Real Effects of Search Frictions in Consumer Credit Markets*

Bronson Argyle† Taylor Nadauld‡ Christopher Palmer§

January 2022

Abstract

We show that search frictions in credit markets affect accepted interest rates and loan sizes and distort consumption. Using data on car loan applications and originations not intermediated by car dealers, we isolate quasi-exogenous variation in both the costs and benefits to searching for credit. After identifying lender-specific policies that price risk discontinuously, we study the differential response to offered interest rates by borrowers in areas with high and low search costs. High-search-cost borrowers are 10% more likely to accept loan offers with higher markups, consequently originating smaller loans and purchasing older and less expensive cars than lower-search-cost borrowers.

Keywords: search, credit markets, auto loans, durable consumption
JEL Codes: D12, D83, E43, G21, L11

*We thank our discussants Paul Calem, Anthony DeFusco, Paul Goldsmith-Pinkham, Steven Laufer, Neale Mahoney, and Johannes Stroebel; seminar, conference, and workshop participants at Berkeley, Boston Federal Reserve, BYU, CFPB, Chicago Federal Reserve, Duke, FDIC, Housing-Urban-Labor-Macro Conference, Imperial, LSF, MIT, NBER Household Finance, NYU, OSU, Oxford, Philadelphia Federal Reserve, Red Rock Finance Conference, San Francisco Federal Reserve, UNC, and the WFA; and John Campbell, Claire Celerier, Maryam Farboodi, Amy Finkelstein, Brigham Frandsen, Peter Ganong, Kyle Herkenhoff, Brad Larsen, Lars Lefgren, Andres Liberman, Brigitte Madrian, Adrien Matray, Gregor Matvos, Adair Morse, Holger Mueller, Hoai-Luu Nguyen, Andrew Paciorek, Jonathan Parker, Brennan Platt, Tobias Salz, Amit Seru, David Sraer, Bryce Stephens, Johannes Stroebel, Stijn Van Nieuwerburgh, Stephen Zeldes, and Jonathan Zinman for helpful conversations. Tommy Brown, Sam Hughes, and Tammy Lee provided excellent research assistance. An anonymous information-technology firm provided the data.

†Brigham Young University; bsa@byu.edu
‡Brigham Young University; taylor.nadauld@byu.edu
§Massachusetts Institute of Technology and NBER; cjpalmer@mit.edu

1 Introduction

Some of the most important questions in household finance center around how various credit-market imperfections affect consumption.[1] In this paper, we demonstrate the special role that search frictions in particular play in both the provision of consumer credit and in determining the equilibrium consumption of durables. We provide evidence that costly search in consumer credit markets can affect extensive- and intensive-margin loan and durable consumption choices. In this sense, search frictions in credit markets are special relative to the search frictions typically studied by the empirical search literature that affect discrete demand for a final good. Because credit demand is continuous, elastic, and an input in final demand, search frictions in credit markets affect not only the distribution of interest rates but also loan sizes and the demand for goods financed with credit (e.g., durables).

Our analysis uses administrative data on 2.4 million used-car loans extended by 326 different financial institutions and 1.3 million loan applications to 41 lenders. The data allow us to exploit large interest-rate discontinuities at various lender-specific credit-score (FICO) thresholds to isolate quasi-exogenous variation in the returns to search.[2] Historical geographic variation in the density of lender branches allows us to develop quasi-exogenous variation in the costliness of searching for credit from nearby lenders. Combining a regression discontinuity strategy with cross-sectional variation in predicted lender proximity, we demonstrate real effects of credit-market search frictions on loan take-up and consumption. Otherwise similar borrowers' propensity to accept loans with higher interest rates is higher in areas likely to have higher search costs even though such areas have weakly more price dispersion.[3] Using this laboratory to study loan and car-purchase characteristics, we consider how the costliness of shopping for credit subsequently impacts consumption. We find that both financing and durable-goods purchasing decisions respond to interest rates such that borrowers facing markups in high-search-cost areas are more likely to substitute towards older and cheaper cars.

We focus on the market for automobile-secured loans for several reasons. The tight link between credit-supply shocks and the demand for cars (Benmelech et al., 2017) gives the car-loan market aggregate importance and makes it a plausible setting to look for credit-market search frictions affecting consumption. Auto debt is the third-largest category of consumer debt in the U.S., with over 114 million outstanding loans (0.89 per U.S. household) totaling \$1.5 trillion (NY Fed, 2021). Most U.S. car purchases are financed (Bartlett, 2013), and vehicles represent over 50% of total assets for U.S. low-wealth households (Campbell, 2006). From an empirical-design standpoint, auto

[1]See, for example, influential work on credit constraints (Zeldes, 1989; Gross and Souleles, 2002), adverse selection in consumer credit markets (Adams, Einav, and Levin, 2009), and credit market concentration (Allen, Clark, and Houde, 2014b; Scharfstein and Sunderam, 2016). We detail the literature on credit frictions in section 1.1.

[2]Credit-score based discontinuities have also been documented in mortgage underwriting (Hutto and Lederman, 2003), mortgage securitization (Keys et al., 2009), mortgage rates (Laufer and Paciorek, 2016), and credit card underwriting (Agarwal et al., 2017).

[3]Consistent with work on other retail financial markets discussed below, we find significant price dispersion—most borrowers in our data could access significantly dominating loan offers if they queried two additional financial institutions.

loans are a relatively homogeneous credit product and can be simply described by their interest rate, term, and amount. Our focus on the large segment of used-car loans *not* intermediated by car dealers allows us to test for credit- and product-market linkages in a non-mechanical setting.[4] Finally, auto loan markets are quite local. The median borrower in our sample originates a loan from a branch that is within a 15-minute drive of her home, contrasted with the median U.S. worker's commute of 26 minutes to work. The stylized fact that direct auto loan markets seem more local than labor markets motivates our inquiry into the distortions that distance-related search frictions might cause in consumer debt markets.

Our empirical strategy features a setting where search costs vary across space and the potential gains to search are high and quasi-randomly assigned in the cross section. First, we document large discontinuities in offered interest rates around FICO thresholds within lending institutions.[5] On average, borrowers just above a FICO discontinuity at their lender are offered loans with interest rates 1.3 percentage points lower than otherwise similar borrowers just below a FICO discontinuity. Importantly, the location of the discontinuities along the FICO spectrum varies across institutions; while some thresholds seem more popular than others, there is no consensus set of thresholds used by a plurality of lenders. We further confirm that borrowers on the high-markup side of a threshold at one institution are more likely to find a significantly better price from another draw of their local price distribution than their fellow borrowers on the low-markup side at the same institution.

This RD laboratory allows us to test how borrowers with a high return to additional search (from being randomly assigned an expensive loan) differentially change their behavior in markets where search costs are high versus low. In particular, these rate discontinuities isolate supply-side interest-rate variation under the assumption that demand-side factors (e.g., preferences, income, financial sophistication) are unlikely to also change discontinuously at FICO thresholds that vary across institutions within the same geography. Intuitively, this is likely to be satisfied given that borrowers are unlikely to know their precise FICO score that will be used in pricing or the location of pricing discontinuities across lenders. We support this identifying assumption of quasi-randomly assigned markup offers with evidence that ex-ante borrower characteristics (including age, gender, ethnicity, application debt-to-income ratio (DTI), application loan size, and the number of loan applications per FICO bin) are balanced around FICO thresholds.[6]

Next, we develop a measure of cross-sectional variation in credit search costs. There are many aspects of loan shopping that may be particularly costly.[7] The process often entails time, hassle, and effort—each of which may be in short supply while simultaneously shopping for a car. To demonstrate the existence of one such distortionary dimension of search costs, we calculate the

[4]See Grunewald et al. (2020) for evidence that auto-loan intermediation by car dealers ("indirect loans" as opposed to the "direct loans" we study here) has important welfare consequences.

[5]We detect lending policies that jump discontinuously at various FICO thresholds at 173 of the 326 lending institutions in our sample.

[6]To compare outcomes conditional on origination, we further demonstrate that ex-ante borrower characteristics and ex-post credit outcomes are also balanced across FICO thresholds in both the origination sample and the application sample.

[7]Consistent with the notion of shopping for credit being costly overall, many borrowers in the 2013 Survey of Consumer Finance report doing very little shopping around for a loan.

number of financial institutions within a 20-minute drive from each borrower address using the physical branch locations of every bank and credit union in the United States. We hypothesize that obtaining multiple loan quotes will be less costly for borrowers with more nearby lenders, and, indeed, the strength of this proxy is borne out in the data in multiple ways.[8] While borrowers with varying lender proximity are presumably different on multiple dimensions, one of the virtues of combining our regression discontinuity (RD) strategy with geographic variation is that our results cannot be driven by fixed differences across high- and low-search cost areas.[9] To account for unobserved factors that could both affect borrower sensitivity to rate markups and be correlated with our physical search cost proxy, we further construct an instrument that isolates exogenous changes to the local branch network using historical spatial variation in bank branching.

Using our proxies for distance-dependent search costs, we show that borrowers on the expensive side of FICO thresholds reject high-interest-rate loans more often when we measure search costs to be relatively low (i.e., when the number of actual or predicted nearby alternative lenders is high). By contrast, borrowers who we predict would 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 most often strongly dominated by nearby alternatives. Finally, we show that our results on price dispersion and search behavior are inconsistent with a simple market-concentration explanation, and we show that our results hold even in relatively low-concentration markets.[10]

Given that borrowers with higher costs of physically shopping for credit are more likely to accept expensive loans, we investigate the real effect of this friction on loan and car-purchase outcomes. On average, borrowers that accept quasi-randomly offered more-expensive credit take out \$550 smaller loans and purchase cars that are two months older, spending an average of \$375 less. The balance of borrower characteristics across FICO thresholds suggests that borrowers quasi-randomly drawing high loan markups would actually also take out larger loans and purchase a more expensive and newer car had they not been offered higher interest rates. To attribute these conditional-on-origination consumption outcomes to search frictions, we rule out selection in loan take-up being correlated with borrower-level demand shocks by further verifying the balance of borrower characteristics and outcomes conditional on origination. The exercise of rejecting significant differences in subsequent credit outcomes also rules out adverse selection as in Stiglitz and Weiss (1981) and Adams, Einav, and Levin (2009), which, if present in our RD setting, would be revealed over time in ex-post higher defaults, etc. Post-origination changes in credit scores and ex-post loan performance do not change differentially across discontinuities for either high- or low-search-cost borrowers. Taken together, the evidence suggests that costly search for credit represents an important market

[8]We observe potential borrowers without many nearby lenders applying for fewer car loans in our data and fewer mortgages in national Home Mortgage Disclosure Act data. Estimating search-cost distributions using the structural search model of Hortaçsu and Syverson (2004) also finds that search costs are higher when the density of nearby lenders is low.

[9]This approach is similar to the difference-in-discontinuities design of Grembi, Nannicini, and Troiano (2016).

[10]The number of sellers in a market directly affects equilibrium pricing in many imperfect competition models. However, our search-cost measure appears to capture variation in the cost of search instead of other variation in market structure given our finding that higher search-cost borrowers indeed search less despite facing, if anything, larger local price dispersion. See Appendix D.1 for further discussion.

friction that ultimately distorts financing and durable consumption.

After contextualizing our work in several related literatures, the remainder of the paper proceeds as follows. Section 2 sketches a conceptual framework to discuss our identification and how credit-market search frictions and elastic demand for both loans and durables could combine to affect consumption. Section 3 details the administrative data we use throughout the paper and documents significant price dispersion in the market for auto loans. Section 4 introduces our regression-discontinuity identification strategy to isolate exogenous variation in the benefits to search using lender pricing rules, and section 5 introduces our measure of search costs. In sections 6 and 7, respectively, we present evidence that consumers' propensity to search is correlated with measures of search costs, and we estimate the effects of costly search on loan and durable-purchase outcomes. Section 8 concludes.

1.1 Related Literature

In this section, we motivate our work in the context of literatures on search, credit frictions, auto loans, and FICO-based discontinuities. While many studies establish the existence of search frictions, our paper uniquely documents the connection between costly search for loans and subsequent outcomes in final-goods markets. Relative to the literature on various credit market imperfections that affect consumption, our specific contribution is to demonstrate and highlight the distortionary effects of search frictions in credit markets on consumption.

Theories of consumer search since Stigler (1961) suggest that under many formulations, when agents find it costly to solicit the full menu of offered prices, equilibrium prices will feature price dispersion. Applying this logic to a credit market, lenders can expect to originate loans at positive markups because a randomly arriving borrower will not exert effort sufficient to find better rates given the equilibrium distribution of interest rates. Baye, Morgan, and Scholten (2006) survey many empirical papers that document a particular product's price dispersion and connect it to evidence that its buyers find search to be costly. In consumer credit markets in particular, Woodward and Hall (2012), Alexandrov and Koulayev (2018), and Bhutta, Fuster, and Hizmo (2020) establish the role of low borrower search intensity in explaining mortgage interest rate dispersion. In a series of papers on the Canadian mortgage market, Allen, Clark, and Houde (2014a, 2014b, 2019) document price dispersion, demonstrate its response to market concentration, and quantify the lost consumer surplus from higher markups. Stango and Zinman (2015) use a self-reported measure of shopping intensity to explain variation in price dispersion in the credit-card market. Relative to this literature on price dispersion and search in consumer credit markets, our setting allows for measurement of the real effects on subsequent consumption quantities that can result from costly search for credit.

Recent work by Agarwal et al. (2020) shows that in the cross section, intensive loan search is counterintuitively correlated with *higher* interest rates. Agarwal et al. (2020) explain this with a model of borrower private information about the returns to search—low credit-worthy borrowers search until they find a lender who offers them an advantageous interest rate, albeit higher than the rates offered without search to (observably) high-quality borrowers. In our setting, the quasi-random assignment of our regression-discontinuity design allows us to abstract away from unobservable private information and contrast financing and consumption outcomes for borrowers with high and low search costs but similar benefits to search.

Traditionally, most search models feature inelastic and discrete demand, inhibiting the model's ability to speak to quantity or welfare effects. By contrast, our focus on costly search for credit highlights the importance of continuous and elastic demand in driving equilibrium outcomes.[11] Modern work on search and differentiated products emphasizes how buyer valuations could change because of the structure of search costs in various settings. For example, Zhou (2014) studies search in multi-product settings where consumers incurring a search cost for each shopping trip they make generates a complementarity between products sold at the same location. Moraga-González, Sándor, and Wildenbeest (2017, 2021) model how price elasticities and substitution patterns in a search model depend on whether consumers have full information about their valuation of differentiated products. Our emphasis on how costly search for credit affects the choice of goods purchased with that credit highlights a new search-cost source of demand elasticity distinct from other characterizations of search with elastic demand.

Our work also contributes to a literature that documents how features of various credit markets drive prices and quantities in other markets (e.g., Zeldes, 1989; Gross and Souleles, 2002; Melzer, 2011; Zaki 2016; Delavande and Zafar, 2019; Benetton, 2021; Aydin, 2022). See, too, work on the interplay between financing and collateral markets (e.g., Stroebel, 2016; Murfin and Pratt, 2019; Benetton, Mayordomo, and Paravisini, 2021), including work on how credit-market conditions drive house prices (Adelino, Schoar, and Severino, 2012; Favara and Imbs, 2015; Di Maggio and Kermani, 2017; Greenwald, 2018). Relative to studies of the connection between credit markets and related markets, we are the first to highlight how search frictions in credit markets can distort durables-goods consumption.

