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.I). We also discuss programs created to increase the availability of affordable capital for U.S. banks (Section A.II), 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.III) and the Targeted Longer-Term Refinancing Operations (TLTRO) at the European Central Bank (Section A.IV).

A.I 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 VI 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.II 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.III 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.IV 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.I 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.[1] 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.”

[1]Since 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)."

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, $z = \{z_1, z_2 \ldots z_K\}$ 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:

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

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 limits to revenue and costs:

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

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.[2]

B.II 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 A.VIII provides a conceptual illustration for how we calculate a lower bound for the scorecard costs, which we denote $\underline{\chi}$, 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 $MP'(CL) = 0$ 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 shaded triangle to the right of \$2,500.

[2] There 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.

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

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

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

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 A.I 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 II, 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, $\overline{x}$ be the cutoff of interest, and $\overline{\overline{x}}$ be the additional cutoff for which we want to control. Assume that the true data generating process is given by

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

Under this process, the outcome is linearly increasing in x_i with a slope of β, and has a jump of δ at threshold $\overline{x}$, and a jump of γ at threshold $\overline{\overline{x}}$. For our Monte Carlo exercise, we assume that x_i takes integer values on $[-10, 10]$. We set $\overline{x} = 0$ and $\overline{\overline{x}} = 5$. 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 A.IX 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 $\overline{x} = 0$ and $\overline{\overline{x}} = 5$ 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 $\mathbb{I}(x_i \geq \overline{\overline{x}})$. Because we exclude the indicator, we estimate "too steep" a slope to the right of $\overline{x} = 0$, generating a significantly downward biased estimates of δ. Panel C shows the distribution of δ_s when we include the indicator for $\overline{\overline{x}} = 5$. 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 V.B, 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:

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

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

Appendix Figure A.X 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 II. 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 A.II shows the result of the decomposition in Equation A6. Panel A reproduces the baseline effect on interest bearing debt from Table V. Panels B and C decompose this effect on the extensive and intensive margins. Consistent with the small effects in Appendix Figure A.X, the decomposition shows that the extensive margin effect is relatively small compared to the intensive margin response. The extensive margin response is relatively more import 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.I 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 VII).

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 interacted 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.[3]

Appendix Figure A.XI 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 VII. 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.II Robustness of Estimates to Size of Credit Limit Jump

As highlighted in Panel A of Figure IV, 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.[4] 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

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

[4] 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.

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.

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

Appendix Figure A.XII 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.III 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 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:

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

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 A.XIII 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.IV 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 A.XIV 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 A.XIV replicate the impulse response plots shown in Figure VII, 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 A.XV 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.V 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 A.XVI 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 A.XVI 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 A.XVII 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.[5] 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 quantitively the main source of heterogeneity in our data.

E.VI 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 A.XI 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.[6] We

[5] We 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.

[6] For 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,

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 A.I 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 A.III 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.

F Linearity Assumption

In Section VII, 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 opti-

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.

mal 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 A.XVIII 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 $MP'_x(CL^*) < MP'_x(CL^*)$. Second, if $AP_x(CL^*) < AP_y(CL^*)$, then by the same logic, $MP'_x(CL^*) > MP'_x(CL^*)$. But then since $AP_x(CL^*) > AP_y(CL^*) \Rightarrow MP'_x(CL^*) < MP'_x(CL^)$ and $AP_x(CL^*) < AP_y(CL^*) \Rightarrow MP'_x(CL^*) > MP'_x(CL^)$, we know that $AP_x(CL^*) > AP_y(CL^*) \iff MP'_x(CL^*) < MP'_x(CL^*)$. □

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 $\tilde{MP}(\tilde{CL}) = \frac{1}{CL^*} MP(CL^* \cdot \tilde{CL})$. The normalized marginal profit function is constructed so that it takes on the optimal value at $CL = 1$. That is, $\tilde{MP}(1) = \frac{1}{CL^*} MP(CL^*) = 0$. By construction, at the optimal values, the slopes of the un-normalized and normalized marginal profit functions are identical: $\tilde{MP}'(1) = \frac{1}{CL^*} MP'(CL^*) \cdot CL^* = MP'(CL^*)$. Define average profits as $\tilde{AP}(\tilde{CL}) = \frac{\int_0^{\tilde{CL}} \tilde{MP}(\tilde{CL})d\tilde{CL}}{\tilde{CL}}$.

Proposition 2. *If the normalized marginal profit functions, $\tilde{MP}_x(\tilde{CL})$ and $\tilde{MP}_y(\tilde{CL})$, 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 $\tilde{AP}_x(1) > \tilde{AP}_y(1) \iff \tilde{MP}'_x(1) < \tilde{MP}'_y(1)$. Second,

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

where the third line involves a change of variables in which we replace $\tilde{CL}$ with $\frac{CL}{CL^*}$. Third, by construction, we know that $\tilde{MP}'(1) = MP'(CL^*)$. 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. □

G Assigning Measures of Profitability to Credit Card Accounts

As discussed in Section VII, 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 VI 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 II 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. (2015) show that these percentages, which Table II 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 A.XIX, taken from Agarwal et al., 2015). 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 A.XX, taken from Agarwal et al., 2015). 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 A.XX, taken from Agarwal et al., 2015). 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 A.XX, taken from Agarwal et al., 2015). 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.

References

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. 2015. "Regulating Consumer Financial Products: Evidence from Credit Cards." *Quarterly Journal of Economics*, 130(1): 111–164.

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.

GAO. 2009. "Government Accountability Office: Rising Interchange Fees Have Increased Costs for Merchants, but Options for Reducing Fees Pose Challenges." *GAO-10-45*.

Imbens, Guido, and Karthik Kalyanaraman. 2011. "Optimal Bandwidth Choice for the Regression Discontinuity Estimator." *Review of Economic Studies*, 79(3): 933–959.

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.

Office of the Comptroller of the Currency. 2015. *Comptroller's Handbook: Credit Card Lending.*

Figure A.I: FICO Score, Population Distribution

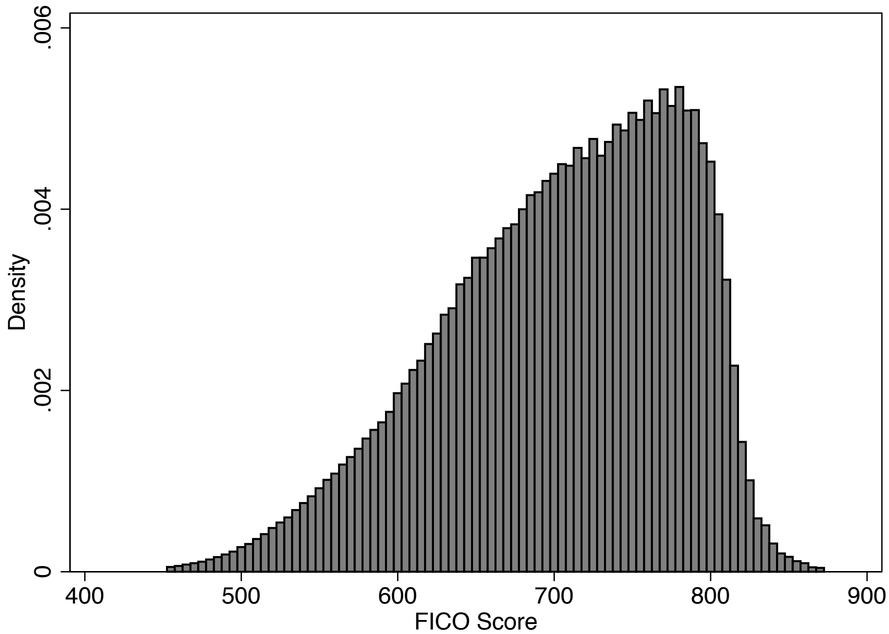

Note: Figure shows the distribution over FICO scores of all credit cards issued by the banks in our sample, averaged over the period from January 2008 to November 2013.

Figure A.II: Initial Borrower Characteristics Around Credit Limit Quasi-Experiments

