Credit Constraints and Search Frictions in Consumer Credit Markets*

Bronson Argyle[†] Taylor Nadauld[‡] Christopher Palmer[§]

February 2017

Abstract

We document consumer credit constraints in the market for used cars and demonstrate that they persist in part because of search frictions in retail loan markets. Using rich microdata from several million auto loans extended by hundreds of financial providers, we isolate plausibly exogenous variation in interest rates due to institution-specific rule-of-thumb pricing rules. We show that these discontinuities are a contributing factor in substantial interest rate dispersion among otherwise similar borrowers in the same local market, distort car-buyer purchasing behavior, and are more consequential in areas we measure as having high search costs. Overall, our results provide an explanation for why both credit constraints and price dispersion persist in household finance.

Keywords: Household finance, credit constraints, price dispersion, search, auto loans

*We thank our discussant Paul Calem; seminar and workshop participants at Berkeley-Haas, BYU, and the 2016 CFPB Research Conference; and John Campbell, Claire Celerier, Anthony DeFusco, Jan Eberly, Brigham Frandsen, Peter Ganong, Lars Lefgren, Andres Liberman, Brigitte Madrian, Adrien Matray, Adair Morse, Hoai-Luu Nguyen, Andrew Paciorek, Brennan Platt, Rodney Ramcharan, David Scharfstein, Aaron Schroeder, Amit Seru, David Sraer, Bryce Stephens, and Jonathan Zinman for helpful conversations. Tommy Brown and Sam Hughes provided excellent research assistance.

†Brigham Young University; bsa@byu.edu

‡Brigham Young University; taylor.nadauld@byu.edu

§University of California, Berkeley; cjpalmer@berkeley.edu

1 Introduction

Modern households frequently depend on consumer credit products such as credit cards, student loans, and mortgages to finance their consumption and investments. In fact, aggregate household debt outstanding exceeds aggregate corporate debt in the United States.[1] However, despite the prevalence of household debt, some of the most important open questions in household finance center around whether credit-market imperfections constrain consumption, including the role of adverse selection in consumer credit markets (Adams, Einav, and Levin, 2009), identifying frictions inhibiting the transmission of monetary policy to the household sector (Agarwal et al., 2015), and why credit constraints persist despite improvements in underwriting technology (Zinman, 2014).

In this paper, we address several such questions by documenting features of retail auto lending markets that act to constrain credit access and distort auto purchasing decisions, offering a novel explanation for why these credit constraints persist despite improvements risk-based pricing technology. Using administrative data on 4 million auto loans extended by 326 different financial institutions in all 50 states, we establish four main empirical facts. First, the segment of the auto lending market we study does not feature pure risk-based pricing; we observe large loan-rate and loan-term discontinuities at various institution-specific FICO thresholds, resulting in significant price dispersion for the same credit product across providers. Second, consumer purchasing decisions are distorted by the resulting interest rate dispersion around these lending thresholds. Third, we show that constrained borrowers could access dominating loan offers if they could costlessly query all nearby financial institutions. Fourth, such search is costly, and borrowers' propensity to search for loans with better terms is lower in areas likely to have higher search costs.[2] Taken together, we argue that this

[1] As of the first quarter of 2016, U.S. aggregate household debt outstanding was \$14.3 trillion and aggregate outstanding non-financial corporate debt was \$8.3 trillion (Federal Reserve Flow of Funds Table D.3).

[2] Nationally representative survey evidence points to the apparent costliness of consumer search in credit markets. According to the 2013 Survey of Consumer Finances, one in five people self-report doing "almost no searching" when taking out a new loan. While such behavior could be driven by expected benefits of non-costly search being low, our results provide evidence that the benefits of search are likely substantial for

evidence supports a search frictions-based explanation for the persistence of credit constraints in the market for auto loans, namely that consumers fail to consistently identify optimal financing terms because of costly search in the retail auto loan market.

What constitutes a credit constraint? Building on Hall (1978), we define a credit constraint as any friction that prevents consumers from borrowing from their future income sufficient to satisfy their intertemporal-substitution Euler equation. A long literature in economics, chronicled by Carroll (2001) and dating at least back to Houthakker's (1958) reply to Friedman's (1957) Permanent Income Hypothesis (PIH), has highlighted constraints that cause consumption to be distorted from first-best levels. Zeldes (1989), for example, defines a borrowing constraint as a cap on a borrower's per-period total indebtedness. Ludvigson (1999) examines limitations on debt-to-income ratios.[3] Carroll (2001) includes a third (non-mutually exclusive) type of credit constraint highlighted by our empirical setting: circumstances where consumers cannot borrow at competitive risk-adjusted interest rates.[4] Facing higher interest rates reduces loan sizes and distorts consumption away from first-best levels by causing payment-to-income ratio constraints to bind at lower debt levels or through the demand elasticity of loan size. Under standard regression-discontinuity identifying assumptions, otherwise identical borrowers in our data are exogenously offered substantially different interest rates, which we show affects their loan sizes and ultimately the quality of the car they purchase. Throughout the remainder of the paper, we refer to the set of borrowers that are quasi-randomly offered high interest rates as credit constrained borrowers.

Our identification strategy relies heavily on our ability to differentiate loan supply from loan demand by exploiting empirically identified discontinuities in offered loan terms around

many borrowers.

[3]Such maximum allowable credit limits or debt-service ratios could also be zero, representing consumers who are unable to borrow any amount from future income. Friedman (1963) describes a related situation wherein consumers are unable to borrow from future wages because they can only access collateralized borrowing.

[4]See also Adams, Einav, and Levin (2009), Ausubel (1991), and Davis, Kubler, and Willen (2006) for empirical and theoretical examples of liquidity constraints arising from high interest rates on borrowing.

FICO thresholds across lending institutions. Lending policies that jump discontinuously at various FICO thresholds appear to exist in 173 of the 326 lending institutions in our sample. Notably, the location of the thresholds along the FICO spectrum varies across institutions; while some thresholds appear more popular than others, there is no consensus set of thresholds used by a plurality of lenders. We document in first-stage results that borrowers just above FICO thresholds are offered longer-maturity loans and lower interest rates. On average, borrowers just above an institution's FICO threshold are offered loans with 1.47 percentage point lower interest rates for 1.4 longer months as compared to otherwise similar borrowers just below a FICO threshold. Figures 1 and 2 provide examples of rate and term discontinuities, respectively, for six different credit unions in our data with detected discontinuities using the lending policy rule estimation procedure described in Section 5.1. As discussed in Section 5 below, the observed FICO thresholds isolate supply-side changes in loan characteristics from demand-driven factors under the assumption that demand-side factors (e.g., preferences, income, financial sophistication) are not likely to also change discontinuously at quasi-random FICO thresholds that vary across institutions even in the same MSA. We support this assumption with evidence that ex-ante and ex-post borrower characteristics (including age, gender, ethnicity, application DTI, application loan size, the number of loan applications per FICO bin, future loan performance and future borrower creditworthiness) are balanced around FICO thresholds.

What impact does sharp variation in loan pricing for otherwise identical borrowers have on borrower outcomes? Borrowers quasi-randomly offered expensive credit on average purchase cars that are 4.8 months older, spending an average of $978.86 less. The similarities mentioned above in borrowers across FICO thresholds suggest that borrowers on the expensive side of an arbitrary FICO threshold have similar preferences to those on the right-hand side of a pricing discontinuity and would thus presumably also like to purchase a more expensive and newer car had they not been assigned higher interest rates. By using below-threshold borrowers as an unconstrained counterfactual for interest-rate–constrained below-cutoff borrowers, we are the first paper to our knowledge to quantify the extent to which this form of credit constraint affects borrowers.

Having established the presence of interest-rate–based credit constraints and the resulting consumption distortion in the market for used cars, we then turn our attention to identifying a new explanation for the persistence of credit constraints, which we view as complementary to traditional adverse-selection explanations for the existence of credit constraints. We first confirm the accessibility of more-attractive loan terms and then address why borrowers are empirically unlikely access these alternatives. In an otherwise frictionless world, borrowers similar in credit risk should be able to obtain similar loan terms. Yet we show using the richness of our loan-level data that loan terms for borrowers constrained by artificially high interest rates are most often strictly dominated by other contemporaneous borrowing opportunities—loans originated within the same MSA at the same time for observationally identical borrowers and collateral. The magnitude of the gap between high-interest-rate borrowers' originated interest rates and available interest rates elsewhere in their MSA is striking; differences in interest rates for similar borrowers average nearly 200 basis points. In other words, there is tremendous price dispersion in the market for auto loans and credit is indeed being offered at terms more favorable to a set of similar borrowers in the same city at the same time.[5] The persistence of variation in loan terms around FICO thresholds implies that frictions do exist leading otherwise similar borrowers to accept loans with economically different prices in equilibrium.

Our proposed search-cost explanation for differences in equilibrium interest rates among similar borrowers relies on the following evidence. We show that borrowers on the expensive side of FICO thresholds reject high-interest-rate loans most often when the number of nearby alternative lenders is high.[6] Using the physical branch locations of every bank and credit

[5]Note that this price dispersion is not simply a case of certain lenders being lowest-cost providers. The identify of the provider with best deal varies across borrower $\times$ collateral types.

[6]Importantly, while loan take-up rates are lower on the expensive-side of FICO thresholds, borrowers do not *apply* for loans at differential rates across the FICO thresholds, bolstering our assumption that demand-side factors do not change at cutoffs.

union in the United States, we calculate the number of financial institutions within a 20 minute drive from each borrower as a proxy for search costs. We find that differences in loan take-up rates across FICO thresholds are smaller for borrowers in high search-cost areas. Borrowers that would presumably have to exert more effort to search for a loan with better terms are more likely to accept the loan pricing they are offered even though these terms are strongly dominated by nearby alternatives. In a set of robustness checks, we verify that borrowers' ex-ante and ex-post observables are not correlated with our measure of search costs. Finally, using a subsample of our data that allows us to link borrowers across loan applications to different lenders, we verify that the incidence of multiple loan applications is negatively correlated with our search-cost measure. Taken together, we interpret these results as suggesting that search costs represent a meaningful market friction that enable interest-rate–based credit constraints to persist by supporting equilibrium price dispersion in the retail auto loan market and preventing consumers from identifying lowest-cost providers.

The remainder of the paper proceeds as follows. After contextualizing our work in several related literatures in Section 2, we fix ideas in Section 3 by providing a brief taxonomy of various forms of credit constraints. Section 4 details the administrative data we use throughout the paper, including an analysis of the its representativeness. Section 5 introduces our regression-discontinuity identification strategy. In Section 6, we present our results detecting discontinuities in lender pricing rules and documenting their role in constraining certain borrowers. In Section 7, we describe our empirical proxies for search costs and present evidence that costly search is an important source of credit constraints. Section 8 concludes.

2 Related Literature

In this section, we connect our work with literatures on credit constraints, search frictions, auto loans, and FICO-based regression discontinuities.

As explained by Carroll (2001), the initial literature on credit constraints arose in response to Friedman (1957), characterizing frictions that could lead to a failure of the PIH.[7] Moving beyond the context of the PIH, two broad types of evidences on credit constraints are most common in the literature: large increases in consumer spending in response to income and wealth shocks and changes in consumption patterns around exogenous changes in financing conditions. Recent papers drawing lessons about liquidity constraints by estimating marginal propensities to consume in response to income and shocks include Johnson, Parker, and Souleles (2006), Agarwal, Liu, and Souleles (2007), Japepelli and Pistaferri (2010), Parker, Souleles, Johnson, and McClelland (2013), and Baker (2015). Papers that explicitly examine changes in access to credit include Gross and Souleles (2002), Adams, Einav, and Levin (2009), Mian and Sufi (2011), and Aydin (2015). For example, Gross and Souleles (2002) identify the presence of credit constraints by documenting consumption responses to credit-card debt limit increases, further showing that the response of consumer debt levels to declining interest rates was the largest among borrowers that carried debt levels near binding credit limits.[8] A reduction in the supply of finance during the financial crisis also appears to have reduced spending among the most ex-ante leveraged borrowers (Mian, Rao, and Sufi, 2013, Mian and Sufi, 2014, and Baker, 2015). Our paper adds to the credit-consumption literature by documenting significant distortions in durable consumption decisions arising from costly access to efficient risk-adjusted interest rates, evidence that borrowers behave as though they are credit constrained via an interest-rate channel.

Borrowing constraints have taken many different forms in the literature, several of which we characterize mathematically in Section 3 below. Owing to the work of Hall (1978), a generation of papers took the failure of an intertemporal Euler Equation to hold as evidence

[7]Although Carroll (2001) argues that most studies purporting to detect credit constraints ought to be viewed instead as measuring consumer impatience, he clarifies that the implications of precautionary savings and liquidity constraints are the same. We remain agnostic on the distinction. Whether the consumption response to the expensiveness of borrowing arises from sharp credit limits, precautionary savings motives, or behavioral factors, we view our setting as a credit-market imperfection distorting consumption and thus a *de facto* credit constraint.

[8]Notably, Aydin (2015) finds the existence of high MPCs even among borrower segments that do not appear to be constrained in any traditional sense, which he interprets as evidence for precautionary savings and investment motives of the unconstrained.

of credit constraints, as discussed in Jappelli, Pischke, and Souleles (1998). Zeldes (1989) operationalizes credit constraints as a limit on the amount a debtor is allowed to borrow each period, similar to the settings studied by Gross and Souleles (2002) and Aydin (2015). Ludvigson (1999) studies limits on debt-service payments as a fraction of income, related to the DTI bunching documented in Adams, Einav, and Levin (2009). While we view the interest-rate–based credit constraints we examine herein as understudied, Adams, Einav, and Levin (2009), Ausubel (1991), Carroll (2001), and Davis, Kubler, and Willen (2006) each discuss borrowers who are constrained because they can only access loans with interest rates well exceeding what would be justified by their individual creditworthiness.

