Administration. Under TALF, the Federal Reserve Bank of New York lent up to \$200 billion (later expanded to \$1 trillion) to holders of certain AAA-rated ABS backed by newly and recently originated consumer and small business loans. The following sources discuss the anticipated impact of this program on the total supply of credit available to the population. They document that an increase in credit availability (and thus borrowing volumes) was a key policy goal of TALF.

- (A) *Board of Governors of the Federal Reserve System, Press Release, November 25, 2008*: "The Federal Reserve Board on Tuesday announced the creation of the Term Asset-Backed Securities Loan Facility (TALF), a facility that will help market participants meet the credit needs of households and small businesses by supporting the issuance of asset-backed securities (ABS) collateralized by student loans, auto loans, credit card loans, and loans guaranteed by the Small Business Administration (SBA). [...] New issuance of ABS declined precipitously in September and came to a halt in October. At the same time, interest rate spreads on AAA-rated tranches of ABS soared to levels well outside the range of historical experience, reflecting unusually high risk premiums. The ABS markets historically have funded a substantial share of consumer credit and SBA-guaranteed small

business loans. Continued disruption of these markets could significantly limit the availability of credit to households and small businesses and thereby contribute to further weakening of U.S. economic activity. The TALF is designed to increase credit availability and support economic activity by facilitating renewed issuance of consumer and small business ABS at more normal interest rate spreads." [\[Link\]](#)

- (B) *Janet L. Yellen, President and CEO, Federal Reserve Bank of San Francisco, Presentation to the Annual AEA/ASSA Conference, January 4, 2009:* "For example, the new Term Asset-Backed Securities Loan Facility (TALF) is a joint program between the Federal Reserve and the Treasury, using TARP funds, and is designed to improve the flow of credit to households and businesses." [\[Link\]](#)
- (C) *Testimony by Elizabeth A. Duke, Member of the Board of Governors of the Federal Reserve, "Credit availability and prudent lending standards," Committee on Financial Services, U.S. House of Representatives, March 25, 2009:* "[T]he Federal Reserve and the Treasury recently launched the Term Asset-Backed Securities Loan Facility (TALF) to facilitate the extension of credit to households and small businesses." [\[Link\]](#)
- (D) *U.S. Department of the Treasury website, "Credit Market Programs." Accessed July 7, 2015:* "The Term Asset-Backed Securities Loan Facility (TALF) is a joint program with the Federal Reserve. The program was launched in March 2009 with the aim of helping to restart the asset-backed securitization (ABS) markets that provide credit to consumers and small businesses. The financial crisis severely impacted these markets. Under this program, the Federal Reserve Bank of New York made non-recourse loans to buyers of AAA-rated asset-backed securities to help stimulate consumer and business lending. Treasury used TARP funds to provide credit support for these loans." [\[Link\]](#)

Similarly, a somewhat more general program – the **Term Auction Facility (TAF)** – was set up to provide short-term collateralized loans to U.S. financial institutions that are judged to be in sound financial condition by their local reserve banks. TAF ran between December 17, 2007 and March 8, 2010. The Fed described the aims of this program as below:

- (A) *Board of Governors of the Federal Reserve System, Press Release, October 6, 2008:* "Consistent with this increased scope, the Federal Reserve also announced today additional actions to strengthen its support of term lending markets. Specifically, the Federal Reserve is substantially increasing the size of the Term Auction Facility (TAF) auctions, beginning with today's auction of 84-day funds. These auctions allow depository institutions to borrow from the Federal Reserve for a fixed term against the same collateral that is accepted at the discount window; the rate is established in the auction, subject to a minimum set by the Federal Reserve. In addition, the Federal Reserve and the Treasury Department are consulting with market participants on ways to provide additional support for term unsecured funding markets. Together these actions should encourage term lending across a range of financial markets in a manner that eases pressures and promotes the ability of firms and households to obtain credit." [\[Link\]](#)

A.2 U.S. programs focused on level and cost of bank capital

In addition to programs aimed at providing liquidity through improving the state of short-term funding markets, U.S. policies also focused on improving the capital position of U.S. banks. Two important

programs with that objective, both using resources of the Troubled Asset Relief Program (TARP), were the Capital Purchase Program (CPP) and the Capital Assistance Program (CAP).

Under the first program, the CCP, nine financial institutions received new capital injections on October 28, 2008, with 42 other institutions participating in the CPP through purchases made on November 14, 2008 and November 21, 2008.

- (A) *U.S. Department of the Treasury website, "Capital Purchase Program."* Accessed August 3, 2015: "The Capital Purchase Program (CPP) was launched to stabilize the financial system by providing capital to viable financial institutions of all sizes throughout the nation. Without a viable banking system, lending to businesses and consumers could have frozen and the financial crisis might have spiraled further out of control. Based on market indicators at the time, it became clear that financial institutions needed additional capital to absorb losses and restart the flow of credit to businesses and consumers. In this context, immediate capital injections into financial institutions were necessary to avert a potential collapse of the system." [\[Link\]](#)
- (B) *Statement by Secretary Henry M. Paulson, Jr. on Capital Purchase Program, October 20, 2008:* "We expect all participating banks to continue to strengthen their efforts to help struggling homeowners who can afford their homes avoid foreclosure. Foreclosures not only hurt the families who lose their homes, they hurt neighborhoods, communities and our economy as a whole. [...] Our purpose is to increase confidence in our banks and increase the confidence of our banks, so that they will deploy, not hoard their capital. And we expect them to do so, as increased confidence will lead to increased lending. This increased lending will benefit the U.S. economy and the American people." [\[Link\]](#)
- (C) *Testimony by Elizabeth A. Duke, Member of the Board of Governors of the Federal Reserve, "Credit availability and prudent lending standards," Committee on Financial Services, U.S. House of Representatives, March 25, 2009:* "The U.S. Treasury, the Federal Deposit Insurance Corporation (FDIC), and the Federal Reserve have taken a number of actions to strengthen the financial sector and to promote the availability of credit to businesses and households. This included injecting additional capital into banks, increasing FDIC deposit coverage, providing guarantees of selected senior bank obligations and noninterest-bearing deposits, and establishing new liquidity facilities to financial markets." [\[Link\]](#)

The Treasury's Financial Stability Plan also included an element to improve the capital position of U.S. banks – the Capital Assistance Program (CAP).

- (A) *Remarks by Treasury Secretary Timothy Geithner, "Introducing the Financial Stability Plan," February 10, 2009:* "First, we're going to require banking institutions to go through a carefully designed comprehensive stress test, to use the medical term. We want their balance sheets cleaner, and stronger. [...] Those institutions that need additional capital will be able to access a new funding mechanism that uses funds from the Treasury as a bridge to private capital. 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." [\[Link\]](#)

A.3 U.K. Funding for Lending Scheme

Programs aimed at increasing lending of banks to households and firms were not limited to the U.S.; in the U.K., the Bank of England's "Funding for Lending Scheme" (FLS) was set up precisely with the purpose of encouraging banks to pass through credit expansions to households and firms:

(A) *Bank of England, News Release, July 13, 2012*: "The FLS is designed to boost lending to the real economy. Banks and building societies that increase lending to UK households and businesses will be able to borrow more in the FLS, and do so at lower cost than those that scale back lending. [...] The FLS is designed to encourage broad participation so that as many institutions as possible have incentives to lend more to the UK real economy through, for example, business loans and residential mortgages, than they otherwise would have. [...] Commenting on the launch of the Scheme, the Governor of the Bank of England said: [...] 'That will encourage banks to make loans to families and businesses both cheaper and more easily available'. The Chancellor of the Exchequer said: 'Today's announcements aim to make mortgages and loans cheaper and more easily available, providing welcome support to businesses that want to expand and families aspiring to own their own home. The Treasury and the Bank of England are taking coordinated action to inject new confidence into our financial system and support the flow of credit to where it is needed in the real economy.'" [\[Link\]](#)

(B) *Spencer Dale, Executive Director, Monetary Policy, and Chief Economist, Bank of England, "Limits of Monetary Policy," September 8, 2012*: "Most recently, the Bank, together with the Government, has launched the Funding for Lending Scheme (FLS), which provides banks with an alternative cheaper source of funding tied to the extent to which they expand lending to the UK real economy. [...] It is bigger and bolder than any scheme tried so far to get the banks lending. In terms of the cost at which funding is being made available, the maturity of that funding and, most importantly, the strong price incentives it provides to banks to expand their lending. By helping to improve the availability of bank lending to companies and households who previously have been effectively starved of credit, it could have a significant effect on demand. Moreover, if some of the recent poor supply side performance of our economy does stem from the constraints on the flow of credit, it may also help to ease that friction." [\[Link\]](#)

A.4 European Central Bank's Targeted Longer-Term Refinancing Operation

More recently, the European Central Bank (ECB) also set up programs to support bank lending to the real economy – the targeted longer-term refinancing operations (TLTRO). In these operations, banks are entitled to borrow from the ECB, conditional on their lending to the private non-financial sector, excluding loans to households for house purchases.

(A) *Mario Draghi, President of the ECB, "Introductory Statement," Hearing at the Committee on Economic and Monetary Affairs of the European Parliament, July 14, 2014*: "[O]ur TLTROs are tailored to incentivise bank lending to the real economy in the euro area. The TLTROs will provide long-term funding to participating banks. This should ease their financing costs, allowing banks to pass on such attractive conditions to their customers. This will ease credit conditions and stimulate credit creation. Moreover, the growth of our balance sheet as a result of a significant take-up in our TLTROs will

put downward pressure on interest rates in the money markets. This will contribute further to lowering the banking sector's funding costs. However, the TLTROs will not merely provide long-term funding. The TLTROs are targeted operations: the stronger the flows of new net lending to non-financial corporations and consumers (relative to a specified benchmark), the higher the amount that banks will be permitted to borrow from the Eurosystem at very attractive terms and conditions over a period of up to four years. Hence, we have built in strong incentives for banks to expand their lending beyond original plans – both banks with a recent record of positive lending and those that have been deleveraging." [\[Link\]](#)

- (B) *Peter Praet, Member of the Executive Board of the ECB, "Current Issues of Monetary Policy," July 3, 2014:* "In this context, the Governing Council decided last month to adopt several credit easing measures – by which I mean, measures aimed at ensuring that the accommodative policy stance is translated into a corresponding easing in credit conditions. In particular, these measures include a series of targeted longer-term refinancing operations (TLTROs) aimed at easing credit conditions. The TLTROs are expected to ease overly tight lending conditions, lower lending rates and stimulate credit volumes through several channels. The first and most important channel is through a reduction in term funding costs for banks. Funding relief, however, does not per se guarantee better credit conditions for banks' customers, unless the supply of loans shifts in parallel and lending mark-ups are kept constant or even pushed down. This is why the targeted nature of the TLTRO is important: by making funding relief conditional on generation of new lending volumes, the TLTRO will encourage a shift outward in the credit supply curve. By simply moving along the demand schedule, this outward shift will reduce the price for lending while increasing new loans. If banks do not manage to exceed a certain benchmark in terms of net lending, they will not benefit from the TLTRO. This shows that the TLTROs are indeed targeted, rather than a broad-based unconditional provision of liquidity as in the case of the earlier 3-year LTROs." [\[Link\]](#)

B Profit Maximization vs. Credit Limit Discontinuities

In this appendix, we further investigate how banks set credit limits. In the first part, we show that the observed step-function relationship between FICO scores and credit limits is consistent with (i) banks maximizing profits given a fixed cost of determining optimal credit limits for a group of observably similar borrowers and (ii) the guidance provided to banks by their regulators. In the second part, we show that not only is the observed step-function relationship *qualitatively* consistent with profit maximization, but the costs of determining optimal credit limits implied by our estimates are also of *similar magnitude* to the costs reported in industry and trade publications.

B.1 Two-Stage Model

We first show that the step-function relationship between FICO scores and credit limits is qualitatively consistent with profit maximization in a simple two-stage model of credit card lending along the lines of the model proposed by Livshits, Macgee and Tertilt (2016). The key component of this model is that, in the first stage, lenders need to pay a fixed cost to develop a *scorecard* for lending to a group of observably similar borrowers. A scorecard is a statistical model that maps consumer characteristics, economic conditions, and contract terms into measures of the profitability of lending to a group of borrowers. The

costs are comprised, for example, of the cost of conducting randomized experiments to determine the sensitivity of borrower behavior to changes in contract terms.⁴¹ In the second stage, lenders use the scorecard to set contract terms for the group of borrowers. Livshits, Macgee and Tertilt (2016) describe the resulting equilibrium in the presence of fixed costs:

“The equilibrium features a finite set of loan contracts, each “targeting” a specific pool of risk types. The finiteness of contracts follows from the assumption that a fixed cost is incurred per contract, so that some “pooling” is necessary to spread the fixed cost across multiple types of borrowers. Working against larger pools is that these require a broader range of risk types, leading to wider gaps between the average default rate and the default risk of the least risky pool members. With free entry of intermediaries, these forces lead to a finite set of contracts for any (strictly positive) fixed cost.”

Indeed, this behavior of setting credit limits separately for different subgroups of the population is highly consistent with industry practice, as described by the Office of the Comptroller of the Currency (2015):

“Scoring systems do not normally consist of a single model. Recognizing that there are differences in available information and behavior patterns, the modeler attempts to segment the group into similarly situated subpopulations. The modeler can then develop individual scorecards for each distinct subpopulation that use the variables most predictive of risk for that particular group, thereby increasing accuracy and precision. [...] The definition of the subpopulations and the determination of how many to use are key components of the model development process.”

In our context, we can think of lenders developing scorecards for consumer segments defined by a partition of the FICO score distribution (e.g., ≤ 620 , 621-660, 661-700, etc). Banks select optimal credit limits for consumers in each group; a step function relationship between FICO scores and credit limits then arises naturally, where the steps will occur at transitions between consumer segments.