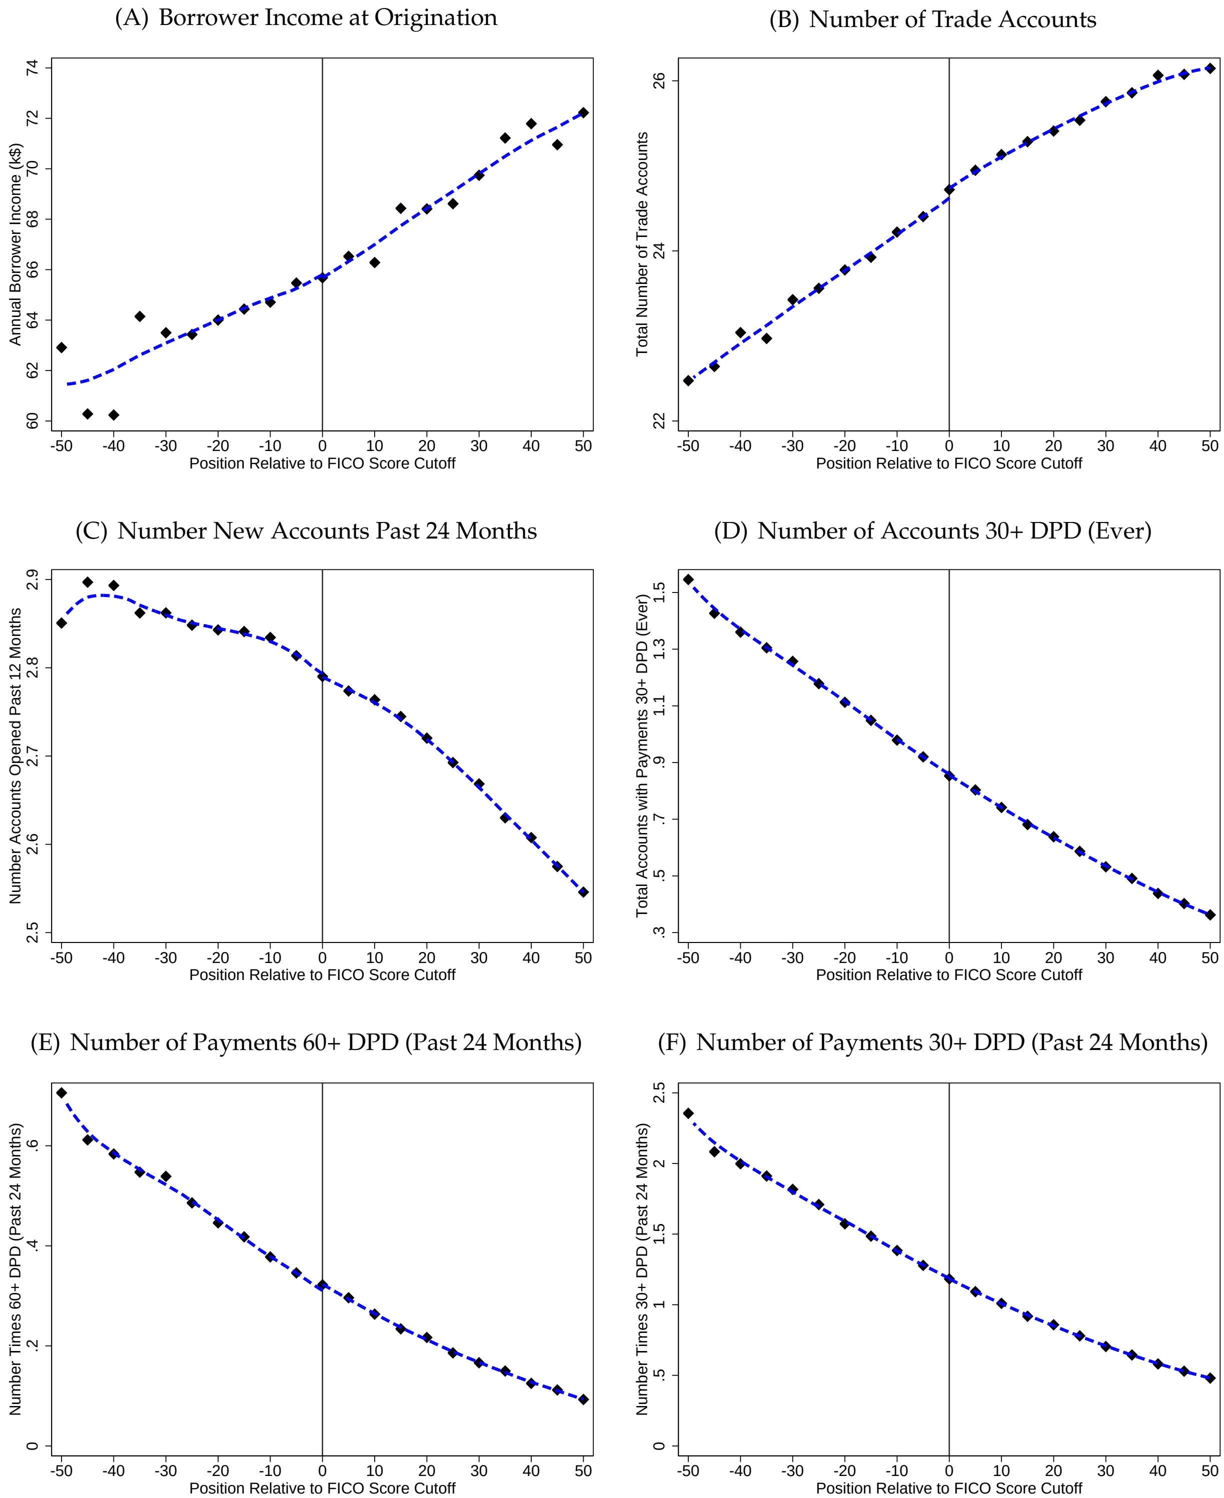

Note: Figure plots average borrower characteristics around our 743 pooled credit limit quasi-experiments. The horizontal axis shows FICO score at origination, centered at the FICO score cutoff. The vertical axis shows reported borrower income at origination (Panel A), total number of trade accounts (Panel 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 A.III: Credit Limits and Cost of Funds in the Time Series

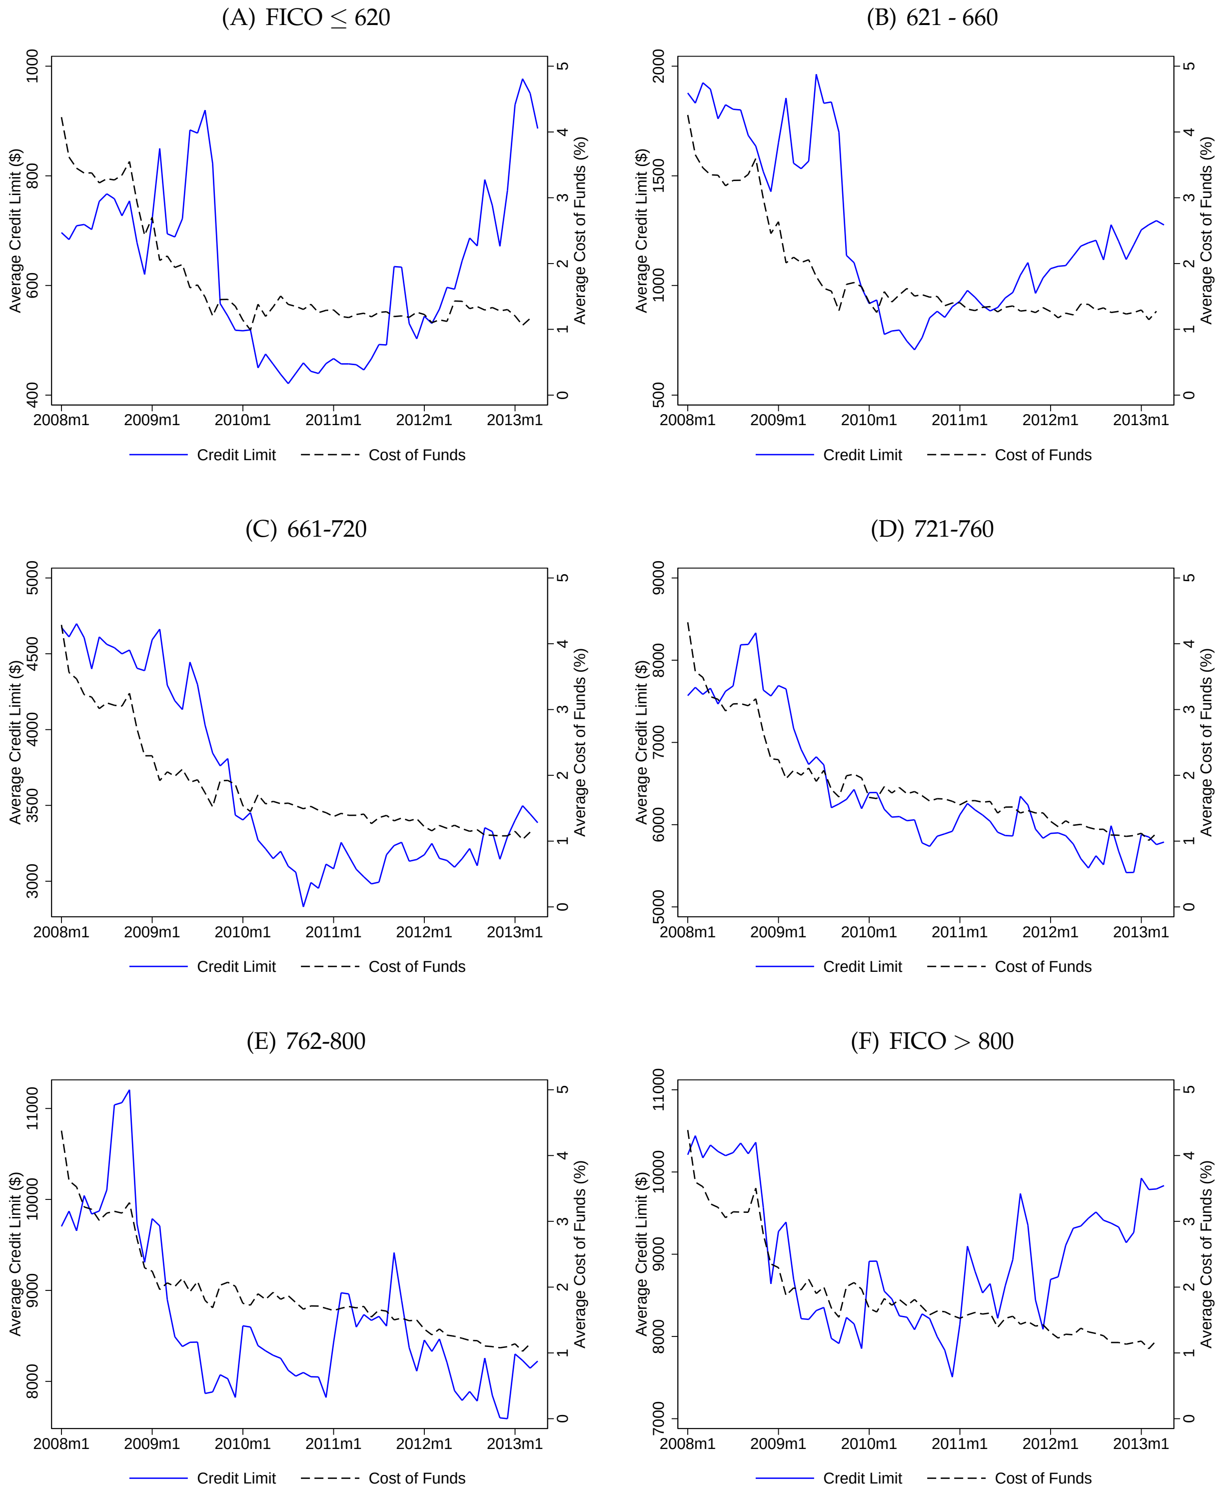

Note: Figure shows average credit limits on newly originated credit cards (solid line) and average cost of funds (dashed line) over time by FICO score group.