We also contribute to a growing literature studying the automobile-loan market and the frictions therein, including Attanasio, Goldberg, and Kyriazidou (2008), Adams, Einav, and Levin (2009), Busse and Silva-Risso (2010), Einav, Jenkins, and Levin (2012 and 2013), and Grunewald et al. (2020). Two other contemporaneous papers use a similar dataset to this paper to answer distinct questions on the importance of loan maturity and budgeting heuristics in the markets for cars and car loans. Argyle, Nadauld, Palmer, and Pratt (2021) (ANPP hereafter) use *vehicle*-level variation (instead of the borrower-level variation we use here) to document that payment-size shocks arising from allowable maturity changes affect bargaining outcomes and are capitalized into transaction prices. Notably, ANPP hold the purchased durable fixed and ask whether consumers pay higher prices for the *same* car when they have access to cheaper financing. By contrast, this paper examines how the cost of searching for better interest rates across lenders changes the borrower's decision of which car to ultimately purchase. By theoretically and empirically illustrating the distortions induced by credit-market search frictions, this paper demonstrates a new dimension of credit-market specialness and extends the vast search literature that neglects any welfare loss from search by

[11] The sequential search model of Reinganum (1979) generates equilibrium price dispersion under elastic demand. Appendix A extends this model to our setting to highlight the linkage between credit search costs and consumption outcomes.

assuming inelastic demand.

Argyle, Nadauld, and Palmer (2020) (ANP hereafter) use FICO-based pricing discontinuities to identify behavioral frictions associated with household debt decision making, showing that even financially unconstrained households bunch at round-number payment sizes and smooth monthly payments when facing payment shocks. While we also use FICO-based pricing discontinuities in this paper, here we focus on the consequences of search costs in credit markets, contrasting the effects of interest rate shocks across high- and low-search cost areas and documenting how interest-rate markups affect final-goods substitution patterns. In at least one respect, the identification strategies of both ANPP and ANP are both rationalized by the search frictions documented in this paper. The various nonlinear lender policies exploited for identification by ANPP and ANP and many other papers using lender-specific rules would have no effect in a perfectly competitive market where consumers are fully informed of all prices because consumers would simply reject any loan offer featuring an arbitrary markup in rate or monthly payment.

Finally, other work also exploits FICO-based discontinuities for identification. See Keys et al. (2009, 2010), Bubb and Kaufman (2014), Laufer and Paciorek (2016), Agarwal et al. (2017), and Aneja and Avenancio-León (2020). Building on this collection of papers that use FICO-based discontinuities as natural experiments or explicitly study their consequences, we are the first to identify credit-score-based discontinuities in loan *pricing* rules and to link those discontinuities to price dispersion, costly consumer search, and effects on final-good consumption.

2 Conceptual Framework

In this section, we develop intuition for how costly search for credit could affect real outcomes including loan take-up, loan sizes, and consumption levels.[12] Our focus is to describe how higher search costs for credit could lead to lower consumption of durables. See Appendix A for a theoretical model of costly search for credit with complementary goods as an example of how this intuition could be formalized.[13]

Let a borrower's demand for durables x and loans q be determined as the solution to

$$(x^*, q^*) = \arg\max u(x, \varepsilon)$$

subject to constraints

$$px \leq w + q \tag{1}$$

$$rq \leq \alpha y \tag{2}$$

where $u(\cdot, \cdot)$ is the utility of consuming x units of durable consumption given idiosyncratic demand shock ε. The first constraint is a budget constraint that the total expenditure on the durable given price per unit of durable consumption p not exceed available funds, given current wealth w and loan size q. The second constraint is an interest coverage constraint that requires that interest payments

[12] Although we study the specific case of costly search in a credit market and consider downstream effects on durable consumption, our results apply more broadly to the demand for any complementary products.

[13] While Appendix A uses a sequential search model to motivate our empirical analysis, the resulting hypotheses apply more broadly to other search formulations.

rq given interest rate r not exceed a maximum fraction α of income y, conceptually similar to so-called front-end debt-to-income ratios commonly imposed by underwriters.

The importance of r in demand is apparent from the constraints. When the budget and underwriting constraints bind, higher interest rates ceteris paribus decrease the maximum loan size q, which decreases the amount of durables x that can be consumed, decreasing utility. In our data, we observe borrowers' loan search decisions after being offered a given interest rate r'. If they choose to not take-up the loan described by r', we assume this means they have chosen to search elsewhere with potential recall of a lower rate. The final interest rate r that governs the choice of x can be written as

$$r = \begin{cases} r' & \text{if accept} \\ \min(r', r^*) & \text{if search} \end{cases}$$

where r^* is another interest rate draw from the distribution of interest rates F_r in the credit market for a particular borrower type.[14] Our empirics below confirm that demand for loans and cars are downward sloping in the final interest rate r.[15]

Potential borrowers will search when their expected net benefit of searching further for credit is positive. The benefit of more search is the increase in indirect utility consumers receive from the interest-rate reduction they expect to realize. The process of searching, in turn, entails incurring search cost k, measured in utility terms. Summarizing our discussion in section 5, key components of the cost of searching for credit include time and hassle costs such as researching potential lenders and submitting signed loan applications. Other frictions captured by the concept of search costs include concerns about perceived credit-score impacts of loan applications (Liberman, Paravisini, and Pathania, 2021) and the need to find a loan quickly to finalize the purchase of a car.

Intuitively, given that the consequence of a higher price paid for one good also includes the utility loss from lower consumption of complementary goods, the search cost needed to rationalize equilibrium markups is larger than would be suggested by the size of accepted markups alone. Many search models assume discrete demand where consumers choose whether to buy a single unit of a final good, allowing researchers to characterize search benefits in price terms (e.g., Hong and Shum, 2006). However, when demand is elastic, the utility loss associated with a given markup will be different from the size of the markup itself. Under continuous and elastic demand, the return to search will include the change in utility from elastic borrowers increasing both their loan amounts and good expenditure when finding a lower interest rate offer. In our specific setting, inferring financing search costs simply from the distribution of markups would fail to account for the fact that the borrower would be paying both the interest rate markup and the disutility of a cheaper and older car.

Combining search costs and expected benefits, borrowers search if the expected utility gain from

[14] Borrowers could subsequently repeat the search process, updating r' and drawing again.

[15] We treat durable consumption as a continuous variable x measuring quality that in our setting can be conceptualized as "car services," with older cars or cars with cheaper predicted ("blue-book") values delivering lower car services, similar to how the housing literature models housing consumption continuously as housing services. These services can be considered to include the dividend flow utility of the durable and its resale value.

another price draw given the current interest rate quote r' exceeds the cost of obtaining another quote. Formally, this condition is

$$\int_{\underline{r}}^{r'} V(r,p)dF(r) - V(r',p) > k \tag{3}$$

where $V(r,p)$ is the indirect utility of a borrower facing interest rate r and price per unit of durables p, $F(r)$ is the CDF of interest rates, $\underline{r}$ is the floor of the interest-rate distribution, and k is a borrower's search cost in utility terms. In our setting, we find that car loans and car services are strong complements, each with finite demand elasticities such that $V(\cdot,\cdot)$ is decreasing in both of its arguments and the cross-price elasticity of demand between loans and cars is negative.

Fixing an outside option interest rate r', more search is likely to be optimal when search costs are lower. Comparing two otherwise identical borrowers with different search costs, the borrower with lower k will be more likely to satisfy condition (3) and therefore more likely to search more and find a lower interest rate $r < r'$. An important implication of costly search for credit then is that borrowers facing high search costs will be more likely to accept higher interest rates. It follows given downward sloping demand and the complementarity between credit and durables that weakly higher accepted interest rates from higher search costs are predicted to lead to smaller loan sizes, lower purchased car services, and less car expenditure.

Identification Taking this prediction to the data to estimate how much loan search costs decrease demand for credit and durable consumption in equilibrium is complicated by the potential correlation between search costs k, interest rates r', and demand shifters ε.[16] Consider a linear probability model of extensive-margin loan demand $takeup_{it}$ of a given loan offer at interest rate r_i for borrower i at time t with search costs k_{it} as

$$takeup_{it} = \beta_0 + \beta_1 r_i + \beta_2 k_{it} + \beta_3 r_i \times k_{it} + \varepsilon_{it}. \tag{4}$$

Given downward sloping demand, we would expect β_1 to be negative. The logic offered above implies that we would further expect β_3 to be positive—but with $\beta_1 + \beta_3 k_{it} < 0$—as borrowers in markets with higher k_{it} would be more likely to accept a given interest rate r because of the costliness of further search. Given the role of the interest rate in constraints (1) and (2), we would then predict that ultimate loan sizes, car services purchased, and car expenditure would be lower for borrowers accepting high r. However, omitted variables could bias β_3 upwards.

The ideal experiment would randomly assign search costs and interest rates for identification to break the correlation between k, r, and ε. Without such exogenous variation, if we observe that accepted interest rates are higher in high-search-cost markets, this could be due to other unobserved demand factors ε potentially correlated with search-cost proxies, such as credit limits, financial literacy, or preferences.[17] To address this challenge, we use a regression discontinuity (RD)

[16] Once we introduce variation in k, (3) implicitly assumes that $F(\cdot)$ is invariant to k. In practice, we will show that $F(\cdot)$ is more dispersed in lower search cost areas, if anything incentivizing search there and biasing our estimates of the effect of search costs downward.

[17] For example, Guiso et al. (2021) demonstrate how financially unsophisticated households may be steered towards more expensive mortgages.

design to isolate quasi-exogenous variation in r that randomly assigns the rate quote borrowers draw to be high or low, as we detail in section 4.

We then combine the RD design with geographic variation in the cost of shopping for credit. As we detail in section 6, our baseline proxy for a measurable component of search costs is based on the number of bank or credit-union branches within a 20-minute drive of each borrower's home. We construct a high-search-cost indicator variable

$$HighSearchCost_{it} = 1(\#\text{ nearby branches}_{it} \leq 10) \tag{5}$$

where the number of nearby branches is determined for each borrower's home address in each year.[18] However, as we discussed, local branch density is not random and is potentially correlated with unobservables that also drive take-up. For example, by this definition, high-search-cost borrowers could be more likely to live in rural areas and could have different financial sophistication, income, and expected future income. Controls for demographic and financial borrower observables can only absorb some of this unobserved heterogeneity. To isolate exogenous variation in $HighSearchCost_{it}$, we develop an instrument that predicts high search costs $High\widehat{Search}Cost_{it}$ using pre-period differences in branch density interacted with national changes to bank branches, as we detail in section 5. The identifying assumption behind this instrument is that conditional on borrower controls and local geographic fixed effects, the local cross-sectional variation interacted with the national time-series variation in the predicted search cost measure is unrelated to time-varying local demand shocks.

Because borrowers just above and below our RD thresholds are similar, the RD step allows us to difference out any differences across borrowers that are correlated with interest rates.[19] Comparing the effect of similarly sized interest-rate markups in markets predicted to have high versus low search costs allows us to focus on how search costs differentially affect the loan demand and consumption decisions of borrowers with high search costs who are more likely to accept higher interest rates. If search costs are the factor driving differential responses to interest-rate discontinuities, then we would expect that being assigned a high r' versus a low r' should be more consequential in high versus low-search-cost markets. This strategy allows us to causally attribute any difference in the difference between high- and low-markup borrowers across low- and high-search-cost areas to search costs.

Summarizing, because credit demand is continuous and elastic, an increase in search costs could lead to higher interest-rate markups being acceptable to borrowers, in turn resulting in smaller loan sizes and less complementary durable consumption being purchased. In low-search-cost markets, we expect ex-ante identical consumers to have relatively similar ex-post credit-market and product-market outcomes. In high-search-cost areas, borrowers will face higher markups, but their higher search costs will discourage them from searching further. Given downward sloping demand for loans and cars, these search-cost-supported higher markups should lead consumers to have lower loan sizes

[18] We use a binary variable specification to capture the nonlinear effect of branch density on search costs. In section 6.2, we vary the cutoff to document this nonlinearity and show robustness to the choice of 10.

[19] Our regression-discontinuity tests in the appendix support the identifying exclusion restriction that borrowers with different initial interest-rate quotes r' on either side of rate discontinuities are similar across a wide set of observables.

and purchase prices.

3 Data

We analyze the loan contract terms and used-car purchasing decisions of 2.4 million individual borrowers in the United States from 326 retail lending institutions between 2005 and 2016. The loan data are provided by a technology firm that provides administrative data warehousing and analytics services to retail-oriented lending institutions nationwide. The majority of the loans in our data (99%) were originated by credit unions ranging between $100 million and $4 billion in asset size, with the remainder originated by non-bank finance companies.[20]

Unlike most studies of secured credit, our dataset contains information capturing all three stages of a loan's life: application, origination, and ex-post performance, although our loan application data is for a subset of our origination data, consisting of approximately 1.3 million loans from 41 different institutions (owing to the smaller number of lenders that share applications data with our provider). The loan application data report borrower characteristics (age, gender, imputed minority status, 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, loan performance, and collateral characteristics. Using Vehicle Identification Numbers (VINs), we observe the make, model and model year of the purchased car. We restrict our sample to direct loans (those not intermediated by a dealer) to address concerns that indirect loans are potentially endogenously steered by sellers to specific financial institutions (perhaps because car dealers become aware of lenders' pricing rules over time).[21] Finally, to measure ex-post loan performance, we observe a snapshot of the number of days each borrower is delinquent, whether individual loans have been charged off, and updated borrower credit scores as of the date of our data extract.

Panels I, II, and III of Table 1 present summary statistics on loan applications, loan originations, and measures of ex-post performance, respectively. As reported in Panel I of Table 1, the median loan application in our data seeks approval for a six-year $20,000 loan at a median interest rate of 4.0%.[22] Borrowers applying for loans in our data have an average credit score of 648 and an average DTI ratio of 26.0%. The percentage of loans approved is 50.2%, with 65% of the approved borrowers subsequently originating a loan (what we refer to as the take-up rate). Panel II of Table 1 reports summary statistics on loan originations. Compared with loan applications, originated loans have smaller average sizes, similar interest rates, shorter terms, and are from more creditworthy and less constrained borrowers. Average monthly payments for originated loans are $324 per month with an

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

[21]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. While indirect loans are more common overall, among used cars financed through credit unions, direct loans are more common. Note that private transactions, a large share of the used-car market, are necessarily financed by direct loans. See Appendix Table A1 for summary statistics for the indirect loans excluded from our estimation sample.

[22]Interest rates in the application data refer to approved loans, regardless of whether they were subsequently originated.

interquartile range of $195.

Panel III tabulates measures of ex-post loan performance. Most loans are current; the 75th percentile of days delinquent is zero, and only 2.1% of loans have been charged off (i.e., accounted as unrecoverable by the lender). Defining default as a loan that is at least 90 days delinquent, default rates average 2.2%.[23] Lending institutions periodically check the credit score of their borrowers subsequent to loan origination, allowing us to monitor borrowers' financial performance across their liabilities and assess the extent of adverse selection at the origination stage. Summary statistics for ΔFICO represent percent changes in borrowers' FICO scores from the time of origination to the lender's most recent (soft) pull of their FICO score.[24] Updated FICO scores indicate that borrowers on average experienced a 1% reduction in FICO score since origination, although borrowers with FICO scores below 600 on average realized a 5.7% increase in FICO score.

Appendix B uses a nonparametric matching exercise combined with our data on interest rates and borrower and loan characteristics to document that significant interest-rate dispersion persists in the market for car loans. We find that 54% of borrowers do not attain the best rate available to nearly identical borrowers in our data at the same time in the same market. Including the borrowers who originate the best available rate given their borrower type, the average borrower in our data pays a 1.3 percentage point higher interest rate than an observationally equivalent borrower. A simulation exercise using the empirical distribution of markups estimates that a borrower would on average need to obtain three price quotes to find an offer within 10 bp of the best available rate for that borrower's type.