As argued by Zinman (2014), compelling explanations in the literature for the existence and persistence of credit constraints remain few.[9] The most notable exception is Adams, Einav, and Levin (2009), who document credit constraints in auto markets and provide evidence of adverse selection and moral hazard as a root cause. Our proposed search-cost explanation need not be mutually exclusive to an adverse-selection explanation. Though we attempt to rule out adverse selection as an explanation for the specific phenomena we detail, we consider search costs as friction number two, alongside adverse selection , on the list of empirically documented frictions contributing to credit constraints. Usury laws are another possible friction contributing to credit constraints. Though theory suggests that prices could and should perform their standard role in clearing credit markets, even for the riskiest of borrowers; in practice, usury laws impose a ceiling on rates that would be required to compensate lenders for the most extreme risk. The credit quality of borrowers in our sample likely fall outside the scope of a usury law explanation.

Although we are the first to tie credit constraints to search frictions, search frictions in retail markets have been well established in many settings. Many of these papers use theoretical results from a rich literature in search theory that seeks to explain real-world

[9]As Zinman (2014) puts it, "The continued prevalence of credit constraints is noteworthy and somewhat puzzling in its own right. For all of the advances in risk-based pricing, mechanism design, nonlinear contracting etc., prices are still quite far from clearing consumer credit markets!"

failures of the Law of One Price and provide conditions under which price dispersion can be sustained in equilibrium.[10] A key result from this literature (e.g., Stahl, 1989) is that price dispersion can persist in equilibrium when there are consumers who must expend costly effort to acquire information on prices. Confirming this finding, multiple empirical papers establish the existence of equilibrium price dispersion (a challenging task that necessitates ruling out product heterogeneity as a driver of price variation) and connect it to positive evidence that consumer search is costly in a given domain. For example, Sorenson (2000) documents dispersion in prices of prescription drugs that are driven by proxies for likely search intensity. In consumer finance, Hortacsu and Syverson (2004) find large dispersion in the fees charged by very similar mutual funds that are driven by information/search frictions, and Stango and Zinman (2015) document price dispersion in the U.S. credit card market, which they connect to variation in shopping intensity (as well as behavioral factors). Relevantly for our setting, in a price-dispersion equilibrium, firms that would like to undercut nearby competitors charging prices above marginal cost are unable to profitably do so because consumers cannot discover this dominating alternative without incurring search costs, explaining why we observe lenders offering (and borrowers accepting) seemingly dominated loan terms.

Finally, we are not the first paper to exploit FICO-based discontinuities in treatment variables. Keys et al. (2009 and 2010) find that the probability of securitization (and thus loan screening) change discontinuously at a FICO score of 620. Bubb and Kaufman (2014) provide evidence for other discrete FICO thresholds in the mortgage underwriting process, including detailing the likely genesis of threshold-based policies. More recently, Agarwal et al. (2015) use sharp FICO-based discontinuities in credit limits to estimate heterogeneity in marginal propensities to borrow, and Laufer and Paciorek (2016) evaluate the consequences of minimum credit-score thresholds for mortgage lending. Building on this collection of papers that either use FICO-based discontinuities as natural experiments or explicitly study their

[10]See, for example, the back and forth (helpfully summarized by Baye, Morgan, and Scholten, 2006) between Stigler (1961), Diamond (1971), Rothschild (1973), Salop and Stiglitz (1982), Burdett and Judd (1983), and Stahl (1989).

consequences, we are the first to identify credit-score–based discontinuities in pricing rules and to link those discontinuities to credit constraints, price dispersion, and costly consumer search.

3 Conceptual Framework

To be more explicit about the interest-rate–based credit constraints we study here and to facilitate comparisons with other forms of credit constraints in the literature, this section provides a brief taxonomy of common credit constraints. In the frictionless world of the PIH, consumers maximize present-discounted utility by borrowing from future income in amounts and at interest rates that obey lifetime budget constraints. We motivate our taxonomy of credit constraints with the standard utility maximization problem wherein consumers chose the consumption path $\{c_t\}_{t=0}^{T}$ to solve the optimization problem

$$\max_{\{c_t\}} E\left(\sum_t \frac{u(c_t)}{(1+\delta)^t}\right) \tag{1}$$

$$s.t. \sum_t \frac{c_t}{(1+r^*)^t} \leq \sum_t \frac{y_t}{(1+r^*)^t} + (1+r^*)A_t \tag{2}$$

where $u(\cdot)$ is the per-period mapping from per-period consumption c_t to utility, δ is the discount rate that makes consumers indifferent to receiving $u(c)$ utils today and $(1+\delta)u(c)$ utils tomorrow, y_t is (uncertain) per-period income, A_t represents total consumer assets, and r^* represents the market-clearing, risk-adjusted cost of lending to each borrower. In this simple model, lending to a given borrower is risky because of individual-specific income risk, leading a price-taking lender earning zero profits to charge each person an interest rate r^* based on that borrower's observable income risk.

Maximizing utility with respect to lifetime consumption yields the standard Euler equation of Hall (1978) describing the trade-off between adjacent periods' marginal utilities for any given period t,

$$u'(c_t) = E\left(\frac{1+r^*}{1+\delta}u'(c_{t+1})\right). \tag{3}$$

Credit constraints impose additional restrictions on borrowing that prevent the Euler Equation from holding as in (3), leading to a wedge $\lambda_t > 0$ such that

$$u'(c_t) = E\left(\frac{1+r^*}{1+\delta}u'(c_{t+1})\right) + \lambda_t \tag{4}$$

holds instead. The implication of the positive wedge λ_t is intuitive: absent sufficient borrowing opportunities, consumption is too low today, resulting in too high of a marginal utility relative to what the consumer would prefer in a world without borrowing constraints.

In their prototypical form (e.g., Zeldes, 1989), credit constraints are explicit borrowing limits on the level of indebtedness. Under the strongest type of credit constraints, this bound is zero and consumers are unable to access credit in any form. That is, borrowers face an additional constraint that requires consumption to be less than the sum of current income and assets such that

$$c_t \leq y_t + A_t \; \forall \; t. \tag{5}$$

A less-extreme version of credit credit constraints more resembles underwriting practices that specify maximum loan sizes such as credit limits for credit cards or because of maximum loan-to-value rules for collateralized borrowing wherein borrowers can access positive amounts of credit but only up to a per-period upper bound B_t. Under these conditions, consumption is constrained to be less than the sum of current income, total assets, and the pre-specified borrowing limit (B_t)

$$c_t \leq y_t + A_t + B_t \; \forall \; t. \tag{6}$$

Borrowing limits could also arise from a maximum debt service-to-income ratio D as opposed to a level credit limit (see Ludvigson, 1999). In this case, this ratio determines the maximum amount that can be borrowed through a limit on the debt load $rB/y \leq D$ such that the credit constraint becomes

$$c_t \leq y_t + A_t + \frac{Dy_t}{r} \; \forall \; t \tag{7}$$

In our setting credit constraints are characterized by a friction that affects the offered interest rate. Equation (3), the utility maximizing first-order condition, requires access to borrowing/saving technologies at the break-even rate of r^*. We define consumers to be credit constrained if they are unable to access credit at r^*. If actual rates offered to a particular borrower are $r > r^*$, then borrowers will not be able to borrow in the amount that satisfies the Euler equation. An offered $r > r^*$ thus constrains the amount of consumption via the interest rate channel in the budget constraint. To see what impact this has on the Euler Equation, we can solve explicitly for the wedge that quantifies the degree to which a given period's marginal utility is too high relative to what it would be with access to borrowing opportunities at r^*:

$$\lambda_t = (r - r^*)\frac{u'(c_{t+1})}{1+\delta} > 0. \tag{8}$$

In other words, the distortion from facing overly expensive interest rates r is most severe when the difference between the available rate r and the efficient rate r^* is large and when the present value of next period's marginal utility is high.

4 Data

We analyze the loan contract terms and auto purchasing decisions of 3.9 million individual borrowers in the United States from 315 retail lending institutions from 2005–2016. The loan data are provided by a technology firm that provides administrative data warehousing and analytics services to retail-oriented lending institutions nationwide. Roughly two thirds of the lending institutions represented in the data set are credit unions ranging between \$100 million and \$4 billion in asset size. The remainder are non-bank finance companies of unknown total asset size, although the vast majority (98.5%) of the loans in our data were originated by credit unions.[11] Borrowers from all 50 states are represented in the data, but the five largest states in the data are Washington (465,553 loans), California (335,584 loans), Texas (280,108 loans), Oregon (208,358 loans), and Virginia (189,857 loans).

[11]Our results are unchanged if we exclude loans from finance companies, which are generally of lower credit quality.

The dataset contains information capturing all three stages of a loan's life: application, origination, and ex-post performance, although we have loan application data for only approximately 1.9 million loans from 41 different institutions. The available loan application data report borrower characteristics (ethnicity, age, gender, FICO scores, and debt-to-income (DTI) ratios at the time of application), whether a loan application was approved or denied, and whether it was subsequently withdrawn or originated. For originated loans, the data additionally include information on loan amounts, loan terms, car purchase prices, and whether the loan came through a direct or indirect origination channel.[12] We restrict our sample to direct loans in an effort to address concerns that indirect loans are potentially endogenously steered to specific financial institutions (perhaps because car dealers become aware overtime of lenders' pricing rules). Finally, to measure ex-post loan performance, we observe a snapshot of the number of days each borrower is delinquent, whether each loan has been charged off, and updated borrower credit scores as of the date of our data extract.

Panels A, B, and C of Table 1 present summary statistics on loan applications, loan originations, and measures of ex-post performance, respectively. As reported in Panel A of Table 1, the average loan application in our data seeks approval for a five-year $18,884 loan at a median interest rate of 4.7%.[13] Borrowers applying for loans in our data have an average credit score of 646 and an average DTI ratio of 28.3%. The percentage of loans approved is 50.2%, with 78.4% of the approved borrowers subsequently originating a loan. Throughout the paper we refer to the number of loans originated divided by the number of applications approved for a particular group as the loan take-up rate. We exploit variation in the loan take-up rate in Section 7.2.

Panel B of Table 1 reports summary statistics on loan originations, revealing several interesting patterns. Compared with loan applications, originated loans have larger average

[12]The terms direct and indirect loans refer, respectively, to whether the borrower applied for a loan directly to the lending institution or through an auto dealership that then sent the loan application to lending institutions on the buyer's behalf.

[13]Application interest rates are strongly right skewed with a mean interest rate of 17.3% and a 75th percentile of 12.7%. These risky outliers appear to be rejected, as they are not in the originated loan sample in Panel B.

sizes, lower interest rates, longer terms, are from more creditworthy and less constrained borrowers, and secure purchases of more expensive cars. Average monthly payments for originated loans are $338 per month with an interquartile range of only $200.

Panel C tabulates measures of ex-post loan performance. While the average loan is 36 days delinquent, most loans are current; the 75th percentile of days delinquent is zero and only 2.1% of loans have been charged-off (accounted as unrecoverable by the lender). Defining default as a loan that is at least 90 days delinquent, default rates average 2.2%. In untabulated results, default rates for borrowers with sub-600 FICOs average 6.8%, compared to a default rate of 2.6% for borrowers with FICOs between 600 and 700 and 1.6% over-700 FICO borrowers. Lending institutions periodically check the credit score of their borrowers subsequent to loan origination, creating a novel feature of our data. Summary statistics for ΔFICO represent changes in borrowers' FICO scores from the time of origination to the lender's most recent (soft) pull of their FICO score.[14] Updated FICO scores indicate that borrowers on average experienced a 1.8% reduction in FICO score since origination, although borrowers with FICO scores below 600 on average realized a 5.7% increase in FICO score.

4.1 Data Representativeness

The bulk of our auto loan data come from credit unions, prompting questions about the representativeness of the data. Popular perception is that credit union usage is concentrated in an older demographic. Our data confirm this fact. Over 41% of borrowers in our sample were between 45 and 65 years old at loan origination. In contrast, census data indicates 34% of the general U.S. population are between the ages of 45–65. Borrowers in our sample are also less racially diverse than the general public. Over 73% of our sample are estimated to be white (as of 2015), compared to a 65% of adults in the general population recorded

[14]The time between FICO queries varies by institution, but institutions that provide updated FICO scores do so at least once a year such that conditional on having an updated FICO score, the amount of time between the original FICO recording and the current FICO is roughly equal to loan age.

by the 2015 American Community Survey.[15] Borrowers in our data report median FICO scores at origination of 715 (Table 1, Panel B) over the full 2005-2016 sample period. The NY Federal Reserve Consumer Credit Panel (CCP), a representative 5% sample of U.S. borrowers, reports median FICO scores for originated auto loans of 695 during the period our sample was collected.[16] Almost 70% of the loans in our sample were originated between 2012 and 2015, with median FICO scores of 714. In comparison, the CCP reports median FICO scores of 696 over the same period. In summary, our sample contains borrowers that are slightly older, less racially diverse, and of a higher average credit quality than national averages. These sample biases should not limit our ability to draw inference given that the biases in our sample likely tilt towards borrowers less likely to be credit constrained.

A second data validity issue involves the distribution of loan originations through time. As reported previously, over 70% of loan originations in our sample occurred between 2012 and 2015, despite a sample period that runs from 2005–2016. The large increase in loans through time reflects the increase in the client base of our data provider through time rather than auto credit origination in general. Auto loan originations in the general population have increased through time, from an aggregate outstanding balance of $725 million in 2005 to just over $1 trillion in 2016, but not at the rate reflected in our database. We view the non-representative time series of our data as less relevant to any inference we attempt to draw given that we rely on a cross-sectional RD approach for identification.

A third data validity issue is whether credit unions capture a meaningful fraction of the auto loan market. Experian data from 2015 indicates that credit unions originated 22% of all used car loan originations and 10% of new car originations in the U.S.. The Experian data do not differentiate direct lending from indirect lending, but of the auto loan data made available to our data provider by its clients, roughly two-thirds are direct loans. Finally, we note that data on the performance of auto loans as reported in the CCP suggests that

[15] Borrowers do not report race at the time of loan origination but most lenders in our sample estimate race ethnicity in an effort to comply with fair lending standards regulations.