Figure A.IV: Credit Card Interest Rates vs. Federal Funds Rate

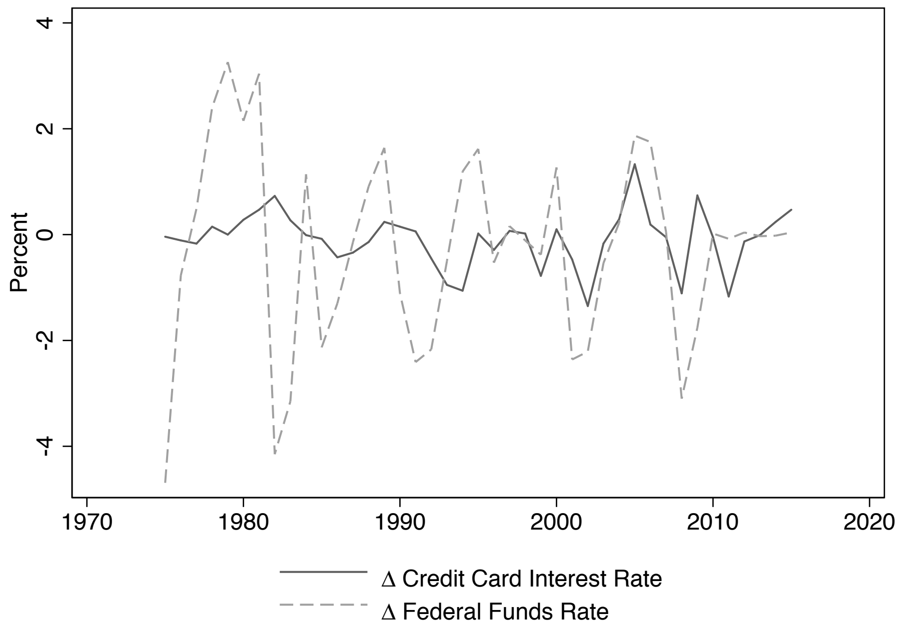

Note: Figure shows the year-on-year change in credit card interest rates and year-on-year change in the Federal Funds Rate between 1974 and 2015. Before 1994, credit card interest rates were those reported in the Federal Reserve's "Quarterly Report of Interest Rates on Selected Direct Installment Loans." 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 A.V: Credit Card Credit Limits vs. Interest Rates

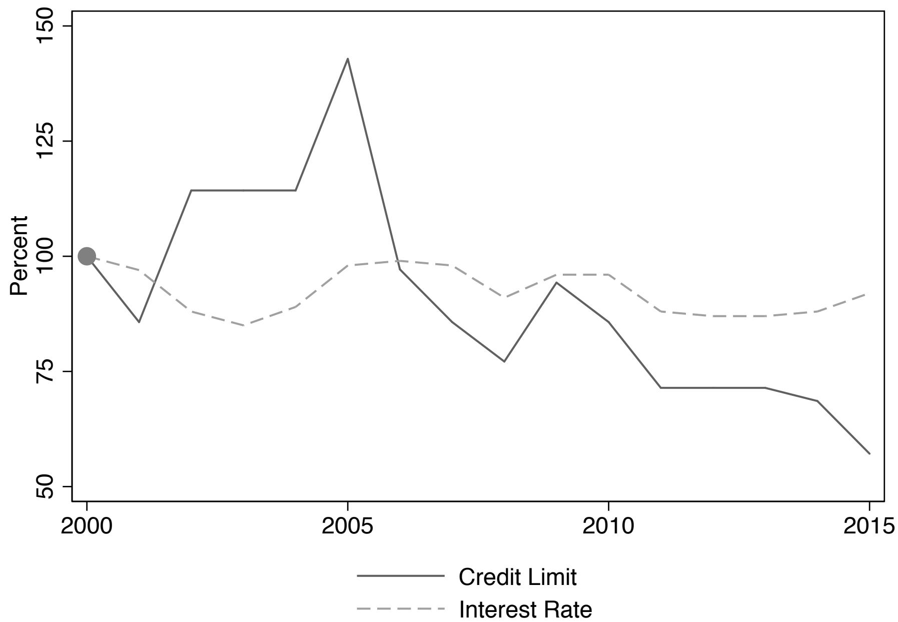

Note: Figure shows credit card credit limits and interest rates between 2000 and 2015, with the values normalized to 100% in the year 2000 for comparability. The interest rates are from Federal Reserve's "Quarterly Report of Credit Card Interest Rates" for those credit card holders incurring interest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

Figure A.VI: Probability of Delinquency at 48 Months After Origination

(A) Probability 60+ DPD

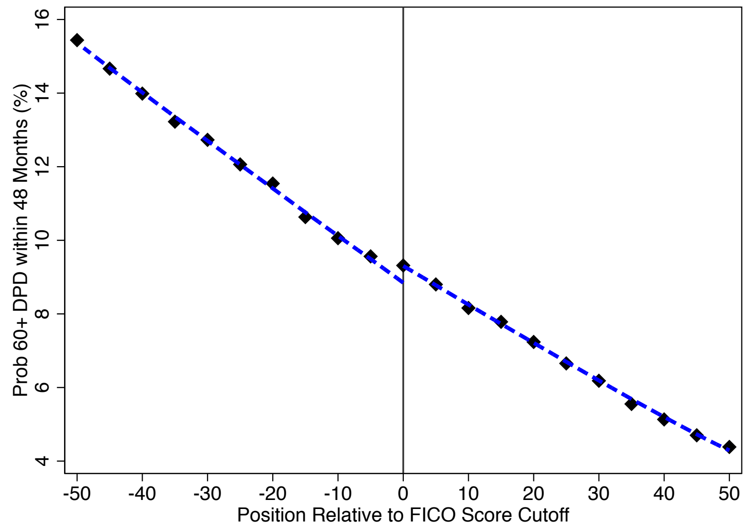

(B) Probability 90+ DPD

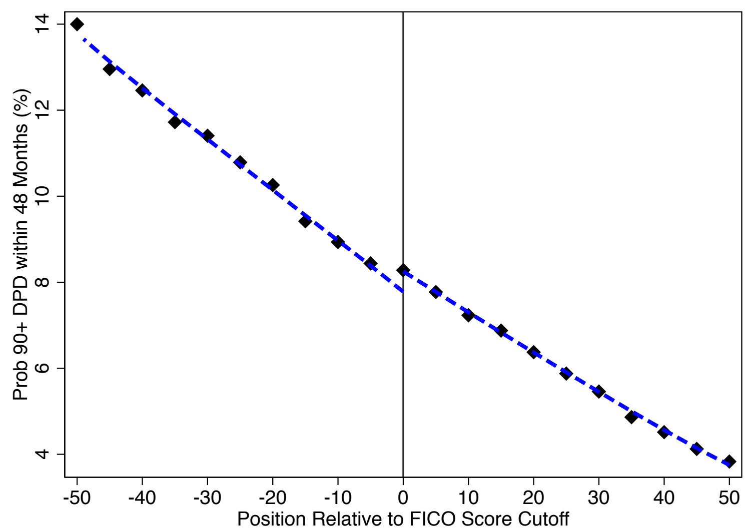

Note: Figure shows the effects of credit limits on the probability of delinquency around our 743 pooled credit limit quasi-experiments. Panel A shows effects on the cumulative probability of an account being more than 60 days past due (60+ DPD); Panel B shows effects on the cumulative probability of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