We next sketch this two-step model more formally. Let x indicate FICO scores, $z = \{z_1, z_2 \dots z_K\}$ be a sequence of increasing numbers that partitions FICO scores into $K - 1$ segments, and N be the number of new accounts the bank expects to originate. In the first stage, the lender chooses a partition to maximize the sum of expected profits for each segment subject to a sunk cost χ of developing a scorecard for each segment:

$$\max_z \left[N \sum_{k=1}^{K-1} \Pr[z_k < x \leq z_{k+1}] \mathbb{E}[\Pi(CL_k | z_k < x < z_{k+1})] \right] - (K - 1) \cdot \chi \quad (\text{A1})$$

In the second stage, the lender chooses a credit limit to maximize profits for each segment k . Because the firm has invested in a scorecard for group k , the lender knows the k -specific function that maps credit

⁴¹Since this sensitivity changes over time, these fixed costs need to be paid repeatedly. As highlighted by the Office of the Comptroller of the Currency (2015): “In simple terms, scoring employs mathematical techniques to predict future behavior based on past performance. Predictive horizons range from six months to two years. The assumption is that the behaviors of the scored population going forward will not change markedly from those of the population used to develop the model. The ability of models to differentiate risk deteriorates with time, however, as a result of shifts in consumer behavior, economic conditions, and bank and industry product terms and marketing. The majority of scoring models rely on statistical regression techniques (linear, logistic, or neural network).”

limits to revenue and costs:

$$\Pi(CL_k | z_k < x < z_{k+1}) = \max_{CL_k} q_k(CL_k)(r_k - c) + \tilde{R}_k(CL_k) - \tilde{C}_k(CL_k) \quad \text{for } k = 1, \dots, K-1 \quad (\text{A2})$$

Note that this is the exact same objective function as Equation 6 in the main body of the paper, except that we now have a k subscript for each group of borrowers.

This model rationalizes the observed step-function relationship between FICO scores and credit limits. Suppose that the scorecard costs are large, or the bank expects to originate relatively few cards. In this case, the bank will find it optimal to segment the FICO score distribution coarsely and so we will observe large jumps in credit limits when we switch from one segment to another, even if credit limits are optimal on average for each FICO score group. On the other hand, if the scorecard costs are small, then the bank will find it optimal to have a large number of small segments. In this case, if underlying borrower characteristics are smooth in the FICO score, we will not observe large jumps in credit limits.

At this point, it is important to note that while jumps imply scorecard costs in our model of profit-maximizing banks, the converse might not be true in a richer model of bank behavior. For instance, if we allowed a bank to choose a function $CL = f_k(x)$ that maps FICO scores to credit limits for each segment k , then we might observe a smooth relationship even with large scorecard costs.⁴²

B.2 Quantitative Estimates of Fixed Cost of Developing Scorecard (χ)

If it were costless to develop a scorecard (i.e., $\chi = 0$), the presence of discontinuous jumps in credit limits would be hard to reconcile with profit maximizing behavior. In particular, accounts to the left of the discontinuity would have credit limits that are too low and accounts to the right of the discontinuity would have credit limits that are too high relative to the profit maximizing levels. In the section above, we showed that the presence of a fixed cost of developing a scorecard can rationalize a step-function relationship between FICO scores and credit limits. In this section, we investigate whether the implied costs of developing scorecards are quantitatively reasonable.

In particular, we use our data to estimate a lower bound for χ for account holders with FICO scores "close" to the discontinuities. Since a bank can pay χ to develop a scorecard, the incremental profits for adjusting credit limits for these accounts provide a lower bound for χ . If χ were lower, banks would be better off paying this cost and setting credit limits in a more granular fashion.

Figure A8 provides a conceptual illustration for how we calculate a lower bound for the scorecard costs, which we denote χ , for a given "jump" in credit limits. Panel A shows an origination group where credit limits are a step function of the FICO score. For our example, we focus on accounts with FICO scores that are "close" to 680, the score at which the credit limit jumps from \$2,000 to \$3,000. Panel B shows marginal profits as a function of the credit limit for these accounts. Because these accounts have FICO scores close to 680, they have approximately the same marginal profit function.

For these accounts, suppose that the optimal credit limit, defined by the point where $MP'(CL) = 0$

⁴²There is another reason that we might not identify cutoffs for some origination groups in our data. Note that our data allow us to consider the credit supply function for origination groups defined as combinations of month-of-origination, bank, loan channel, and product type. This was the most granular way our data allow us to define a "credit card type." However, in practice, banks could still use different lending functions for cards within an origination group. In particular, some banks might issue a number of different credit cards in an origination group that we cannot tell apart using our data. If these credit supply functions had different cutoffs, and if there were a substantial number of different lending functions used, then we would not be able to observe these cutoffs at the origination group level.

is \$2,500, half-way between the \$2,000 and \$3,000 credit limits on either side of the discontinuity. Since total profits are the integral under the marginal profit curve, for accounts that are just to the left of the threshold, increasing credit limits to \$2,500 would raise total profits by the area of the shaded triangle to the left of \$2,500. For accounts that are just to the right, credit limits are too high, and lowering credit limits to \$2,500 would reduce losses by the area of the shaded triangle to the right of \$2,500.

Given our assumptions of linear marginal profits, the areas of both of these triangles are identical and equal to:

$$\underline{\chi} = -\frac{1}{2} \underbrace{\left(\frac{\Delta CL}{2}\right)}_{\text{Base}} \underbrace{\left(\frac{\Delta CL}{2} MP'(CL^*)\right)}_{\text{Height}} \quad (\text{A3})$$

$$= -\frac{1}{8} (\Delta CL)^2 MP'(CL^*), \quad (\text{A4})$$

which is just the standard "one-half base times height" formula for the area of a triangle. Intuitively, $\underline{\chi}$ is increasing in the slope of marginal profits (or equivalently the curvature of total profits), and the size of ΔCL .

Table A1 implements this formula using our estimates for each of our four FICO groups. We define accounts just to the left as those with a FICO score 1-5 points below the cutoff, and those to the right as accounts with a FICO score 1-5 points above. Column 1 shows the average jump in credit limits within each FICO group (ΔCL), column 2 shows the effect of a \$1,000 increase in credit limits on marginal profits ($1,000 \times MP'(CL^*)$) over our baseline 48 month horizon (i.e., the "slope" of marginal profits), and column 3 calculates $\underline{\chi}$ using Equation A4 above. Scorecard costs per account average between \$1 and \$11 across FICO score groups. These costs are much larger for low FICO score groups due to the fact that marginal profits are much more steeply sloped for these accounts, and therefore deviations from the frictionless-optimal credit limits are more costly. Intuitively, the higher estimates of $\underline{\chi}$ could be rationalized by the fact that underwriting is more difficult for these types of accounts.

We conduct two exercises to put these estimates in context. First, column 5 shows these estimates as a percentage of average total profits over the same 48 month horizon as we used to calculate the slopes. For these accounts, the scorecard costs are a fairly modest 1 to 11 percent. Second, we calculate the aggregate scorecard costs, defined as the product of the per-account costs multiplied by the number of accounts originated in our window of +/- 5 FICO score points of the threshold. The resulting estimates range from \$3,500 to \$18,000. The estimates are below the industry estimates of developing a scorecard of between \$40,000 and \$100,000 that are discussed in Livshits, Macgee and Tertilt (2016). Our estimates are, of course, a lower bound, and would be somewhat larger if we included the (smaller) gains from adjusting credit limits for accounts further away from the threshold. In addition, banks might use the same scorecard for originating accounts for a number of months, allowing them to spread the fixed cost over a larger number of accounts.

Overall, we conclude that our estimates suggest that the step-functions we observe are consistent with profit maximization by banks under fixed costs of the order of magnitude reported in industry and trade publications cited by Livshits, Macgee and Tertilt (2016).

C Econometric Approach: Controlling for Additional Cutoffs

As illustrated in Panels C and D of Figure 2, some of our quasi-experiments have additional cutoffs within the +/- 50 FICO score window we use for our regression discontinuity specifications. In our analysis, we control for the presence of these additional cutoffs with an indicator variable that is equal to 1 for all FICO scores above any additional cutoff beyond the one under investigation. In this section, we show the results of a Monte Carlo exercise that confirms that this approach to controlling for the presence of additional cutoffs allows us to recover the true treatment effect of interest.

Let y be the outcome, x be the running variable, \bar{x} be the cutoff of interest, and $\bar{\bar{x}}$ be the additional cutoff for which we want to control. Assume that the true data generating process is given by

$$y_i = \beta(x_i - \bar{x}) + \delta \mathbb{I}(x_i \geq \bar{x}) + \gamma \mathbb{I}(x_i \geq \bar{\bar{x}}) + \epsilon_i \quad \text{where } \epsilon_i \sim N(0,1). \quad (\text{A5})$$

Under this process, the outcome is linearly increasing in x_i with a slope of β , and has a jump of δ at threshold \bar{x} , and a jump of γ at threshold $\bar{\bar{x}}$. For our Monte Carlo exercise, we assume that x_i takes integer values on $[-10, 10]$. We set $\bar{x} = 0$ and $\bar{\bar{x}} = 5$. We also set $\beta = 1$, $\delta = 10$, and $\gamma = 10$. For each simulation, denoted by s , we randomly draw $n = 200$ values of x_i and ϵ_i , and calculate y_i according to the data generating process. Panel A of Figure A9 plots the average values of y for each value of x from a single simulation with $n = 200$ draws. The jumps of approximately 10 at $\bar{x} = 0$ and $\bar{\bar{x}} = 5$ are visible.

For each simulation s , we then estimate a single value of the coefficient of interest δ_s using the same locally linear regression specification we use for the main results in the paper (see Equation 3). We repeat this process 200 times, and examine how the distribution of δ_s compares to the true value of δ . Panel B shows the distribution of δ_s when we estimate a specification that *excludes* an indicator for $\mathbb{I}(x_i \geq \bar{\bar{x}})$. Because we exclude the indicator, we estimate “too steep” a slope to the right of $\bar{x} = 0$, generating a significantly downward biased estimates of δ . Panel C shows the distribution of δ_s when we include the indicator for $\bar{\bar{x}} = 5$. The average of δ_s across simulations is 9.993, which is very close to the true estimate of $\delta = 10$. This indicates that our econometric approach allows us to recover an unbiased estimate of the true effect.

D MPB: Extensive versus Intensive Margins

In Section 4.2, we document a large response of household borrowing to an increase in credit limits. This effect combines an extensive margin response, whereby some households that would not otherwise borrow start to borrow, and an intensive margin response, whereby households that already borrow increase their borrowing.

To quantify the relative importance of the extensive versus intensive margins, we conduct a simple decomposition of our baseline MPB estimates. Since the extensive margin is conceptually most clear for interest bearing debt, we focus on this measure of borrowing. Let q denote interest bearing debt and CL denote the credit limit. The average effect of extra credit on expected interest bearing debt can be decomposed into:

$$\frac{d\mathbb{E}[q]}{dCL} = \underbrace{\frac{d\Pr(q > 0)}{dCL} \mathbb{E}[q|q > 0]}_{\text{Extensive Margin}} + \underbrace{\frac{d\mathbb{E}[q|q > 0]}{dCL} \Pr(q > 0)}_{\text{Intensive Margin}}, \quad (\text{A6})$$

where the first term captures the extensive margin response and second term captures the intensive margin response.

Appendix Figure A10 shows the effects of a \$1,000 increase in credit limits on the probability of having cumulative positive interest bearing debt by FICO score group and at different time horizons. For reference, the average probability of having cumulative positive interest bearing debt is shown in Table 2. The increase in credit limits has a measurable but economically small impact on extensive margin borrowing. For the lowest FICO score group, a \$1,000 increase in credit limits raises the probability of borrowing within 12 months by 3 percentage points on a base of 58%, with a smaller effect at longer time horizons. For the highest FICO score group, a \$1,000 increase in credit limits raises the probability of borrowing by 1.1 percentage points on a base of 27%, with the effect similarly tailing off at longer time horizons. These 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 an automobile downpayment).

Appendix Table A2 shows the result of the decomposition in Equation A6. Panel A reproduces the baseline effect on interest bearing debt from Table 5. Panels B and C decompose this effect on the extensive and intensive margins. Consistent with the small effects in Appendix Figure A10, the decomposition shows that the extensive margin effect is relatively small compared to the intensive margin response. The extensive margin response is relatively more important at short time horizons, accounting for approximately one-quarter of the effect at 12 months after origination. At longer time horizons, the extensive margin effect becomes much less important, accounting for less than 5% of the overall effect at time horizons of 36 months and longer.

E Robustness Checks and Additional Heterogeneity

A key objective of this paper is to explore how the 743 regression discontinuity (RD) estimates vary by FICO score. In this Appendix, we provide additional robustness checks of the main results discussed in the paper. We also explore heterogeneity in our RD estimates along dimensions other than the FICO score. Finally, we analyze the sensitivity of our results to the distribution of where we observe the credit limit quasi-experiments.

E.1 Non-Parametric Relationship between RD Estimates and FICO Scores

In the main text, we examined heterogeneity in the 743 regression discontinuity estimates by projecting these estimates onto indicators for four different FICO score groups (≤ 660 , 661-700, 701-740, > 740) and controls. See Equation 5 for more details. The FICO score groups were chosen to partition the distribution of originated credit cards into four approximately equal-sized groups. This approach allowed us to show heterogeneity across FICO scores in the impulse responses plots (see, for example, Figure 7).

In this Appendix, we explore the extent to which our results are sensitive to this specific partition of the FICO score distribution. We do this by showing binned scatter plots using the `binscatter` command in Stata. The plots are constructed in two steps. In the first step, the command calculates residuals from regressions of the RD estimates and FICO scores on the control variables, adding back the sample mean of each variable to aid interpretation. As in the baseline regressions (Equation 5), we include 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. We also include fully inter-

acted loan channel by "zero initial APR" fixed effects. These controls partial out any heterogeneity in treatment effects that might come from quasi-experiments occurring, for example, at different points in time. In the second step, the residualized FICO scores are grouped into 50 buckets, and the mean residualized treatment effects and FICO scores within each bin are calculated. The `binscatter` command plots these data points and the quadratic line of best fit.⁴³

Appendix Figure A11 shows these binned scatter plots for the key outcome variables analyzed in the main body of the paper. Panels A to D show plots for MPB outcomes (ADB, interest bearing debt, balances across all cards, and cumulative purchase volume), the same outcomes shown in Figure 7. We show outcomes at a horizon of 12 months, but the results are similar if we consider outcomes at other horizons. All measures of marginal borrowing and spending decline smoothly and monotonically in FICO score, indicating that our finding of a declining MPB is robust to our definition of FICO score groups.

Panels E and F show binned scatter plots of the key determinants of banks' MPL: cumulative marginal chargeoffs and the slope of cumulative marginal profits, both at a 48 month time horizon. Again, we find a smooth, monotone relationship between the RD estimates and FICO scores. In particular, the slope of marginal profits becomes flatter (less negative) as the FICO score increases, indicating that our finding of a larger MPL for higher FICO score borrowers is also robust to how we define the FICO score groups.