[16] The CCP data report quarterly median FICO scores over our sample period. The reported 695 median FICO is actually the median of the quarterly medians that span our sample period.

auto loans originated by credit unions and banks have substantially lower default rates as compared to loans originated by auto finance companies.[17]

5 Detecting Discontinuities

Lending institutions make underwriting decisions about whether to approve a loan application using a combination of hard and soft information on borrower credit quality. Hard information generally consists of quantifiable credit metrics provided by credit bureaus or verified with paystubs and tax statements such as FICO scores, debt-to-income ratios, bankruptcy history, and annual earnings. Soft information, loosely defined as information that cannot be easily quantified related to the likelihood of a borrower's future willingness or ability to repay a loan, is by definition unobservable to the econometrician.[18] Any econometric analysis that specifies loan outcomes as the dependent variable is subject to the critique that equilibrium loan outcomes are influenced by unobservable soft information, complicating inference seeking to isolate factors causing an outcome of interest. Our setting is no exception. While our dataset consists of millions of equilibrium lending outcomes, our ability to draw inference is hindered by the possibility that unobserved soft information plays a role in jointly determining selection into application and origination, observed loan terms, and subsequent loan performance. Because our sample consists of direct auto loans, soft information in our setting would most likely be generated from the relationship between credit unions and their long-term customers, observable to a loan officer.

We address this possibility, and other potential omitted variables, with a regression-discontinuity design that exploits observed discontinuities in offered loan terms across several FICO thresholds. Unlike the 620 FICO heuristic in mortgage *underwriting* first exploited by Keys et al. (2009 and 2010) that affects screening at both origination and securitization (Bubb and Kaufman, 2014), we focus on discontinuities in loan *pricing*, i.e., the interest rate

[17] The CCP does not separate auto loans made by credit unions from those made by banks.

[18] See Petersen (2004) for a careful treatment of hard and soft information in financial markets.

offered to a borrower conditional on having a loan application approved by underwriting. Moreover, no industry standard set of thresholds exists in auto lending as in mortgage lending. Still, while auto-loan lending institutions do not adhere to a common set of FICO cutoffs, the use of a given threshold at some point across the FICO spectrum is prevalent for most lenders in our data. Multiple lending institution executives have confirmed to us in private conversations that their pricing functions explicitly incorporate discrete FICO thresholds to set interest rates and loan terms.[19] Also in contrast to Keys et al. (2010), FICO thresholds observed in our data have little to do with secondary markets given that many auto loans are retained by the lending institutions in our dataset. Rather than reflecting demand for securitization or a loan's subsequent marketability on a secondary market, FICO discontinuities may have been incorporated into software systems as a holdover from a time when pricing was done via rate sheets instead of automated algorithms.[20]

To illustrate the effect of FICO thresholds on equilibrium interest rates, we estimate lender-specific interest-rate and loan-term policies nonparametrically. For each lender c in our data, we characterize their lending policies across FICO bins with a set of parameters $\{\psi_{ck}\}$ where k indexes FICO bins denoted $\mathcal{F}_k$. Pooling loan-level data from individuals i, we estimate ψ by regressing an origination outcome y (interest rates or loan terms) on a set of indicator variables for each 5-point FICO bin $\mathcal{F}_k$

$$y_{ic} = \sum_k \psi_{ck} \mathbb{I}(FICO_i \in \mathcal{F}_k) + \varepsilon_{ic} \tag{9}$$

where ε_{ic} includes all other factors that influence loan pricing. The 5-point FICO bins begin at a FICO score of 501 where the first bin includes FICO scores in the 501-505 range, the

[19] As an example, one executive pointed to a FICO score of 610 as the explicit cutoff that determines the loan terms offered to prospective borrowers at that executive's credit union. Applicants with a FICO score just below 610 were offered higher rates and loan terms below 60 months in contrast to applicants with FICO scores above 610.

[20] In the mortgage industry, Bubb and Kaufman (2014) write that "Though [Automated Underwriting Systems] calculate default risk using smooth functions of FICO score, they also employ a layer of 'overwrites' which trigger a 'refer' recommendation when borrowers fall into certain categories—for instance, borrowers with FICO scores below 620." See Hutto & Lederman (2003) for a history of the incorporation of discrete credit score cutoffs into automated underwriting systems for mortgage lending, such as those created by Fannie Mae and Freddie Mac.

second bin includes 506-510, etc., up through FICO scores of 800. The estimated coefficients on each FICO bin represent the average interest rate for loans originated to borrowers with FICO scores in that bin relative to the estimated constant (the omitted category is loans outside this range—we focus on relative magnitudes for this exercise).

Figure 1 presents interest-rate plots for three different financial institutions. The estimated $\hat{\psi}$ point estimates represent how that lender's pricing rules appear to vary with borrower FICO score, and the accompanying 95% confidence intervals provide a sense of how reliant on FICO scores was each lender's pricing rule. Panel A of Figure 1, estimated on one institution in our data with approximately 12,000 borrowers (rounded to preserve lender anonymity), illustrates breaks in average interest rates for borrowers with FICO scores around FICO cutoffs at 600, 660, and 700. The breaks in interest rates at the FICO cutoffs are large (representing jumps of over 2 percentage points). Average interest rates for borrowers in the 595-599 FICO bin are 2.5 percentage points higher than the average interest rate for borrowers in the 600-604 FICO bin, and the difference in average interest rates between the two bins are statistically significant at the .001 level. Panels B and C illustrate similar rule-of-thumb FICO breaks for unique institutions with approximately 6,000 and 25,000 loans, respectively. One important observation arising from these anecdotal plots is the fact that the breaks occur at different FICO scores across different institutions, consistent with our understanding that the discontinuities are reflective of idiosyncratic pricing policies across institutions.

In addition to interest rates, loan terms (the number of months until a loan matures) also often change discontinuously across FICO thresholds. In Figure 2, we plot point estimates and confidence intervals from estimating equation (9) for a set of loan term, 5-point FICO bin regressions. The institution represented in Panel A, with approximately 162,000 loans in the portfolio, has loans in its portfolio with longer terms for borrowers just to the right of the 630 FICO threshold. Panels B and C illustrate similar loan term discontinuities around different rule-of-thumb FICO thresholds for institutions originating approximately 42,000 and 36,000 loans, respectively. The plots presented in Figures 1 and 2 are illustrative of rule-of-thumb thresholds that exist at different institutions throughout our sample.

In order to standardize our analysis to include every institution that employs discontinuous pricing rules, we empirically identify the existence of discontinuities at each institution (if they exist at all) in our sample through the following criteria. We first estimate the interest-rate FICO bin regressions following equation (9) for each institution in our sample separately. To establish the existence of a economically and statistically significant interest-rate discontinuity, we require that interest rate differences across consecutive bins be larger than 50 basis points and be estimated with p-values that are less than 0.10. We further refine the set of discontinuities by requiring that an identified discontinuity not lie within 20 FICO points of another identified discontinuity within the same institution. This restriction limits any potential contamination that could occur if borrowers simultaneously fall into a treated sample at one observed threshold but serve as a control for a sample at a different threshold. We further examine each potential threshold visually to ensure that the identified discontinuities are well behaved around the candidate thresholds. Finally, in an effort to maximize the statistical power in our RD design, we require that each candidate threshold contain 100,000 loans within the span of 38 FICO points around the candidate threshold, forming a discontinuity sample (Angrist and Lavy, 2004). The 38 FICO points represent 19 points on either side of a threshold that do not bump up against a different threshold that could exist within 20 FICO points. Implementing each of these restrictions ultimately results in large and meaningful discontinuities in interest rates and loan terms at FICO scores of 600, 640, and 700 across 173 institutions and 489,993 loans.[21] Table 2 reports summary statistics for our ultimate estimation sample (the set of loans within 19 points of one of our thresholds). A comparison of the full sample summary statistics (Table 1) with the

[21] We reiterate that not all institutions have thresholds at 600, 640, and 700—these are merely the most popular detected discontinuities satisfying our criteria. Relaxing the requirement of 100,000 loans within 38 FICO points around the threshold results in a larger set of identified thresholds. The two most populated thresholds outside of our selected three thresholds are at 680 and 660 which contain approximately 90,000 and 80,000 loans, respectively.

threshold-specific sample (Table 2) reveals that the threshold sample is similar to the full sample along observables. All of the estimates reported in the paper use the discontinuity sample.

5.1 First-Stage Results

To validate our Regression-Discontinuity (RD) design, we present a series of diagnostics designed to test whether our data meet the two main identifying assumptions underlying RD estimation. First, the RD approach assumes that the probability of borrower treatment (i.e., offered interest rates) with respect to loan terms is discontinuous at FICO thresholds of 600, 640, and 700. Second, valid RD requires that any borrower attribute (observed or unobserved) that could influence loan outcomes change only continuously at interest-rate discontinuities. This smoothness condition requires that borrowers on either side of a FICO threshold are otherwise similar, such that borrowing outcomes on either side of a threshold would be continuous absent the difference in treatment induced by policy differences at the threshold.

In our remaining specifications, we normalize FICO scores to create a running variable $\widetilde{FICO}_{ict}$ that measures distance from a interest-rate discontinuity. For example, for loans near the 600 FICO score threshold, $\widetilde{FICO}_{ict} = FICO_{ict} - 600$. Panel A of Figure 3 plots average interest rates against normalized borrower FICO scores for a sample restricted to loans with borrower FICO scores between 581 and 619. The plots demonstrate smoothness in the conditional expectation function except for the points corresponding to a FICO score of 599 and 600, where interest rates jump discontinuously. We repeat the plot using similar 38 point FICO ranges for the 640 and 700 FICO thresholds in panels B and C of the same figure. These plots confirm the existence of large interest-rate discontinuities at these thresholds. The magnitude of the discontinuities appears to be smaller at higher FICO thresholds, which might arise from smaller relative differences in credit quality at high FICO score levels.

To establish statistical significance and introduce our RD design, we estimate first-stage regressions of the form

$$y_{ict} = \beta_1 \widetilde{FICO}_{ict} + \beta_2 \mathbb{I}(\widetilde{FICO}_{ict} \geq 0) + \beta_3 \widetilde{FICO}_{ict} \cdot \mathbb{I}(\widetilde{FICO}_{ict} \geq 0) + \alpha_c + \delta_t + \varepsilon_{ict} \quad (10)$$

where y_{ict} is the outcome for loan i originating from lending institution c in quarter t, $\mathbb{I}(\widetilde{FICO}_{ict} \geq 0)$ is in indicator variable equal to one if the normalized FICO score $\widetilde{FICO}_{ict}$ is above the threshold, and α_c and δ_t are lender and quarter fixed effects, respectively. In practice, we conservatively estimate equation (10) using the Robust RD estimator of Calonico, Cattaneo, and Titiunik (2014), estimating the effect of the running variable $\widetilde{FICO}$ above and below the cutoff at $\widetilde{FICO} = 0$ using local linear regression (as opposed to the unweighted linear specification we provide for intuition in equation (10)) and a local quadratic bias correction.[22] Our baseline regression specification pools each of the three discontinuities into one dataset using the FICO normalization described above. We cluster our standard errors by normalized FICO score.

Table 3 presents results of this exercise. Interest rates for borrowers with FICO scores immediately above one of our thresholds are estimated to be 1.47 percentage points lower than borrowers just below (column 1). Column 2 reports that loan terms for borrowers just above a FICO threshold are 1.38 months longer than otherwise similar borrowers below the threshold. Given an average interest rate in our estimation sample of 6.7% (Panel B of Table 2, the magnitude of this coefficient is economically meaningful and shows that landing on the so-called wrong side of a interest rate discontinuity has material consequences on the cost of credit.

5.2 Testing Exogeneity Assumption

To test whether other observables besides the treatment variables (interest rate and loan term) also change discontinuously at our detected FICO thresholds, in Figure 4, we pool loans in the neighborhood of all three FICO thresholds and plot the average value of other

[22]While our reported results use a uniform kernel with a bandwidth of 19, our results are robust to alternative kernels and a wide range of bandwidths.

borrower characteristics around these FICO thresholds. Importantly, these graphs are constructed with loan application data in order to ensure that borrowers are similar at FICO thresholds along characteristics at the time of application. Panels A–E plot borrower debt-to-income ratios, loan amounts, borrower age in years, borrower gender (an indicator for male), and borrower ethnicity (an indicator for white), respectively. These plots indicate smoothness in ex-ante borrower characteristics around FICO thresholds. Borrowers on either side of FICO thresholds do not appear meaningfully different in terms of their debt capacity, their willingness to borrow, or along demographics. Finally, Panel F plots the number of applicants within each normalized FICO bin, showing that borrowers don't appear to select into applying for a loan based on their FICO score.[23] Such manipulation of the running variable—a discontinuity in the propensity to apply for a loan at a FICO threshold—would raise selection concerns but would be difficult to accomplish given the uncertainty applicants face about their own credit scores (owing to the volatility of FICO scores and uncertainty about which credit bureau(s) a lender will query) and the low likelihood that prospective borrowers are aware of the precise thresholds used by a given lender.

Table 4 reports more formal tests of the smoothness condition using the loan-application data, available for a subset of lending institutions. The estimates indicate no statistical difference in requested loan amounts for borrowers on either side of the threshold (column 1). In column 2 we present estimates of differences in debt-to-income ratios around the thresholds. Ex-ante debt-to-income ratios of borrowers on either side of the thresholds are statistically indistinguishable. Finally, we count the number of applications received from borrowers of each normalized FICO score and examine these counts at the FICO-score level using our RD estimator. Column 3 shows that the number of borrowers applying for loans on either side of as threshold are also not statistically different.