Figure A.VII: Total Revenue, Total Cost, and Components

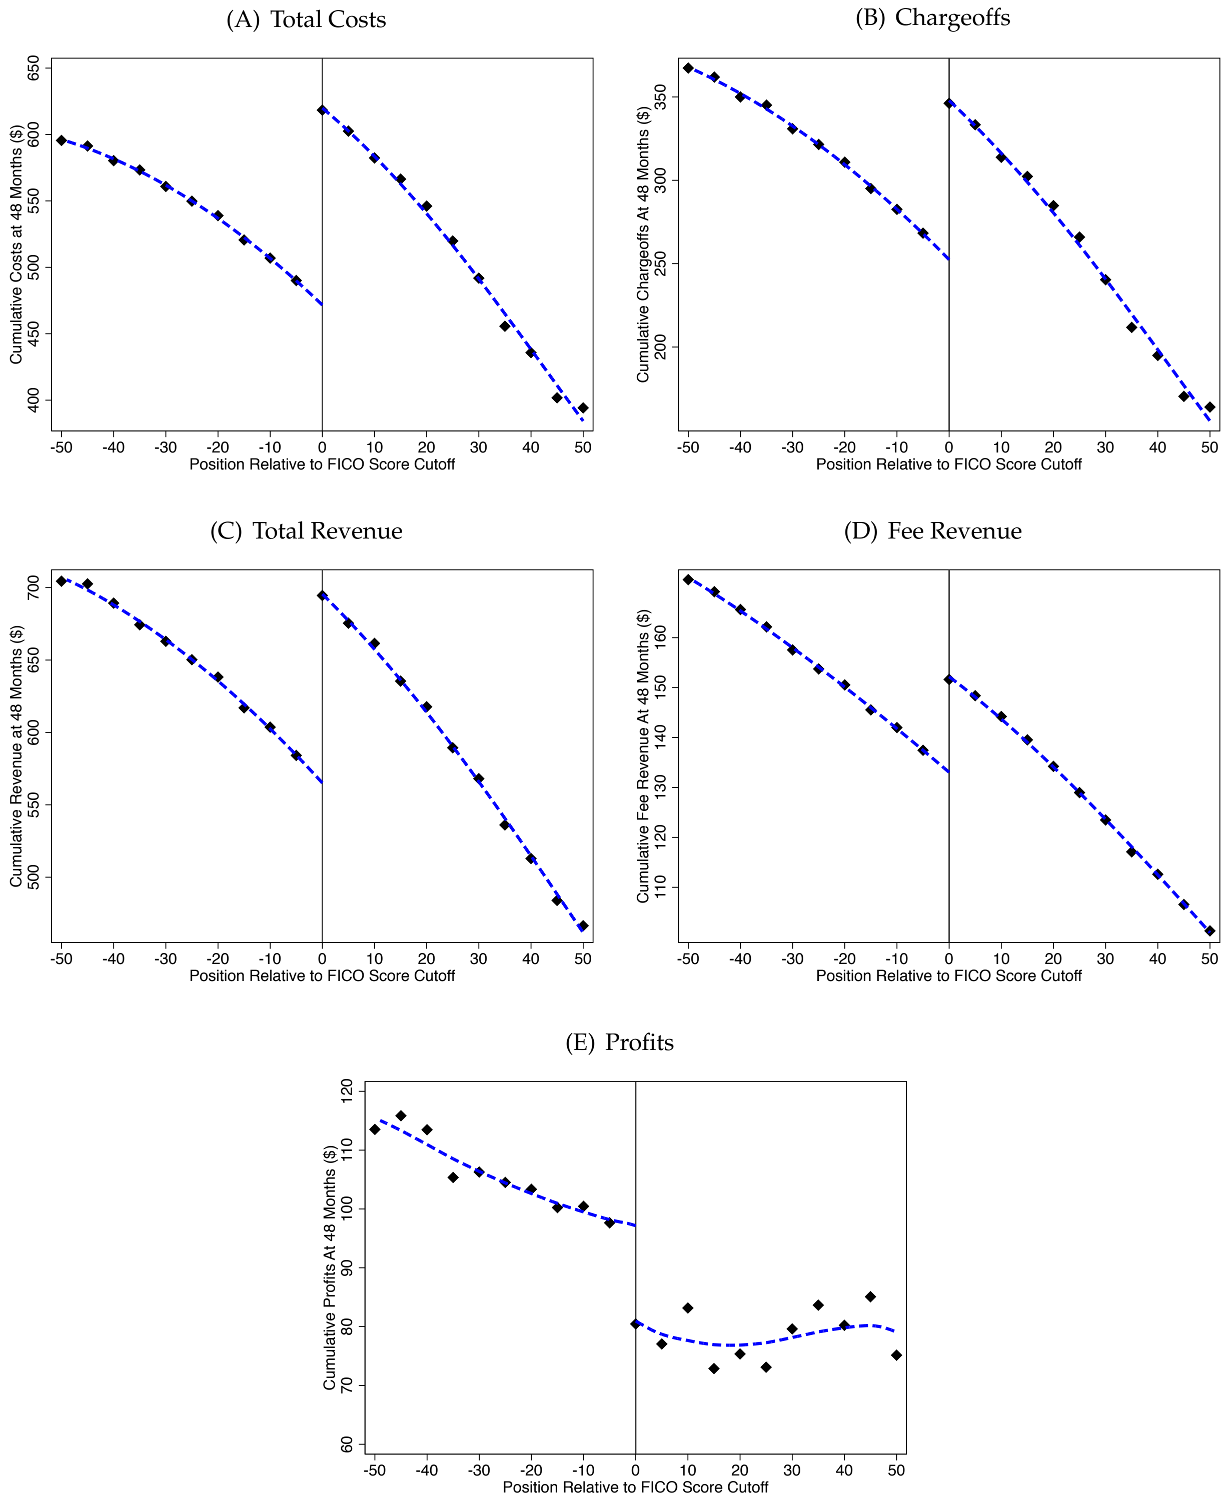

Note: Figure shows the effects of credit limits on cumulative total costs (Panel A), cumulative chargeoffs (Panel B), cumulative total revenue (Panel C), cumulative fee revenue (Panel D), and cumulative profits (Panel E), all measured over the first 48 months after account origination. These plots pool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

Figure A.VIII: Scorecard Costs (χ): Conceptual Framework

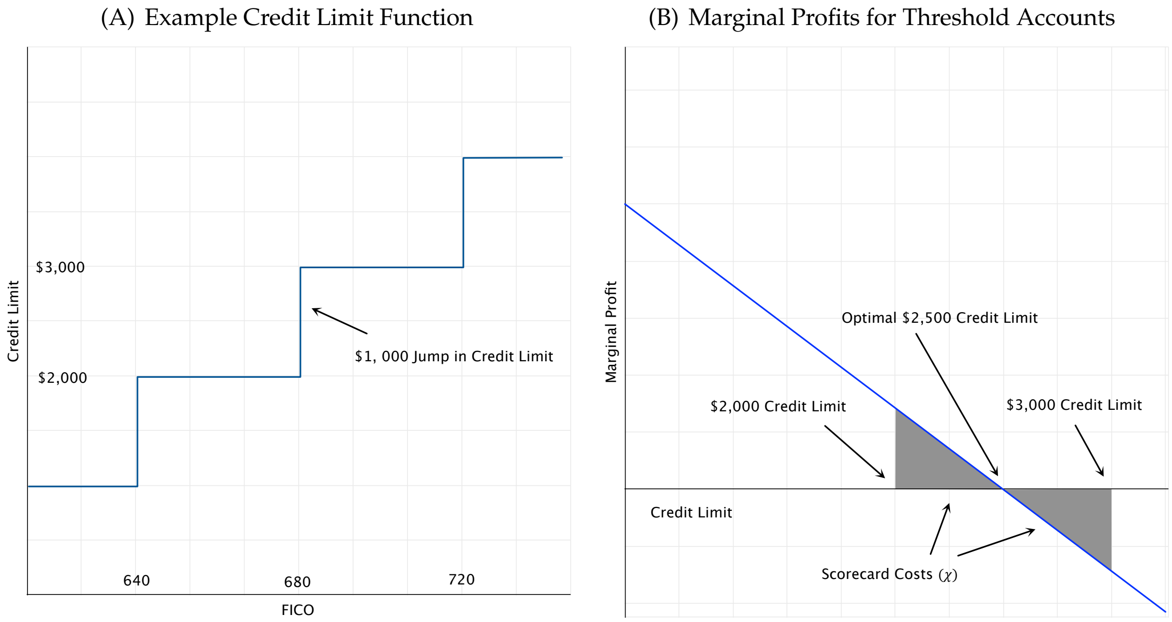

Note: Figure illustrates how we calculate a lower bound for the scorecard costs χ for a given "jump" in credit limits. Panel A shows an origination group where credit limits are a step function of FICO score. We focus on accounts with FICO scores that are "close" to 680 where the credit limit jumps from $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 A.IX: Monte Carlo Simulations: Controlling for Other Experiments

(A) Mean Outcome from a Single Simulation

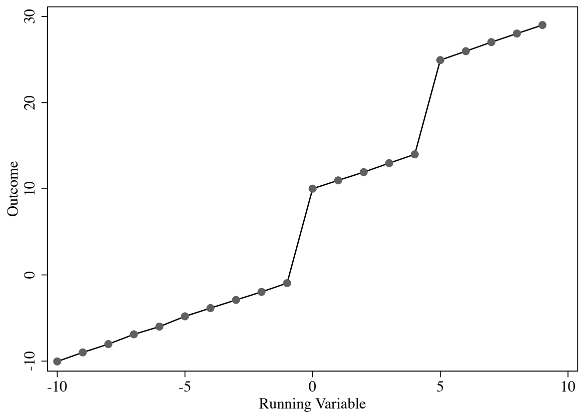

(B) No Indicator for Additional Cutoff

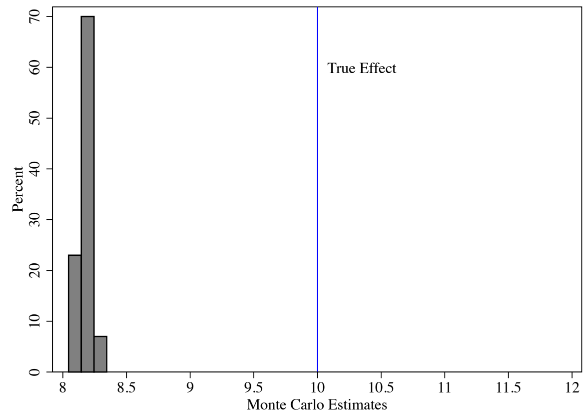

(C) Indicator for Additional Cutoff

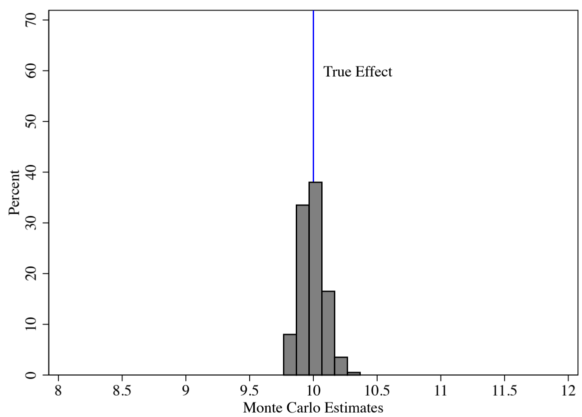

Note: Figure shows output from a Monte Carlo exercise that assesses whether we can control for the presence of additional cutoffs with an indicator variable that is equal to 1 for all FICO scores above this additional cutoff value. Panel A plots the average value of the outcome from a single simulation 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 Carlos 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 A.X: Effect of $1K Increase in Credit Limits on Probability of Positive Interest Bearing Debt

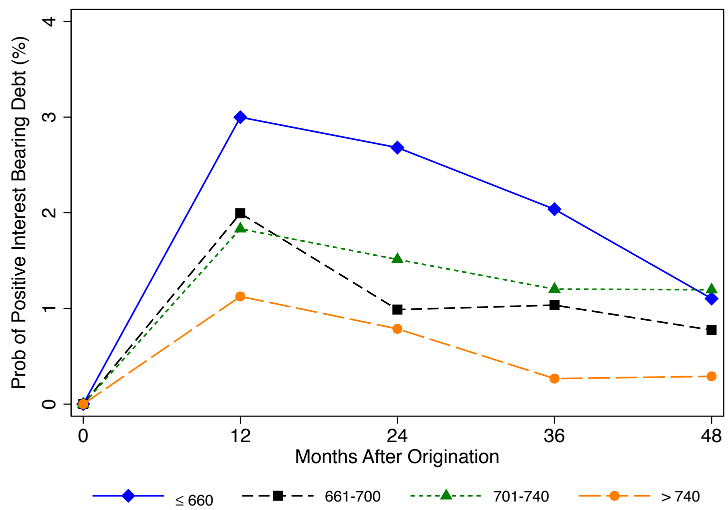

Note: Figure shows the effects of a $1,000 increase in credit limits on the cumulative probability of positive interest bearing debt for different FICO score groups and different time horizons after account origination. FICO score groups are determined by FICO score at account origination.