E.2 Robustness of Estimates to Size of Credit Limit Jump

As highlighted in Panel A of Figure 4, the size of the first stage increase in credit limits varies across our 743 quasi-experiments. In particular, the size of the jump ranges from \$214 to \$5,463, with a 10-90 percentile range of \$627 to \$2,635, and an interquartile range of \$883 to \$1,939.⁴⁴ In this Appendix, we analyze whether the RD estimates depend on the size of the credit limit jump. In particular, for each of our key outcome variables, we estimate a version of the baseline regression (Equation 5) where we fully interact the FICO score bucket fixed effects with an indicator for quasi-experiments with jumps larger than the median jump size of \$1,282.

$$\tau_j = \left(\sum_{k=1}^4 \beta_k FICO_{j,k} \times \mathbf{1}_{SmallJump,j} \right) + \left(\sum_{k=1}^4 \beta_k FICO_{j,k} \times \mathbf{1}_{LargeJump,j} \right) + X'_j \delta_X + \epsilon_j. \quad (A7)$$

Appendix Figure A12 shows the results of this analysis. In particular, we show estimates for the MPB outcomes (ADB, interest bearing debt, balances across all cards, and cumulative purchase volume) and key MPL outcomes (cumulative chargeoffs and slope of marginal profits) across the FICO score buckets separately for experiments with small and large credit limit jumps. For all of the outcome variables, we find economically and statistically similar effects across jumps of different magnitudes, indicating that our effects are robust to the size of the credit limit jump.

E.3 Heterogeneity of Borrowing Response by Loan Channel

One interesting question is whether the estimated MPBs differ across credit card originations that were initiated by the bank versus originations that were initiated by the consumer. Conceptually, it is unclear

⁴³ See [here](#) for more details on the `binscatter` command, and Chetty, Friedman and Rockoff (2014) for more details on the binned scatter plot methodology.

⁴⁴ Interestingly, while the average size of the credit limit jump is slightly larger for quasi-experiments at higher FICO scores, the size of the credit limit jumps in our four FICO score buckets come from largely overlapping distributions.

which group should be more responsive to credit expansions. On the one hand, consumers presumably have private information about their (future) demand for credit. If consumers who know they will need to borrow in the future are more likely to initiate credit card applications, we might expect to observe larger MPBs for consumer-initiated accounts. On the other hand, banks likely target consumers who they think will be profitable, and who are presumably also more likely to borrow in the future, resulting in a high MPB. Moreover, since consumers still need to respond to bank-initiated credit card offers, the sample of bank-initiated cards may also reflect some degree of private information held by consumers.

To analyze heterogeneity along this dimension, we obtain separate estimates of the key MPB effects for accounts that were initiated by the bank and accounts that were initiated by the consumers:

$$\tau_j = \left(\sum_{k=1}^4 \beta_k FICO_{j,k} \times \mathbf{1}_{ConsumerInitiated,j} \right) + \left(\sum_{k=1}^4 \beta_k FICO_{j,k} \times \mathbf{1}_{BankInitiated,j} \right) + X_j' \delta_X + \epsilon_j. \quad (\text{A8})$$

where we define "bank-initiated" accounts as those originated through a pre-approved mailing or an invitation to apply and "consumer-initiated" accounts as those originated through either a "take one" branch application or in response to an internet or magazine advertisement.

Appendix Figure A13 shows the results of this analysis. There are no clear patterns across the four MPB outcome variables, with the estimates for bank-initiated originations being moderately larger for some outcomes and moderately smaller for others. The estimated differences across the bank-initiated and consumer-initiated credit cards are small from a statistical perspective (relative to the standard errors) and small economically relative to the across-FICO score group differences in the estimates.

E.4 Heterogeneity of Borrowing Response by Income and by Utilization

Most of the analysis in the main body of the paper focuses on heterogeneity in the MPB by FICO score. Given that banks set credit limits based on FICO scores, this is the appropriate dimension of heterogeneity to consider for analyzing the pass-through of credit expansions. However, our data allow us to examine heterogeneity along other borrower characteristics, which might be the relevant dimensions of heterogeneity in other settings. In this Appendix, we examine the heterogeneity in the MPB on two other dimensions: self-reported borrower income and the utilization rate across all credit cards, both defined at account origination.

To set up the analysis, Panel A of Appendix Figure A14 shows a binned scatter plot (see description above) of borrower income against FICO score. The plot shows that these measures exhibit only a modest positive correlation, indicating that there is information in income that is not contained in the FICO score. Panels B to E of Appendix Figure A14 replicate the impulse response plots shown in Figure 7, except that we show heterogeneity by income group instead of by FICO score group. Interestingly, we find borrowing and spending behavior to be very similar across different income groups, indicating that self-reported borrower income is not a strong predictor of subsequent borrowing and spending behavior.

Appendix Figure A15 shows analogous plots that examine heterogeneity by the utilization rate across all credit cards, calculated as the ratio of total balances to total credit limits in the credit bureau data. Panel A shows that utilization and FICO score are highly negatively correlated, consistent with utilization being one of the key input variables in the FICO score formula. Consistent with the strong correlation between FICO score and utilization, the patterns of heterogeneity by utilization in the

borrowing and spending plots mirror those that examine heterogeneity by FICO score.

E.5 Heterogeneity of Estimates by Time of Credit Card Origination

Our data cover credit cards originated between January 2008 and November 2013 – stretching from prior to the start of the Financial Crisis through much of the recovery period. In our main analysis, we pool credit cards originated in different years while removing time fixed effects, to ensure that the heterogeneity of RD estimates across FICO scores is orthogonal to a possible heterogeneity over time. In this Appendix, we explore whether there is heterogeneity in the RD effects for credit cards originated at different points in time. One important caveat to this analysis is that any observed heterogeneity combines effects of heterogeneity over time in the unobserved characteristics of new borrowers and heterogeneity in the treatment effects of credit expansions holding the borrowers' unobservable type fixed. In addition, the limitations of our sample period prevent us from conducting an analysis of the full business cycle variation of our estimates.

Appendix Figure A16 examines the MPB across accounts originated in different years of our sample. Panel A shows the effect on ADB at different time horizons after origination by year of origination, controlling for the FICO score of the quasi-experiment. Because we only observe outcomes through December 2014, for accounts that were originated in later years, we are only able to examine outcomes for shorter time horizons. The plot shows fairly similar effects across origination years. In particular, at twelve months after origination, the MPB is between 35% and 40% across accounts originated in different years. At 24 months after origination, the MPB ranges between about 25% and 35% across origination years. To explore heterogeneity across both time and FICO score groups, Panel B of Appendix Figure A16 shows effects by FICO score group and year of origination, restricting the plot to the MPB at 12 months after origination. The plot again shows very little heterogeneity and no obvious pattern by origination year.

The sufficient statistics for pass-through are the MPB and the slope of marginal profits. To complement the MPB analysis discussed above, we also examine heterogeneity in the slope of marginal profits over time. For this analysis, a key question is whether the slope of marginal profits is different for accounts that were originated before the onset of the Financial Crisis, when the degree of asymmetric information in the market may have been different.

Appendix Figure A17 shows the effect of a \$1,000 increase in credit limits on cumulative marginal profits over 48 months (i.e., the slope of marginal profits) for accounts originated in Q1 and Q2 of 2008 (before the onset of the Financial Crisis) or in Q3 of 2008 and later.⁴⁵ The plot shows fairly similar slopes of marginal profits before and after the start of the Financial Crisis. For both time periods, the slope of marginal profits is more steeply negative for accounts with lower FICO scores and the variation across time periods is much smaller than the variation across FICO score groups. Thus, while we do not interpret this analysis as indicating that there is no over time variation in the slope of marginal profits, we do interpret the estimates as indicating that the patterns we document are qualitatively robust, and that the heterogeneity by FICO score we focus on is quantitatively the main source of heterogeneity in our data.

⁴⁵We are unable to explore the same range of across-time heterogeneity as for the MPB because we need to observe at least 48 months of post-origination data to construct the measure of marginal profits used in the main body of the paper. Along a number of indicators, such as house prices, January 2008 was already a crisis period. This means that we cannot rule out heterogeneity in the slope of marginal profits across the full peak-to-trough range of the business cycle.

E.6 Sensitivity of Results to the Distribution of Quasi-Experiments

In our baseline analysis, we examine heterogeneity in treatment effects across the FICO score distribution by including indicator variables for whether the cutoff for our quasi-experiment falls into one of four different FICO score groups (see Equation 5). In this section, we explore the sensitivity of our results to partitioning the FICO score distribution in this manner, as well as to the distribution of FICO score cutoffs at which we observe the quasi-experiments.

This sensitivity analysis proceeds in three steps. First, we estimate a non-parametric relationship between our 743 quasi-experimental estimates and the FICO score cutoffs at which these experiments occur. In particular, Figure A11 shows binned scatter plots of our estimates against the FICO score cutoffs, along with the best fit second-order polynomial. Like our baseline specification (Equation 5), these binned scatter plots partial out fully interacted controls for origination quarter, bank, and zero initial APR, and additively separable fully interacted controls for loan channel and zero initial APR. Across the different outcome variables, the estimates trend quite smoothly in FICO score.

In the second step, we use these estimated non-parametric relationships to project an effect for every possible FICO score, including FICO scores where we do not observe quasi-experiments. Conceptually, this approach entails interpolating effects for FICO scores where we do not observe quasi-experiments, but which are within the range of FICO scores where we do observe quasi-experiments; we also extrapolate effects for FICO scores outside of the range of quasi-experiments we observe.⁴⁶ We observe quasi-experiments at FICO scores that run from 630 to 785, a range that covers 67.8 percent of the distribution of all credit card accounts originated during our sample (Figure A1 shows this distribution). Because the relationship between the estimates and FICO scores is well approximated by a second-order polynomial, we think these interpolations and extrapolations are reasonable.

In the third step, we use the estimated non-parametric relationship to assess the sensitivity of our baseline estimates to the population distribution of credit card accounts. Specifically, we calculate the mean of the predicted effect for each of our four FICO score groups, weighting the estimates by the distribution of FICO scores in the population. Table A3 shows these population-weighted estimates and the unweighted baseline estimates for reference. The differences due to reweighting are small, both relative to the precision of the estimates and relative to the across-FICO score group variation in magnitudes. Because the population distribution has more mass in the tails, the population-weighted estimates for the lowest and highest FICO score groups tend to be more extreme than the baseline estimates. Since the population-weighted estimates exhibit slightly larger differences across FICO score groups, accounting for the population weighting would modestly increase the negative correlation between MPB and MPL. Thus, relative to population reweighting, our approach provides a somewhat conservative lower bound to the mismatch between MPB and MPL.

⁴⁶For example, if we estimate an MPB of 0.5 for FICO scores of 660, an MPB of 0.6 for FICO scores of 680, and assume a linear functional form (for expositional purposes), then we will interpolate an MPB of 0.55 for FICO scores of 670. Similarly, if we estimate an MPB of 0.65 for FICO scores of 640, and an MPB of 0.7 for FICO scores of 630, and again assume a linear functional form (for expositional purposes), we would predict an MPB of around 0.75 for FICO scores of 620, even though we do not observe any experiments for FICO scores below 630.

F Linearity Assumption

In Section 6, we parameterized the marginal profit curve using a linear functional form, and showed that the slope of marginal profits was the largest for the lowest FICO score borrowers. In this Appendix, we show our results are qualitatively robust to a wide class of functional forms for the marginal profit function. 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 $AP(CL)/CL$ has a larger value. Since we find in our data that $AP(CL)/CL$ is larger for lower FICO score borrowers, this implies that our finding that the slope of marginal profits is steeper for lower FICO score borrowers holds for any functional form that satisfies this condition and does not depend on our choice of functional form.

First consider two marginal profits functions, $MP_x(CL)$ and $MP_y(CL)$, that have the same optimal credit limit: $MP_x(CL^*) = 0$ and $MP_y(CL^*) = 0$. Define average profits as $AP(CL) = \frac{\int_0^{CL} MP(CL)dCL}{CL}$.

Proposition 1. *If $MP_x(CL)$ and $MP_y(CL)$ satisfy a strict single crossing property, then $AP_x(CL^*) > AP_y(CL^*) \iff MP'_x(CL^*) < MP'_y(CL^*)$.*

That is, higher average profits is a necessary and sufficient condition for marginal profits to have a steeper (negative) slope at the optimum. Figure A18 provides an illustration of this phenomenon.

Proof. Since the marginal profit functions (i) have the same optimal credit limits CL^* and (ii) satisfy a strict single-crossing property, the marginal profit functions cross at CL^* . First, if $AP_x(CL^*) > AP_y(CL^*)$, then $MP_x(CL) > MP_y(CL)$ for $CL < CL^*$ and $MP_x(CL) < MP_y(CL)$ for $CL > CL^*$. But then, since MP_x crosses MP_y "from above", we know that $MP'_x(CL^*) < MP'_y(CL^*)$. Second, if $AP_x(CL^*) < AP_y(CL^*)$, then by the same logic, $MP'_x(CL^*) > MP'_y(CL^*)$. But then since $AP_x(CL^*) > AP_y(CL^*) \implies MP'_x(CL^*) < MP'_y(CL^*)$ and $AP_x(CL^*) < AP_y(CL^*) \implies MP'_x(CL^*) > MP'_y(CL^*)$, we know that $AP_x(CL^*) > AP_y(CL^*) \iff MP'_x(CL^*) < MP'_y(CL^*)$. \square

Now consider two marginal profit functions that have different optimal credit limits: $MP_x(CL_x^*) = 0$ and $MP_y(CL_y^*) = 0$, where $CL_x^* \neq CL_y^*$. Define the normalized marginal profit function as $\tilde{M}P(\tilde{C}L) = \frac{1}{CL^*}MP(CL^* \cdot \tilde{C}L)$. The normalized marginal profit function is constructed so that it takes on the optimal value at $CL = 1$. That is, $\tilde{M}P(1) = \frac{1}{CL^*}MP(CL^*) = 0$. By construction, at the optimal values, the slopes of the un-normalized and normalized marginal profit functions are identical: $\tilde{M}P'(1) = \frac{1}{CL^*}MP'(CL^*) \cdot CL^* = MP'(CL^*)$. Define average profits as $\tilde{A}P(\tilde{C}L) = \frac{\int_0^{CL} \tilde{M}P(\tilde{C}L)d\tilde{C}L}{\tilde{C}L}$.