Data Representativeness The bulk of our auto loan data come from credit unions and seem broadly representative of that lending segment. Experian data from 2015 indicates that credit unions originated 22% of all used-car loan originations and 10% of new car originations in the United States.[25] In the auto-loan data made available to our data provider by its clients, roughly half of our clients' loans are direct loans. Data from the New York Federal Reserve Consumer Credit Panel (CCP) suggests that auto loans originated by credit unions and banks have substantially lower default rates compared to loans originated by auto finance companies. We discuss issues associated with online lending in Appendix D.2.

Credit-union borrowers differ somewhat from the average U.S. adult. Our sample contains borrowers who are slightly older, less racially diverse, and of a higher average credit quality than national averages. Over 41% of borrowers in our sample were between 45 and 65 years old at origination whereas 34% of the adult U.S. population is between the ages of 45–65. Roughly 73% of our sample is estimated to be white, compared to an estimated 65% of adults in the 2015 American

[23]We find that the default rate for sub-600 FICO borrowers is 6.8%, compared to a default rate of 2.6% for borrowers with FICOs between 600 and 700 and 1.6% for over-700 FICO borrowers.

[24]The time between FICO queries varies by institution, but the amount of time between the original FICO recording and the current FICO is roughly equal to the loan's age.

[25]While most new-car buyers finance their purchase through the car dealer, used-car purchases (including from dealers, private transactions, and independent used-car sellers) are more frequently financed elsewhere.

Community Survey.[26] Borrowers in our data report median FICO scores at origination of 711 over the full 2005-2016 sample period. The CCP reports median FICO scores for originated auto loans of 695 during the period our sample was collected. The representativeness of our sample should not limit our ability to make causal inference given that we rely on an RD design that leans crucially on an assumption of smoothness in borrower demographics across discontinuities at a given institution. However, it is possible that the search behavior of the segment we study differs from other segments of the population. That said, while borrowers in our data may have different search costs than non-credit-union borrowers, our data still represents a very substantial set of auto-loan borrowers.

4 Isolating Variation in the Benefits to Search

In this section, we introduce an empirical strategy designed to identify exogenous variation in the benefits to search, which we will use below to estimate the impact of costly search on equilibrium outcomes in both credit and durables markets. As noted, individual-level heterogeneity in transacted prices could be correlated with unobserved heterogeneity in search costs or taste shocks that could plausibly be correlated with other outcomes or durable-good product characteristics. To address this, we exploit quasi-exogenous within-lender markup variation in our data that serves as a laboratory where the potential gains to search are quasi-randomly assigned across borrowers.

Our regression-discontinuity (RD) design assigns otherwise nearly identical borrowers to high or low offered interest rates. Facing a high initial quote r' should be relatively inconsequential for borrowers in markets with low search costs. In low-search-cost markets, borrowers with high initial quotes r' should have similar loan and durable-consumption outcomes as borrowers with low initial quotes because borrowers drawing high initial quotes should be willing to search further.[27] As such, rate discontinuities should have stronger effects on search behavior and borrowing and consumption outcomes in high versus low-search-cost markets.

4.1 Detecting Discontinuities

Lenders make underwriting and loan pricing decisions based on observable ("hard") and unobservable ("soft") information (Liberti and Petersen, 2018). An econometrician's ability to make inference from lending outcomes is complicated by the possibility that unobservables play a role in jointly determining selection into application and origination, observed loan terms, and subsequent loan performance. We address this possibility, and other potential omitted variables, with an RD design leveraging 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 the origination and securitization stages (Bubb and Kaufman, 2014), we focus on discontinuities in loan *pricing*, i.e., the interest rate offered to a borrower conditional on having a loan application approved by underwriting. Moreover, unlike in mortgage

[26]Borrowers do not report race at the time of loan origination, but most lenders in our sample estimate minority status to document compliance with fair lending standards.

[27]Similar logic applies in a fixed-sample-size search model; in low-search-cost markets, borrowers should be less affected by one high draw r' because they should have a larger set of quotes to choose from.

lending, no industry standard set of thresholds exist in auto lending, and credit unions were prohibited from securitizing auto loans until 2017 (after our sample ends). This heterogeneity in FICO pricing discontinuities across lenders likely makes each given lender's FICO bin locations unknown to most potential borrowers. While auto-loan lending institutions do not adhere to a common set of FICO cutoffs, the use of discontinuous pricing at some point across the FICO spectrum is prevalent for more than half of the lenders in our data.[28] See Bubb and Kaufman (2014), Matějka and McKay (2015), Livshits, MacGee, and Tertilt (2016), and Agarwal et al. (2017) for models of credit risk processing and Al-Najjar and Pai (2014) for a model of overfitting that each rationalize binning risk types in pricing decisions. 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 and could persist in part because costly consumer search prevents more accurately risk-based pricers from gaining market share.[29]

To detect FICO-based discontinuities for each lender, we estimate lender-specific interest-rate policies nonparametrically. For each lender l in our data, we characterize its lending policies across FICO bins with a set of parameters $\{\psi_{kl}\}$ where k indexes FICO bins denoted $\mathcal{F}_k$. Separately for each lender l, we estimate ψ by regressing individual interest rates r_{il} on a set of indicator variables for each 5-point FICO bin $\mathcal{F}_k$

$$r_{il} = \sum_k \psi_{kl} 1(FICO_i \in \mathcal{F}_k) + v_{il}. \tag{6}$$

The five-point FICO bins begin at a FICO score of 500 where the first bin includes FICO scores in the 500-504 range, 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 at a given lender relative to the estimated constant (the omitted category is loans outside this range—we focus on relative magnitudes for this exercise).[30]

Appendix Figure A1 presents interest-rate plots for two different financial institutions. The point estimates $\hat{\psi}$ represent how each lender's pricing rules vary with borrower FICO score, and the accompanying 95% confidence intervals provide a sense of the precision with which we estimate ψ, owing both to how reliant on FICO scores each lender's pricing rule is and how many loans are in each bin. Panel I of Appendix Figure A1, estimated on one institution in our data with approximately 12,000 borrowers, illustrates breaks in average interest rates for borrowers with FICO scores around cutoffs at 600, 660, and 700. The breaks in interest rates at the FICO cutoffs are large, representing jumps of over two 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

[28]In principle, variation across lenders in the use of pricing discontinuities could introduce selection on unobservables into our estimation sample that uses only lenders with discontinuities. In terms of external validity, the similarity between our estimation sample and our overall sample mitigates this concern. For internal validity, our RD design allows for lender × discontinuity fixed effects to absorb any such fixed differences in borrower unobservables across lenders.

[29]See Hutto and 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.

[30]Our results are unchanged if we estimate ψ conditional on other borrower characteristics that affect loan pricing such as DTI.

bin, and the difference in average interest rates between the two bins is statistically significant at the 0.001 level. Panel II of Appendix Figure A1 illustrates similar rule-of-thumb FICO breaks for a lender with approximately 6,000 borrowers. Note that the breaks occur at different FICO scores across different institutions, consistent with our understanding that the discontinuities are reflective of institution-level idiosyncratic pricing policies.

In order to restrict our analysis to only institutions that employ discontinuous pricing rules, we empirically identify the existence and location of discontinuities at each institution in our sample through the following criteria. We first estimate the interest-rate FICO bin regressions following equation (6) for each institution in our sample separately. To establish the existence of an economically and statistically significant interest-rate discontinuity, we require interest rate differences across consecutive bins to be larger than 50 basis points and to be estimated with p-values that are less than 0.001. We refine the set of discontinuities by requiring smoothness (p-value of at least 0.1) for the difference in the pair of coefficients leading and the pair of coefficients following a potential discontinuity.[31] Finally, we examine each potential threshold visually to ensure that the identified discontinuities are well behaved around the candidate thresholds. Our estimation sample, termed the discontinuity sample by Angrist and Lavy (1999), consists of 514,834 loans within 19 points of one of our lender-specific detected thresholds—see Appendix Table A2 for summary statistics.[32] In our empirical specifications, we normalize FICO scores to create a running variable $\widetilde{FICO}$ that measures distance from a given interest-rate discontinuity. For example, for loans near the 600 FICO score threshold, $\widetilde{FICO}_i = FICO_i - 600$.

4.2 The Effect of Pricing Discontinuities on Interest Rates

To validate our RD design, we present a series of diagnostics designed to test whether our data meet the two main identifying assumptions required of valid RD estimation. Our objective is to establish that borrowers with FICO scores just above and below discontinuities in a lender's pricing function are quasi-randomly assigned different interest-rate markups. By contrasting the borrowing and consumption outcomes of such consumers in high- and low-search-cost environments, we can assess the distortions caused by search frictions.

First, the RD design assumes that the probability of borrower treatment (i.e., offered interest rates) with respect to loan terms is discontinuous at detected FICO thresholds. Second, valid RD requires that borrower attributes (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. We provide evidence that the smoothness condition is satisfied in Appendix D.3 and Appendix Figure A2. Across a large set of observables, we do not see any statistically or visually significant changes at a FICO discontinuity, which we take to be strong

[31] See Appendix C of Agarwal et al. (2017) for a discussion of estimation with overlapping cutoffs.

[32] Our results are robust to detecting the discontinuities using a hold-out sample and using the remaining sample for estimation.

evidence that unobservables are also not changing across FICO discontinuities.

Panel I of Figure 1 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 700 FICO thresholds in panel II of Figure 1. These plots confirm the existence of large interest-rate discontinuities at these thresholds. The narrow confidence intervals in Appendix Figure A1 and Figure 1 also indicate that interest rates in this market seem to be strongly determined by FICO. If there were substantial residual variation after controlling for FICO scores nonparametrically, the confidence intervals would be much larger.

To estimate the average magnitude of the interest-rate discontinuity across all detected thresholds, we estimate RD regressions. For intuition, we first introduce the RD estimating equation in the context of a single threshold and with linear controls for the running variable as

$$r_{iglt} = \pi_1 \widetilde{FICO}_i + \delta \cdot 1(\widetilde{FICO}_i \geq 0) + \pi_2 \widetilde{FICO}_i \cdot 1(\widetilde{FICO}_i \geq 0) + \alpha_{gt} + \gamma_l + v_{iglt} \tag{7}$$

where r_{iglt} is the interest rate of loan i originating in Commuting Zone g from lending institution l in quarter t, $1(\widetilde{FICO}_i \geq 0)$ is in indicator variable equal to one if the normalized FICO score $\widetilde{FICO}_i$ is above the threshold, and α_{gt} and γ_l are Commuting Zone (or zip code) $\times$ quarter and lender fixed effects, respectively. In this specification, δ is the key RD coefficient and estimates how interest rates r change discontinuously at a policy threshold while allowing the running variable $\widetilde{FICO}$ gradient to also change at the threshold.

In practice, there are two differences between (7) and our actual estimating equation. First, we allow for the effect of the running variable $\widetilde{FICO}$ above and below the cutoff at $\widetilde{FICO} = 0$ to be quadratic. Second, to deal with loans that may be within 19 FICO points of multiple discontinuities as in Agarwal et al. (2017), we sum across discontinuities d from the set of discontinuities $\mathcal{D}$ to estimate

$$r_{iglt} = \sum_{d \in \mathcal{D}} 1(il \in \mathcal{D}_d) \left(\delta \cdot 1(\widetilde{FICO}_{id} \geq 0) + f(\widetilde{FICO}_{id}; \pi) + \varphi_{dl} \right) + \alpha_{gt} + v_{iglt} \tag{8}$$

where $1(il \in \mathcal{D}_d)$ is an indicator for whether loan i is within a bandwidth of 19 FICO points of a discontinuity at lender l, φ_{dl} are discontinuity $\times$ lender fixed effects to allow for each lender to have a different selection of borrowers around each threshold, and the function $f(\cdot;\cdot)$ is defined as

$$f(x;\pi) = \pi_1 x + \pi_2 x^2 + 1(x \geq 0) \left(\pi_3 x + \pi_4 x^2 \right) \tag{9}$$

to allow for a smooth but nonlinear effect of the running variable that potentially changes shape discontinuously at the threshold.[33] Standard errors are double clustered by lender and FICO score, and the sample used to estimate (8) is the discontinuity sample described in Appendix Table A2.[34]

Table 2 presents results of this exercise for varying levels of fixed effects. Using the stringent zip

[33] The specification in (8) also allows us to accommodate loans on the left of one threshold and the right of another, similar to Agarwal et al. (2017).

[34] 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.

code × quarter and discontinuity × lender fixed effects in column 4, interest rates for borrowers with FICO scores immediately above a detected threshold are an average of 1.37 percentage points lower than borrowers just below. Given an average interest rate in our estimation sample of 6.0% (panel II of Appendix Table A2), the magnitude of these effects is economically meaningful, amounting to a $488 higher present value and a $9 higher monthly payment for otherwise identical loans taken out by borrowers on the expensive side of a FICO discontinuity. Conveniently, this magnitude is comparable with the average spread between an originated interest rate and the best available rate for a similar borrower in our data (see Appendix B). In the context of our conceptual framework, drawing a rate quote from the expensive side of an interest-rate discontinuity constitutes a significantly higher initial interest rate r', and, depending on whether searching for financing is costly, could have material consequences on the ultimate cost of credit for such borrowers. As we discuss in section 6 after introducing our measure of financing search costs, column 5 shows that the average FICO discontinuity has similar markup effects in high- and low-search-cost areas.

4.3 Measuring Potential Gains to Search Across Lenders

By quasi-randomizing initial interest-rate quotes r' to each borrower, these discontinuities represent exogenous variation in markups *within* lender. However, if search *across* lenders were costless (or demand inelastic), a high r' quote would have no effect on transaction outcomes. We next document empirically that borrowers who find themselves on the expensive side of a pricing threshold at one lender could reasonably expect to find a lower interest rate (all else equal) were they to search across lenders.[35] Figure 2 provides visual evidence of differentially higher returns to search across lenders for left-of-threshold borrowers by plotting the density of the spread to the lowest available rate for left- and right-of-threshold borrowers using the matching strategy of Appendix B. Dotted and solid lines in each plot are for borrowers just below and above a given threshold, respectively. For both thresholds, average spreads to the lowest available rate and the variance of those spreads are larger for left-of-threshold borrowers, implying that price dispersion (and thus the returns to search) are higher for those offered exogenously higher rates. For borrowers with FICO scores between 595 and 599, for example, there was on average a loan with a 3.8 pp lower interest rate originated to someone with the same FICO and DTI in the same CZ at the same time and used to secure a similarly priced car (see Appendix Table A3 for details). Taken together, these plots confirm that price dispersion is largest for borrowers exogenously offered higher interest rates and that such borrowers are more likely to draw a much lower interest rate from an additional search.[36]

5 Isolating Variation in the Costs of Search

The existence of high and dispersed markups in equilibrium—especially given their detection using quasi-exogenous pricing discontinuities in section 4—is prima facie evidence that at least some

[35]Note that this would not be the case if every institution shared the same FICO cutoffs.