Figure A.XI: Main Results: Binned Scatter Plots by FICO Score

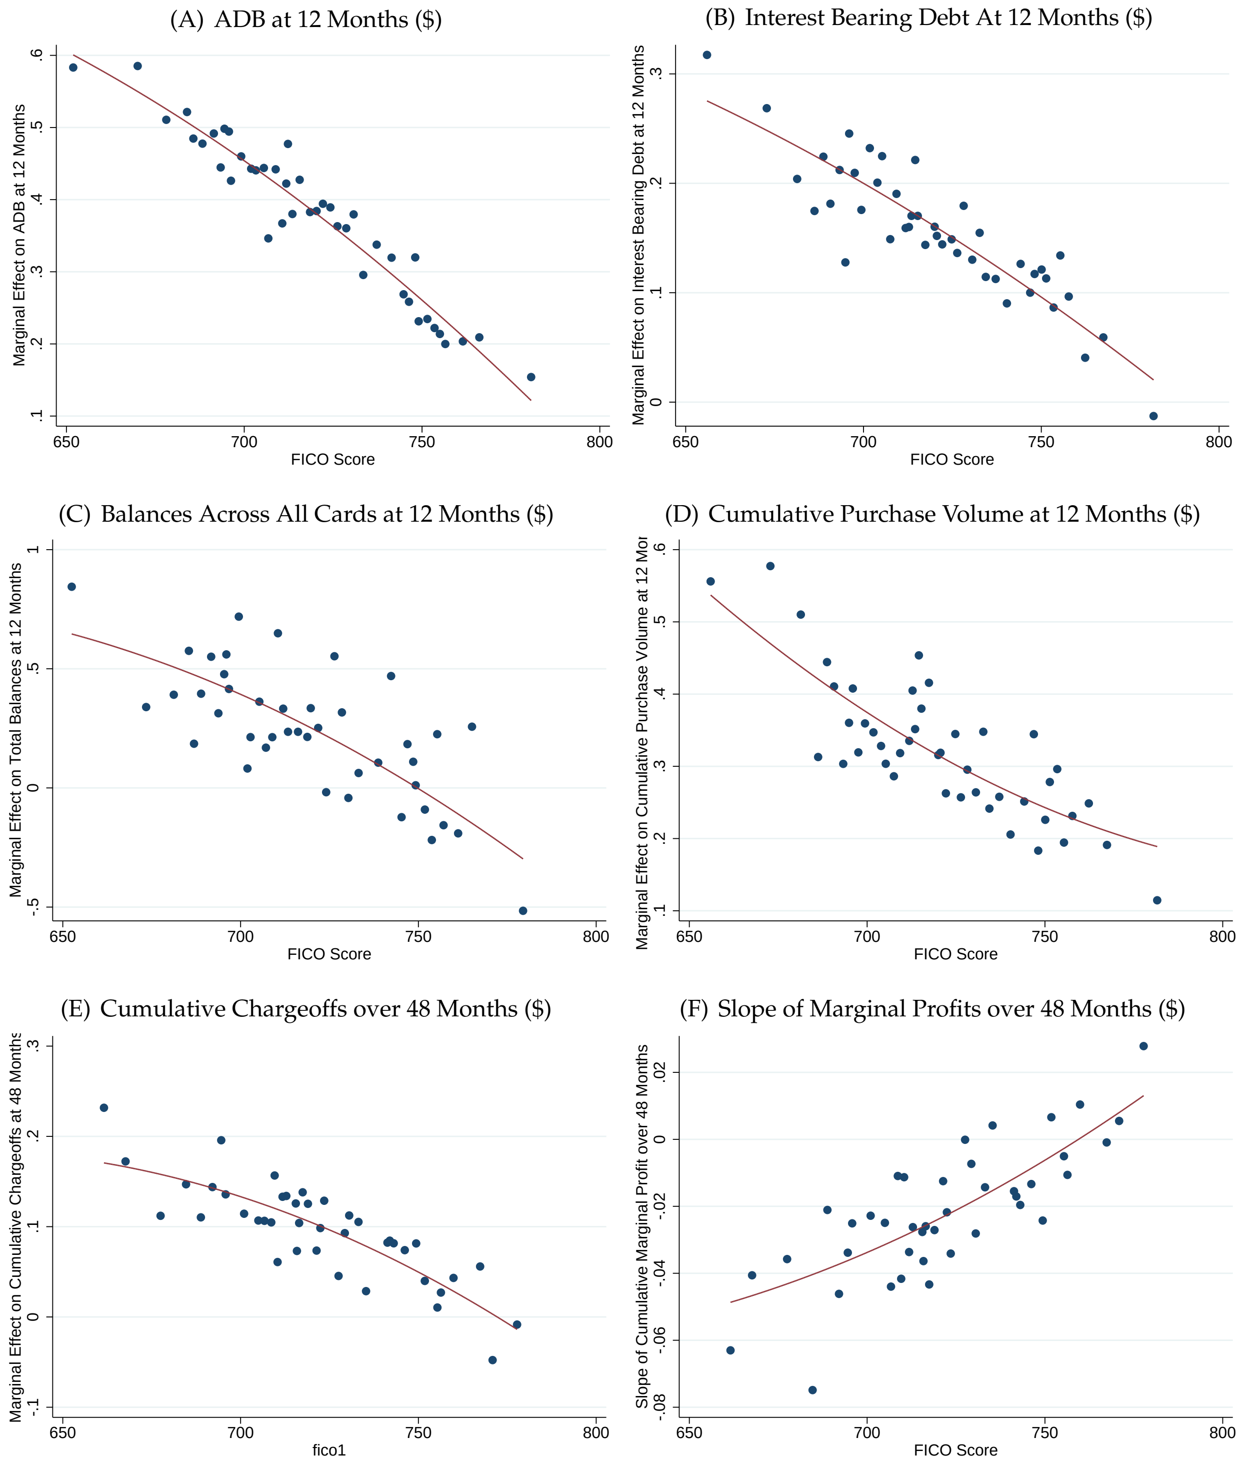

Note: Figure shows binned scatter plots of the relationship between FICO score and regression discontinuity estimates of the effect of a $1 increase in credit limits. We residualize the x-variable and y-variable on controls, before binning and plotting. The controls are fully interacted dummies for origination quarter, bank, and a "zero initial APR" indicator, as 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 A.XII: Main Results: By Size of Credit Limit Jump