Proposition 2. *If the normalized marginal profit functions, $\tilde{M}P_x(\tilde{C}L)$ and $\tilde{M}P_y(\tilde{C}L)$, satisfy a strict single crossing property, then $\frac{AP_x(CL_x^*)}{CL_x^*} > \frac{AP_y(CL_y^*)}{CL_y^*} \iff MP'_x(CL_x^*) < MP'_y(CL_y^*)$.*

That is, a higher value for $AP(CL)/CL$ is a necessary and sufficient condition for marginal profits to have a steeper (negative) slope at the optimum.

Proof. First, by Proposition 1, we know that $\tilde{A}P_x(1) > \tilde{A}P_y(1) \iff \tilde{M}P'_x(1) < \tilde{M}P'_y(1)$. Second,

$$\begin{aligned}\tilde{A}P(1) &= \int_0^1 \tilde{M}P(\tilde{C}L) d\tilde{C}L \\ &= \int_0^1 \frac{1}{CL^*} MP(CL^* \cdot \tilde{C}L) d\tilde{C}L \\ &= \int_0^{CL^*} \frac{1}{CL^*} MP(CL) \frac{1}{CL^*} dCL \\ &= \frac{AP(CL^*)}{CL^*}\end{aligned}$$

where the third line involves a change of variables in which we replace $\tilde{C}L$ with $\frac{CL}{CL^*}$. Third, by construction, we know that $\tilde{M}P'(1) = MP'(CL^*)$. Therefore, substituting in, we have $\frac{AP_x(CL_x^*)}{CL_x^*} > \frac{AP_y(CL_y^*)}{CL_y^*} \iff MP'_x(CL_x^*) < MP'_y(CL_y^*)$, as desired. \square

G Assigning Measures of Profitability to Credit Card Accounts

As discussed in Section 6, in order to estimate banks' MPL, we need to measure the profitability of credit card lending at the account level. In this Appendix, we describe the profitability components we are able to observe in our various data sets, and how we assign them to the individual credit card accounts.

Before we do so, however, it is worth re-emphasizing that, from a conceptual standpoint, any component of profits (both revenue and cost) that does not vary with credit limits has no effect on the degree of pass-through of credit expansions. Instead, what matter for pass-through are the variable components of profits that are affected by credit limits. Section 5 formalizes this insight. We next discuss how we measure the various components of variable profits in the data, and assign them to individual accounts. There are three classes of variables.

Variables observed at the account level. The first category of variables are already reported by the banks to the OCC at the account level. These variables include interest charges, fees, and chargeoffs. Table 2 shows that interest charges and fee revenue make up approximately 90% of all cumulative revenue over the first 48 months (the number is slightly smaller, at 83%, for the highest FICO score group, where interchange income plays a somewhat larger role). On the cost side, chargeoffs constitute about 60% of total cumulative costs. Agarwal et al. (2015b) show that these percentages, which Table 2 reports for our experimental sample, are very similar to the percentages in the overall credit card portfolios.

Variables observed in portfolio data. A second category of variables are reported by the banks to the OCC at the monthly level for their entire credit card portfolio. We are required to manually apportion these variables to each account. We do this in the following way.


- **Interchange income.** Interchange fees are paid by a merchant's acquiring bank to a cardholder's issuing bank as part of an electronic payment card transaction. Interchange income makes up the remaining 10% of observed bank revenue from managing a credit card portfolio. In the portfolio data, we find that total interchange income is a very constant fraction of about 2% of total purchase volume across the credit cards in the portfolio data (see Appendix Figure A19, taken from Agarwal et al., 2015b). This is highly consistent with evidence from industry sources, which cite an average

interchange fee of about 2% of purchase volume (GAO, 2009). We therefore construct account level interchange income as 2% of account level purchase volume.

- **Cost of funds.** These are the costs for the bank to fund their credit card receivables, and are reported in the portfolio data. Together with data on total credit card receivables at the portfolio level, we can calculate a monthly cost of funding \$1 of credit card receivables. The resulting series closely tracks other aggregate cost of funds measures (see Panel A of Appendix Figure A20, taken from Agarwal et al., 2015b). We apportion this portfolio-level monthly cost to the account level based on each account's ADB. This allows us to allocate another 5% to 6% of total cost.
- **Rewards and fraud expenses.** From an institutional perspective, rewards and fraud expenses are likely to scale with purchase volume: rewards such as airline miles and cash-back are usually paid as a fraction of the amount spent on the credit card, and the probability of credit card fraud rises with spending. Indeed, in the portfolio data we find these expenses to correspond to a constant fraction of 1.4% of purchase volume (see Panel B of Appendix Figure A20, taken from Agarwal et al., 2015b). We construct account-level values by applying this percentage to account-level purchase volume. This allows us to allocate another 10% to 15% of total costs of credit card lending to the account level.
- **Operational costs.** These are the costs for marketing and acquisition, collections, servicing, cardholder billing, processing payments, and card issuing and administration. They average at between 3.5% and 4.5% of ADB in the portfolio data (see Panel C of Appendix Figure A20, taken from Agarwal et al., 2015b). Some of these costs, such as the cost of debt collection, naturally scale with ADB, and thus with credit limits. Others, such as marketing and customer acquisition costs, could be modeled as fixed costs that do not scale with credit limits, and therefore should be excluded from our measure of variable profits. We choose to assign the majority, but not all of the operational costs to the account level, taking them as 3.5% of ADB. Since these costs are fairly small, the results are not particularly sensitive to the exact choice of their assignment.

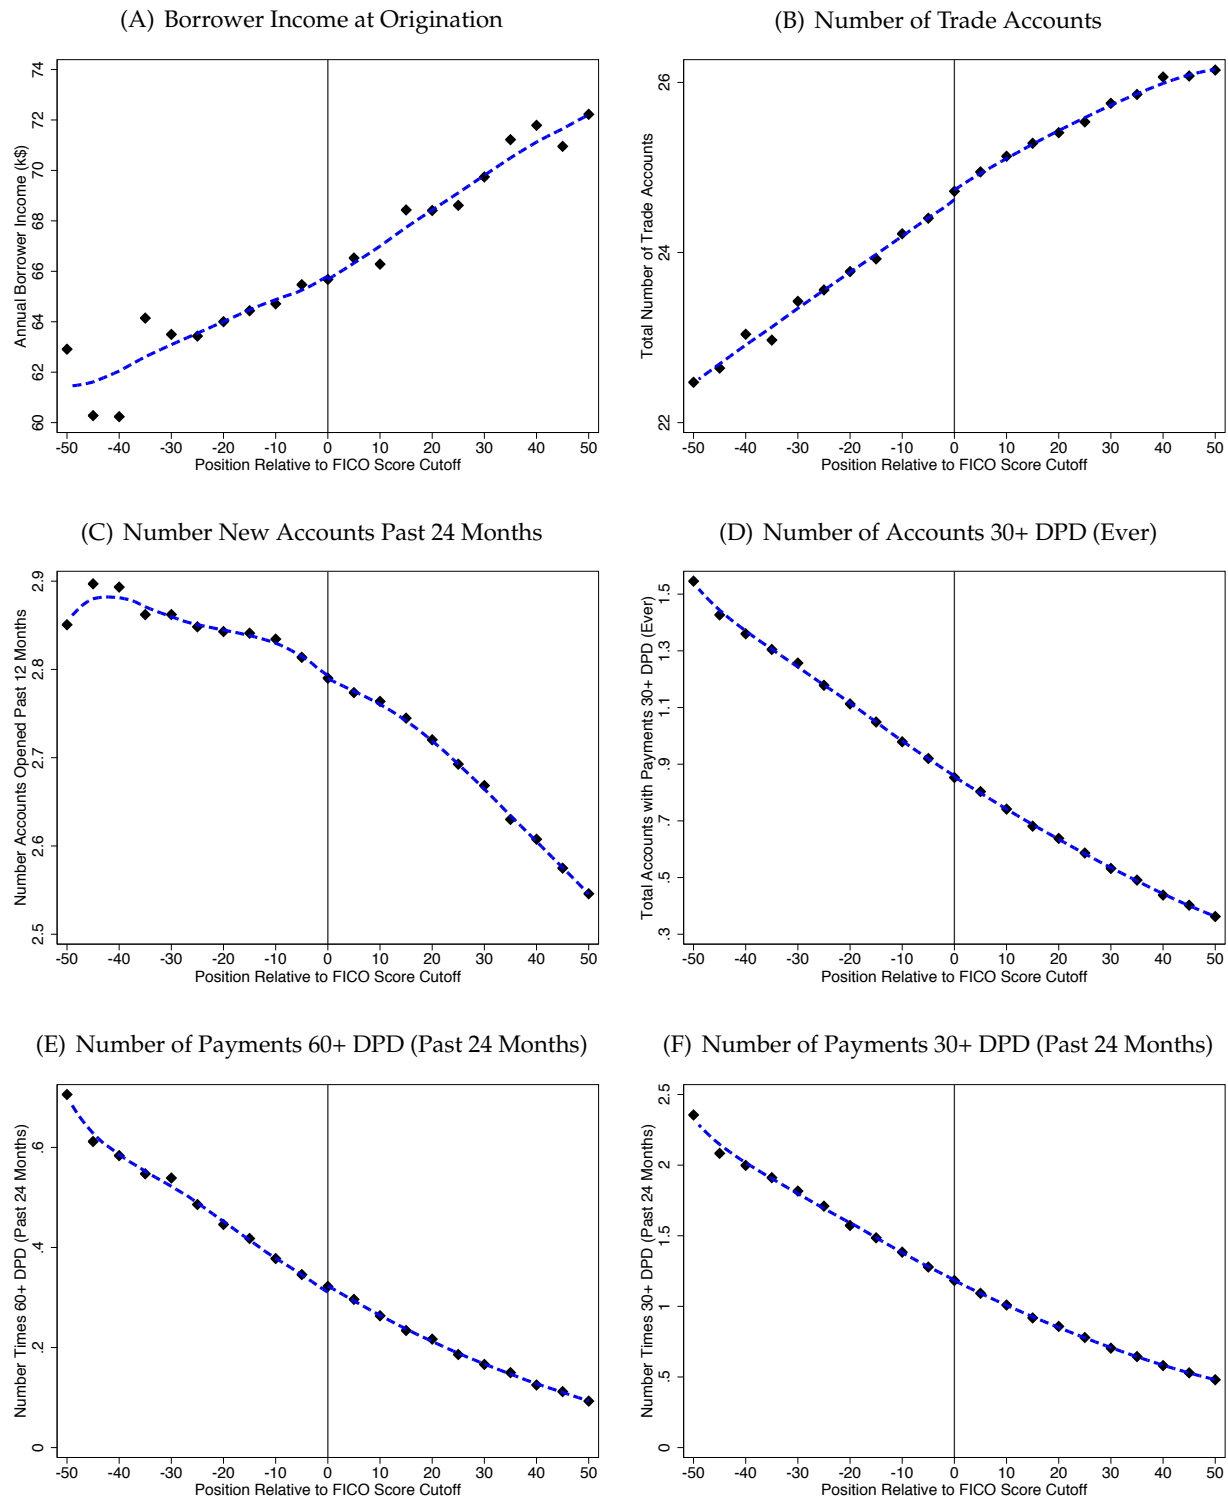
Variables not observed. The third category of variables includes any items that are not observed in the OCC data. The primary variable in this category is the cross-selling benefits that banks obtain by being able to sell other products to their credit card customers. However, since the benefit of cross selling is unlikely to scale with credit limits, even if we did observe it, we would not want to use it in our construction of variable profits. We cannot think of any other major component of variable cost or revenue of credit card lending that is not reported in our data.