[36]While markups would also distort quantities in an oligopolistic market, too, simple market concentration explanations for our results fail to predict the price dispersion we observe. For completeness, we show in Appendix D.1 that the contrast between high- and low-search-cost borrowers holds in both concentrated and less concentrated markets.

borrowers find soliciting the full menu of interest rates costly. Otherwise, we would expect borrowers to search until finding the lowest available rate and that the resulting competition would drive dispersion to zero. Moreover, the results of Appendix D.3 support the initial assignment of markups to borrowers on either side of a FICO discontinuity being as good as random. However, attributing variation in subsequent conditional-on-origination outcomes to search frictions raises selection concerns. If borrowers who accept high markups are different on other dimensions from borrowers accepting lower markups then simple comparisons of outcomes conditional on markup size are biased, even when markups have been randomly assigned. In this section, we introduce a geography-based measure of search costs, along with exogenous variation therein. In the next section, we use our RD apparatus to demonstrate that borrowers with high search-cost proxies are more likely to accept higher markups, consistent with our intuition in section 2. Importantly, our RD strategy and market-by-time and lender-by-discontinuity fixed effects absorb any fixed differences in borrowers across high- and low-search-cost areas. The remaining threat to identification is only that unobserved heterogeneity correlated with our search-cost proxy also changes how borrowers respond to FICO discontinuities, motivating our development of quasi-exogenous variation in search costs to isolate the effect of costly search on loan take-up.

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 report doing "almost no searching" when taking out a new loan. Potential borrowers face a variety of non-monetary costs when shopping for a car loan. While many car buyers—perhaps precisely because of financing search costs—choose to finance their purchase through a lender vertically integrated with a dealer, used-car buyers frequently finance their purchase from a separate source.[37] Such borrowers may purchase used cars from a seller that does not have a financing arm (e.g., private transactions), seek loan preapproval before negotiating with the seller over purchase price to refine their own budget, or seek to avoid double marginalization (Busse and Silva-Risso, 2010; Grunewald et al., 2020).

As car-loan pricing is specific to the credit risk of each individual, obtaining rate quotes in the direct market most often entails filling out a loan application, undergoing a credit check, and potentially verifying assets and income. For measurability and the potential to isolate shocks thereto, we focus on the dimension of search costs that scales with time and distance, such as the time and hassle required to travel to a branch and physically sign financial paperwork or the cost of ascertaining the choice set of potential lenders.[38] However, we note that there are many other dimensions that we do not measure over which search is costly, for example the disutility of filling out financial paperwork, the effort required to become informed about price dispersion, and potential concerns that additional credit-registry queries negatively impact credit scores (Liberman, Paravisini, and Pathania, 2021).[39] Given the many contributors to the reduced-form concept of

[37]Recall that our sample consists of direct auto loans originated through a lending institution as opposed to indirect auto loans where dealers broker an electronic search across multiple lenders at the time of the auto purchase and mark up the resulting loan offer (Grunewald et al., 2020).

[38]We discuss the option borrowers have to search for loans online in Appendix D.2.

[39]Note that we do not consider several other plausible proxies for search costs in our data because of their likely

search costs, we view our results as providing a lower bound on the role of search and information frictions in affecting consumer borrowing and consumption.

To proxy for distance-based search costs, we use FDIC and NCUA data to identify the precise 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 the number of nearby financial institutions by calculating the driving-time lender density for each borrower. To do so, we geocode and count the number of physical branch locations within a 20-minute drive of the borrower's address.[40] This density measure is designed to capture the effort (proxied by time and distance) for each borrower to shop for an additional interest-rate quote from a lending institution that is within a reasonable distance from their home.[41] Supporting this search-cost proxy, Degryse and Ongena (2005) find evidence of the important role of transportation costs in local credit markets.[42]

Borrowers in the 25th percentile of driving distance live less than a 20-minute drive from 23 lending institutions compared to 168 institutions for borrowers in the 75th percentile. Our baseline results categorize borrowers as having high search costs if their home address is within a 20-minute drive of at most 10 lending institutions, although our results are robust to alternative definitions (see section 6.2). This definition classifies roughly 15% of borrowers as living in high search-cost areas and is designed to capture the diminishing effect of an additional nearby lender on search costs. Figure 3 geocodes the borrowers in our our data in the Jacksonville, Florida metropolitan area to illustrate the spatial distribution of our search-cost classification. Relative to the blue dots for high-search-cost borrower locations, the red dots for borrowers with more than ten nearby lenders are concentrated closer to downtown Jacksonville.

To provide a more direct test of whether search behavior varies with our search cost measure, Appendix C shows that borrowers in low-search-cost areas indeed submit a higher number of loan applications to lenders covered by our data than borrowers in high-search-cost areas. Panel I of Appendix Table A8 uses HMDA data on the near-universe of mortgage applications to show that borrowers in tracts in which our search-cost measure is on average high submit 0.18 fewer applications than borrowers in tracts where our average search-cost measure is low. Panel II of Appendix Table A8 uses our car-loan application data to show a similar result; applicants in high-search-cost areas submit 0.07 fewer applications, on average, than applicants in low-search-cost

nonmonotonic mapping to search costliness. For example, borrowers with high FICO scores or older borrowers may have both better financial literacy and a higher opportunity cost of time.

[40]Our driving-time calculations rely on posted speed limits along current driving routes and do not incorporate traffic conditions or changes to the road network between the time of loan origination and 2016 (the date of our driving-time data). For each borrower, we use only those institutions that existed at the time of that borrower's loan origination.

[41]While distance can also proxy for soft-information producing relationships (see Nguyen, 2019; Granja, Leuz, and Rajan, 2021), auto loans are not a particularly relationship-intensive credit product. Consistent with this, we find a lack of adverse selection around discontinuities and a high R^2 in our interest-rate regressions based on lender pricing rules. While brand loyalty effects could still affect take-up (Allen, Clarke, and Houde, 2019), our RD design allows within-lender analysis. Appendix D.1 also presents findings inconsistent with the number of nearby lenders directly affecting outcomes through market concentration.

[42]See also Moraga-González et al. (2021), who use the density of nearby automobile dealers to proxy for search costs in the car-buying market.

areas. The lower coefficient is unsurprising, given that we do not observe applications to institutions outside of our data. Overall, Appendix Table A8 supports our interpretation of the number of nearby lenders affecting the costliness of shopping for credit.

5.1 Exogenous Variation in Search Costs

People who live within 20 minutes of more or less than 10 lenders may be different from each other on other dimensions besides search costs. Our RD setup mostly accounts for even unobservable differences across high- and low-search-cost areas by comparing borrowers within lender, commuting zone, and time who are randomly assigned to above or below pricing discontinuities. Applicants in high- and low-search-cost areas also look similar on observables. For example, as we discuss below, two measures of creditworthiness tell mixed but muted stories about applicant differences across our search cost measure. On average, applicants in high-search-cost areas have slightly lower FICO scores (0.3 points) but also lower DTI ratios (2.6 points). To address any unobserved heterogeneity correlated with both our search-cost proxy and borrowers' response to rate discontinuities, we employ a measure of predicted search costs that isolates exogenous variation in the number of nearby lenders. In Appendix D.1, we further address the possibility that the number of local lenders may directly affect the level of competition in ways unrelated to consumer search.

To consider how unobserved heterogeneity could affect our contrasting of effects across high- and low-search-cost areas, consider the following example. If a given location has few nearby lenders because lenders anticipate local economic conditions to be poor, then our proxy for high search costs may be correlated with borrowers in such areas being constrained and more likely to accept high-markup loan offers. Such local endogeneity of nearby lender density could generate the results of section 6 even if the number of nearby lenders had remained high.

We address possible bias from such time-varying omitted variables through an instrument $High\widehat{Search}Cost_{it}$ that predicts the high-search-cost status of a given borrower's location using only historical spatial variation in lender density and national trends in bank branches. Specifically, we calculate the number of nearby bank and credit union branches for each borrower in our sample using the density of nearby financial institutions as of 1990 using NETS historical data on the location of every financial institution in the U.S. We then grow the 1990 lender density measure using the national growth rate in financial institutions from 1990 through the year of each observation in our sample period. Mathematically, the predicted high-search-cost indicator is defined as

$$High\widehat{Search}Cost_{it} = 1(\#\ \text{near}\widehat{\text{by bra}}\text{nches}_{it} \leq 10). \tag{10}$$

The predicted number of nearby branches for each borrower i in each year t is calculated as the product of the number of branches within a 20-minute drive of borrower i's home location in 1990 and the ratio of the current total number of bank and credit union branches nationwide to the total number of such branches nationwide in 1990:

$$\#\ \text{near}\widehat{\text{by bra}}\text{nches}_{it} = \#\ \text{nearby branches}_{i,1990} \times \frac{\text{nationwide branches}_t}{\text{nationwide branches}_{1990}}. \tag{11}$$

Variation in this predicted high-search-cost measure is driven by local branching concentration in

1990 and aggregate variation in national branching trends. Appendix Table A4 demonstrates that $High\widehat{Search}Cost_{it}$ is a strong predictor of $HighSearchCost_{it}$, with partial F-statistics ranging from 12 to 83, depending on the controls.

Importantly, we expect neither object to be correlated with time-series variation in local demand shocks during our sample period of 2005–2016. If the local composition of borrowers affects their response to interest-rate discontinuities in a way that is correlated with the nearby branch density, replacing our search cost proxy with $High\widehat{Search}Cost$ will test whether the effect of search costs that we measure is directly driven by search or spuriously driven by unobserved geographic heterogeneity. Table 3 tests whether borrower characteristics in our origination data (panel I) or Census tract-level characteristics from the American Community Survey (panel II) are predicted by our High Search Cost measure (column 1) or our predicted High Search Cost measure (column 2). Each reported coefficient is the estimated effect from a separate regression of a binary search cost measure on the dependent variable in a given row, conditional on zip code-by-quarter fixed effects. This tests the identifying assumption supporting estimating (4) that any observed effect of search costs on outcomes is attributable to search costs and not other correlates of our proxy measure of search costs.

For example, since our search costs proxy depends on the number of financial institutions near a borrower's home address, if this borrower attribute is predictive of credit constraints, then it's possible that our estimated effects of search costs are confounded with the effect of credit constraints. Instead, Table 3 shows that our search cost measure is generally unrelated to credit characteristics. The statistically and economically insignificant credit score coefficient of -0.267 in column 1 indicates that borrowers in high search cost areas had FICO scores that were 0.3 points lower than borrowers in low search cost areas. Both search costs and predicted search costs are similarly economically and statistically insignificant predictors of other credit attributes in panel I.

The results of panel II highlight some of the benefits of examining robustness of our take-up results to using our predicted search cost measure. Three census tract characteristics are correlated with our search cost measure in column 1. Tracts with higher average search costs have \$35 lower rent, 11.5 log points higher wage growth, and 0.9 log points higher job growth. While the sign of these effects generally alleviate concerns that high search cost areas are distressed or rural areas with latently lower demand for loans or cars, key for our purposes is that the magnitude and significance of these correlations disappear in column 2 when we use our predicted search cost measure. Although having only a few nearby lenders is statistically related to some local characteristics, predicting search costs using the 1990 branch network breaks these correlations.

Overall, the results of Table 3 support our use of predicted search costs to ensure that our results on loan demand are driven by costly search. Combining these results with those in Appendix Table A8 rules out most forms of unobserved heterogeneity that could confound our interpretation of the effect of nearby lenders on outcomes. For an unobserved correlate of our search-cost measure to drive our results instead of search costs, it would have to affect borrowers on one side of a pricing discontinuity differently than borrowers on the other, be correlated with the number of mortgage and car loan applications submitted by the average borrower, and yet be uncorrelated with average borrower and borrower tract characteristics.

6 Effects of Search Frictions on Loan Take-up

Can costly search explain why many borrowers randomly assigned expensive interest rates do not avail themselves of better credit terms available elsewhere? Using application FICO scores, we estimate differences in loan take-up rates around FICO thresholds. If the nearby lender density captures a dimension of search costs that influences the propensity to search, the effect of similarly sized pricing discontinuities on loan take-up-rates should be larger in areas with fewer nearby lenders. In particular, we predict that applicants who are both below-threshold and have fewer nearby lenders (and thus face higher search costs and high markups) will be less likely to reject unfavorable loan offers and less likely to search for better terms elsewhere.

In a difference-in-differences spirit, our empirical specification measures how differences in loan take-up rates around FICO thresholds vary with search costs.[43] This allows us to exploit the quasi-random assignment of borrowers in the neighborhood of a FICO discontinuity to high and low markups to control for any unobserved differences in borrowers across search-cost categories. Even if borrowers in high- and low-search-cost areas do vary meaningfully on some unobservable dimension, the conditional variation in the FICO-discontinuity indicator variable allows us to difference out any demand-side variation common to borrowers at the same lender, FICO threshold, market, and time period and then compare the resulting responses across search-cost areas.

To see this argument statistically, consider that our conceptual framework's object of interest is whether high-search-cost borrowers are more likely than low-search-cost borrowers to accept a high-interest-rate loan offer. Using a binary measure of search costs $k \in \{k_{high}, k_{low}\}$ and interest rates $r \in \{r_{high}, r_{low}\}$, the expected estimated difference in loan take-up likelihood can be written as

$$E(takeup_{iglt}|r_{high}, k_{high}, X_{iglt}) - E(takeup_{iglt}|r_{high}, k_{low}, X_{iglt}), \tag{12}$$

where conditioning on $r = r_{high}$ in both terms represents that we are interested in the differential acceptance of high interest rates by high- and low-search-cost borrowers conditional on characteristics X_{iglt} measured for borrower i in geography g applying to lender l at time t.[44] While we expect this take-up effect to vary with search costs, endogeneity bias arises when unobservables ε vary with interest rate offers or search costs. Normalizing $k_{high} - k_{low} = 1$ and letting r_{high} be an indicator for above-threshold interest rates, the expected estimated take-up effect in (12) would be equal to

$$\begin{aligned} &E(takeup_{iglt}|r_{high}, k_{high}, X_{iglt}) - E(takeup_{iglt}|r_{high}, k_{low}, X_{iglt}) \\ &= \beta_2 + \beta_3 + E(\varepsilon_{iglt}|r_{high}, k_{high}, X_{iglt}) - E(\varepsilon_{iglt}|r_{high}, k_{low}, X_{iglt}) \end{aligned}$$

[43] See Grembi, Nannicini, and Troiano (2016) for a related difference-in-discontinuities identification strategy that uses discontinuities to identify a time difference instead of the spatial difference we study here.

[44] Borrower characteristics X consist of the additional controls included in (8): flexible controls for normalized FICO scores, and discontinuity × lender and geography × time fixed effects. Note that conditioning on X in the neighborhood of a discontinuity allows us to interpret r_{high} as a higher markup instead of merely a higher interest rate.

where β_2 and β_3 are the coefficients in (4) on k and $r \times k$, respectively. To isolate β_3 from fixed spatial differences in takeup rates correlated with our search-cost measure that load onto β_2, we will turn to a difference-in-discontinuities strategy below. More concerning is the endogeneity bias coming from the third and fourth terms; $E(\varepsilon|r_{high}, k, X)$ might not be zero because in general, borrowers with high interest-rate offers or high search costs may be selected in some unobservable way correlated with their likelihood to take up a loan. For example, if high-search-cost borrowers are more likely to be credit constrained, then they may be more likely to accept high interest-rate loans for reasons unrelated to retail credit-market search frictions.