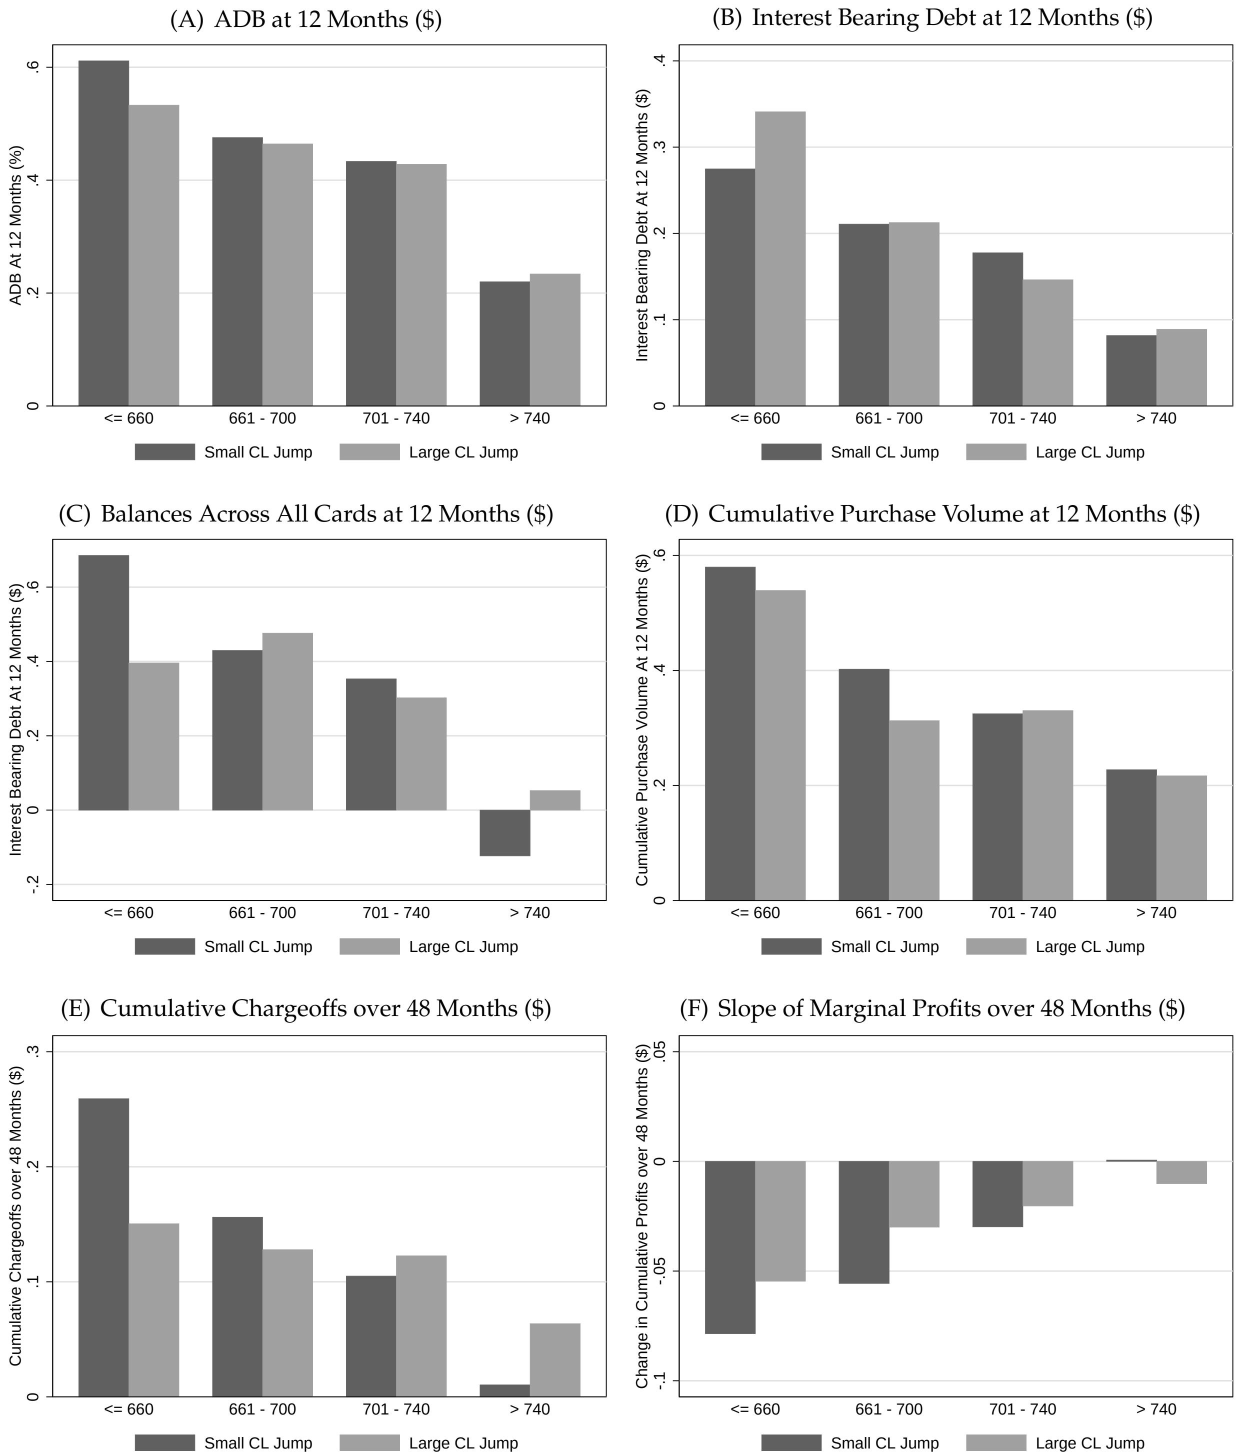

Note: Figure shows main results, splitting quasi-experiments by whether the size of the credit limit jump at the discontinuity is above or below the median value. We show regression discontinuity estimates of the effect of a $1 increase in credit limits. Panel A shows effects on average daily balances 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 A.XIII: Borrowing Behavior by Source of Credit Card Origination

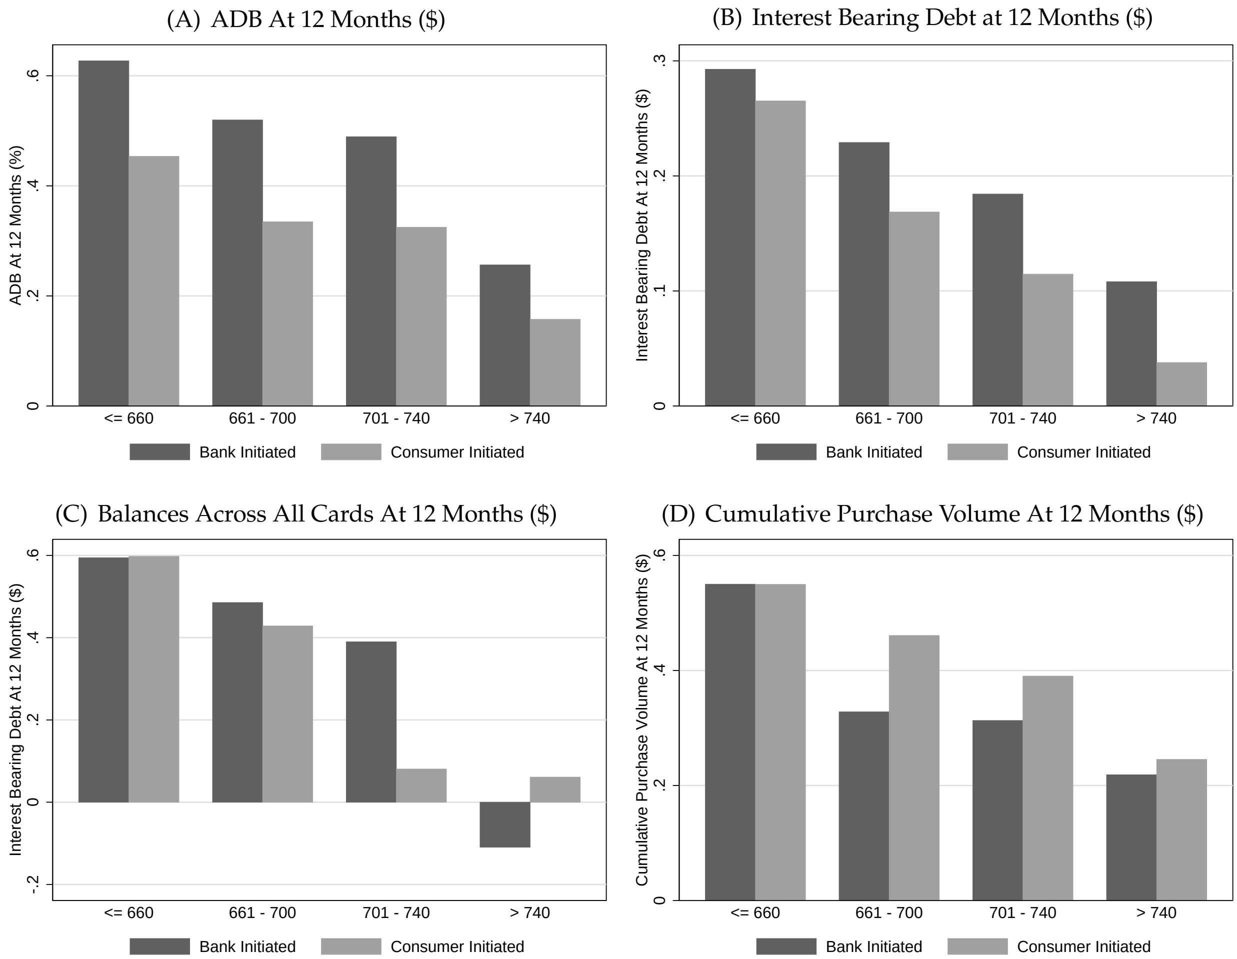

Note: Figure shows main results on borrowing and spending behavior, splitting quasi-experiments by whether the bank or the consumer initiated the origination of the credit card. We show regression discontinuity estimates of the effect of a $1 increase in credit limits. Panel A shows effects on average 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 A.XIV: Borrowing Behavior by Reported Income at Origination