Figure A1: FICO Score, Population Distribution



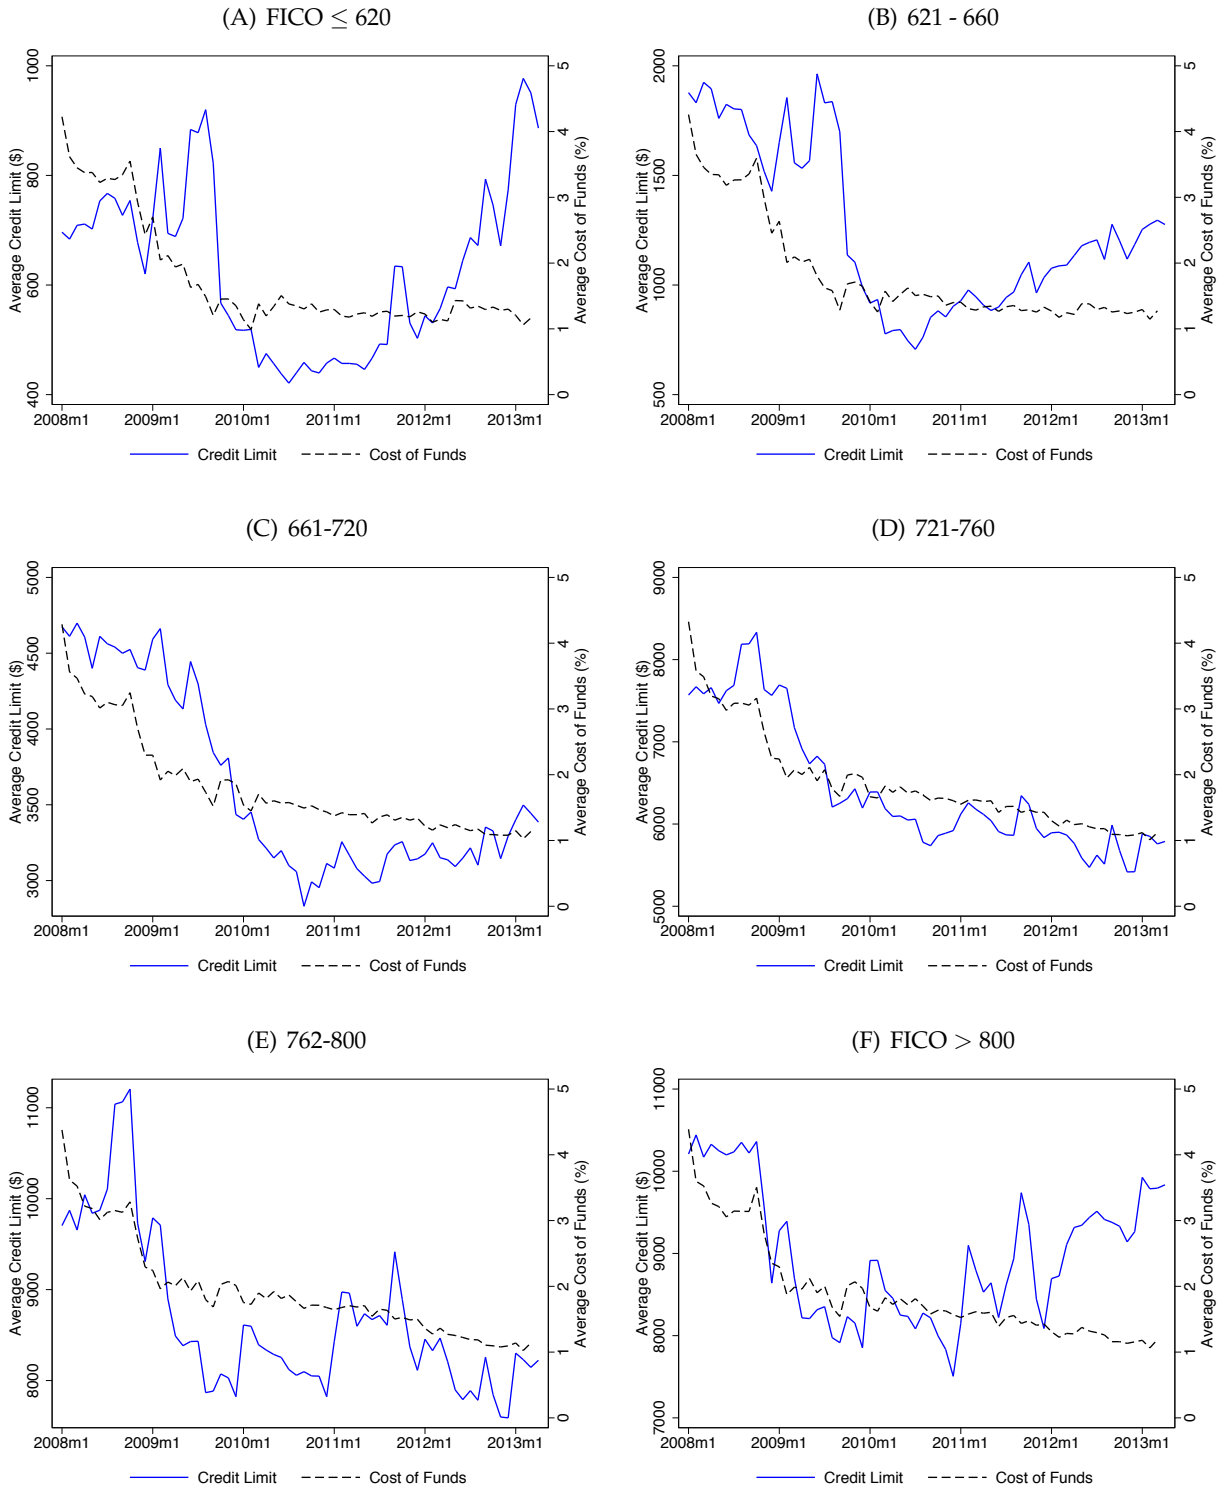
Note: Figure shows the distribution over FICO scores of all credit cards issued by the banks in our sample, averaged over the period from January 2008 to November 2013.

Figure A2: Initial Borrower Characteristics Around Credit Limit Quasi-Experiments




Note: 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 reported borrower income at origination (Panel A), total number of trade accounts (Panel B), total number of new trade accounts over the past 24 months (Panel C), number of accounts that were ever 30+ days past due (Panel D), number of payments that were 60+ days past due in last 24 months (Panel E), and the number of payments that were 30+ days past due in last 24 months (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. Blue lines show predicted values from locally linear regressions estimated separately on either side of the cutoff using the Imbens and Kalyanaraman (2011) optimal bandwidth.

Figure A3: Credit Limits and Cost of Funds in the Time Series




Note: Figure shows average credit limits on newly originated credit cards (solid line) and average cost of funds (dashed line) over time by FICO score group.

Figure A4: Credit Card Interest Rates vs. Federal Funds Rate




Note: Figure shows the year-on-year change in credit card interest rates and year-on-year change in the Federal Funds Rate between 1974 and 2015. Before 1994, credit card interest rates were those reported in the Federal Reserve's "Quarterly Report of Interest Rates on Selected Direct Installment Loans." From 1994 onwards, credit card interest rates are from the Federal Reserve's "Quarterly Report of Credit Card Interest Rates" for those credit card holders incurring interest charges. The full-sample time-series correlation is 0.166.

Figure A5: Credit Card Credit Limits vs. Interest Rates



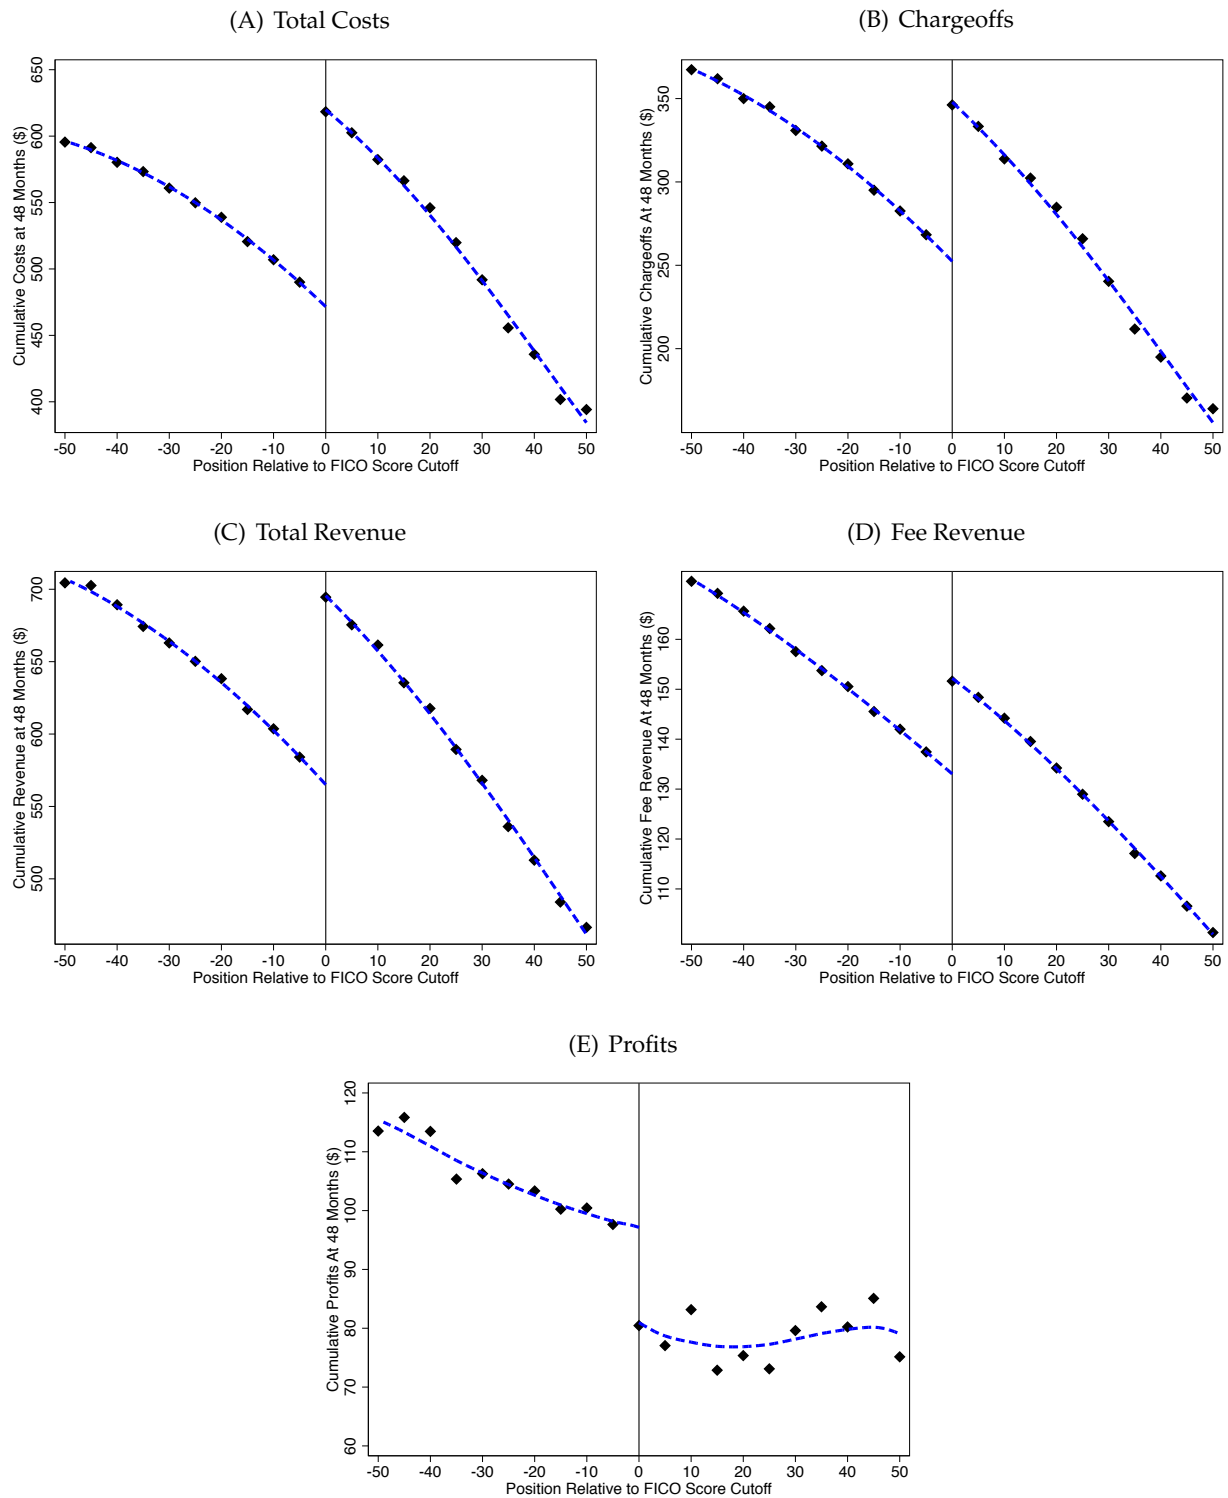
Note: Figure shows credit card credit limits and interest rates between 2000 and 2015, with the values normalized to 100% in the year 2000 for comparability. The interest rates are from Federal Reserve's "Quarterly Report of Credit Card Interest Rates" for those credit card holders incurring interest charges. The credit limits are calculated using a random sample of credit reports from TransUnion between 2000 to 2015.

Figure A6: Probability of Delinquency at 48 Months After Origination



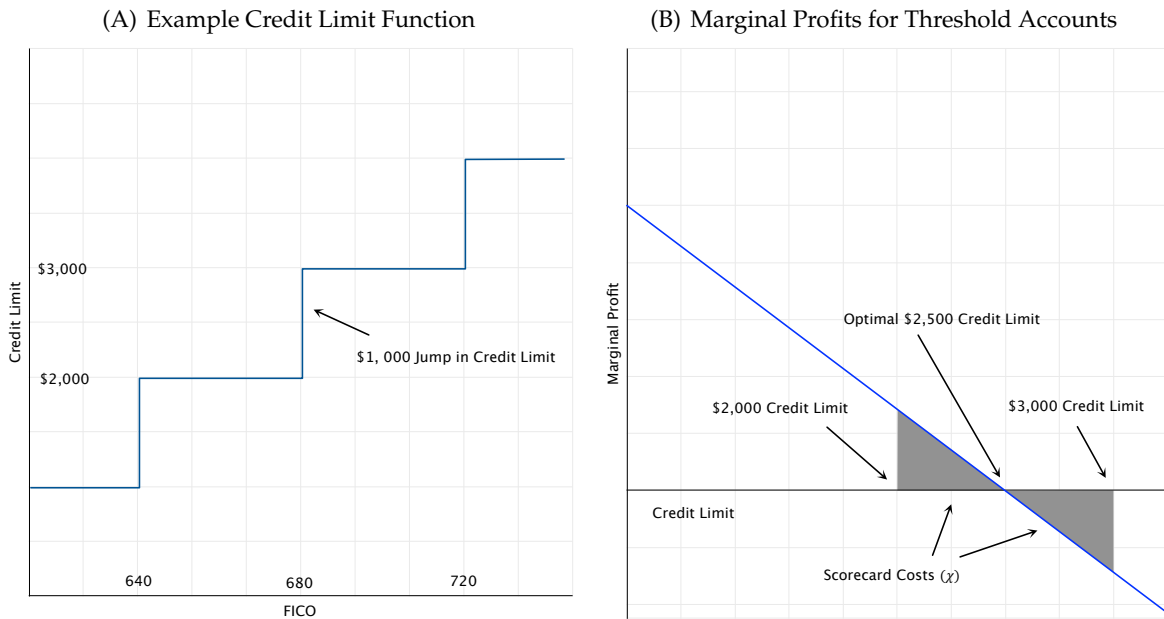
Note: Figure shows the effects of credit limits on the probability of delinquency around our 743 pooled credit limit quasi-experiments. Panel A shows effects on the cumulative probability of an account being more than 60 days past due (60+ DPD); Panel B shows effects on the cumulative probability of being more than 90 days past due (90+ DPD). These plots are constructed as described in Figure 3.