Taking stock, our empirically detected discontinuities in loan pricing at specific FICO thresholds are large (nearly 150 basis points) and are unaccompanied by similar discontinu-

[23] We do, however, note one unusual spike in applications at FICO = 600 that does not show up as statistically significant in the RD specifications of Table 4.

ities in borrower composition supporting our reliance on a regression discontinuity design.

6 Documenting Credit Constraints

To establish whether the discontinuities in interest rates discussed above act as *de facto* credit constraints, we next establish that being treated with a higher interest rate causally distorts consumption decisions. Whether a given credit-market imperfection is a binding constraint on optimal consumption is difficult to ascertain because it requires estimating counterfactual consumption in the absence of the alleged constraint. However, we can determine credit-constraint existence by evaluating the auto purchasing decisions of borrowers on either side of the documented FICO thresholds. Given the empirical result that, ex-ante, borrowers are similar around FICO thresholds, we start with the null hypothesis that borrowers around FICO thresholds would also have similar demand for cars, conditional on obtaining the same set of financing terms. Exploiting a useful feature of the data, namely our ability to observe the exact amount that each borrower spent on a car, we test whether borrowers spend differently around the observed FICO thresholds.

Figure 5 plots car purchase amounts around the normalized FICO threshold. Purchase amounts are smooth leading up to the FICO threshold and then jump discontinuously at the threshold. Using the same RD design used in our first-stage analysis above, we formally test for statistical differences in purchase amounts. As before, we estimate (10) by controlling for lending institution fixed effects, quarter-of-origination fixed effects, allowing for a local linear function of the running variable, bias-correcting using the local quadratic approach of Calonico et al. (2014), and using a bandwidth of 19 around the normalized FICO threshold with a uniform kernel. Column 1 of Table 5 presents these reduced-form results. Borrowers quasi-randomly offered more expensive loans spend an average of \$979 less on the cars they purchase. Column 2 presents results with loan amounts as the dependent variable. Realized loan sizes decrease by an average of \$1,480 on the expensive side of a detected FICO discontinuity. The fact that loan sizes increase by larger amounts around the threshold than purchase amounts indicates that, ex-post, borrowers on the right side of the cutoff are allowed and choose higher loan-to-value ratios. This result is confirmed in column 3 of Table 5, which indicates that ex-post LTV ratios are slightly higher, 2.7 percentage points on average, for borrowers to the right of FICO thresholds.

Detailed data on loan amounts and loan terms allow us to calculate the implied monthly payment of borrowers on either side of the thresholds. In column 4 of Table 5 we test whether ex-post monthly payments are different around the thresholds. On average, monthly payments increase by $9.67 per month for borrowers to the right of a interest rate threshold. Shorter terms and higher interest rates cause the relatively constrained borrowers just below FICO thresholds to purchase less expensive cars and use less financing in their purchase than unconstrained borrowers, essentially purchasing less car and less credit for only a $10 reduction in monthly payment.

Still, higher ex-post LTV ratios and slightly higher monthly payments, implying higher ex-post DTI ratios, warrant further consideration. One concern is the possibility that borrowers on either side of the FICO thresholds are different ex-ante in their ability to service debt, violating the smoothness/exogeneity conditions required for valid RD. We interpret these results differently. Given that ex-ante DTI ratios in the loan application data are continuous around the thresholds (Table 4), we interpret these results as further evidence of the easing of credit terms for borrowers on the right side of FICO thresholds. That is, ex-post, borrowers on the right side of thresholds are offered lower rates, longer terms, and apparently allowed higher ex-post LTV and DTI ratios.

This evidence of otherwise similar borrowers spending different amounts on the cars they purchase as a result of the financing terms they are offered is consistent with our hypothesized interest-rate channel of credit constraints. An alternative explanation is the possibility that dealers somehow price discriminate and exploit borrowers' ability to service larger debt amounts by charging more for the exact same car than otherwise similar borrowers on the expensive side of FICO thresholds purchase.[24] We address this possibility by controlling for year-make-model (e.g. 2013 Honda Accord) fixed effects in our RD regressions.[25] Column 1 of Table 6 reports results when controlling only for make-model fixed effects. Even within a make and model category, borrowers quasi-randomly assigned expensive credit continue spend $887 less on cars, suggesting that the purchasing behavior we observe in Table 5 is not driven by people choosing to purchase different model cars as a result of their assigned credit. Contrasting the coefficients in columns 1 and 2 provides indirect evidence as to the exact nature of the substitution patterns in this market. When we include year-make-model fixed effects in column 2, we find no statistical differences in the car purchase price, suggesting that it's not that borrowers with more affordable credit are paying more for the same car by choosing extra add-on features or simply being overcharged by the seller. Reconciling these two results, column 3 provides direct evidence with vehicle age at purchase in months as the dependent variable (and naturally controlling for make-model fixed effects since age would be collinear with year-make-model fixed effects). Borrowers with easier access to credit purchase 4.8 months newer cars, on average.

How do borrowers respond to being arbitrarily offered more expensive credit than their creditworthiness would warrant? The evidence presented in Tables 5 and 6 indicate that borrowers offered expensive credit spend less on their car purchases by selecting an older car than they would have otherwise, originating smaller loans in the process and having slightly smaller monthly payments. The lack of a difference in car pricing outcomes with year-make-model fixed effects is also evidence that there is little private information or selection around buyers' savvy or skill in car buying decisions based on assignment to one side of a FICO threshold. We view this as evidence of lenders' nonlinear pricing rules distorting consumption away from efficient levels.

[24]Again, our decision to restrict our sample to direct loans mitigates this concern somewhat, although buyers may share details of their financing terms with dealers.

[25]Year-make-model fixed effects are made possible by vehicle identification numbers (VINs) provided in our data set.

6.1 Evaluating Alternative Explanations: Adverse Selection

In this section we address the possibility that borrowers that originate loans at interest rates on the expensive side of FICO thresholds are different on unobservable dimensions from car buyers borrowing at above-threshold, lower interest rates. While we have already demonstrated that borrowers seem similar on observable dimensions at the ex-ante application stage, an alternative explanation for our results is that (unobservably) high credit-quality borrowers who are arbitrarily offered expensive interest rates withdraw their loan applications and look elsewhere for credit. In this story, borrowers who follow through originating expensive loans are those who know they are of poor credit quality and unlikely to do better given their unfavorable soft attributes. Lending institutions could also recognize that borrowers that choose to accept the unfavorable terms are indeed lemons, as anticipated, and so the arbitrary thresholds reinforce an equilibrium that separates high credit-quality borrowers from low credit-quality borrowers, with the appropriate pricing differences offered to each borrower type.

We test for the possibility that adverse selection drives the observed equilibrium outcomes in our data by comparing ex-post borrower performance around the FICO thresholds. If an unobservable selection process guides differences in who accepts expensive loan offers, this should be revealed by ex-post credit outcomes as lower credit-quality borrowers eventually default relatively more. To test this hypothesis, we first specify as a dependent variable in our RD setting the number of days a borrower is subsequently delinquent on their car loan. The coefficient in column 1 of Table 7 estimates that above-threshold, cheaper-credit borrowers are an average of four fewer days delinquent than constrained borrowers, indicating that borrowers on either side of the threshold do not exhibit economically meaningful or statistically significant differences in delinquency. Similarly, constrained borrowers are 0.08% more likely to have their loan charged off (written off as a loss by the lender) and 0.2% more likely to be in default (over 90 days past due), both of which estimates we view as relatively precise zeroes.[26]

A novel feature of our dataset allows for a second test of adverse selection as an explanation for our observed results. As a means of monitoring borrowers, many lending institutions in our dataset pull credit scores on borrowers after loan origination. Ex-post credit score queries occur as frequently as every six months, and, in a few cases, as infrequently as once post-origination. The most common convention for the subset of institutions that pull credit ex-post is to pull credit scores once a year. Ex-post credit scores allow us to calculate changes in credit scores over time, capturing broad changes in borrower distress and financial responsibility. Any unobserved heterogeneity driving selection into loan take-up should impact credit scores over time if low credit-quality borrowers for whom the below-threshold expensive interest rate is "fair" are the only ones to originate such loans. Using the sub-sample of institutions that collect updated FICO scores after origination, we use the percentage change between credit scores at origination and the most recently observable credit score as the dependent variable in our RD framework. Results presented in column 4 of Table 7 show no meaningful differences in credit score changes for borrowers around the threshold.

Taken together, the evidence on borrower delinquency, defaults, and ex-post changes in credit scores indicate that borrowers to the left of FICO thresholds do not represent meaningfully different credit risks as compared to otherwise similar borrowers to the right of thresholds. While adverse selection is undoubtedly a motivator of many features of retail car loan markets (Adams, Einav, and Levin, 2009), information asymmetries do not appear to be a primary determinant of the acute differences in lending behavior around the observed FICO thresholds.[27]

[26]The sample size differs across columns in Table 7 because of inconsistent data coverage of all variables across lenders.

[27]Adverse selection is not the only alternative explanation for our observed results around thresholds. For example, FICO thresholds could promote the steering of financially unsophisticated borrowers into higher rate loans. However, as any such borrower naïveté is not manifest in differences at loan application, more expensive car purchase prices paid, differential ex-post default rates, or differences in ex-post credit scores, it is likely not a driving factor for the phenomena we document here.

7 Credit Constraints and Search Costs

While we have demonstrated that discontinuities in offered loan terms act as *de facto* credit constraints that alter consumption decisions, the equilibrium persistence of credit constraints remains a puzzle. In frictionless, competitive credit markets, borrowers of similar risk types should be offered similar rates. While adverse selection and moral hazard have not disappeared from retail loan markets, the modern era of Big Data has promised to help mitigate asymmetric information. We offer an additional explanation for ongoing credit constraints, motivated by recent work in household finance suggesting that borrowers are reluctant to shop for loans. As mentioned in the introduction, questions on search intensity in the Survey of Consumer Finance indicate that many borrowers self-report doing very little shopping around for a loan. Woodward and Hall (2012) document that mortgage borrowers overpay for mortgage origination services due to a reluctance to shop for mortgages. Zinman & Stango (2015) document price dispersion in the credit card market that varies with borrower shopping intensity. In this section we explore whether a search cost explanation for persistence in credit constraints has support in the data.

As discussed in Section 2, theories of search costs (e.g., Stahl, 1989) suggest that when there is heterogeneity in the costliness of consumer search, many agents find it too costly to solicit the full menu of offered prices. As a result, equilibrium prices reflect the distribution of offered prices and the random draw that agents get from the offered price distribution. Consider a financial institution that offers an interest rate on auto loans that is high relative to competitors, conditional on borrower quality. If search is costly, consumers that arrive randomly to solicit a loan are more likely to accept the offered rate, despite the existence of better available rates. Lenders can expect to make loans in the presence of search costs despite not offering the lowest rates among their competitors because of the possibility that a randomly arriving customer will not exert the effort required to find better rates (see Sorensen, 2000, for a recent discussion). Similarly, entrants cannot profitably undercut overpriced competitors because of entrants' inability to attract consumers. Lowering search costs will result in lower price dispersion as consumers increase their propensity to search, thus tracing out a more complete distribution of available prices. In a similar spirit, if consumer search costs are reduced, lenders will be forced to offer more competitive rates knowing that consumers are more likely to search out competitors rates.

Empirical evidence on the prevalence of price dispersion in many markets that should otherwise follow the Law of One Price was a prime motivation behind many theories of search. In the following sections we bring these ideas to our auto-lending data in two ways. First, we document the existence of price dispersion in the auto loan market and test whether this dispersion interacts with discontinuities in lender pricing rules. Doing so allows us to assert whether constrained borrowers could have found better loans. Second, we evaluate indirectly (using take-up data) and directly (using borrower-linked application data across lenders within a market) whether borrowers' propensity to search for better loan terms is correlated with measures of search costs.

7.1 Are loans with better terms available to constrained borrowers?

To assess whether interest-rate constrained borrowers (on the left side of FICO thresholds) could have found a loan with more favorable loan terms, we first check for the presence of price dispersion. Diagnosing a market with dispersed prices requires ruling out any product differentiation, i.e., that differences in prices truly represent identical goods being sold for different prices in the same market. We first estimate the spread between a loan note's rate and the lowest available interest rate at another lender in our data for borrowers with FICO scores just to the left of observed FICO thresholds in each market. To calculate this spread, we group borrowers in the same MSA, six-month window, five-point FICO bin, \$1,000 purchase-price bin, five percentage-point DTI bin. The \$1,000 auto purchase bins are non-overlapping, beginning from \$2,000 to \$2,999, up to a maximum purchase amount of \$100,000. We consider loans originated to borrowers within the same MSA $\times$ time $\times$ price $\times$ FICO $\times$ DTI cell to be effectively identical. Owing to the strictness of this criteria, many borrowers in our data are in their own cell, limiting our ability to calculate a spread. Moreover, because we do not observe interest-rate offers from lenders that are not clients of our data provider, these spreads are lower bounds (having the universe of interest rates offered to a given cell could only weakly decrease the best available rate).[28] However, albeit incomplete, because of the richness of our data coverage, and the opportunity in this exercise to use data from all originators not just those in our discontinuity sample, we have hundreds or thousands of cells with multiple borrowers for each FICO range we consider.

Table 8 tabulates summary statistics of the spread to the best available rate for constrained borrowers. The average spread for borrowers with FICO scores from 595 to 599, 635–639, and 695–699 is 3.4 percentage points (pp), 2.5 pp, and 1.6 pp, respectively. That is, borrowers with FICO scores between 595 and 599 that took out a loan from a lender with a pricing discontinuity at 600, there was a loan with a 3.4 percentage point lower interest rate originated to someone with the same FICO and DTI in the same MSA at the same time and to secure a similarly priced car. The standard deviation of the spread across these cells is 2.5%, 2.1%, and 1.4% with an average number of borrowers in the cell of 2.46, 2.99, and 3.99, respectively.