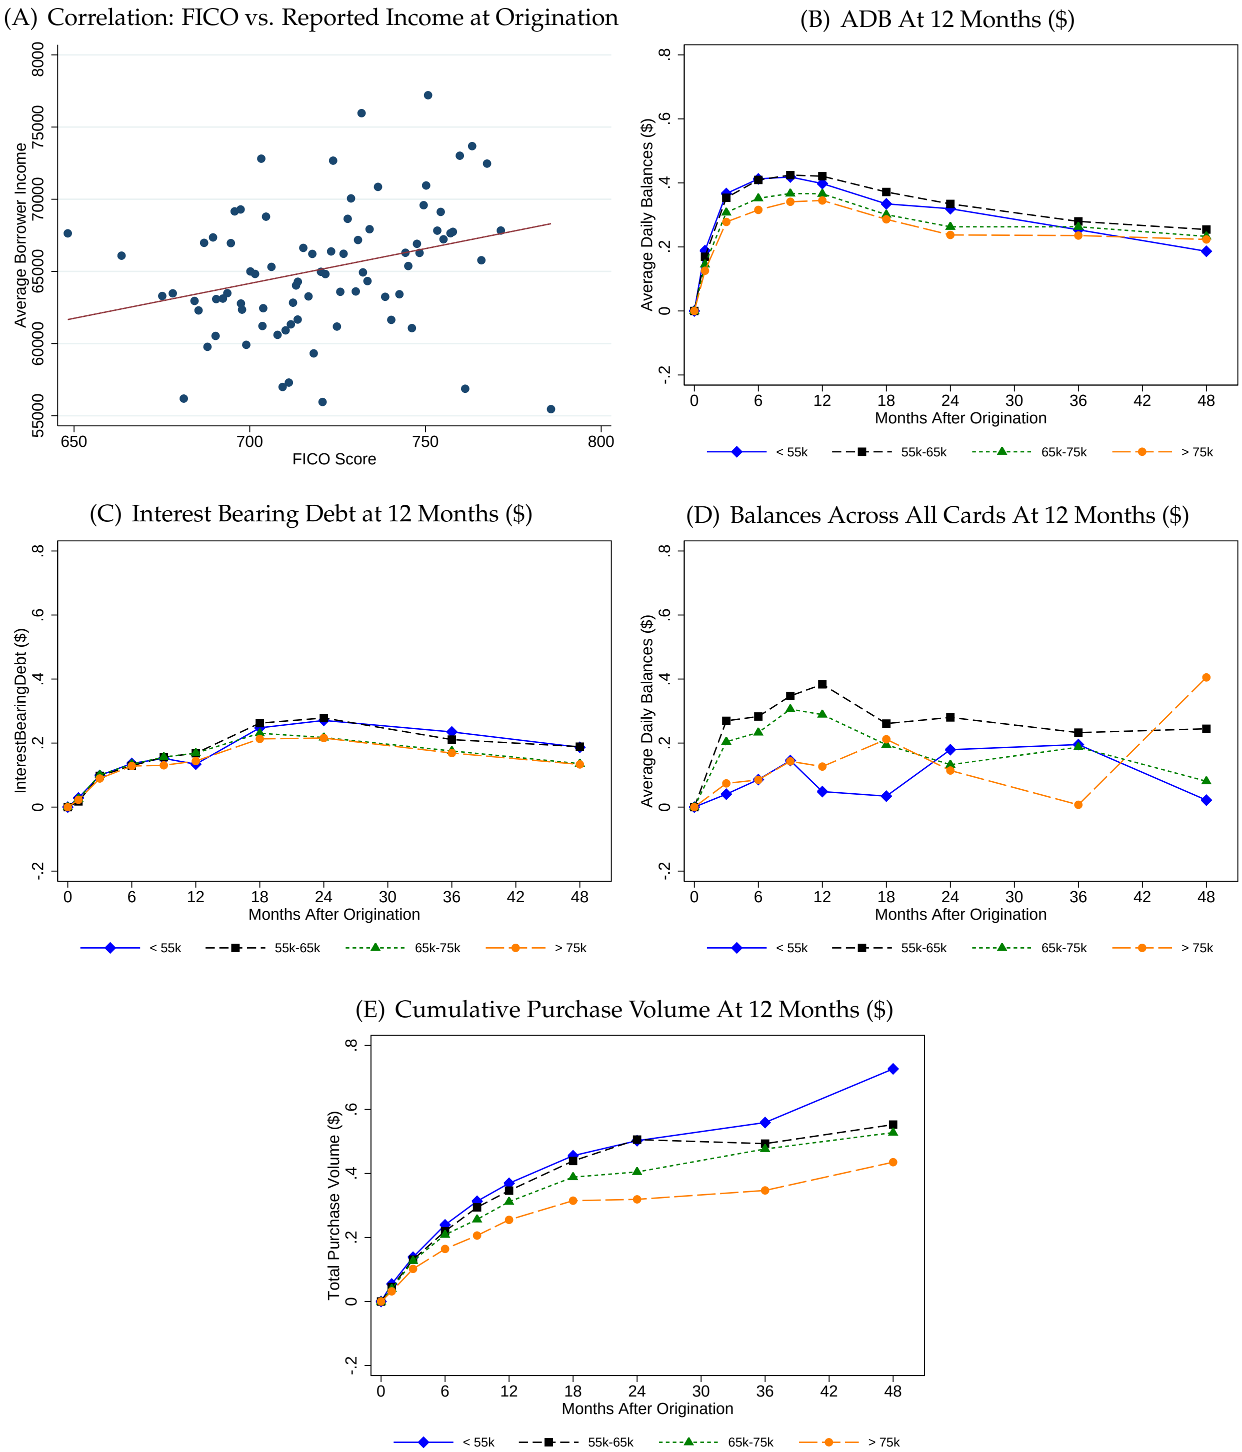

Note: Figure shows the effects of credit limits on borrowing and spending. We show regression discontinuity estimates of the effect of a $1 increase in credit limits for different groups of reported income at origination. Panel A shows the correlation between FICO score and reported income across our 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 A.XV: Borrowing Behavior: By Utilization at Origination

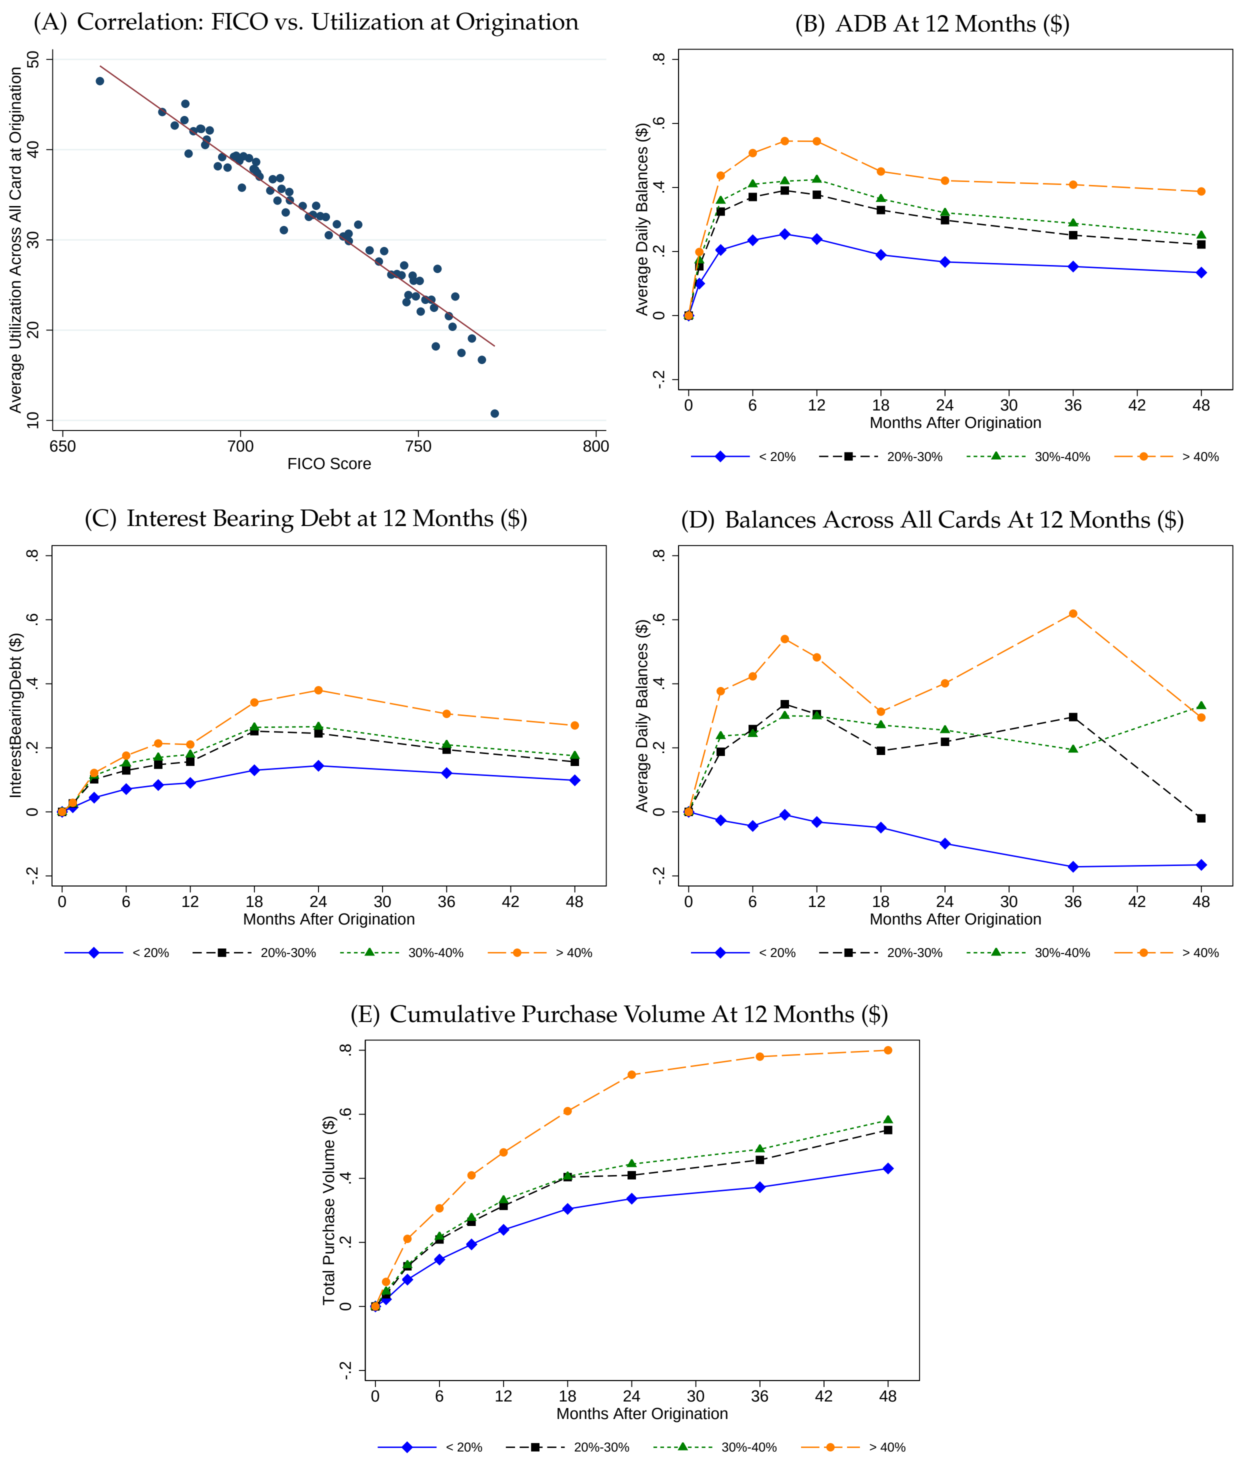

Note: Figure shows the effects of credit limits on borrowing and spending. We show regression discontinuity estimates of the effect of a $1 increase in credit limits for different groups of credit card utilization across all cards at origination. Utilization is determined as the ratio of total balances 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 A.XVI: Marginal Effect on ADB Over Time