Figure A7: Total Revenue, Total Cost, and Components



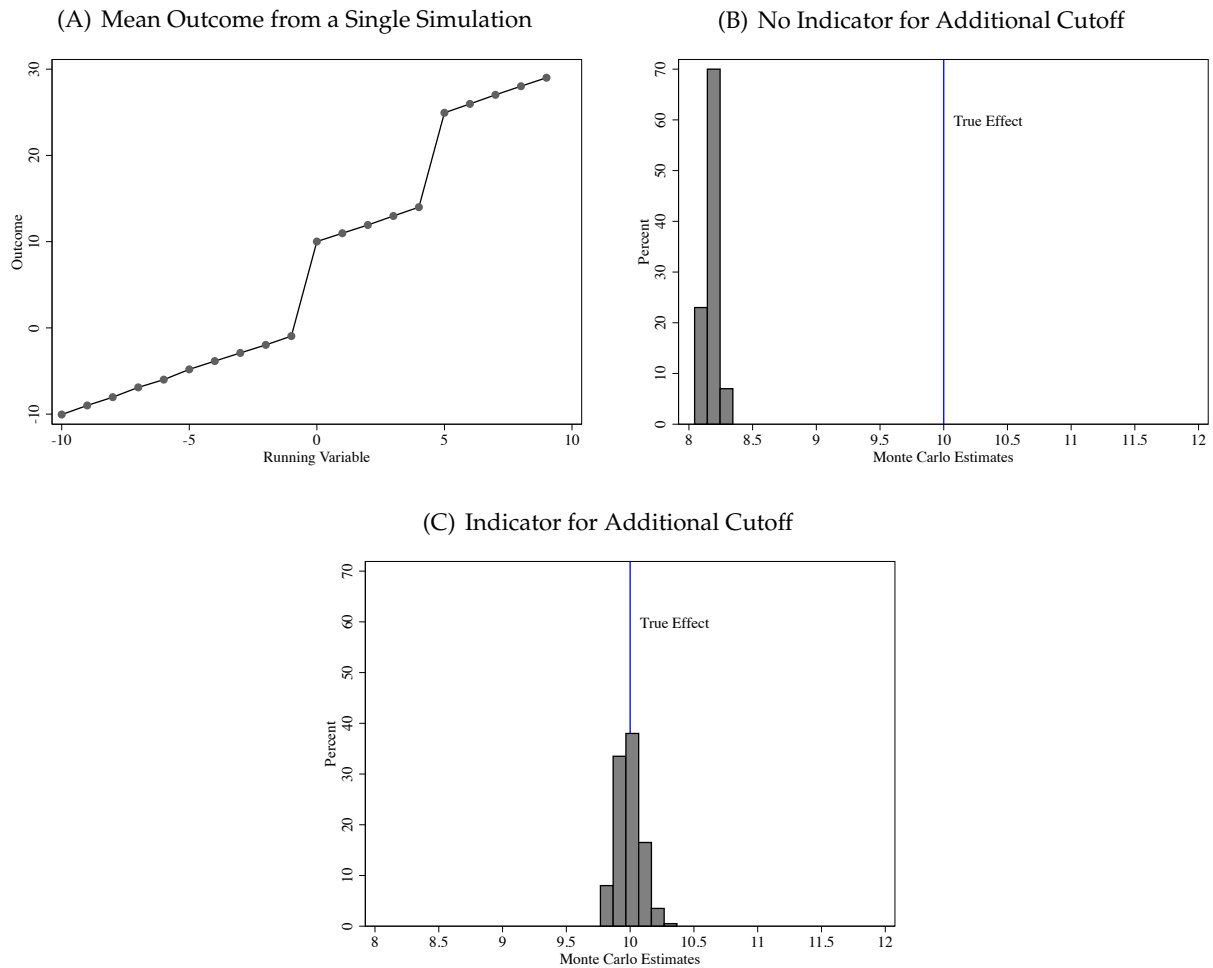
Note: Figure shows the effects of credit limits on cumulative total costs (Panel A), cumulative chargeoffs (Panel B), cumulative total revenue (Panel C), cumulative fee revenue (Panel D), and cumulative profits (Panel E), all measured over the first 48 months after account origination. These plots pool across our 743 credit limit quasi-experiments, and are constructed as described in Figure 3.

Figure A8: Scorecard Costs (χ): Conceptual Framework




Note: Figure illustrates how we calculate a lower bound for the scorecard costs χ for a given “jump” in credit limits. Panel A shows an origination group where credit limits are a step function of FICO score. We focus on accounts with FICO scores that are “close” to 680 where the credit limit jumps from \$2,000 to \$3,000. Panel B shows marginal profits as a function of the credit limit for these accounts, where the optimal credit limit is assumed to be \$2,500. The area of the shaded triangles indicates the incremental profits that could be achieved from adjusting the credit limits to the optimal level, which is a lower bound for the scorecard costs χ . See text for more details.

Figure A9: Monte Carlo Simulations: Controlling for Other Experiments



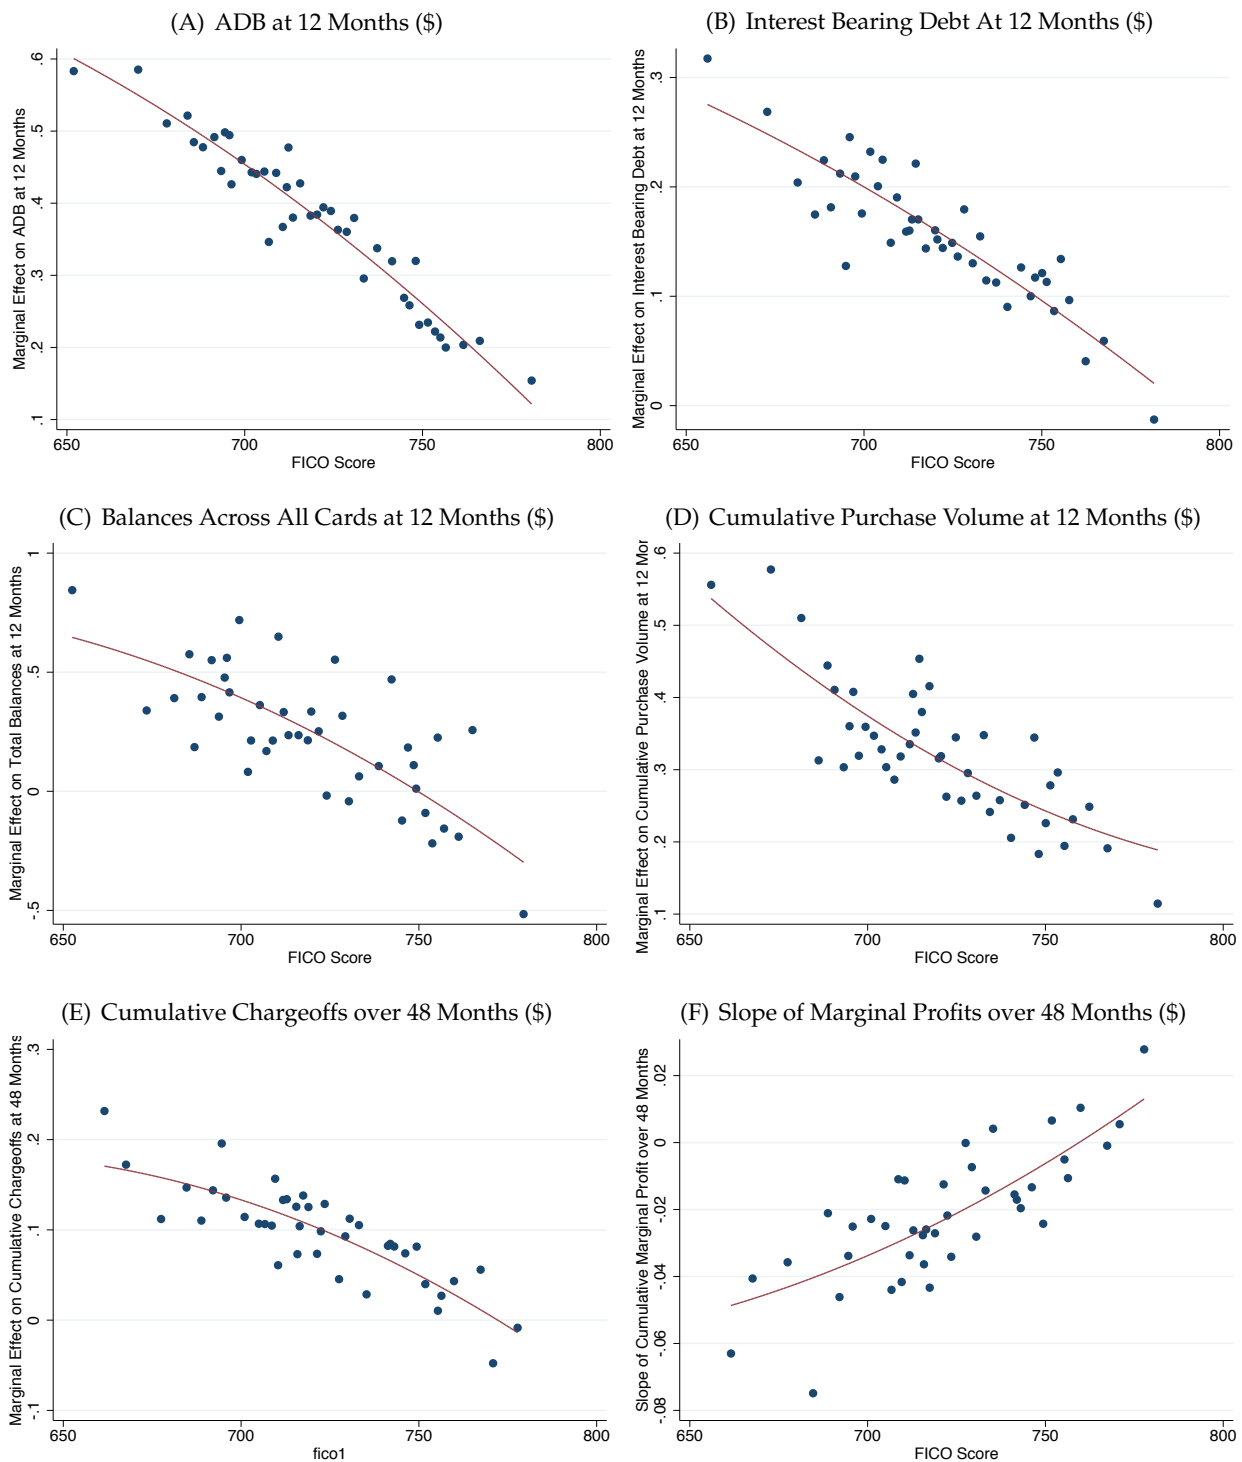
Note: Figure shows output from a Monte Carlo exercise that assesses whether we can control for the presence of additional cutoffs with an indicator variable that is equal to 1 for all FICO scores above this additional cutoff value. Panel A plots the average value of the outcome from a single simulation with $n = 200$ draws, a cutoff at $x = 0$ and an additional cutoff at $x = 5$. See text for a description of the data generating process. Panel B shows the distribution of δ_s from $s = 200$ Monte Carlo simulations in which we estimate a specification that *excludes* an indicator for the additional cutoff. Panel C shows the distribution of δ_s when we include the indicator for the additional cutoff. In Panels B and C, the true effect at the cutoff is indicated with a vertical line.

Figure A10: Effect of \$1K Increase in Credit Limits on Probability of Positive Interest Bearing Debt



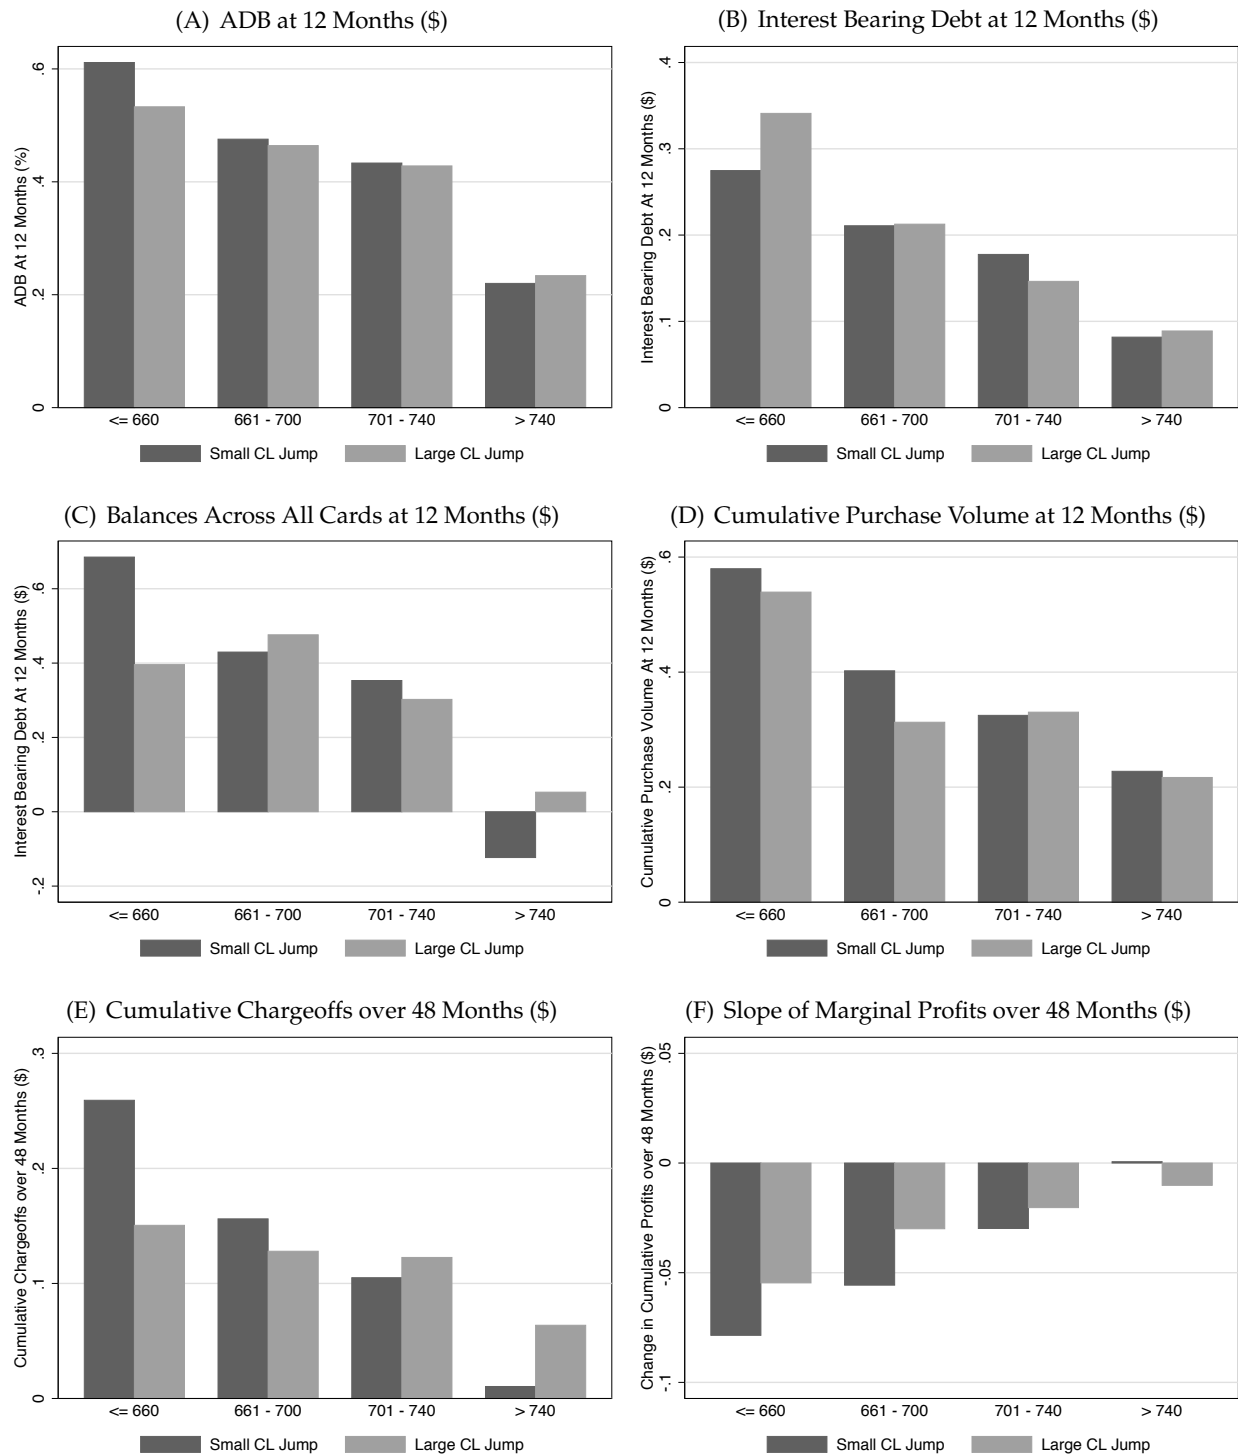
Note: Figure shows the effects of a \$1,000 increase in credit limits on the cumulative probability of positive interest bearing debt for different FICO score groups and different time horizons after account origination. FICO score groups are determined by FICO score at account origination.