To be clear, we do not consider FICO-based pricing discontinuities as the main driver of price dispersion. Costly search in the market for retail auto loans would lead to price dispersion even if pricing policies were completely smooth functions of credit-risk metrics. Figure 6 plots kernel densities for the spread to the lowest available rate for borrowers below and above pricing discontinuities at FICO scores of 600, 640, and 700 in panels A–C, respectively. The solid red line shows that there is significant variation in prices even for basically identical borrowers in the control group (those just above FICO discontinuities) who were offered relatively competitive interest rates. Comparing the red line across panels, price dispersion seems to be more acute at lower FICO scores, consistent with lower creditworthy borrowers having a harder time searching for credit perhaps because of a more binding

[28]We discuss the particular case of digital lenders in section 7.4 below.

underwriting process for that population. Contrasting the densities of the below- and above-threshold borrowers (dashed blue and solid red lines, respectively), there is clearly more dispersion for borrowers facing expensive prices. On average, they are foregoing much more attractive opportunities than borrowers offered competitive rates.[29]

7.1.1 Exclusivity of Credit Unions

By definition, a credit union is member-owned cooperative financial institution that requires membership to receive financial services. Often, credit unions' membership requirements restricting eligibility to well-defined groups. Because most of the loans in our sample were originated by credit unions, one concern is whether a given borrower could have joined the credit union with the best available rate used to demonstrate the existence of dominating loan opportunities for that particular cell. For example, if the lowest available interest rate that we assume could have been obtained by a borrower was offered by a firefighters credit union, then borrower search costs would not only involve the effort required to find the low rate but also the effort required to become a firefighter. To address this concern, our primary estimation sample is comprised entirely of credit unions whose primary membership requirement is residence in a specified geographic area. In other words, all borrowers in our MSA-based matched portfolios are eligible to become a member at any of the credit unions included in their cell by virtue of living in the same MSA as others in their cell. In practice, because we also exclude institutions that do not have a sufficient number of loans around discontinuities, this restriction removes little of the data—smaller, niche credit unions do not generally have enough members to meet our sample requirements. We also note that the finance companies in our sample have no membership requirements.

[29] We readily acknowledge that some of the thick right tails of these densities could be driven by unobserved heterogeneity not captured by FICO or our other matching covariates. Still, our earlier results suggest that adverse selection is not a on average. Moreover, the difference in means remains significant even after censoring the extreme values of each distribution.

7.2 Measures of loan search

The summary statistics tabulated in Table 8 and the kernel densities in Figure 6 confirm that borrowers to the left of lending thresholds had significantly better interest rates available to them. Earlier results indicated that ex-post borrower and loan outcomes are indistinguishable around FICO thresholds, indicating that adverse selection is not the primary cause for differences in offered rates around thresholds. Why, then, did borrowers treated with expensive rates not avail themselves of better lending opportunities? In this section we evaluate more directly whether borrowers' propensity to search is correlated with search costs.

This analysis requires the construction of two measures; a measure of search propensity and a measure of search costs. We primarily measure search propensity as the fraction of borrowers that reject an offered loan under the assumption that borrowers who applied for and then reject an offered loan reject the loan in favor of searching for a different loan (see below for a more direct measure of search using a limited subset of our data). While borrowers could reject loans for different reasons, including the decision not to originate any loan, we view differences in the decision to reject an offered loan around lending thresholds as a reasonable proxy for differences in borrowers' propensity to engage in further loan search. Using reported FICO scores at the time of loan application, we estimate differences in take-up rates around FICO thresholds.

To proxy for search costs, we use FDIC and NCUA data to identify the physical location of every bank branch and credit union branch in the United States for each year in our application data. We then create a measure of driving-time density for an individual borrower by geocoding and counting the number of physical branch locations within a 20-minute drive of the borrower, similar to the approach employed by Degryse and Ongena (2005). The calculation of actual driving times relies on posted speed limits along current driving routes and abstracts from traffic conditions and any changes to the road network between the time of loan origination and 2016 (the date of our driving-time data). We calculate driving distances and trip durations for only those institutions which existed at the time of loan origination.

The driving-time density measure is designed to capture the effort, proxied by travel time and physical distance, for each borrower to search out a lending institution that is within a reasonable distance from their home. Clearly, there are many other dimensions over which search is costly besides time and distance, for example the disutility of filling out financial paperwork and potential concerns that applying for too much credit negatively impacts credit scores (see Liberman, Paravisini, and Pathania, 2016, for a discussion of the impact of this characteristic of credit scoring). If our search-costs proxy influence the propensity to search, differences in loan take up-rates around FICO thresholds should be higher in areas with higher driving-time density. If a borrower lives in an area with higher driving-time density, borrowers in that area to the left of FICO thresholds, i.e. those offered unfavorable loan terms, would be more likely to reject the unfavorable loan terms in favor of searching for better terms elsewhere. Importantly, in a differences-in-differences spirit, our empirical specification measures *differences* in loan take-up rates around FICO thresholds. We then compare differences in loan take-up rates for borrowers in high versus low driving-time density areas.

The median borrower in our application data lives within a 20 minute drive of 63 lending institutions. Borrowers in the 25th percentile of driving distance live less than a 20-minute drive from 20 institutions, as compared to 148 institutions for borrowers in the 75th percentile. Just under 4% of applicants in our sample live in an area with one or fewer lending institutions within a 20-minute drive. In contrast, 66% of applicants live within a 20-minute drive of at least 100 different lending institutions.

Using our standard RD framework described by equation 10, we specify a loan take-up indicator as the dependent variable and estimate differences in the take-up rate around FICO thresholds. We then estimate RD regressions for the full sample and separately for borrowers in the top and bottom half of the search cost distribution and report results in Table 9. For the full sample, column 1 documents differences in take-up rates around the threshold of 16 percentage points; that is, borrowers assigned expensive interest rates are 16.4 percentage points less likely to take out the loan. In high search cost areas, defined as locations below the median number of lenders in a 20-minute drive, treated borrowers (with expensive interest rates) are only 13 percentage points less likely to accept a loan offer than control-group borrowers. In contrast, differences in take-up rates around FICO thresholds are 21 percentage points in low search cost areas (column 3). These results indicate that borrowers in high search cost areas are more likely than borrowers in low search cost areas to accept dominated loan terms that are less favorable than they might otherwise obtain given the relatively higher cost associated with finding better terms in such an area.

Proxying for search costs with driving-time density may not uniquely measure borrower search costs. Driving-time density, as constructed, might also be a correlate of other local factors such as the degree of price competition among lenders. Indeed, costly search is a fundamental source of imperfect competition. In an effort to differentiate between search costs and a pure competition story, we construct empirical measures of lending competition within MSAs. We calculate the share of originated mortgage loans by each HMDA lenders within a given MSA and use the origination shares to construct an MSA-level Herfindahl index and divide loans into low and high competition areas based on our constructed Herfindahl index. We then reestimate the specification of Table 9 for all four combinations of high and low search cost areas crossed with high and low competition areas.

The results of this exercise in Table 10 highlight that even within a competition category, there are statistically significant differences by search costs in the difference of take-up rates across FICO thresholds. In other words, even for areas with a highly competitive banking sector, car-loan borrowers in high search cost areas are much more likely to accept dominated loan terms. For low-competition MSAs, in low search-cost areas, the difference in take-up rates around lending thresholds is 25 percentage points (treated borrowers are 25 pp more likely to walk away from an expensive loan than control-group borrowers). In comparison, borrowers in high search-cost areas in the same competition bin are only 11 percentage points more likely to walk away when offered an expensive loan. For the high competition bin, we find similar results with higher search costs (lower search costs) resulting in a take-up differential of 8 percent (13 percent). This result show that regardless of the level of competition, borrowers in areas we expect to have high search costs are much less sensitive to interest rates in their extensive-margin loan take-up decision. Combined with our prior results that those borrowers exogenously offered expensive loans are effectively constrained away from their first-best car purchase, we conclude that interest-rate–driven credit constraints are more pernicious when search is costly.

7.3 Direct Evidence of Loan Search

In this section, we present a more direct measure of loan search by attempting to link loan applications across financial institutions for the same borrower. We then evaluate whether the propensity of an individual borrower to search for an auto loan, as measured by the number of filed loan applications, varies with our proxy for search costs. We assume that loan applications originating from the same nine-digit zip code from a borrower with the same birthdate are from the same individual and that loan applications occurring within 90 days of each other are for the same prospective purchase. We divide borrowers into search cost quartiles using the driving-time density as a proxy. Table 11 reports that applicants in high search cost areas (column 1) apply for an average of 1.08 loans per vehicle purchase. In contrast, applicants in low search cost areas (column 2) apply for 1.101 loans/vehicle purchase, on average. Column 3 shows that this difference is statistically significant at a 1% significance level. Although small, these averages are a lower bound as we do not observe applications to any lender not in our data (the vast majority of potential lenders). In untabulated results, we find that borrowers in areas with fewer than 2 lending institutions within a 20-minute drive apply for 1.06 loans per purchase, compared to 1.10 for borrowers with at least 100 institutions within 20 minutes. Regressions of loan applications per purchase on the count of lending institutions within a 20-minute drive confirms a positive and significant relationship. These results are consistent with the more comprehensive indirect evidence in section 7.2 above that borrowers facing high search costs search less and accept worse rates than borrowers facing relatively low search costs.

7.4 Digital Search

Many consumers now search for loans on the internet (including using such information aggregators as Bankrate.com), potentially limiting the relevance of lender density and driving distances as a proxy for 21st-century search costs. We view our results with respect to lender density and driving times as even more noteworthy given the increased propensity for borrowers to search for loans online. One potential explanation for the ability of physical search measures to explain variation in loan search propensity is the fact that our sample is skewed towards older borrowers. Another possibility is that, although borrowers can be easily pre-approved on the internet, the actual closing of loans (signing documents, transfer of title, etc.) still most frequently occurs at physical branch or dealer locations, even for direct loans.

8 Conclusion

Mounting evidence indicates that credit constraints continue to be a central topic in household finance and that the supply of finance does indeed influence consumer purchasing decisions. In this paper, we made two main points. First, we presented evidence that many borrowers in retail auto loan markets behave as though they are credit constrained by overly expensive interest rates. The quasi-random assignment of expensive loan terms causes car buyers to spend less on their car purchase by selecting an older car than they would have at competitive risk-adjusted interest rates. We rule out alternative explanations for the statistically significant response of tightening car buyers' credit constraints using data on ex-ante loan application behavior and ex-post loan and borrower outcomes.

Second, we proposed search costs as a new explanation for the persistence of both credit constraints and equilibrium price dispersion in the auto loans market. Because the arbitrary pricing schedules vary across lending institutions within the same MSA, borrowers on the expensive side of FICO discontinuities in loan pricing at one institution would likely find themselves on the favorable side of a pricing threshold at a different institution. Absent search frictions, borrowers would never accept dominated loan terms. However, we find that for constrained borrowers in our sample, a lower interest rate loan was frequently originated by an equally creditworthy borrower purchasing a similarly priced car on a similar date, suggesting that borrowers are either unwilling or unable to search for more favorable loan terms. Measuring the costliness of acquiring information about loan terms using the driving-time density of borrowers to the nearest lending institutions, we show that in areas with higher search costs, borrowers are more likely to accept inferior loan terms, constraining their access to efficient prices in credit markets. Although we have focused on probing the interaction between our detected discontinuities in pricing rules and search costs, the evidence suggests that the entire retail car loan market is subject to costly search and the resulting price dispersion.

Even with a well-developed financial sector including secondary markets for many forms of consumer debt, households are still constrained by their access to credit. While this result is well-known from prior empirical work, we provide a novel answer to an unresolved follow-up question in household finance—why such constraints continue to persist even in an era of adverse-selection–mitigating Big Data. Even with the possibility of shopping for interest rates online, searching for consumer credit products remains an opaque, local, and costly process for many borrowers. This relationship between costly search and distortionary credit market imperfections extends our understanding of equilibrium price dispersion to credit markets and could motivate both the provision of pricing information as a public good and extra regulatory attention on so-called banking deserts.

References

Adams, W., L. Einav, and J. Levin (2009). "Liquidity constraints and imperfect information in subprime lending." *The American Economic Review* 99.1, pp. 49–84.

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2015). "Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?" NBER Working Paper No. 21567.

Agarwal, S., L. Chunlin, and N. S. Souleles (2007). "The Reaction of Consumer Spending and Debt to Tax Rebates-Evidence from Consumer Credit Data." *Journal of Political Economy* 115.6, pp. 986–1019.

Angrist, J. D. and V. Lavy (1999). "Using Maimonides' rule to estimate the effect of class size on scholastic achievement." *The Quarterly Journal of Economics* 114.2, pp. 533–575.

Ausubel, L. M. (1991). "The Failure of Competition in the Credit Card Market." *The American Economic Review* 81.1, pp. 50–81.

Aydin, D. (2015). "The Marginal Propensity to Consume Out of Credit: Evidence From Random Assignment of 54,522 Credit Lines." Working Paperi.

Baker, S. R. (2015). "Debt and the Consumption Response to Household Income Shocks." Working Paper.

Baye, M. R., J. Morgan, and P. Scholten (2006). "Information, search, and price dispersion." *Handbook on Economics and Information Systems*. Ed. by T. Hendershott. Vol. 1. Citeseer, pp. 323–77.

Bubb, R. and A. Kaufman (2014). "Securitization and moral hazard: Evidence from credit score cutoff rules." *Journal of Monetary Economics* 63, pp. 1–18.

Burdett, K. and K. L. Judd (1983). "Equilibrium price dispersion." *Econometrica*, pp. 955–969.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82.6, pp. 2295–2326.

Carroll, C. D. (2001). "A theory of the consumption function, with and without liquidity constraints." *The Journal of Economic Perspectives* 15.3, pp. 23–45.