(A) ADB

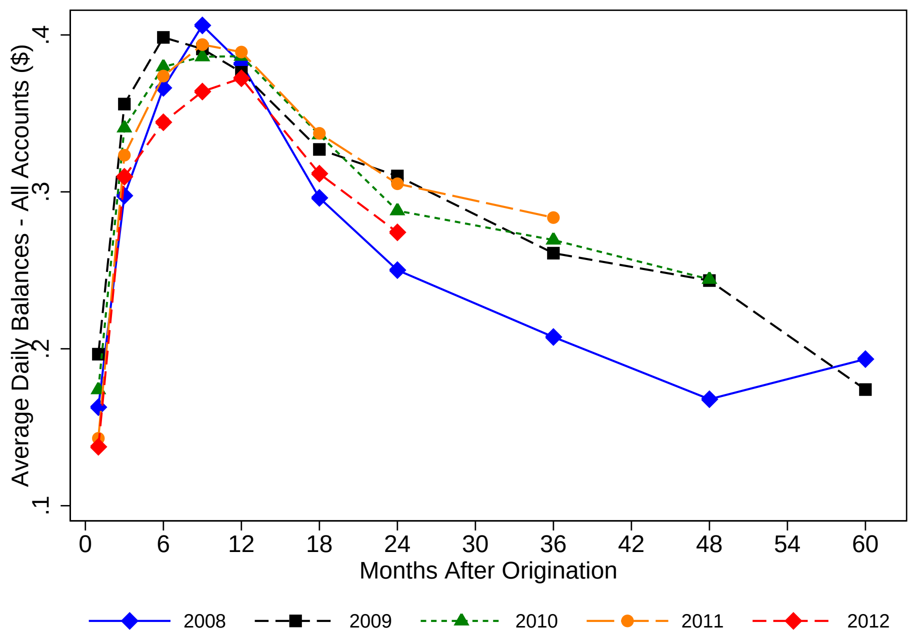

(B) ABD at 12 Months After Origination

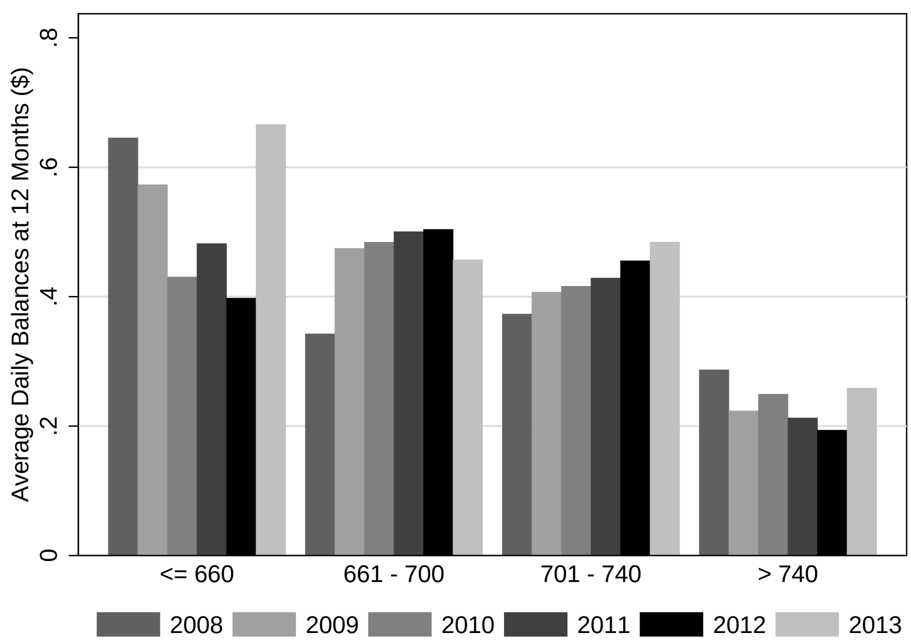

Note: Figure shows the effects of credit limits on borrowing and spending. Panel A shows regression discontinuity estimates of the effect of a $1 increase in credit limits on average daily balances (ADB) for quasi-experiments originated in different years and different time horizons after account origination. 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 A.XVII: Slope of Marginal Profit Over Time

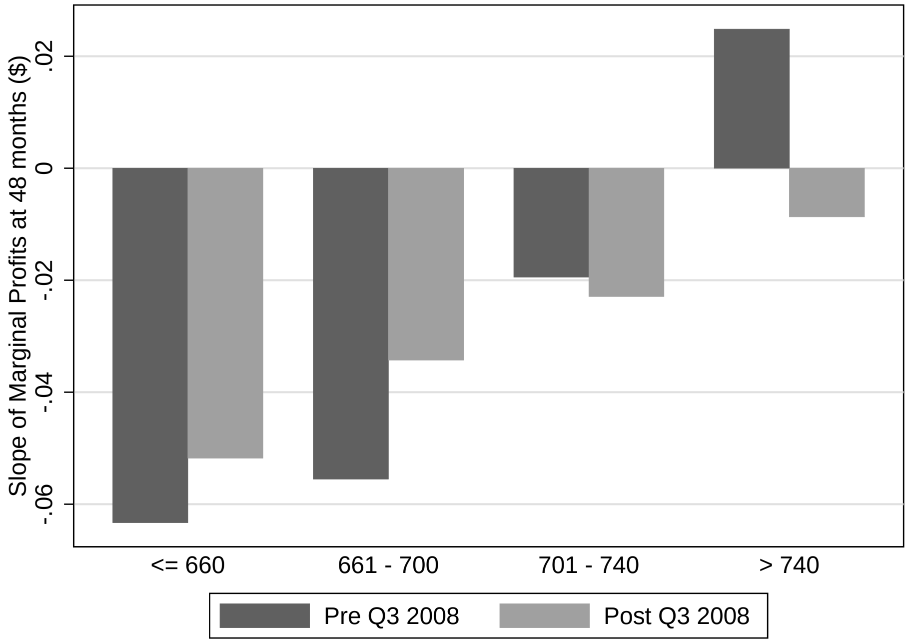

Note: Figure shows the effect of a $1,000 increase in credit limits on cumulative marginal profits after 48 months, separately for accounts originated in Q1 or Q2 of 2008 (pre Financial Crisis) and accounts originated in Q3 of 2008 or later. See Figure IX for more details.

Figure A.XVIII: Graphical Illustration of Proposition 1.

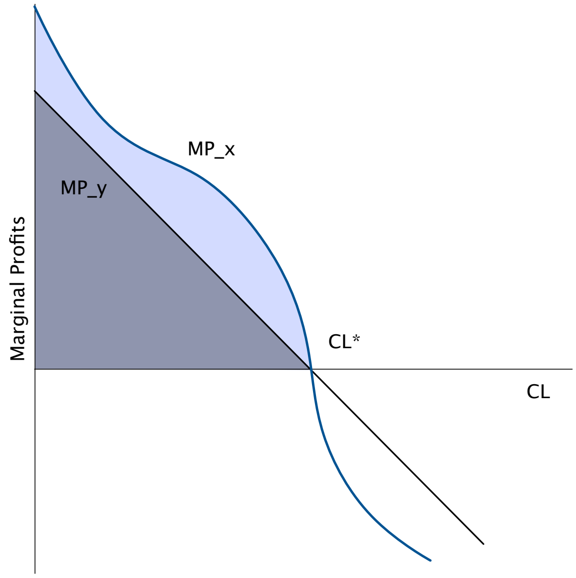

Note: Figure provides a graphical illustration of Proposition 1. The figure shows two marginal profit functions, $MP_x(CL)$ and $MP_y(CL)$, that have the same optimal credit limit: $MP_x(CL^*) = 0$ and $MP_y(CL^*) = 0$. Average profit is defined as $AP(CL) = \frac{\int_0^{CL} MP(CL)dCL}{CL}$. AP_x is 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 A.XIX: Interchange Income at Portfolio Level

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

Figure A.XX: Cost Components at Portfolio Level

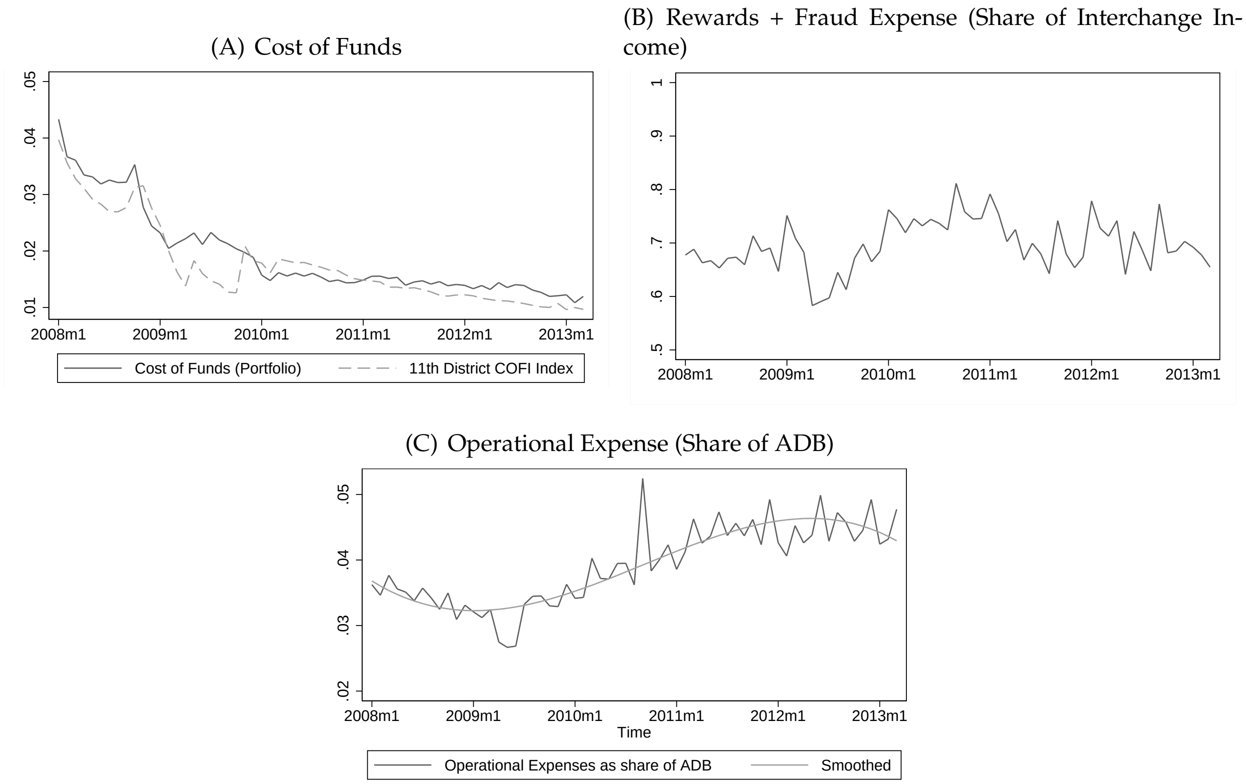

Note: Figure corresponds to Figure A.I in Agarwal et al. (2015). Panels show plots of cost components at the portfolio level by month. Panel A shows the cost of funds, calculated as the annualized interest expense ("total interest expense accrued for the month to fund credit card receivables") as a share 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 A.I: Scorecard Costs: Lower Bounds

	ΔCL	1,000 x MP'(CL*)	Per Account Scorecard Costs (χ)	Average Total Profit	% of Average Total Profits	Number of Accounts	Aggregate 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 A.II: 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				
≤660	30%	46%	40%	35%
661-700	21%	34%	30%	28%
701-740	16%	27%	23%	21%
>740	8%	13%	12%	12%
Panel B: Extensive Margin Effect				
FICO				
≤660	4.1%	2.1%	1.0%	0.3%
661-700	4.9%	1.9%	1.1%	0.4%
701-740	4.4%	3.0%	1.3%	0.7%
>740	2.4%	1.3%	0.2%	0.1%
Panel C: Intensive Margin Effect				
FICO				
≤660	25.9%	43.8%	39.1%	34.8%
661-700	16.4%	32.5%	28.6%	27.5%
701-740	11.5%	23.7%	21.8%	20.1%
>740	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 V. 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 A.III: 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						
FICO						
≤660	0.58	0.30	0.59	0.56	0.216	-0.068
661-700	0.47	0.21	0.46	0.35	0.136	-0.037
701-740	0.43	0.16	0.32	0.33	0.119	-0.024
>740	0.23	0.08	-0.05	0.22	0.037	-0.004
Panel B: Population Reweighted Estimates						
FICO						
≤660	0.698	0.311	0.759	0.825	0.181	-0.082
661-700	0.508	0.227	0.498	0.436	0.149	-0.042
701-740	0.379	0.161	0.245	0.315	0.104	-0.022
>740	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 V and Table VII. 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.