To address this potential bias, we first break the correlation between ε and r by using the random assignment of high r in the neighborhood of the discontinuity by differencing out the effect of search costs on take-up rates for low-interest-rate borrowers. If we difference out the selection effect by considering the difference in take-up rates for *low* interest-rate markup borrowers in high- and low-search-cost areas, then our difference-in-discontinuities estimate for $\hat{\beta}_3$ is

$$\begin{aligned} E(\hat{\beta}_3) = & [E(takeup_{iglt}|r_{high}, k_{high}, X_{iglt}) - E(takeup_{iglt}|r_{low}, k_{high}, X_{iglt})] \\ & - [E(takeup_{iglt}|r_{high}, k_{low}, X_{iglt}) - E(takeup_{iglt}|r_{low}, k_{low}, X_{iglt})] \\ = & \beta_3 + [E(\varepsilon_{iglt}|r_{high}, k_{high}, X_{iglt}) - E(\varepsilon_{iglt}|r_{low}, k_{high}, X_{iglt})] \\ & - [E(\varepsilon_{iglt}|r_{high}, k_{low}, X_{iglt}) - E(\varepsilon_{iglt}|r_{low}, k_{low}, X_{iglt})] \\ = & \beta_3 \end{aligned} \tag{13}$$

The last line follows because the random assignment of interest rates in the neighborhood of each FICO discontinuity implies that the RD smoothness of unobservables assumption holds. Formally, this can be stated as

$$\begin{aligned} E(\varepsilon|r_{high}, k, X) - E(\varepsilon|r_{low}, k, X) &= \lim_{\widetilde{FICO} \to 0^-} E(\varepsilon|r(\widetilde{FICO}), k, X) - \lim_{\widetilde{FICO} \to 0^+} E(\varepsilon|r(\widetilde{FICO}), k, X) \\ &= 0, \end{aligned}$$

which is supported by our analysis in Appendix D.3 showing the smoothness of observables around FICO thresholds.

To implement this difference-in-discontinuities strategy, we augment our RD specification (8) with two additional controls: a High Search Cost indicator that proxies for k_{high} and an interaction between this dummy and the discontinuity indicator $1(\widetilde{FICO}_{id} \geq 0)$ that proxies for $r_{high} = 0$:

$$\begin{aligned} takeup_{iglt} = \sum_{d \in \mathcal{D}} 1(il \in \mathcal{D}_d) \Big(& \beta_1 \cdot 1(\widetilde{FICO}_{id} > 0) \\ & + \beta_2 \cdot HighSearchCost_{it} \\ & + \beta_3 \cdot 1(\widetilde{FICO}_{id} \geq 0) \times HighSearchCost_{it} \\ & + f(\widetilde{FICO}_{id}; \theta) + \lambda_{dl} \Big) + \eta_{gt} + \varepsilon_{iglt}. \end{aligned} \tag{14}$$

The difference-in-discontinuities coefficient β_3 captures how having high search costs affects loan take-up for borrowers who are quasi-randomly assigned high interest rates. As before, the RD function $f(\cdot;\cdot)$ captures a flexible function of the running variable, and geography $\times$ time fixed effects η_{gt} absorb differences in take-up rates arising from shocks at the commuting zone or zip code $\times$ quarter level. The summation over the set of discontinuities $\mathcal{D}$ allows us to pool all discontinuities in our estimation, accounting for potential differences across lenders or discontinuities in takeup rates with lender-by-discontinuity fixed effects λ_{dl}. The estimation sample for take-up regressions is the subset of the applications data within 19 FICO points of a pricing discontinuity, approved for a loan offer, and with non-missing address and FICO score data; these many restrictions result in a subsample of roughly 30,000 observations.

Table 4 reports estimates of (14), beginning with η_{gt} defined as commuting zone $\times$ quarter fixed effects. Column 1 reports an 11.4 pp difference in take-up rates across FICO discontinuities pooling borrowers with high and low search costs; that is, borrowers quasi-randomly drawing high markups from the distribution of interest rates are 11.4 percentage points less likely to accept the offered loan. Column 2 adds controls for high search costs and the interaction term of high interest rates and high search costs, defined as borrowers with at most 10 lending institutions within a 20-minute drive. The coefficient β_2 on the high-search-cost indicator is statistically and economically insignificant. That lower-interest-rate borrowers have quite similar take-up rates in our specification supports our identifying assumption of the comparability of borrowers in these two areas conditional on our RD controls and fixed effects. While low-search-cost borrowers are $\hat{\beta}_1 = 12.9$ pp more likely to accept a loan offer when they are just above a pricing discontinuity than below, high-search-cost borrowers are equally likely to accept loans on either side of a discontinuity—we cannot reject that $\beta_1 + \beta_3 = 0$ in column 2. Coupled with FICO pricing discontinuities being indistinguishable for high- and low-search-cost addresses, this suggests that borrowers who have higher search costs are much more willing to accept high interest rates than borrowers with lower search costs.[45] Column 3 shows that this result is robust to using within zip code $\times$ quarter variation.[46]

Given the mean take-up rate of 0.51 (Appendix Table A2), being in a high-search-cost has a 20-30% effect on the likelihood a borrower will accept an expensive loan. These differences in take-up rates are consistent with our conjectured mechanism where borrowers with high search costs are more likely to accept high interest-rate markups rather than search, potentially adjusting their consumption on other margins as discussed below.[47]

[45]A threat to our interpretation of Table 4 is that the magnitude of the interest-rate discontinuity ($r_{high} - r_{low}$ in (13)) may be different in high- and low-search-cost areas, naturally leading to differences in responses to the discontinuities. Note, however, that for this to explain our results, low-search-cost areas would have to have *larger* discontinuities in rates. Column 5 of Table 2 finds a small and insignificant difference in the size of the average interest-rate change at a FICO threshold across our measure of search costs, confirmed by graphical evidence in panel I of Appendix Figure A3. Similarly, we demonstrate in Appendix B that differences in the benefits to search across areas cannot explain our results given high-search-cost borrowers if anything can expect higher benefits to search from higher price dispersion.

[46]See panel II of Appendix Figure A3 for graphical evidence paralleling column 3 that the discontinuity in take-up rates is only present for low-search-cost borrowers.

[47]These take-up results also suggest that, as would be expected with costly search, individual lenders in both high- and low-search-cost areas face residual demand curves that are not infinitely elastic.

6.1 Estimates Using Quasi-Exogenous Variation in Search Costs

Our empirical setting generates exogenous variation in the benefits to search across lenders by quasi-randomizing within-lender interest-rate markups to borrowers around pricing discontinuities. Borrowers appear balanced across FICO discontinuities at the application stage, which is intuitive given the likely low level of precise awareness borrowers have about their own FICO score or the discontinuous mapping from FICO scores to loan pricing. Moreover, low-markup borrowers with high and low search costs have very similar take-up rates. However, there could be several time-varying or time-invariant factors correlated with our search-cost proxy that could differentially affect borrowers offered high markups. We emphasize that our RD results are unaffected by any differences between high- and low-search-cost areas that affect borrowers on both sides of a discontinuity. For example, differences in income, creditworthiness, or financial literacy are differenced out by our RD specifications before comparing across search-cost areas. Instead, such factors bias our estimates of β_3 only to the extent low-search-cost-area borrowers respond to pricing discontinuities differently from high-search-cost-area borrowers for reasons other than differences in search costs.

Consider that despite the ability of our difference-in-discontinuities estimator to hold many unobservables fixed, in practice we cannot be fully nonparametric in our conditioning on r, k, and X (especially their interactions) because our treatment of interest is at the $r \times k$ level.[48] Although unobserved heterogeneity correlated with r or k is addressed by our regression discontinuities by virtue of their quasi-random assignment, takeup factors that are unique to high search costs but which only affect high-interest-rate borrowers would be problematic. For example, if borrowers that are nearby fewer lenders are more likely to differ in their financial literacy, such effects are accommodated by our RD design because any such unobservables are balanced across the discontinuities within a fixed level of search costs. However, it would be problematic for our RD strategy if borrowers not nearby many lenders have lower financial literacy *and* financial literacy *only* affects takeup for high-interest-rate borrowers (e.g., if most borrowers accept low-interest-rate loan offers and only borrowers with low financial literacy accept high-interest-rate loans). Such a factor would survive conditioning on the random assignment to one side or the other of a discontinuity. Mathematically, let such an effect be $W_{it} \times r_{high}$, that has an effect γ on takeup rates, with W_{it} an indicator for a latent factor positively correlated with k.[49] If such an unobserved effect at the $r \times k$ level is included in ε, and if r and k are controlled for linearly, then the estimated coefficient $\hat{\beta}_3$ on $r \times k$ will be

[48]This identifying assumption is analogous to the usual difference-in-differences concern that unobserved shocks correlated with treatment $\times$ post lead to biased estimates.

[49]For example, W_{it} could be an indicator for credit constraints or low financial literacy. Such factors could lead to geographic differences in demand elasticities not driven by a causal effect of search frictions. If demand is less elastic in high-search-cost areas for reasons unrelated to search, then borrowers from such areas could be less sensitive to loan pricing, which would manifest only for borrowers offered high rates.

biased because

$$\begin{aligned}
&[E(\varepsilon_{iglt}|r_{high}, k_{high}, X_{iglt}) - E(\varepsilon_{iglt}|r_{low}, k_{high}, X_{iglt})] \\
&- [E(\varepsilon_{iglt}|r_{high}, k_{low}, X_{iglt}) - E(\varepsilon_{iglt}|r_{low}, k_{low}, X_{iglt})] \\
&\qquad\qquad\qquad\qquad = \gamma \frac{Cov(W_{it}, k_{high}|X_{iglt})}{Var(k_{high}|X_{iglt})} \neq 0,
\end{aligned}$$

where the bias fails to equal zero in the last line if and only if the latent demand factor W is correlated with high search costs. In this case, (13) evaluates to $E(\hat{\beta}_3) = \beta_3 + \gamma\theta$, where $\theta = \frac{Cov(W_{it}, k_{high}|X_{iglt})}{Var(k_{high}|X_{iglt})}$ is the coefficient on k_{high} in a regression of W on k_{high} and controls X. The potential presence of unobserved shocks positively correlated with $r \times k$ could lead to upward biased estimates of the effect of search costs on accepting a given loan offer.

To separately identify β_3 from the effect of any such unobservables, section 5.1 introduces an instrument for search costs based on historical variation in lender branch density combined with time-series variation in the nationwide total number of lender branches. By exploiting quasi-exogenous time-varying variation in our search-cost proxy, we can hold fixed geographic differences in unobservables correlated with current search costs $k_{it} = HighSearchCost_{it}$. As discussed above, Table 3 demonstrates that for the few candidate characteristics we find correlated with our search cost measure, these characteristics are unrelated to our predicted search costs measure.

Columns 4-5 of Table 4 repeat the specifications of columns 2-3 but replace our high search cost indicator with a predicted measure of search costs to address the possibility that omitted variables correlated with $r \times k$ could bias our estimates of β_3. To do so, we reestimate (14) instead conditioning on predicted search costs $\hat{k}_{it} = High\widehat{Search}Cost_{it}$ as defined in (10). While an unobservable like $lowliteracy \times r_{high}$ may be correlated with $k_{high} \times r_{high}$, such unobservables are conceptually unlikely to be correlated with $\hat{k}_{high} \times r_{high}$, supported empirically by Table 3. Again, the rationale behind this identifying assumption used in columns 4-5 is that the current unobserved demand factors that might affect high- and low-markup borrowers differently in low- versus high-search-cost areas are unlikely to be correlated with the 1990 local lender density used to define $\hat{k}$. As before, borrowers predicted to have higher search costs are significantly less sensitive to interest-rate discontinuities, with the point estimates for β_3 in columns 4-5 suggesting they are 7-10 pp less likely to accept a higher-markup offer. Borrowers in low-search-cost areas exhibit strong and statistically significant reactions to loan offers. The estimated main effect β_2 of higher search costs is again statistically and economically insignificant, suggesting that our specification makes low-markup borrowers in these areas roughly comparable. These results suggest that differences in loan take-up rates across borrowers predicted to have high and low search costs are not driven by the endogeneity of the local branching network to time-varying economic shocks. Finally, combining these results with Table 3 also rules out some sort of persistent spatial variation in credit-demand elasticities being correlated with our search-cost measure.

Overall, we conclude that borrowers in low-search-cost areas are less likely to accept a loan if they are assigned a high instead of a low interest-rate markup, whereas high-search-cost borrowers are more likely to accept loan terms regardless of where their offers fall in the distribution of markups.

6.2 Robustness to Definition of High-Search-Cost Area

To probe whether our take-up results are sensitive to the choice of cutoff for our high-search-cost area definition, we vary the cutoff used for the definition of the High Search Cost dummy and plot the results in Figure 4. Each point plots the coefficient on the interaction between $HighSearchCost_{it} \times 1(\widetilde{FICO}_{id} \geq 0)$ from a separate regression. Along the horizontal axis, we indicate the number of proximate lenders within a 20-minute drive used to define High Search Cost for the corresponding point estimate and confidence interval, varying the maximum number of nearby institutions to qualify as high search cost to be 5, 10, 15, etc. When the max number of nearby lenders on the horizontal axis is 10, the point estimate corresponds to column 2 of Table 4.

Several interesting patterns emerge from the plotted coefficients. First, when the contrast between high- and low-search-cost areas is meaningful (i.e., when the maximum number of nearby lenders on the x-axis is below 25), the point estimates are negative and largest when the number of nearby lenders is small. Recall that the main effect of a FICO discontinuity on take-up rates is positive; applicants just above a FICO discontinuity are much more likely to accept an offered loan because they are being offered much more favorable interest rates (an average of 137 bp lower) than just-below applicants. Being in a high-search-cost area undoes all of the main effect of an interest rate markup on take-up. We interpret this as high-search-cost borrowers being insensitive to their offered interest rate r' because they would find shopping for a better rate particularly costly relative to low-search-cost borrowers. Second, while borrowers with five or fewer nearby lenders are rare and so the resulting take-up discontinuity is estimated imprecisely, the other estimates for borrowers who have less than 25 nearby lenders are significantly less than zero. Third, estimates for high-search-cost definitions near our baseline cutoff of 10 nearby lenders are similar, suggesting that our conclusions are robust to the exact definition of a high-search-cost area. Finally, the point estimates become insignificant as the number of nearby lenders used in the definition of high search cost grows. This is consistent with the intuition that from a search-cost perspective, the value of an additional nearby lender diminishes (i.e., each lender matters more when there are only a handful of nearby lenders compared to when there are, for example, over 40 lenders nearby). From an identification standpoint, too, it is reassuring that there is no difference in take-up rates across a FICO discontinuity when the number-of-lenders cutoff is too large to reasonably be capturing something related to search.

7 Effects on Loan and Consumption Quantities

To directly demonstrate how search frictions in credit markets can have real effects on consumption, we next establish that being treated with a higher offered interest rate affects subsequent loan and purchase decisions. Whether a given credit-market imperfection constrains consumption is usually challenging to identify empirically because it requires estimating counterfactual consumption in the absence of the alleged friction. However, our RD setup allows us to test for quantity effects by using the borrowing and purchasing decisions of borrowers on one side of a FICO threshold as a counterfactual for borrowers on the other. Given the empirical results in Appendix D.3 that borrowers are ex-ante similar around FICO thresholds, we assume that borrowers around FICO thresholds would have similar demand for loans and cars if quoted the same set of financing terms. Exploiting our ability to observe the exact amount borrowed and spent on a car, we test whether borrowers spend differently around the observed FICO thresholds and whether the composition of borrowers accepting loans changes across thresholds. Combined with higher interest rates on the expensive side of FICO thresholds being much more likely to be accepted by high-search-cost borrowers, these results demonstrate that search costs for credit affect subsequent consumption.