Davis, S. J., F. Kubler, and P. Willen (2006). "Borrowing costs and the demand for equity over the life cycle." *The Review of Economics and Statistics* 88.2, pp. 348–362.

Degryse, H. and S. Ongena (2005). "Distance, lending relationships, and competition." *The Journal of Finance* 60.1, pp. 231–266.

Diamond, P. A. (1971). "A model of price adjustment." *Journal of Economic Theory* 3.2, pp. 156–168.

Friedman, M. (1957). *A Theory of the Consumption Function*. Princeton: Princeton University Press.

— (1963). “Windfalls, the ‘Horizon,’ and Related Concepts in the Permanent-Income Hypothesis.” *Measurement in Economics: Studies in Mathematical Economics and Econometrics in Memory of Yehuda Grunfeld*. Ed. by C. F. Christ. Stanford, California: Stanford University Press, pp. 3–28.

Gross, D. B. and N. S. Souleles (2002). “Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data.” *Quarterly Journal of Economics* 117.1, pp. 149–185.

Hall, R. E. (1978). “Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence.” *Journal of Political Economy* 86.6, pp. 971–987.

Hortaçsu, A. and C. Syverson (2004). “Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds.” *The Quarterly Journal of Economics* 119.2, pp. 403–456.

Houthakker, H. S. (1958). “The permanent income hypothesis: reply.” *The American Economic Review* 48.5, pp. 991–993.

Hutto, G. W. and J. Lederman (2003). *Handbook of Mortgage Lending*. Mortgage Bankers Association of America.

Jappelli, T., J. Pischke, and N. S. Souleles (1998). “Testing for liquidity constraints in Euler equations with complementary data sources.” *Review of Economics and Statistics* 80.2, pp. 251–262.

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

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). “Household Expenditure and the Income Tax Rebates of 2001.” *The American Economic Review* 96.5, pp. 1589–1610.

Keys, B. J., T. Mukherjee, A. Seru, and V. Vig (2009). “Financial regulation and securitization: Evidence from subprime loans.” *Journal of Monetary Economics* 56.5. Carnegie-Rochester Conference Series on Public Policy: Distress in Credit Markets: Theory, Empirics, and Policy November 14-15, 2008, pp. 700–720.

Keys, B. J., T. Mukherjee, A. Seru, and V. Vig (2010). “Did Securitization Lead to Lax Screening? Evidence from Subprime Loans.” *The Quarterly Journal of Economics* 125.1, pp. 307–362.

Laufer, S. and A. Paciorek (2016). “The Effects of Mortgage Credit Availability: Evidence from Minimum Credit Score Lending Rules.” FEDS Working Paper No. 2016-098.

Liberman, A., D. Paravisini, and V. S. Pathania (2016). “High-Cost Debt and Borrower Reputation: Evidence from the UK.” Working Paper.

Ludvigson, S. (1999). "Consumption and credit: a model of time-varying liquidity constraints." *Review of Economics and Statistics* 81.3, pp. 434–447.

Mian, A., K. Rao, and A. Sufi (2013). "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128.4, pp. 1687–1726.

Mian, A. and A. Sufi (2011). "House prices, home equity-based borrowing, and the US household leverage crisis." *The American Economic Review* 101.5, pp. 2132–2156.

— (2014). "House price gains and US household spending from 2002 to 2006." NBER Working Paper No. 20152.

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013). "Consumer spending and the economic stimulus payments of 2008." *The American Economic Review* 103.6, pp. 2530–2553.

Petersen, M. A. (2004). "Information: Hard and Soft." Working Paper.

Rothschild, M. (1973). "Models of market organization with imperfect information: A survey." *Journal of Political Economy* 81.6, pp. 1283–1308.

Salop, S. and J. E. Stiglitz (1982). "The theory of sales: A simple model of equilibrium price dispersion with identical agents." *The American Economic Review* 72.5, pp. 1121–1130.

Sorensen, A. T. (2000). "Price Dispersion in Retail Markets for Prescription Drugs." *Journal of Political Economy* 108.4, pp. 833–850.

Stahl, D. O. (1989). "Oligopolistic pricing with sequential consumer search." *The American Economic Review*, pp. 700–712.

Stango, V. and J. Zinman (2015). "Borrowing high versus borrowing higher: price dispersion and shopping behavior in the US credit card market." *Review of Financial Studies*, hhv072.

Stigler, G. J. (1961). "The economics of information." *Journal of Political Economy* 69.3, pp. 213–225.

Woodward, S. E. and R. E. Hall (2012). "Diagnosing consumer confusion and sub-optimal shopping effort: Theory and mortgage-market evidence." *The American Economic Review* 102.7, pp. 3249–3276.

Zeldes, S. P. (1989). "Consumption and liquidity constraints: an empirical investigation." *Journal of Political Economy* 97.2, pp. 305–346.

Zinman, J. (2014). "Consumer Credit: Too Much or Too Little (or Just Right)?" *The Journal of Legal Studies* 43.S2, S209–S237.

Figure 1: Examples of FICO-Based Discontinuities in Interest-Rate Policies

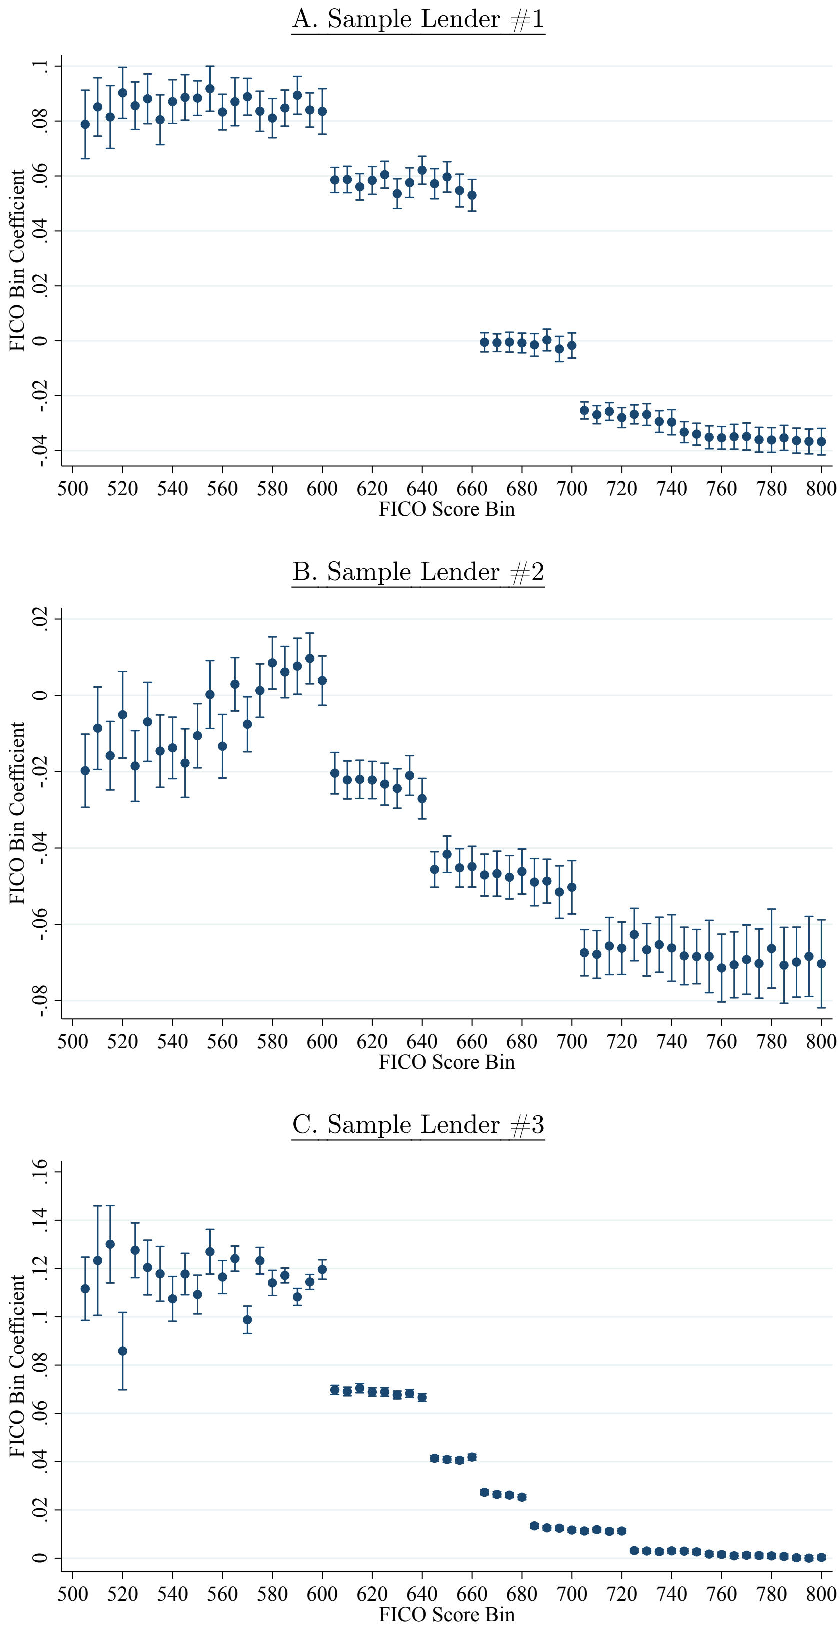

Notes: Each panel plots estimated interest-rate rules (with 95% confidence intervals) for a lender in our sample. Loan rates in percentage points are regressed on 5-point FICO bin indicators as in equation (9).

Figure 2: Examples of FICO-Based Discontinuities in Loan-Term Policies

A. Sample Lender #4

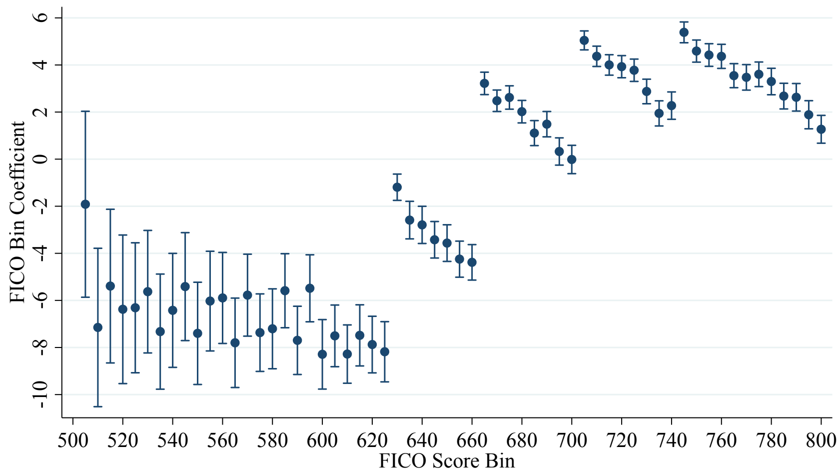

B. Sample Lender #5

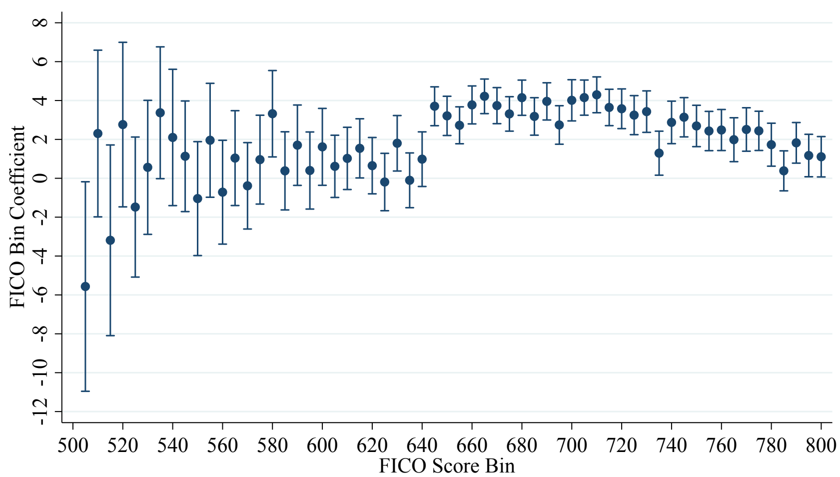

C. Sample Lender #6

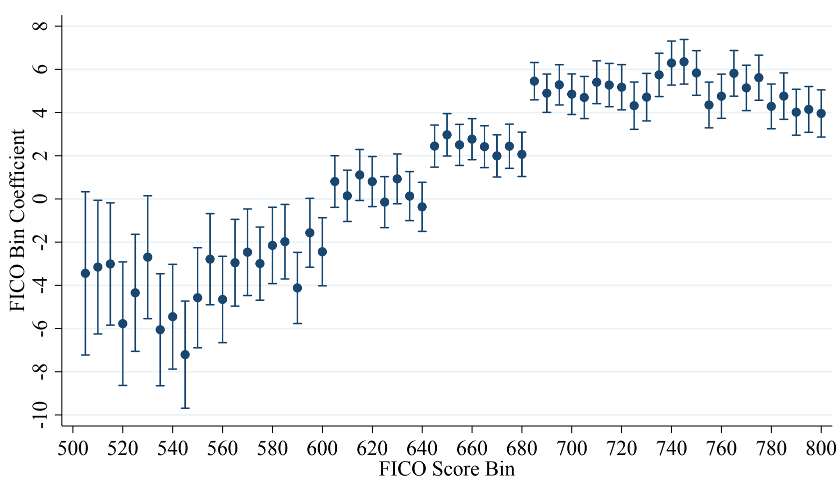

Notes: Each panel plots estimated loan-term rules (and 95% Confidence Intervals) for a lender in our sample. Loan terms (in months) are regressed on 5-point FICO bin indicators as in equation (9).