Figure A11: Main Results: Binned Scatter Plots by FICO Score



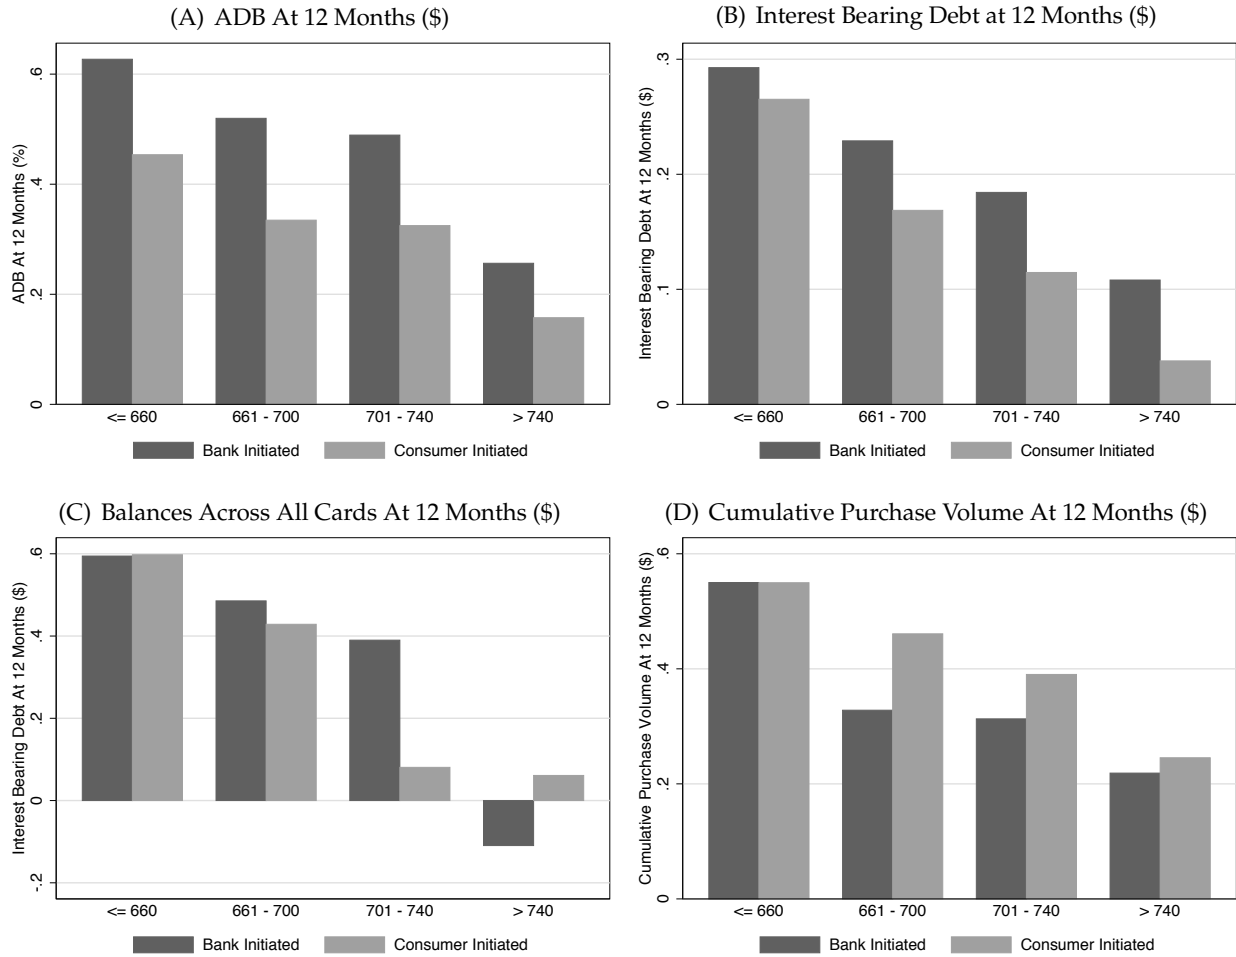
Note: Figure shows binned scatter plots of the relationship between FICO score and regression discontinuity estimates of the effect of a \$1 increase in credit limits. We residualize the x-variable and y-variable on controls, before binning and plotting. The controls are fully interacted dummies for origination quarter, bank, and a "zero initial APR" indicator, as well as fully interacted dummies for loan channel and a "zero initial APR" indicator. Panel A shows effects on average daily balances on the treated credit card after 12 months. Panel B shows effects on interest bearing debt on the treated card after 12 months. Panel C shows effects on total balances aggregated across all credit cards held by the account holder after 12 months. Panel D shows effects on cumulative purchase volume on the treated card after 12 months. Panel E shows effects on cumulative chargeoffs after 48 months. Panel F shows effects on cumulative marginal profits after 48 months of increasing credit limits by \$1,000 from their equilibrium value (i.e., the slope of marginal profits).

Figure A12: Main Results: By Size of Credit Limit Jump



Note: Figure shows main results, splitting quasi-experiments by whether the size of the credit limit jump at the discontinuity is above or below the median value. We show regression discontinuity estimates of the effect of a \$1 increase in credit limits. Panel A shows effects on average daily balances on the treated credit card after 12 months. Panel B shows effects on interest bearing debt on the treated card after 12 months. Panel C shows effects on total balances aggregated across all credit cards held by the account holder after 12 months. Panel D shows effects on cumulative purchase volume on the treated card after 12 months. Panel E shows effects on cumulative chargeoffs after 48 months. Panel F shows effects on cumulative marginal profits after 48 months of increasing credit limits by \$1,000 from their equilibrium value (i.e., the slope of marginal profits).


Figure A13: Borrowing Behavior by Source of Credit Card Origination




Note: Figure shows main results on borrowing and spending behavior, splitting quasi-experiments by whether the bank or the consumer initiated the origination of the credit card. We show regression discontinuity estimates of the effect of a \$1 increase in credit limits. Panel A shows effects on average daily balances on the treated credit card after 12 months. Panel B shows effects on interest bearing debt on the treated card after 12 months. Panel C shows effects on total balances aggregated across all credit cards held by the account holder after 12 months. Panel D shows effects on cumulative purchase volume on the treated card after 12 months.

Figure A14: Borrowing Behavior by Reported Income at Origination


(A) Correlation: FICO vs. Reported Income at Origination




(B) ADB At 12 Months (\$)




(C) Interest Bearing Debt at 12 Months (\$)



(D) Balances Across All Cards At 12 Months (\$)

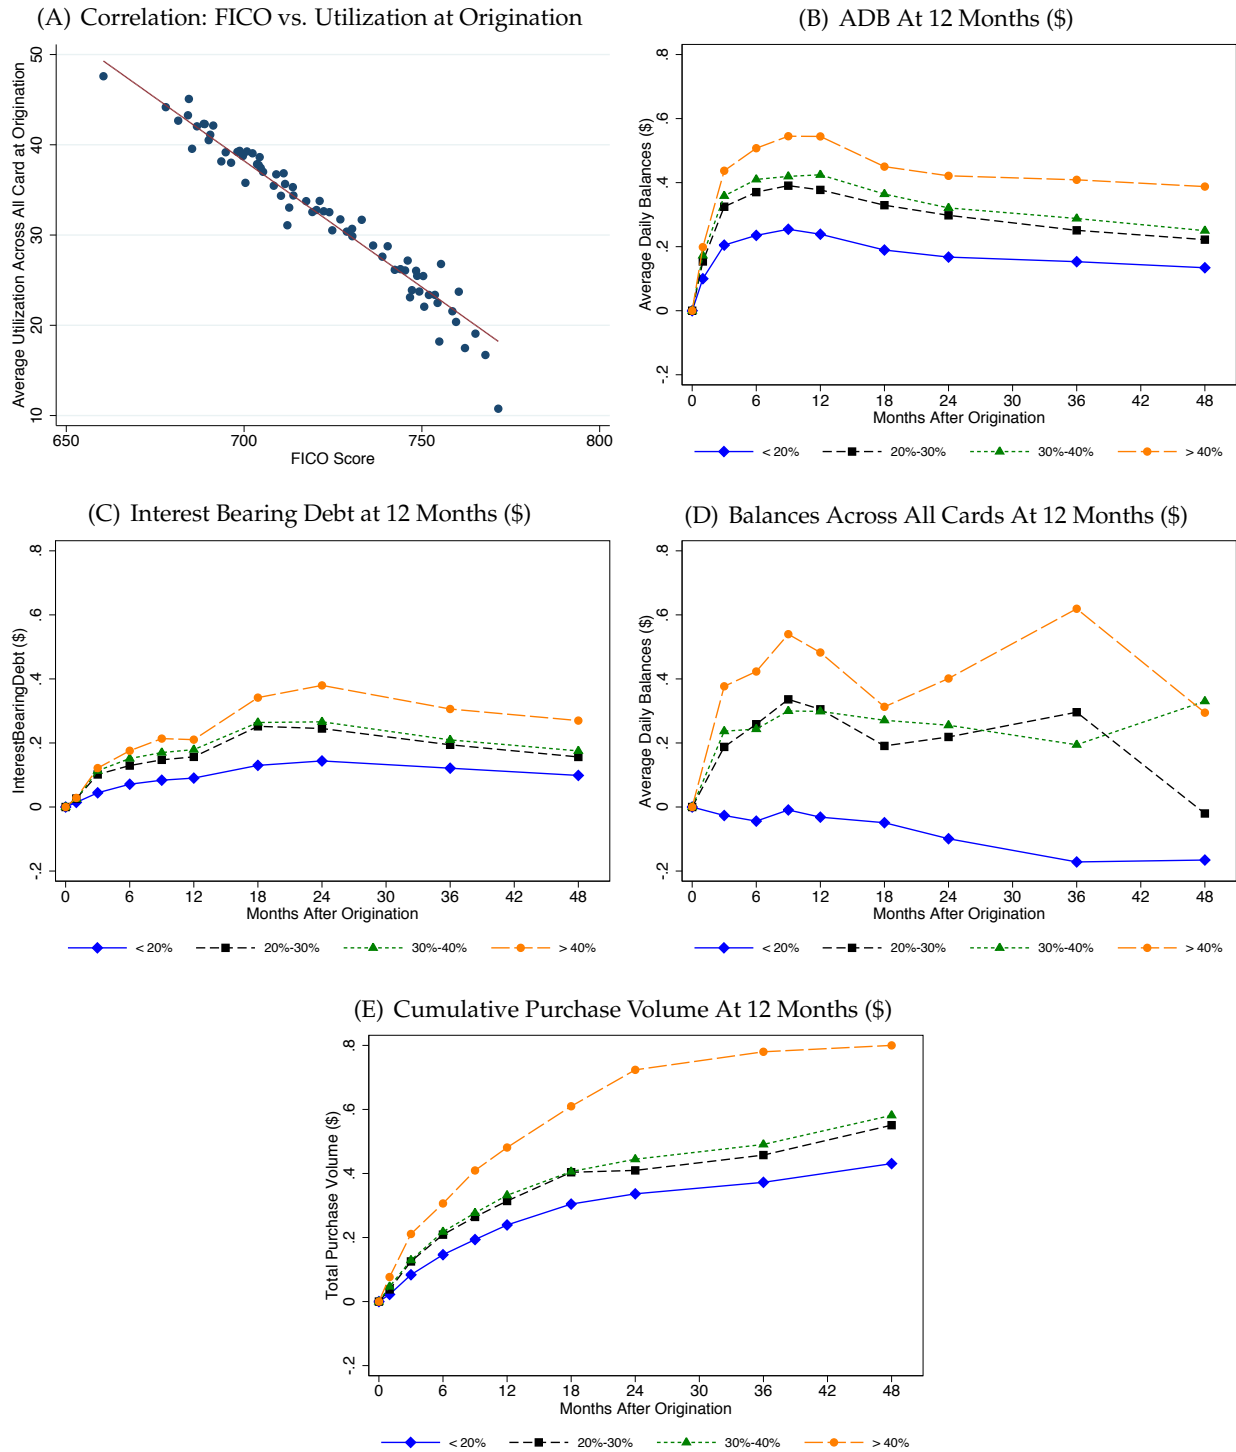


(E) Cumulative Purchase Volume At 12 Months (\$)