Summarizing our main result graphically, Figure 5 plots car purchase expenditures around the normalized FICO threshold. Purchase amounts are smooth leading up to the FICO threshold and then jump down discontinuously at the threshold. Using the same RD design used in our pricing discontinuity analysis above, we formally test for statistical differences in purchase amounts. As before, we estimate equation (8) by controlling for commuting-zone $\times$ quarter-of-origination fixed effects and discontinuity $\times$ lender fixed effects and allow for a quadratic function of the running variable using a bandwidth of 19 around the normalized FICO threshold with a uniform kernel.[50]

Table 5 presents these reduced-form results. Borrowers quasi-randomly offered more expensive loans spend an average of $375 less on the cars they purchase (a 2.4% effect evaluated at the mean purchase price). Column 2 presents results with loan amounts as the dependent variable. Originated loan sizes are an average of $557 (4%) lower on the expensive side of a detected FICO discontinuity.[51] 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 approved for and take up higher loan-to-value (LTV) ratios. Column 3 of Table 5 indicates that ex-post LTV ratios are an average of 1.4 percentage points higher for borrowers to the right of FICO thresholds. Given that ex-ante DTI ratios in the loan application data are continuous around the thresholds (Appendix Table A5), we interpret these results as further evidence of the easing of credit terms for above-threshold borrowers. Borrowers with FICO scores just above a pricing discontinuity are offered lower rates, longer terms, and allowed higher ex-post LTV ratios.

Microdata on loan amounts and loan terms allow us to calculate the implied monthly payment of each borrower 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 decrease by a statistically and economically insignificant $0.08 for above-threshold borrowers. Shorter terms and higher interest rates lead below-threshold borrowers to purchase less expensive cars and use less financing in their purchase than above-threshold borrowers, essentially purchasing less car and using less credit to keep the same monthly payment.[52] Decomposing how high-search-cost borrowers accommodate high interest rates instead of searching, roughly 50% of accepted higher loan costs is

[50]Because this exercise necessarily conditions on origination, we are able to use our much larger origination sample for this set of results.

[51]The lower demand for loans by borrowers facing higher interest rates further argues against an adverse selection explanation, which would imply simultaneously higher demand and higher default (Finkelstein and Poterba, 2004). Instead, we observe similar ex-ante applications, lower ex-post takeup and loan sizes, and similar default by high-interest-rate borrowers.

[52]See Argyle et al. (2020) for related evidence on borrowers' monthly payment targeting.

offset purely by lower LTVs, 20% is offset by buying cheaper cars, and the remainder is offset by the combination of the two.

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 a quantity effect of search frictions. Absent search frictions, we would expect all borrowers to find the lowest-price provider leading to infinitely elastic residual demand curves for a given lender. Instead, these results suggest each lender faces a finite residual demand elasticity, consistent with search frictions giving direct auto-loan lenders market power. Given the balance of borrowers across discontinuities, we would also expect to observe similar purchasing decisions for above- and below-threshold borrowers absent their differential interest-rate draws.

One concern with this interpretation is that purchase-price effects may not represent quantity effects if borrowers only *pay* different amounts but actually purchase the same cars they would have otherwise, deriving the same flow utility from their purchase. For example, if dealers can use financing terms to price discriminate, they may exploit above-threshold borrowers' increased marginal willingness to pay by charging more for the exact same car than otherwise similar borrowers with more expensive financing.[53] We test for this possibility by controlling for year-make-model (e.g., 2013 Honda Accord) fixed effects in our RD regressions. Column 1 of Table 6 reports results when controlling for make-model fixed effects. Even within a make and model category, borrowers quasi-randomly assigned expensive credit continue to spend $344 less on cars, suggesting that the bulk of 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 on the nature of the substitution patterns in this market. When we include *year*-make-model fixed effects in column 2, we find a much smaller change in purchase price at the discontinuity of $72. Because fixing the model year of a car has such large explanatory power on the effect of an interest-rate markup, we conclude that much of the effect in column 1 is explained by substitution within a model and across model years.[54]

Reconciling the strong effect on purchase prices within make-models and the relatively weaker effect on purchase prices within make-model-years, column 3 provides direct evidence with vehicle age at purchase in months as the dependent variable (controlling for make-model fixed effects since vehicle age would be collinear with year-make-model and time fixed effects). Borrowers with access to easier credit purchase cars that are on average 1.8 months newer, suggesting that roughly one in seven borrowers respond to a high interest-rate markup by buying a car that is one model year older, keeping their monthly payments roughly constant. These car-age effects are consistent with some borrowers preferring to purchase older cars over searching for better financing. Appendix Figure A4 uses data from the National Highway Transportation Survey (U.S. Department of Transportation, 2017) to show the average relationship between car age and mileage. On average, every additional year of car age in the NHTS data is associated with 8,000 more miles. Evidence in Busse et al.

[53]See Argyle et al. (2021) for evidence that individual-level used-car prices capitalize loan maturities.

[54]The small estimated treatment effect on prices within make-model-year stems from the *borrower* being treated in contrast to the vehicle-level treatment in Argyle et al. (2021) that has larger effects on price.

(2013) suggests that controlling for make, model, model year, and trim, the used-car market values every additional 8,000 miles of mileage at -$960, implying a valuation of a 1.8-month age effect of $144.

These results on car expenditures and car age are consistent with search frictions distorting the quantity of car services purchased. These changes in durable purchases affect consumers' utility in two ways, by lowering flow utility of their purchased car and the vehicle's durability and resale value. The estimated differences in purchases across interest-rate discontinuities—combined with the higher propensity of higher-search-cost borrowers to accept quasi-randomly higher interest rates—are sufficient to imply real (as opposed to purely financial) effects of financing search frictions. Overall, high financing search costs induce borrowers facing high rate markups to take out higher-interest-rate and smaller loans and purchase older and cheaper cars.

Appendix Table A6 divides the sample into high- and low-search-cost borrowers and shows that FICO pricing discontinuities affect loan sizes and purchase prices in both samples, as expected.[55] Again, the mechanism described in section 2 is that borrowers with high financing search costs will purchase cheaper cars because they face weakly higher interest rates. High-search-cost borrowers have these larger quantity distortions because they are more likely to accept high-markup loans, and on average, any borrower who accepts a high-markup loan will take out a smaller loan and purchase a cheaper car (Tables 4-5).

Taking stock, how do borrowers respond to being arbitrarily offered high financing markups? The evidence presented in Tables 5 and 6 indicates that many borrowers accepting expensive credit adjust their loan- and car-purchasing behavior to keep their monthly payments the same despite the higher interest rates. Such borrowers spend less on their car purchases by selecting an older car than they would have otherwise, originating smaller loans at higher loan-to-value ratios. Combined with borrowers in high-search-cost areas being more likely to accept arbitrarily expensive loans, these overall quantity distortions are more consequential when search is more costly. We view this as evidence that borrowers' inability to costlessly identify the best available loan terms distorts consumption away from the equilibrium quantities that would prevail absent search frictions.

7.1 Assessing Selection into Origination

Limiting our sample to direct loans eliminates the possibility of borrower selection driven by car dealers steering of borrowers to lenders via so-called indirect loans. Moreover, borrowers are unlikely to be aware of their precise FICO score as calculated by the credit bureau queried by the lender pricing their loan application, and even less aware of that lender's or alternative lenders' FICO discontinuities. However, another possibility is that borrowers *who accept loans* on either side of FICO thresholds might differ systematically (even if borrowers are balanced at the application stage), violating the smoothness condition required for valid RD inference of second-stage effects.[56]

[55] Borrowers with high financing search costs seem more affected by a similarly sized loan markup than low-search-cost borrowers in Appendix Table A6, but the difference is statistically insignificant.

[56] For example, perhaps borrowers who accept high loan markups are particularly inelastic or lenders may change their reliance on soft information around FICO discontinuities.

In this section, we address the possibility that borrowers who take up below-threshold, high-markup loan offers are different on unobservable dimensions from above-threshold, low-markup car buyers.

Theories of adverse selection in credit markets such as Stiglitz and Weiss (1981) suggest that the population of borrowers who accept high interest-rate markups is disproportionately composed of borrowers with private information that they are more likely to default. Such a selection story could be at play in our setting, too, where conditional on origination, the sample of borrowers on the expensive side of FICO discontinuities is somehow selected relative to the sample of borrowers with similar FICOs who apply but then withdraw after observing the interest rate. However, the premise of the Stiglitz and Weiss (1981) unraveling argument is that although such borrower heterogeneity is unobserved by the credit market at the origination stage, it reveals itself in higher defaults later. Our setting enables us to follow borrowers after origination and observe outcomes for borrowers on both sides of discontinuities. Using several ex-post creditworthiness measures, we find no evidence that borrowers taking up loans on either side of discontinuities are systematically different ex-ante.

The evidence in Appendix D.3 demonstrates that interest-rate markups seem quasi-randomly assigned (i.e., at the ex-ante application stage, borrowers are similar on observable dimensions). An alternative explanation for our results is that rate-insensitive borrowers may accept high loan markups and face other constraints that lead them to demand lower loan sizes and spend less on their car purchases. Similarly, perhaps (unobservably) high credit-quality borrowers who are arbitrarily offered expensive interest rates withdraw their loan applications and look elsewhere for credit. Under this private-information scenario, borrowers who accept expensive loans are those who know they are of poor credit quality and unlikely to do better given their unfavorable soft attributes. If lenders recognize that borrowers who choose to accept unfavorable terms are riskier, ex-ante arbitrary FICO discontinuities could reinforce an equilibrium that separates high- and low-credit-quality borrowers with the appropriate pricing differences offered to each borrower type.

To test for the possibility that some form of selection drives the observed equilibrium outcomes in our data, we compare the balance around the FICO discontinuities of borrower characteristics and ex-post borrower performance for the accepted loan applications in our applications data. If some correlated selection process guides differences in who accepts expensive loan offers, this should be revealed by an imbalance of offer-accepting borrower characteristics or ex-post credit outcomes. Appendix Figure A5 repeats the exercise of Appendix Figure A2 by checking for the smoothness of borrower and loan characteristics around FICO discontinuities.[57] For the four borrower attributes that should not respond to the thresholds if there is not selection into origination based on borrower demand, we find smoothness. Although the estimates are noisier in this applications-originations merged sample due to smaller sample sizes, age, DTI, gender, and ethnicity have economically small and statistically insignificant differences across FICO discontinuities, indicating that borrowers who accept high markups seem similar to those above FICO discontinuities that accept low markups. Panels II and VI show that loan sizes and the fraction of borrowers accepting loans do

[57] We also merge our application data to our origination data to demonstrate that the same smoothness in ex-ante borrower attributes that we saw at the application stage persists conditional on origination.

respond to interest-rate markups, consistent with the extensive- and intensive-margin effects of the discontinuities we estimate in Tables 4 and 5, respectively. The results of the Appendix Figure A5 exercise suggest that the observed borrower responses to FICO discontinuities are causal effects of the search-cost-induced interest-rate markups and not explainable by selection into which borrowers accept higher markups.

As any selection on credit quality should eventually manifest itself in the average ex-post credit performance of selected borrowers, we further test for selection by specifying as a dependent variable in our RD setting various ex-post credit outcomes. We also interact our discontinuity indicator with our high search cost indicator to verify that the selection into origination is not different for high search cost borrowers. The coefficient in column 1 of Table 7 estimates that above-threshold borrowers are an average of 3.5 more days delinquent than below-threshold borrowers, indicating that borrowers on either side of the threshold do not exhibit economically meaningful or statistically significant differences in delinquency.[58] Similarly, above-threshold borrowers are 0.4 percentage points more likely to have their loan charged off (written off as a loss by the lender, column 2) and 0.1 percentage points more likely to be in default (over 90 days past due, column 3), which we consider precise zeroes. For each of these outcomes in columns 1-3, the magnitude of the discontinuity in ex-post credit outcomes is statistically indistinguishable for borrowers with high financing search costs.[59] At least one reason for why we do not observe any evidence of ex-post differences in default is likely the monthly payment smoothing behavior of borrowers documented in Table 5. Given that borrowers adjust loan sizes and expenditure amounts to keep monthly payments roughly constant, there is less scope for differential defaulting by borrowers with higher markups.

A novel feature of our dataset allows for a second test of private information on creditworthiness 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.[60] Ex-post credit scores allow us to calculate changes in credit scores over time, capturing broad changes in borrower distress and financial responsibility incorporating other credit products beyond the given auto loan in question. 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 actually reflects their riskiness) are the only ones to originate such loans. This omnibus test is particularly valuable given that the effect of markups on equilibrium monthly payments is close to zero, possibly resulting in negligible effects on the performance of the auto loan itself. Using the subsample 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 (0.2 percentage points) in credit score changes for borrowers around the

[58]The sample size varies across columns in Table 7 because of inconsistent data coverage of all monitoring variables across lenders.

[59]Appendix Figure A6 presents these results graphically.

[60]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.

threshold, with small and statistically insignificant differences for borrowers with high financing search costs.

While adverse selection undoubtedly motivates 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 and purchasing behavior around the observed FICO pricing discontinuities. Of course, selection into take-up correlated with ex-ante borrower characteristics or ex-post credit worthiness 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, any explanation such as borrower naïveté would have to not be manifest in differences across thresholds at loan application or origination, significantly higher prices paid for the same make-model-year, differential ex-post default rates, or differences in ex-post credit scores. Given this set of outcomes showing borrowers offered high and low interest rate markups to be otherwise quite similar, we find it unlikely our results are driven by a missing factor that drives the differential response of above- and below-discontinuity borrowers in high- and low-search-cost areas.

8 Conclusion

Mounting evidence indicates that a variety of credit-market imperfections influence household debt and consumption outcomes. A parallel empirical search literature establishes the consequences of the costliness of learning prices in a wide variety of markets. We present empirical evidence connecting these two literatures, demonstrating the effect of financing search frictions—a market featuring elastic and continuous demand—on credit-market and consumption outcomes.

The average borrower in our data pays 137 bp more than the best rate available to observationally similar borrowers—$488 in present-value terms for the average loan. In this setting where the gains to search are high, we show that borrowers' acceptance of dominated loan terms is related to measures of the cost of searching for retail credit. Because discontinuous pricing schedules vary across lenders within the same commuting zone, borrowers on the expensive side of FICO discontinuities in loan pricing at one institution would be more likely to find favorable pricing at a different institution. Absent search frictions, borrowers are unlikely to accept seemingly dominated loan terms. Moreover, under the assumption that borrowers on either side of pricing discontinuities are otherwise identical, we would expect them to ultimately find similar financing opportunities and purchase similar cars absent search frictions. Proxying for the costliness of loan shopping with the density of nearby lenders, we show that borrowers with higher search costs face weakly more dispersed prices, are more likely to accept quasi-randomly offered dominated loan terms, and apply for fewer loans. Robustness exercises show this result is distinct from a simple market-concentration explanation and holds using alternative identifying variation. Reinforcing our interpretation, we show that to bias our estimates, any omitted variables correlated with our search-cost proxy would have to also have differential effects for high- and low-markup borrowers and be uncorrelated with the characteristics of both borrower and their Census tracts.

Next, we confirm that accepting a high-markup loan (an effect of search costs) has material intensive-margin effects on loan quantities and durable consumption. Borrowers quasi-randomly offered high-markup loans on average borrow \$560 less, spend \$375 less on their car purchases, and buy 1.8 months older cars than otherwise similar borrowers offered and accepting lower rates. This downward sloping continuous demand for cars and car loans, combined with high-search-cost borrowers being more likely to accept high-markup loans, highlights the importance of well-functioning consumer credit markets in determining durable goods consumption patterns. Moreover, relative to traditional search models with inelastic discrete demand where dispersed prices for final goods have no associated deadweight loss and just represent a transfer from buyers to sellers, there appear to be aggregate welfare consequences of costly search for credit in the real world. When financing search costs are non-negligible, consumers facing firm-specific markups for credit may adjust the quantity or characteristics of both that good and its complements away from first-best levels.