Figure 3: FICO-Based Lending Policies - Interest Rates

A. Interest Rates Around FICO = 600 Discontinuities

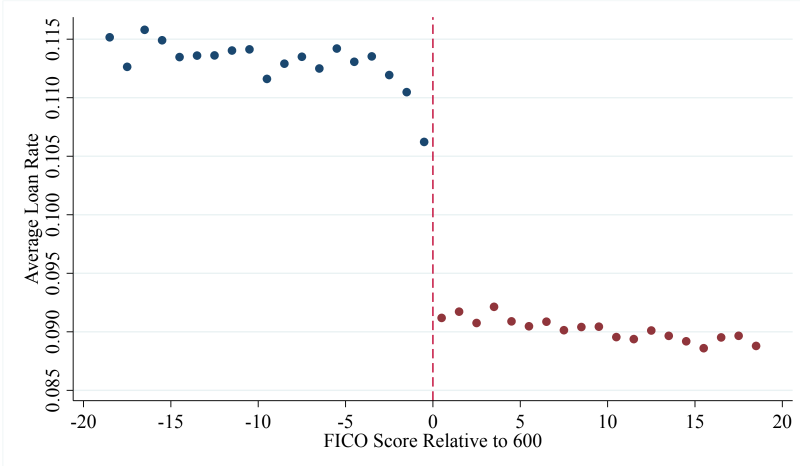

B. Interest Rates Around FICO = 640 Discontinuities

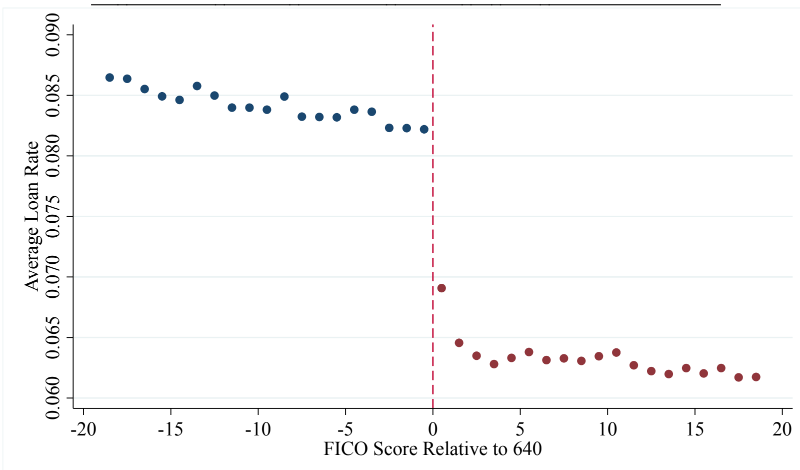

C. Interest Rates Around FICO = 700 Discontinuities

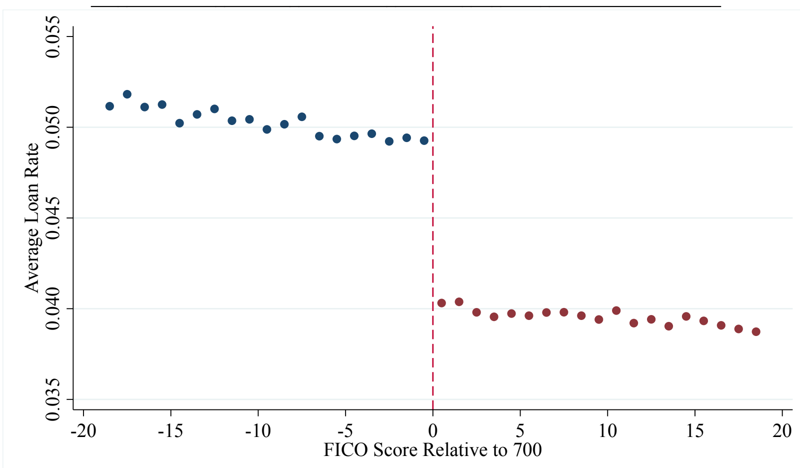

Notes: Figures plot average interest rates on the vertical axis against borrower FICO scores normalized to each threshold along the horizontal axis for institutions with pricing discontinuities detected at FICO scores of 600, 640, and 700, respectively.

Figure 4: Balance of Borrower Characteristics Across FICO Thresholds

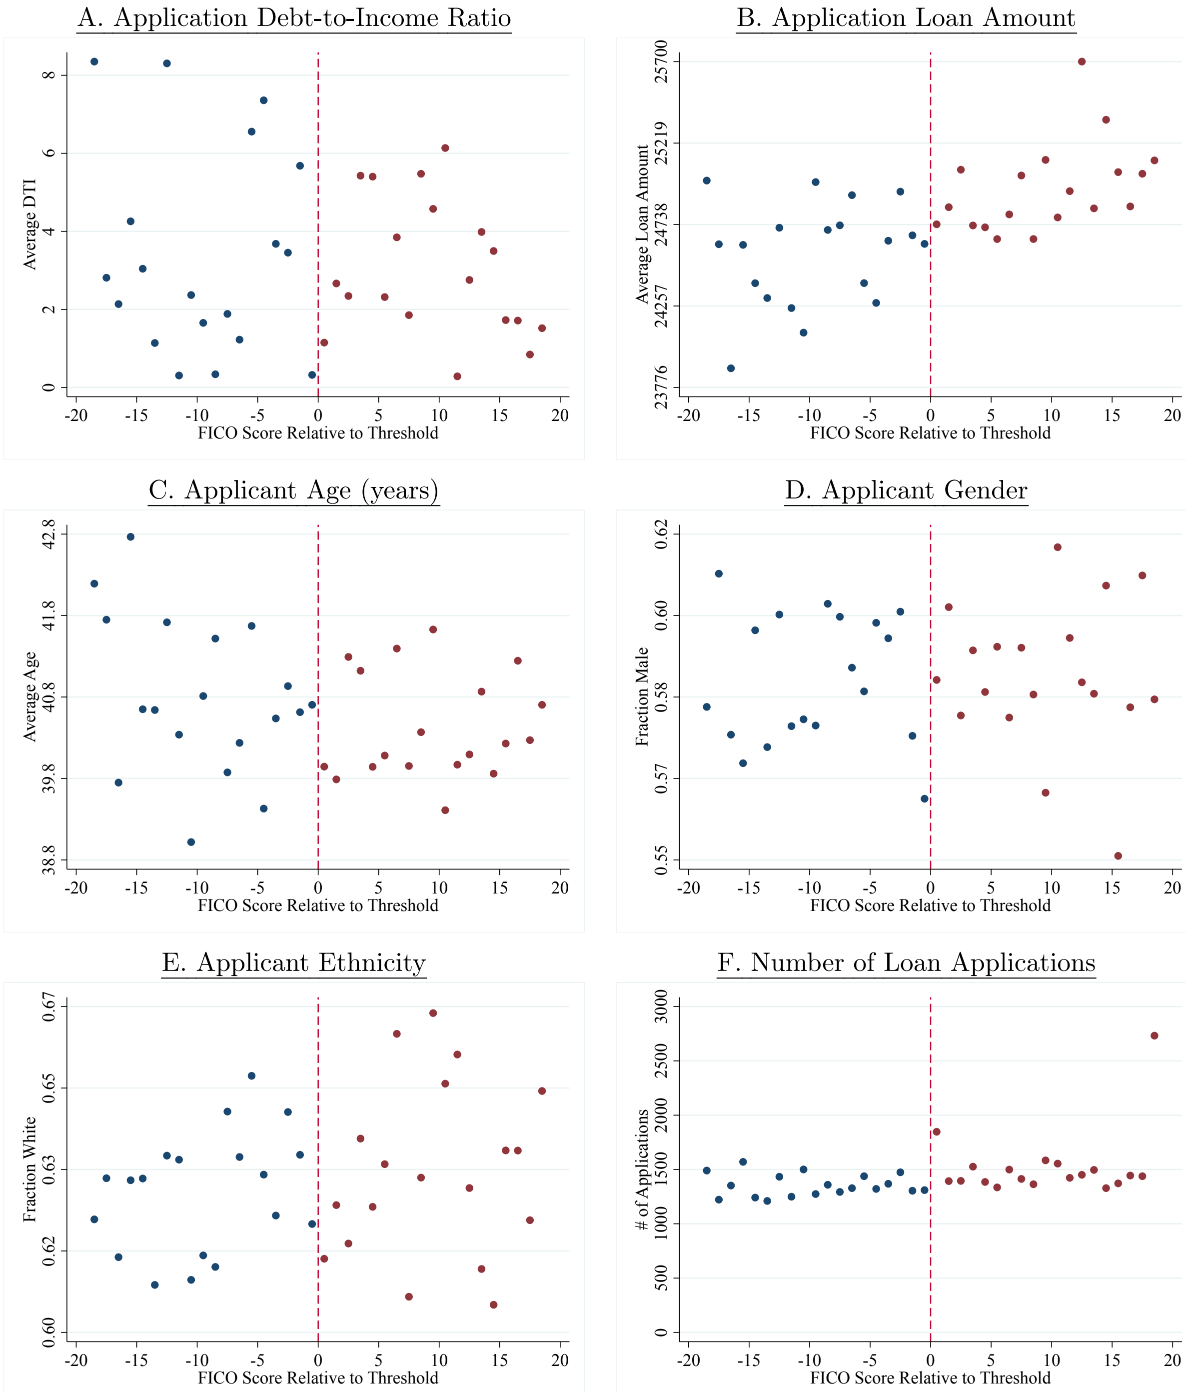

Notes: Figures plot average values of ex-ante borrower characteristics around FICO thresholds for institutions with detected discontinuities. Applicant gender in panel D is an indicator for male, and ethnicity in panel E is an indicator for whether the applicant is estimated as white by the lender. Panel F plots the number of applicants within each normalized FICO bin.

Figure 5: Effect of FICO Threshold on Value of Car Purchased

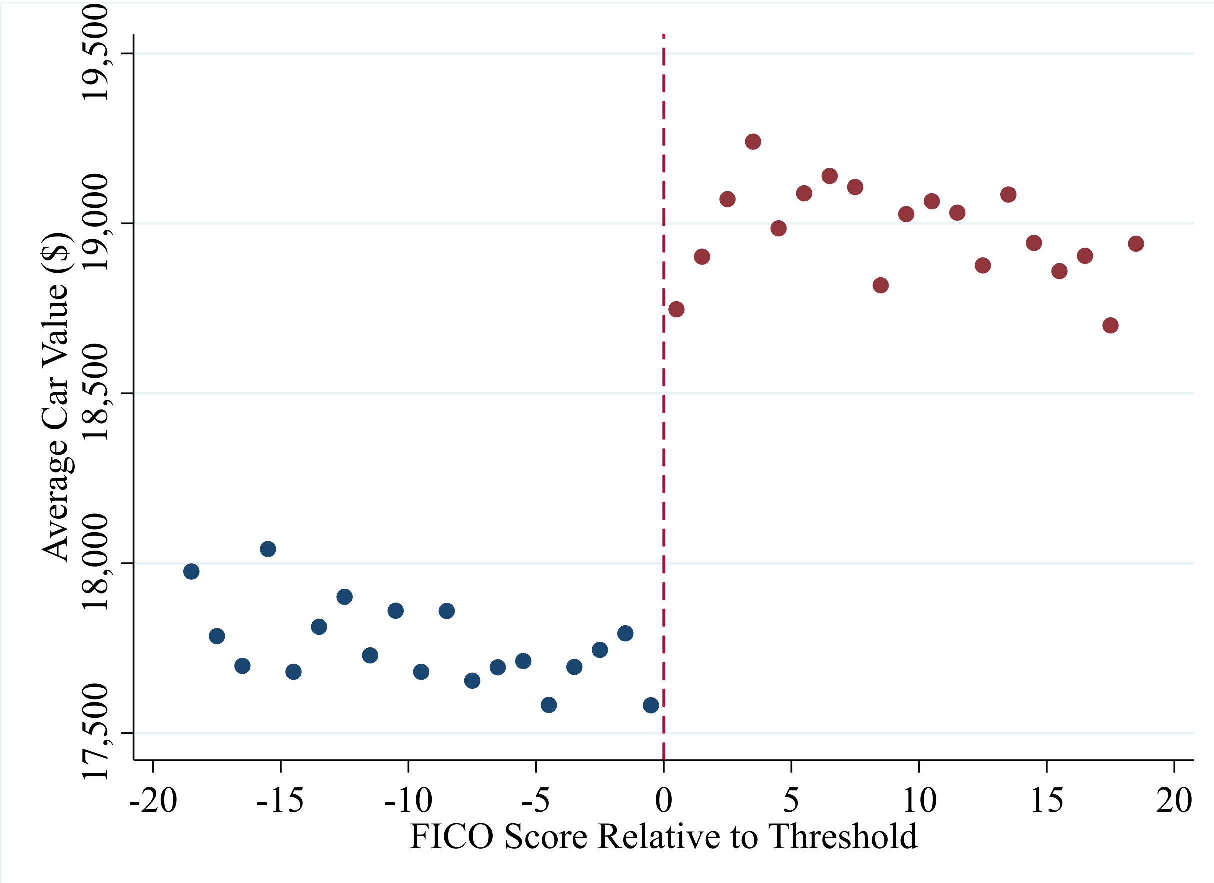

Notes: Figure plots average car value by normalized FICO score for loans originated by lenders with detected interest discontinuities at FICO thresholds of 600, 640, and 700.