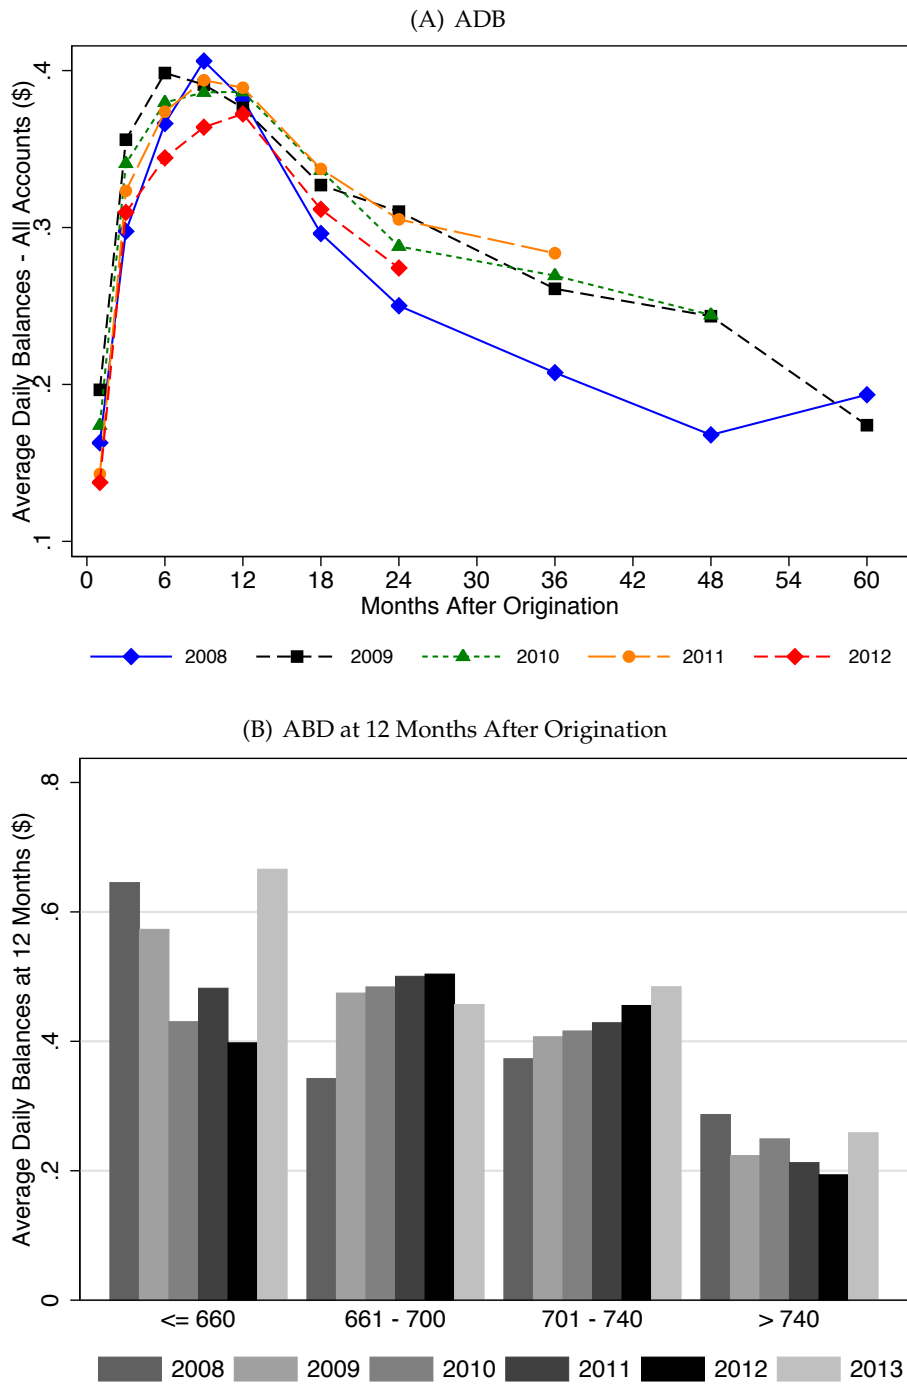
Note: Figure shows the effects of credit limits on borrowing and spending. We show regression discontinuity estimates of the effect of a \$1 increase in credit limits for different groups of reported income at origination. Panel A shows the correlation between FICO score and reported income across our experiments. Panel B shows effects on average daily balances on the treated credit card after 12 months. Panel C shows effects on interest bearing debt on the treated card after 12 months. Panel D shows effects on total balances aggregated across all credit cards held by the account holder after 12 months. Panel E shows effects on cumulative purchase volume on the treated card after 12 months.

Figure A15: Borrowing Behavior: By Utilization at Origination




Note: 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 groups of credit card utilization across all cards at origination. Utilization is determined as the ratio of total balances to total credit limits in the credit bureau data. Panel A shows the correlation between FICO score and utilization across our experiments. Panel B shows effects on average daily balances on the treated credit card after 12 months. Panel C shows effects on interest bearing debt on the treated card after 12 months. Panel D shows effects on total balances aggregated across all credit cards held by the account holder after 12 months. Panel E shows effects on cumulative purchase volume on the treated card after 12 months.

Figure A16: Marginal Effect on ADB Over Time




Note: Figure shows the effects of credit limits on borrowing and spending. Panel A shows regression discontinuity estimates of the effect of a \$1 increase in credit limits on average daily balances (ADB) for quasi-experiments originated in different years and different time horizons after account origination. Panel B shows regression discontinuity estimates of the effect of a \$1 increase in credit limits on average daily balances (ADB) after 12 months for different FICO score groups and for quasi-experiments originated in different years.

Figure A17: Slope of Marginal Profit Over Time




Note: Figure shows the effect of a \$1,000 increase in credit limits on cumulative marginal profits after 48 months, separately for accounts originated in Q1 or Q2 of 2008 (pre Financial Crisis) and accounts originated in Q3 of 2008 or later. See Figure 9 for more details.

Figure A18: Graphical Illustration of Proposition 1.



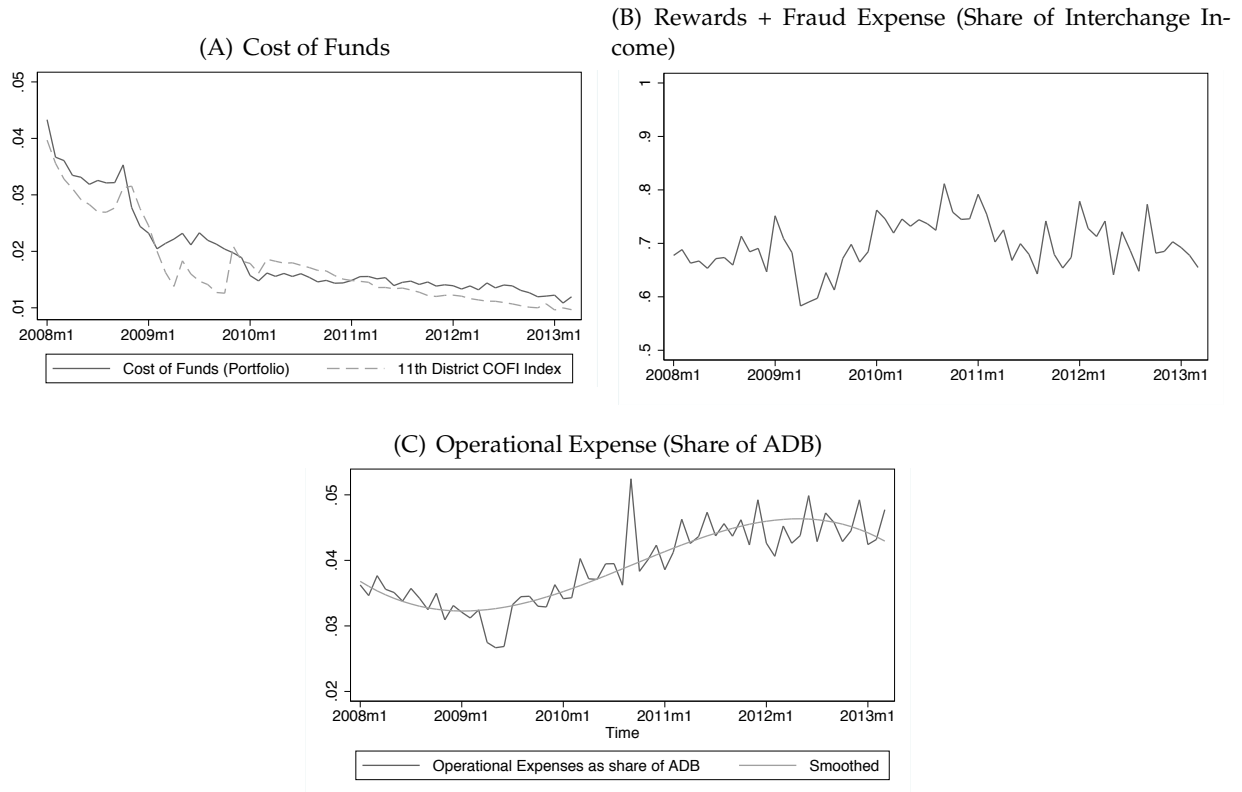
Note: Figure provides a graphical illustration of Proposition 1. The figure shows two marginal profit functions, $MP_x(CL)$ and $MP_y(CL)$, that have the same optimal credit limit: $MP_x(CL^*) = 0$ and $MP_y(CL^*) = 0$. Average profit is defined as $AP(CL) = \frac{\int_0^{CL} MP(CL)dCL}{CL}$. AP_x is the shaded area under the curve MP_x divided by CL^* , and AP_y is the shaded area under the curve MP_y divided by CL^* .

Figure A19: Interchange Income at Portfolio Level



Note: Figure corresponds to Figure A.II in Agarwal et al. (2015b). It shows the ratio of interchange income to purchase volume at the portfolio level.

Figure A20: Cost Components at Portfolio Level



Note: Figure corresponds to Figure A.I in Agarwal et al. (2015b). Panels show plots of cost components at the portfolio level by month. Panel A shows the cost of funds, calculated as the annualized interest expense (“total interest expense accrued for the month to fund credit card receivables”) as a share of average daily managed receivables for that month. It also shows the 11th District Cost of Funds Index (COFI). Panel B shows the share of rewards and fraud expenses as a ratio of the interchange income. These figures are constructed using the monthly general purpose credit card portfolio-level data. Numbers are averages across banks. Panel C shows the share of annualized operational expenses (including marketing and acquisition, collections, servicing, cardholder billing, processing payments, and card issuing and administration) as a share of average daily managed receivables.

Table A1: Scorecard Costs: Lower Bounds

	Per Account					Aggregate	
	ΔCL	1,000 x $MP'(CL^*)$	Scorecard Costs (χ)	Average Total Profit	% of Average Total Profits	Number of Accounts	Scorecard Costs (χ)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
FICO							
≤ 660	\$1,124	-0.068	\$10.74	\$365	2.9%	1,675	\$17,992
661-700	\$1,363	-0.037	\$8.59	\$126	6.8%	807	\$13,868
701-740	\$1,434	-0.024	\$6.17	\$55	11.1%	1,065	\$13,144
> 740	\$1,666	-0.004	\$1.39	\$75	1.9%	1,256	\$3,486

Note: Table shows lower bounds of the the implied scorecard costs χ that rationalize the observed step-function relationship between FICO scores and credit limits. The analysis focuses on accounts within +/- 5 FICO score points of the discontinuities. Columns 1 and 2 show the inputs into the scorecard cost function: the jump in credit limits (ΔCL) and the effect of a \$1,000 increase in credit limits on marginal profits ($MP'(CL^*)$) over our baseline 48 month horizon (i.e., the "slope" of marginal profits). Column 3 shows the estimated scorecard cost calculated using Equation A3. Column 4 shows average profits for these accounts over a 48 month horizon and column 5 shows scorecard costs as a percentage of average profits. Column 6 shows the average number of accounts within the +/- 5 FICO score point band and column 7 shows aggregate scorecard costs.

Table A2: Effect on Interest Bearing Debt: Extensive versus Intensive Margins

	Months After Account Origination			
	12	24	36	48
Panel A: Effect on Interest Bearing Debt				
FICO				
<i>≤660</i>	30%	46%	40%	35%
<i>661-700</i>	21%	34%	30%	28%
<i>701-740</i>	16%	27%	23%	21%
<i>>740</i>	8%	13%	12%	12%
Panel B: Extensive Margin Effect				
FICO				
<i>≤660</i>	4.1%	2.1%	1.0%	0.3%
<i>661-700</i>	4.9%	1.9%	1.1%	0.4%
<i>701-740</i>	4.4%	3.0%	1.3%	0.7%
<i>>740</i>	2.4%	1.3%	0.2%	0.1%
Panel C: Intensive Margin Effect				
FICO				
<i>≤660</i>	25.9%	43.8%	39.1%	34.8%
<i>661-700</i>	16.4%	32.5%	28.6%	27.5%
<i>701-740</i>	11.5%	23.7%	21.8%	20.1%
<i>>740</i>	6.0%	11.3%	11.5%	11.5%

Note: Table decomposes the effect on interest bearing debt into extensive and intensive margin responses. Panel A shows the baseline regression discontinuity estimates of the effect of a \$1 increase in credit limits on interest bearing debt, reproduced from Table 5. Panels B and C decompose this effect on the extensive and intensive margins. See text for more details on the decomposition. 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.

Table A3: Sensitivity to Population Reweighting

	Average Daily Balances	Interest Bearing Debt	Total Balance Across All Cards	Cumulative Purchase Volume	Chargeoffs	Slope of Marginal Profits*
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Baseline Estimates						
<i>FICO</i>						
<i>≤660</i>	0.58	0.30	0.59	0.56	0.216	-0.068
<i>661-700</i>	0.47	0.21	0.46	0.35	0.136	-0.037
<i>701-740</i>	0.43	0.16	0.32	0.33	0.119	-0.024
<i>>740</i>	0.23	0.08	-0.05	0.22	0.037	-0.004
Panel B: Population Reweighted Estimates						
<i>FICO</i>						
<i>≤660</i>	0.698	0.311	0.759	0.825	0.181	-0.082
<i>661-700</i>	0.508	0.227	0.498	0.436	0.149	-0.042
<i>701-740</i>	0.379	0.161	0.245	0.315	0.104	-0.022
<i>>740</i>	0.135	0.027	-0.301	0.205	-0.027	0.008

Note: Table shows the effects of population reweighting on the main parameter estimates. Panel A reproduces the baseline estimates from Table 5 and Table 7. Panel B constructs estimates for each FICO group where we reweight by the distribution of FICO scores in the population within each FICO group. See text for more details.

* Effect of \$1,000 increase in credit limits on marginal effect.