Even with a well-developed financial sector including secondary markets for many forms of consumer debt, household consumption still appears distorted by costly search for credit as an important credit market imperfection. At least one answer to Zinman's (2014) query as to why efficient risk-based pricing is still not ubiquitous in the era of big-data-based credit modeling appears to be demand-side obstacles to finding the lowest available interest rates. Even with the possibility of shopping for interest rates online, searching for consumer credit products currently remains an opaque, local, and costly process for many borrowers. This relationship between costly search for financing and consumption outcomes broadens the consequences of search frictions, especially in credit markets, and could motivate the extra regulatory attention paid to so-called banking deserts where the density of lenders is particularly low.

References

Adams, Paul, Stefan Hunt, Christopher Palmer, and Redis Zaliauskas (July 2021). "Testing the Effectiveness of Consumer Financial Disclosure: Experimental Evidence from Savings Accounts." *Journal of Financial Economics* 141.1, pp. 122–147.

Adams, William, Liran Einav, and Jonathan Levin (2009). "Liquidity Constraints and Imperfect Information in Subprime Lending." *American Economic Review* 99.1, pp. 49–84.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino (Mar. 2012). "Credit Supply and House Prices: Evidence from Mortgage Market Segmentation." Working Paper. NBER Working Paper No. 17832.

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel (2017). "Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?" *The Quarterly Journal of Economics* 133.1, pp. 129–190.

Agarwal, Sumit, John Grigsby, Ali Hortaçu, Gregor Matvos, Amit Seru, and Vincent Yao (2020). "Searching for Approval." NBER Working Paper No. 27341.

Alexandrov, Alexei and Sergei Koulayev (2018). "No Shopping in the U.S. Mortgage Market: Direct and Strategic Effects of Providing Information." CFPB Working Paper No. 2017-01.

Allen, Jason, Robert Clark, and Jean-François Houde (2014a). "Price Dispersion in Mortgage Markets." *The Journal of Industrial Economics* 62.3, pp. 377–416.

— (Oct. 2014b). "The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry." *American Economic Review* 104.10, pp. 3365–96.

— (2019). "Search Frictions and Market Power in Negotiated-Price Markets." *Journal of Political Economy* 127.4, pp. 1550–1598.

Al-Najjar, Nabil I and Mallesh M Pai (2014). "Coarse Decision Making and Overfitting." *Journal of Economic Theory* 150, pp. 467–486.

Aneja, Abhay P. and Carlos F. Avenancio-León (Feb. 2020). "No Credit For Time Served? Incarceration and Credit-Driven Crime Cycles." Working Paper.

Angrist, Joshua D and Victor 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.

Argyle, Bronson, Taylor Nadauld, and Christopher Palmer (Jan. 2020). "Monthly Payment Targeting and the Demand for Maturity." *The Review of Financial Studies* 33.11, pp. 5416–5462.

Argyle, Bronson, Taylor Nadauld, Christopher Palmer, and Ryan Pratt (2021). "The Capitalization of Consumer Financing into Durable Goods Prices." *The Journal of Finance* 76.1, pp. 169–210.

Attanasio, Orazio P, Pinelopi Koujianou Goldberg, and Ekaterini Kyriazidou (2008). "Credit Constraints In the Market for Consumer Durables: Evidence from Micro Data on Car Loans." *International Economic Review* 49.2, pp. 401–436.

Aydin, Deniz (Jan. 2022). "Consumption Response to Credit Expansions: Evidence from Experimental Assignment of 45,307 Credit Lines." *American Economic Review* 112.1, pp. 1–40.

Bartlett, Jeff (Sept. 2013). "Consumers rely on car financing more than ever." *Consumer Reports.*

Baye, Michael R, John Morgan, and Patrick Scholten (2006). "Information, Search, and Price Dispersion." *Handbook on Economics and Information Systems*. Ed. by T. Hendershott. Vol. 1, pp. 323–77.

Benetton, Matteo (Aug. 2021). "Leverage Regulation and Market Structure: A Structural Model of the UK Mortgage Market." *The Journal of Finance*.

Benetton, Matteo, Sergio Mayordomo, and Daniel Paravisini (2021). "Credit Fire Sales: Captive Lending as Liquidity in Distress." SSRN Working Paper No. 3780413.

Benmelech, Efraim, Ralf R. Meisenzahl, and Rodney Ramcharan (Nov. 2017). "The Real Effects of Liquidity During the Financial Crisis: Evidence from Automobiles." *The Quarterly Journal of Economics* 132.1, pp. 317–365.

Bhutta, Neil, Andreas Fuster, and Aurel Hizmo (2020). "Paying Too Much? Price Dispersion in the US Mortgage Market."

Bubb, Ryan and Alex Kaufman (2014). "Securitization and Moral Hazard: Evidence from Credit Score Cutoff Rules." *Journal of Monetary Economics* 63, pp. 1–18.

Busse, Meghan R., Nicola Lacetera, Devin G. Pope, Jorge Silva-Risso, and Justin R. Sydnor (May 2013). "Estimating the Effect of Salience in Wholesale and Retail Car Markets." *American Economic Review* 103.3, pp. 575–79.

Busse, Meghan R. and Jorge M Silva-Risso (2010). "'One Discriminatory Rent' or 'Double Jeopardy': Multicomponent Negotiation for New Car Purchases." *American Economic Review* 100.2, pp. 470–474.

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

Campbell, John Y (2006). "Household Finance." *The Journal of Finance* 61.4, pp. 1553–1604.

Chetty, Raj, John N Friedman, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter (Oct. 2018). "The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility." NBER Working Paper No. 25147.

DeGroot, Morris (1970). *Optimal Statistical Decisions*. New York: McGraw-Hill.

Degryse, Hans and Steven Ongena (2005). "Distance, Lending Relationships, and Competition." *The Journal of Finance* 60.1, pp. 231–266.

Delavande, Adeline and Basit Zafar (Oct. 2019). "University Choice: The Role of Expected Earnings, Nonpecuniary Outcomes, and Financial Constraints." *Journal of Political Economy* 127.5, pp. 2343–2393.

Di Maggio, Marco and Amir Kermani (June 2017). "Credit-Induced Boom and Bust." *The Review of Financial Studies* 30.11, pp. 3711–3758.

Diamond, Peter A (1971). "A Model of Price Adjustment." *Journal of Economic Theory* 3.2, pp. 156–168.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl (2017). "The Deposits Channel of Monetary Policy." *The Quarterly Journal of Economics* 132.4, pp. 1819–1876.

Einav, Liran, Mark Jenkins, and Jonathan Levin (2012). "Contract Pricing in Consumer Credit Markets." *Econometrica* 80.4, pp. 1387–1432.

— (2013). "The Impact of Credit Scoring on Consumer Lending." *The RAND Journal of Economics* 44.2, pp. 249–274.

Ellison, Glenn and Sara Fisher Ellison (2009). "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica* 77.2, pp. 427–452.

Favara, Giovanni and Jean Imbs (Mar. 2015). "Credit Supply and the Price of Housing." *American Economic Review* 105.3, pp. 958–992.

Federal Highway Administration (2017). *National Household Travel Survey*. Washington, DC.

Federal Reserve Bank of New York (Aug. 2021). "Quarterly Report on Household Debt and Credit."

Finkelstein, Amy and James Poterba (2004). "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market." *Journal of Political Economy* 112.1, pp. 183–208.

Granja, João, Christian Leuz, and Raghuram Rajan (Oct. 2021). "Going the Extra Mile: Distant Lending and Credit Cycles." NBER Working Paper No. 25196.

Greenwald, Daniel (2018). "The Mortgage Credit Channel of Macroeconomic Transmission." MIT Sloan Working Paper No. 5184-16.

Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano (2016). "Do Fiscal Rules Matter?" *American Economic Journal: Applied Economics* 8.3, pp. 1–30.

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

Grunewald, Andreas, Jonathan Lanning, David Low, and Tobias Salz (2020). "Auto Dealer Loan Intermediation: Consumer Behavior and Competitive Effects." Working Paper.

Guiso, Luigi, Andrea Pozzi, Anton Tsoy, Leonardo Gambacorta, and Paolo Emilio Mistrulli (2021). "The cost of steering in financial markets: Evidence from the mortgage market." *Journal of Financial Economics*.

Hong, Han and Matthew Shum (2006). "Using Price Distributions to Estimate Search Costs." *The RAND Journal of Economics* 37.2, pp. 257–275.

Hortaçsu, Ali and Chad 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.

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

Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig (2009). "Financial Regulation and Securitization: Evidence from Subprime Loans." *Journal of Monetary Economics* 56.5, pp. 700–720.

— (2010). "Did Securitization Lead to Lax Screening? Evidence from Subprime Loans." *The Quarterly Journal of Economics* 125.1, pp. 307–362.

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

Liberman, Andres, Daniel Paravisini, and Vikram Pathania (July 2021). “High-Cost Debt and Perceived Creditworthiness: Evidence from the UK.” *Journal of Financial Economics*.

Liberti, José María and Mitchell A Petersen (2018). “Information: Hard and Soft.” *Review of Corporate Finance Studies* 8.1, pp. 1–41.

Liebersohn, Jack (2017). “Does Local Bank Competition Still Matter? Quasi-Experimental Evidence from Bank Mergers.” Working Paper.

Lippman, Steven A and John J McCall (1976). “Job Search in a Dynamic Economy.” *Journal of Economic Theory* 12.3, pp. 365–390.

Livshits, Igor, James C MacGee, and Michele Tertilt (2016). “The Democratization of Credit and the Rise in Consumer Bankruptcies.” *The Review of Economic Studies* 83.4, pp. 1673–1710.

Matějka, Filip and Alisdair McKay (2015). “Rational inattention to discrete choices: A new foundation for the multinomial logit model.” *American Economic Review* 105.1, pp. 272–98.

McCrary, Justin (2008). “Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test.” *Journal of Econometrics* 142.2, pp. 698–714.

Melzer, Brian T. (Feb. 2011). “The Real Costs of Credit Access: Evidence from the Payday Lending Market.” *The Quarterly Journal of Economics* 126.1, pp. 517–555.

Moraga-González, José Luis, Zsolt Sándor, and Matthijs R Wildenbeest (2017). “Prices and Heterogeneous Search Costs.” *The RAND Journal of Economics* 48.1, pp. 125–146.

— (Aug. 2021). “Consumer Search and Prices in the Automobile Market.” Working Paper.

Murfin, Justin and Ryan Pratt (2019). “Who Finances Durable Goods and Why It Matters: Captive Finance and the Coase Conjecture.” *The Journal of Finance* 74.2, pp. 755–793.

Nguyen, Hoai-Luu Q (2019). “Are Credit Markets Still Local? Evidence from Bank Branch Closings.” *American Economic Journal: Applied Economics* 11.1, pp. 1–32.

Reinganum, Jennifer F (1979). “A Simple Model of Equilibrium Price Dispersion.” *Journal of Political Economy* 87.4, pp. 851–858.

Scharfstein, David S and Adi Sunderam (2016). “Market Power in Mortgage Lending and the Transmission of Monetary Policy.” Working Paper.

Stahl, Dale O (1989). “Oligopolistic Pricing with Sequential Consumer Search.” *American Economic Review*, pp. 700–712.

Stango, Victor and Jonathan Zinman (2015). “Borrowing High versus Borrowing Higher: Price Dispersion and Shopping Behavior in the US Credit Card Market.” *Review of Financial Studies*.

Stigler, George J. (1961). “The Economics of Information.” *Journal of Political Economy* 69.3, pp. 213–225.

Stiglitz, Joseph E and Andrew Weiss (1981). “Credit Rationing in Markets with Imperfect Information.” *American Economic Review* 71.3, pp. 393–410.

Woodward, Susan E and Robert E Hall (2012). “Diagnosing Consumer Confusion and Sub-optimal Shopping Effort: Theory and Mortgage-Market Evidence.” *American Economic Review* 102.7, pp. 3249–3276.

Zaki, Mary (July 2016). “Access to Short-term Credit and Consumption Smoothing within the Paycycle.” FEEM Working Paper No. 007.2016.

Zeldes, Stephen P (Apr. 1989). “Consumption and Liquidity Constraints: An Empirical Investigation.” *Journal of Political Economy* 97.2, pp. 305–346.

Zhou, Jidong (Sept. 2014). “Multiproduct Search and the Joint Search Effect.” *American Economic Review* 104.9, pp. 2918–39.

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

Figure 1: Interest-rate FICO Regression Discontinuity Plots

I. Interest Rates Around FICO = 600 Discontinuities

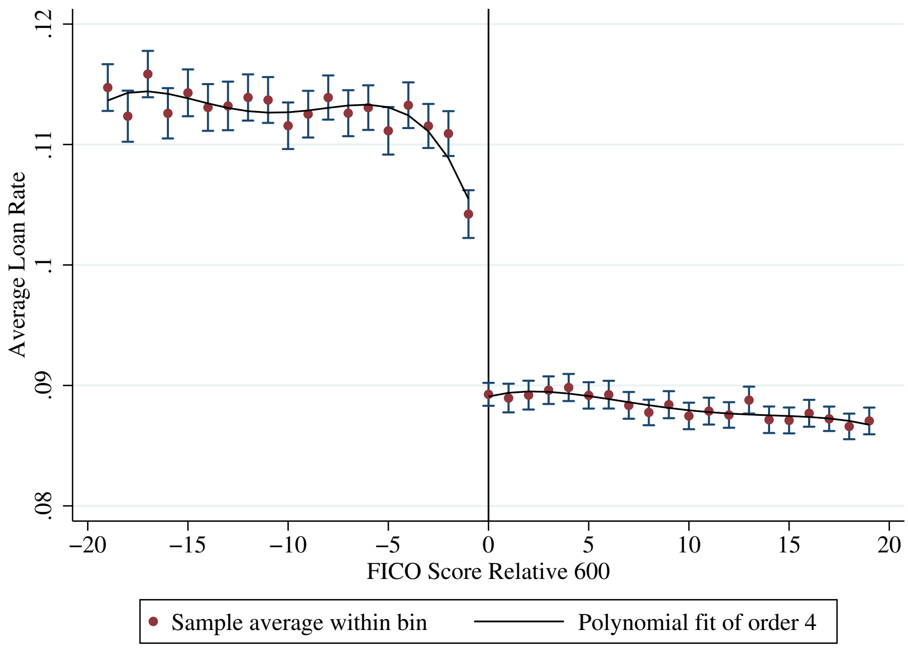

II. Interest Rates Around FICO = 700 Discontinuities

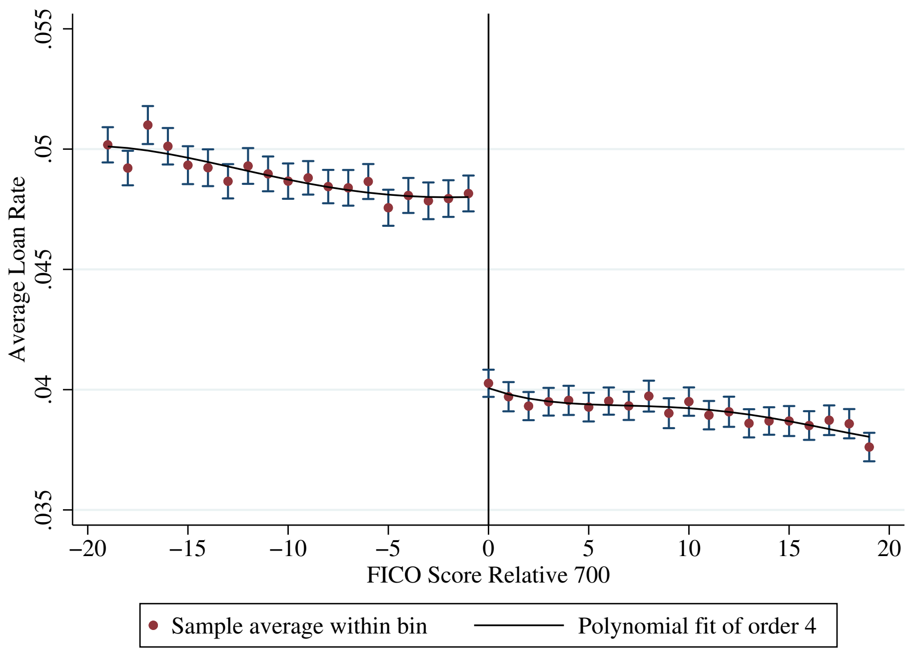

Notes: Figures plot average interest rates against borrower FICO scores normalized to pricing discontinuities. 95% confidence intervals are double clustered by lender and FICO score. Plotted RD functions are estimated using the Calonico et al. (2014) robust RD estimator with fourth-degree polynomials for institutions with pricing discontinuities detected at FICO scores of 600 and 700 in panels I and II, respectively.