Figure 6: Density of Spread to Lowest Available Interest Rate

A. Borrowers Around a FICO = 600 Threshold

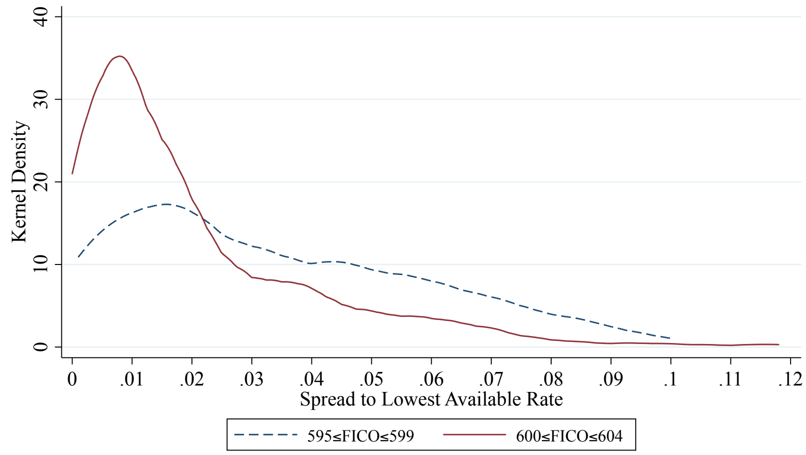

B. Borrowers Around a FICO = 640 Threshold

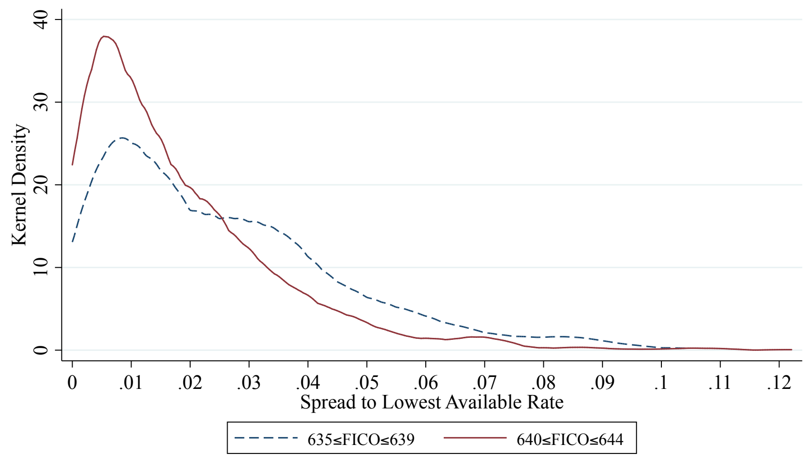

C. Borrowers Around a FICO = 700 Threshold

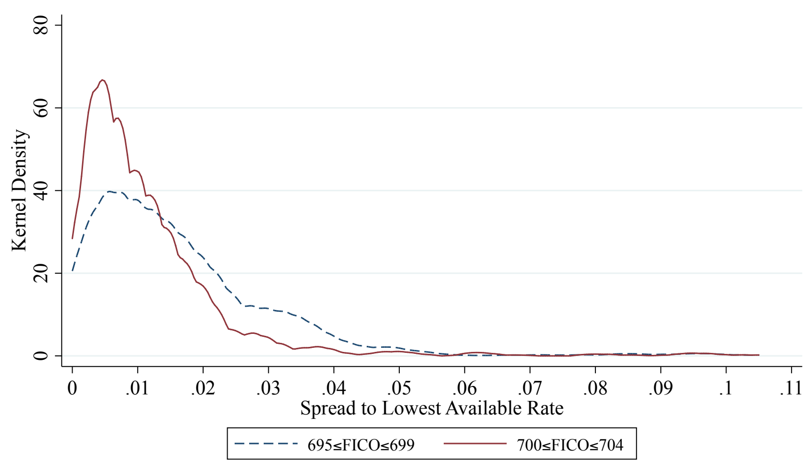

Notes: Figure reports the kernel densities of the spread (in percentage points) to the lowest available rate for borrowers with FICO scores just above just to the left of a threshold that borrowed from institutions with lending thresholds of 600, 640, and 700, respectively.

Table 1: Summary Statistics

	Count	Mean	Std. Dev.	Percentile 25th	Percentile 50th	Percentile 75th
	A. Loan Applications					
Loan Rate	1,716,853	.173	.330	.026	.047	.127
Loan Term (months)	1,666,141	59.80	25.71	42	60	72
Loan Amount ($)	1,916,177	18,884.4	12,242.6	9,415.7	17,100	26,193
FICO	1,297,642	646.4	110.7	594	656	718
Debt-to-Income	1,266,666	.283	.205	.131	.272	.401
	B. Originated Loans					
Loan Rate	3,968,914	.049	.029	.028	.040	.060
Loan Term (months)	3,968,914	64.52	20.45	60	61	72
Loan Amount ($)	3,968,914	19,441	11,196	11,211	17,393	25,366
FICO	3,526,409	712.25	73.4	664	715	769
Debt-to-Income (%)	1,862,758	.264	.357	.089	.300	.370
Collateral Value ($)	3,968,914	20,697	11,071	12,781	18,650	26,405
Monthly Payment ($)	3,968,914	337.74	162.49	222.8	310.5	422.2
	C. Ex-Post Loan Performance Measures					
Days Delinquent	2,848,807	35.63	1054.2	0	0	0
Charged-off Indicator	3,968,914	.021	.142	0	0	0
Default Indicator	3,968,914	.022	.147	0	0	0
Current FICO	2,759,442	704.42	83.53	653	713	771
%ΔFICO	2,727,173	-.010	.089	-.044	-.002	.032

Note: Panels A–C respectively report summary statistics for loan applications, originated loans, and ex-post loan performance. "Loan Term" is the term (in months) of the loan. Loan Rate is the annual interest rate of the loan. Debt-to-Income is the ratio of debt service payments to income. Collateral Value is the value of the car at origination. Current FICO is an updated FICO score for each borrower as of the date of our data extract. ΔFICO is the change in FICO score since origination as a fraction of the FICO score at origination. Days Delinquent is the number of days that a borrower has missed one or more monthly payments. Charged-off Indicator is a dummy for whether a loan has been written off the books of the lending institution. Default is an indicator for whether a borrower has been delinquent for at least 90 days.

Table 2: Summary Statistics for Estimation Sample with Identified FICO Discontinuities

				Percentile		
	Count	Mean	Std. Dev.	25th	50th	75th
	A. Loan Applications					
Loan Rate	38,236	.052	.054	.038	.042	.060
Loan Term (months)	42,568	70.6	16.3	61	72	84
Loan Amount ($)	73,516	24,923	11,924	16,000	23,474	32,294
FICO	54,715	663	42.1	623	683	700
Debt-to-Income	51,257	.270	.195	.130	.270	.390
	B. Originated Loans					
Loan Rate	489,315	.067	.031	.043	.060	.085
Loan Term (months)	489,315	63.7	21.5	50	61	72
Loan Amount ($)	489,315	18,018	10,709	10,149	16,018	23,691
FICO	489,315	655	37.5	627	649	692
Debt-to-Income	278,160	.252	.309	.014	.257	.374
Collateral Value	489,315	18,524	10,074	11,425	16,675	23,625
Monthly Payment	489,315	328	154	219	303	408
	C. Ex-Post Loan Performance Measures					
Days Delinquent	336,961	46.8	289	0	0	0
Charged-off Indicator	489,315	.034	.181	0	0	0
Default Indicator	489,315	.033	.178	0	0	0
Current FICO	369,679	649	69.2	610	654	695
%ΔFICO	369,679	-.011	.093	-.054	.000	.044

Note: Table reports summary statistics for the discontinuity sample (restricted to a 19-point bandwidth around detected FICO discontinuities in lender pricing rules). Panels A, B, and C describe loan applications, loan originations, and ex-post loan performance, respectively. See notes to Table 1 for further details.

Table 3: First-Stage Effects of FICO Discontinuity on Loan Rate and Loan Term

	(1)	(2)
	Loan Rate	Loan Term
Discontinuity Coefficient	-0.0147***	1.378***
	[-29.74]	[5.12]
Institution Fixed Effects	✓	✓
Quarter Fixed Effects	✓	✓
Number of Observations	489,315	489,315

Notes: Table reports regression discontinuity estimates of equation (10), pooling the three discontinuities shown in Figure 3 by normalizing FICO scores around each threshold and using the estimator of Calonico et al. (2014). All specifications include lending institution fixed effects and quarter-of-origination fixed effects. Robust t-statistics reported in brackets are clustered by normalized FICO score.

Table 4: Loan Application Covariate Balance Regressions

	(1) Loan Amount	(2) Debt-to-Income	(3) Number of Loan Applications
Discontinuity Coefficient	-110.090	-2.685	176.212
	[-0.56]	[-1.09]	[1.22]
Institution FE	✓	✓	✓
Quarter FE	✓	✓	✓
Number of Observations	54,715	51,256	39

Notes: Table reports reduced-form RD results for the subset of institutions for which we have detailed loan application data. See notes to Table 3 for more details. Each observation in the data used for column 3 represents a normalized FICO score. Robust t-statistics reported in brackets are clustered by normalized FICO score.

Table 5: Reduced-Form Effects of FICO Discontinuity on Origination Outcomes

	(1) Purchase Price	(2) Loan Amount	(3) LTV	(4) Monthly Payment
Discontinuity Coefficient	978.857***	1,479.674***	0.027***	9.670***
	[11.86]	[13.63]	[5.03]	[6.28]
Institution FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Number of Observations	489,315	489,315	489,315	489,315

Notes: Table reports reduced-form RD estimates of equation (10) using the estimator of Calonico et al. (2014). Columns 1, 2, and 4 are measured in dollars. Robust t-statistics reported in brackets are clustered by normalized FICO score. See notes to Table 3 for more details.

Table 6: Reduced-Form Effects Robustness to Vehicle Heterogeneity

	(1) Purchase Price	(2) Purchase Price	(3) Car Age (months)
Discontinuity Coefficient	887.69***	84.62	-4.80***
	[10.84]	[1.56]	[-20.86]
Institution FE	✓	✓	✓
Quarter FE	✓	✓	✓
Make-Model FE	✓		✓
Year-Make-Model FE		✓	
Number of Observations	448,017	448,017	448,017

Notes: Table reports reduced-form RD estimates of equation (10) on car purchase prices (columns 1–2) and car age in months (column 3). Columns 1 and 3 include make × model fixed effects and column 2 includes year × make × model fixed effects. Robust t-statistics reported in brackets are clustered by normalized FICO score. See notes to Table 3 for more details.

Table 7: Effect of FICO Discontinuities on Ex-Post Credit Outcomes

	(1) Days Delinquent	(2) Charge-off	(3) Default	(4) %ΔFICO
Discontinuity Coefficient	-3.76	-.0008	-.002	.0004
	[-1.12]	[-.64]	[-1.17]	[.18]
Institution FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Number of Observations	336,961	489,315	489,315	369,679

Notes: Table reports RD estimates of equation (10) on ex-post loan and borrower outcomes. Days delinquent is the number of days a borrower is delinquent as of our data extract. Charge-off is an indicator for whether a loan has been written off the books of the lending institution. Default is an indicator for whether a borrower has been delinquent for at least 90 days. ΔFICO is the change in FICO score since origination as a fraction of the FICO score at origination for the sub-sample of institutions that report credit scores after loan origination. Robust t-statistics reported in brackets are clustered by normalized FICO score. See notes to Table 3 for more details.

Table 8: Spread to Lowest Available Rate Summary Statistics

FICO Range	# of Cells	Mean # Borrowers in Cell	Mean	Std. Dev.	Percentile 25th	Percentile 50th	Percentile 75th
$595 \leq FICO \leq 599$	124	2.46	.034	.025	.01	.03	.05
$635 \leq FICO \leq 639$	1,006	2.99	.025	.021	.01	.02	.04
$695 \leq FICO \leq 699$	671	3.92	.016	.014	.006	.01	.02

Notes: Table reports summary statistics for the spread between an above-discontinuity borrower's interest rate and the best available interest rate for borrowers in the same cell. Cells are defined as borrowers with in the same MSA, 5-point FICO bin, $1,000 purchase-price bin, 5 percentage point DTI bin, who take out loans in the same six month window. Within each of the matched bins, we then calculate the average difference between the lowest interest rate in the portfolio and each individual loan in the portfolio. Summary statistics are reported for only those cells that contain at least 2 borrowers.

Table 9: Effect of Search-Cost Proxies on Loan Offer Take-up Decisions

Search Costs	Full	High	Low
	(1)	(2)	(3)
Discontinuity Coefficient	0.164	0.130	0.205
	[8.59]	[4.11]	[7.21]
Institution FE	✓	✓	✓
Quarter FE	✓	✓	✓
Number of Observations	22,735	9,235	9,231

Notes: Table reports results for reduced-form RD regressions of loan take-up (conditional on being offered a loan) separately for the full sample (column 1) and for borrowers in areas with above- and below-median search costs in columns 2 and 3, respectively, using the specification in equation (10). Search costs are estimated using the number of lending institutions within a 20-minute drive. Robust t-statistics reported in brackets are clustered by normalized FICO score. See notes to Table 3 for estimation details.

Table 10: Effects of Search Costs and Competition on Take-up Decisions

		Competition	
		LOW	HIGH
Search Costs	LOW	0.25 [5.33]	0.13 [3.44]
	HIGH	0.11 [2.29]	0.08 [1.65]

Notes: Table reports results for reduced-form RD regressions of loan take-up for borrowers in each combination of areas with above- and below-median search costs and above- and below-median competition. Search costs are estimated using the number of lending institutions within a 20-minute drive. Competition is measured using MSA-level lender mortgage market shares in HMDA data. All regressions include lending institution fixed effects and quarter fixed effects. Robust t-statistics reported in brackets are clustered by normalized FICO score. See notes to Table 3 for estimation details.

Table 11: Number of Loan Applications per Vehicle Purchase by Search Cost Group

	High Search Costs (1)	Low Search Costs (2)	Difference (1) - (2)
Mean	1.08	1.10	-.015
Standard Deviation	(.499)	(.546)	[14.12]
Institution FE	✓	✓	
Quarter FE	✓	✓	
Number of Observations	487,560	482,137	

Notes: Table reports average number of applications per vehicle purchase for applications with reported birthdates and nine-digit addresses. Standard deviations are reported in parentheses. Column 3 calculates the difference in means, along with the robust t-statistic in brackets for the statistical significance of the difference between columns 1 and 2.