Figure 2: Variation in the Returns to Search Around FICO Discontinuities

I. Borrowers Around a FICO = 600 Threshold

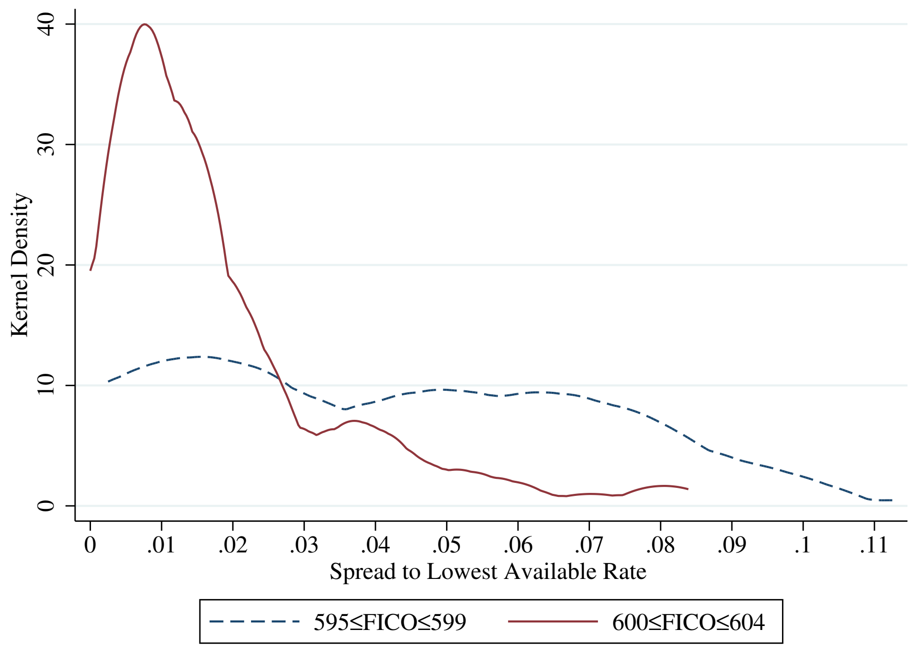

II. Borrowers Around a FICO = 700 Threshold

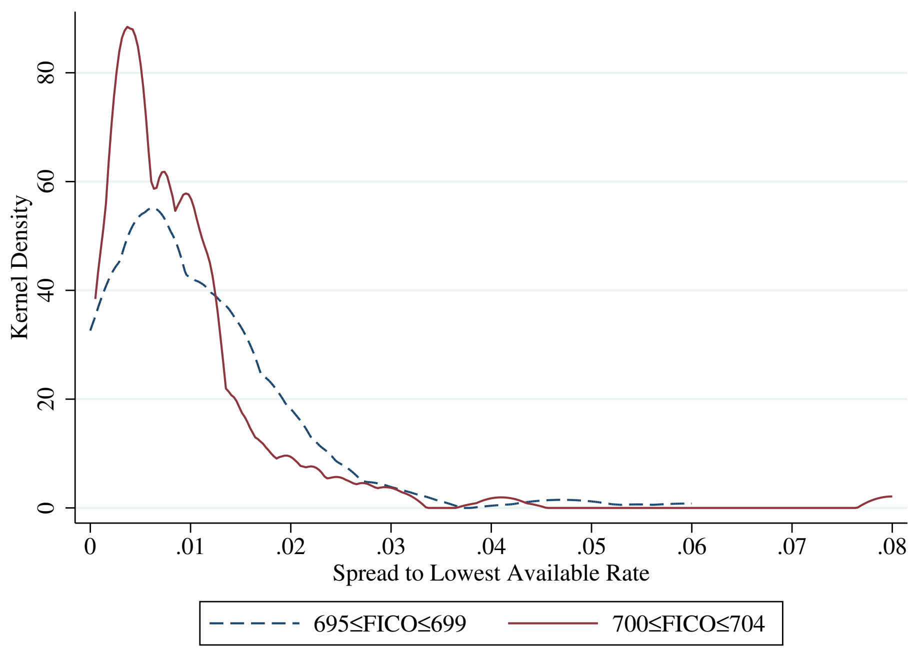

Notes: Figures plot the kernel densities of the spread between the originated interest rate for borrowers at lenders with FICO discontinuities at 600 and 700 in panels I and II, respectively, and the lowest rate available to similar borrowers in the same market. Dashed blue lines and solid red lines, respectively, plot the densities for borrowers with FICO scores just below and above each FICO threshold.

Figure 3: Map of Jacksonville, Florida Borrower Locations by Proximity to Financial Institutions

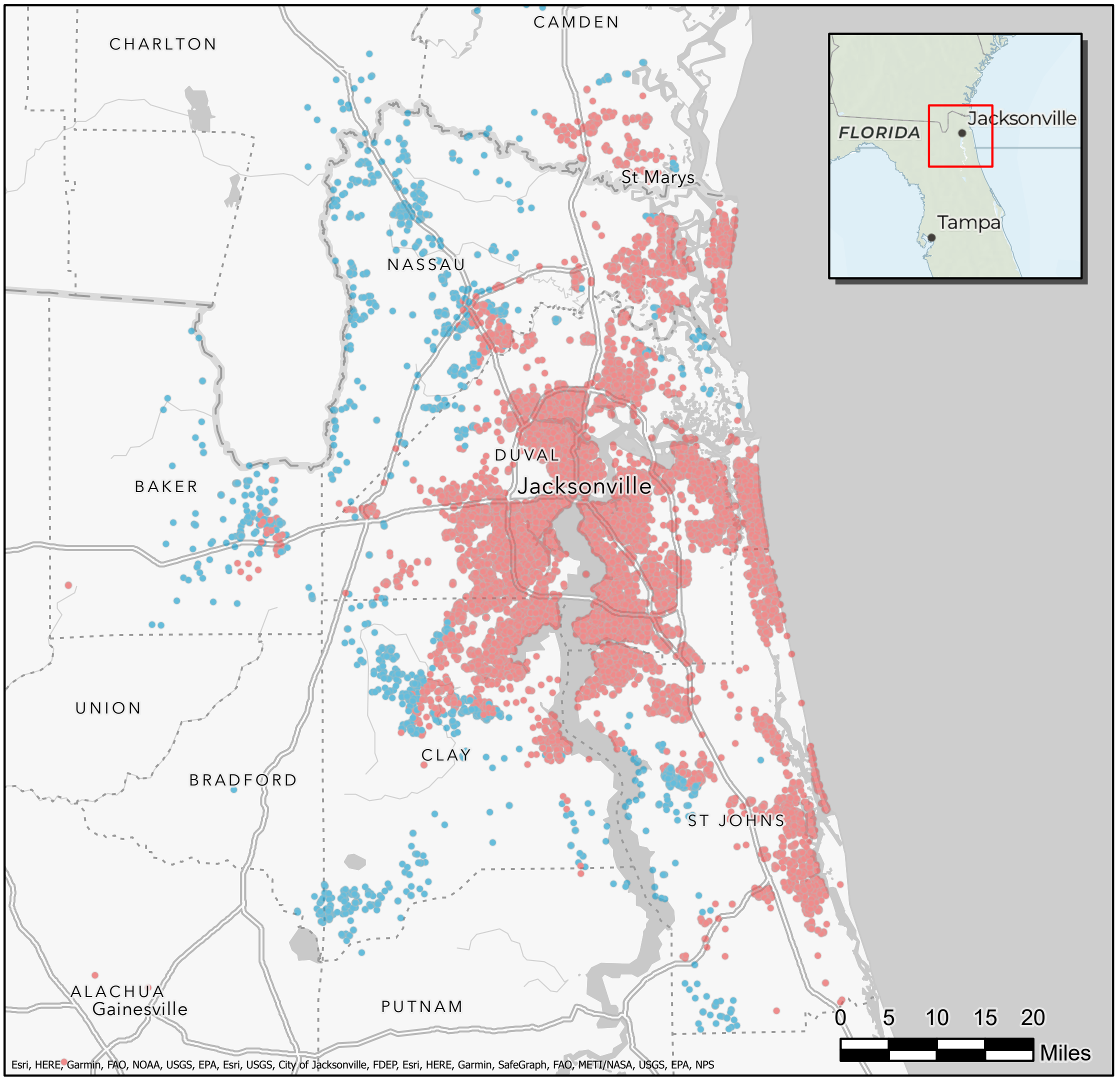

Notes: Figure plots geocoded borrower locations in our data in the Jacksonville, Florida Metropolitan Statistical Area. Red dots denote areas that have more than 10 banks or credit unions within a 20-minute drive, and blue dots denote locations with at most 10 banks or credit unions within a 20-minute drive.

Figure 4: Robustness of Take-up Rate Differentials to High Search Cost Definition

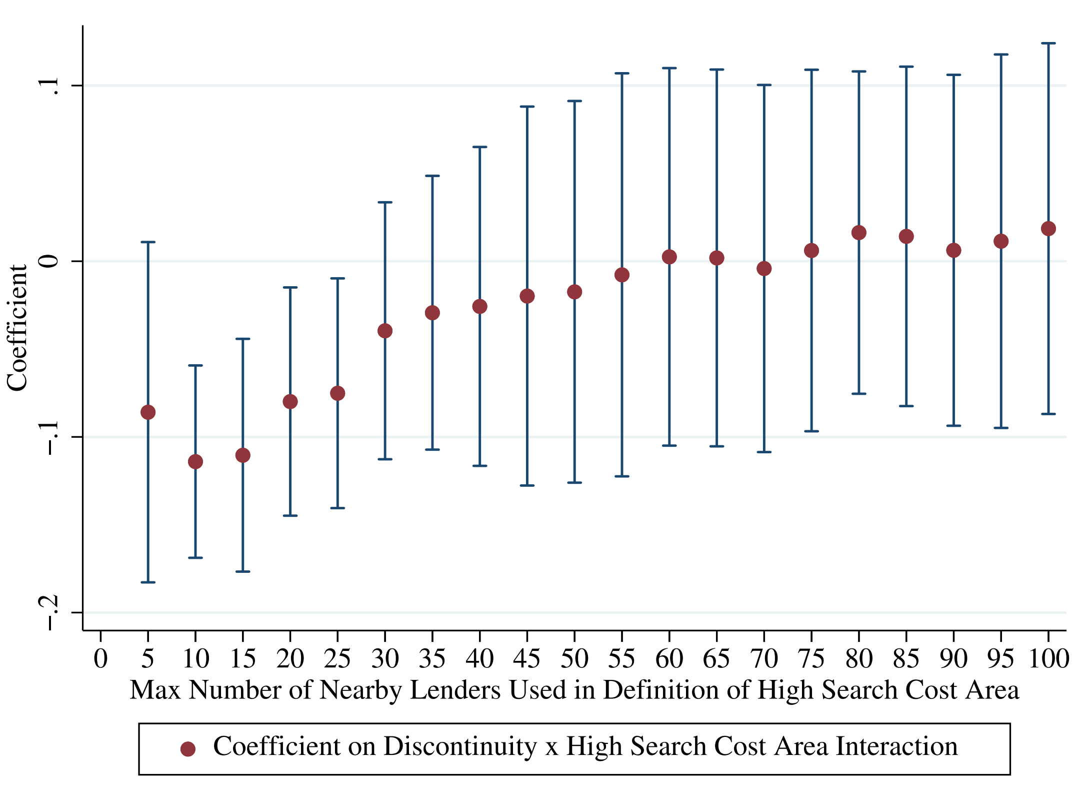

Notes: Figure reports estimates of the coefficient on the interaction of the discontinuity indicator ($\widetilde{FICO} \geq 0$) and a High Search Cost indicator in an RD regression of loan take-up using the specification in (14). High Search Cost is defined as an individual loan applicant having at most the number of lenders indicated on the x-axis within a 20-minute drive of her home at the time of loan application. Estimates condition on discontinuity $\times$ lender and Commuting Zone $\times$ origination-quarter fixed effects. 95% confidence intervals are double clustered by lender and FICO score.

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

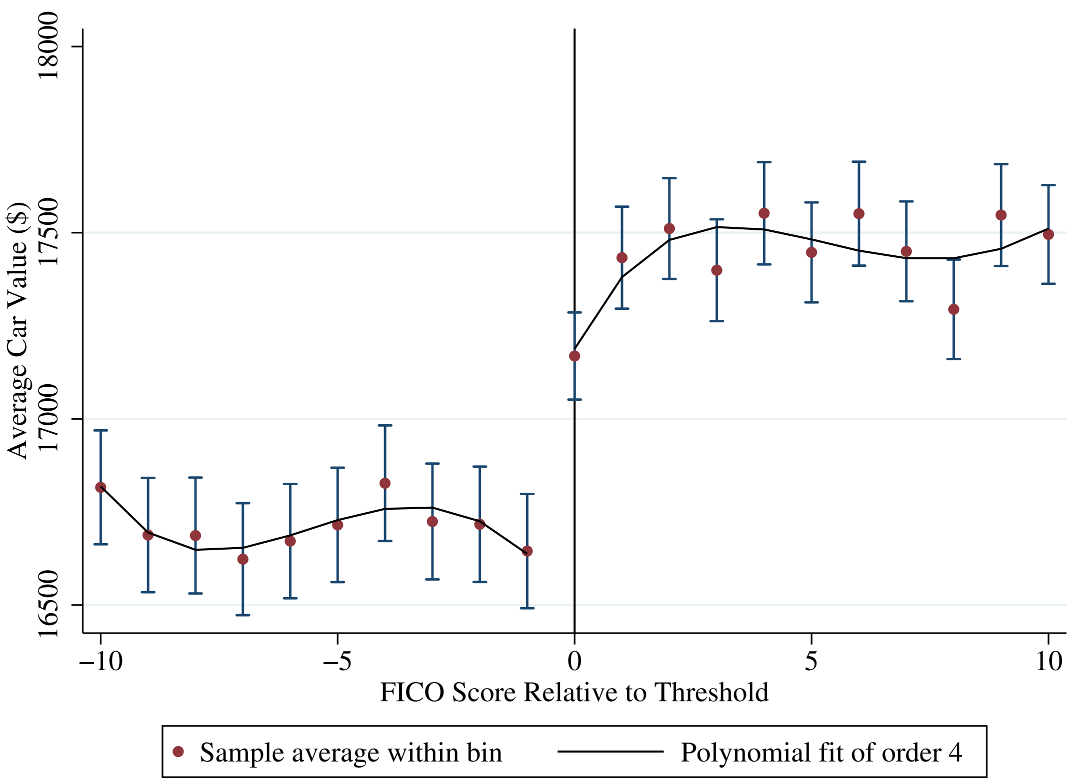

Notes: Figure plots average car transaction prices by FICO scores normalized to detected pricing discontinuities. 95% confidence intervals are double clustered by lender and FICO score. Plotted RD functions are estimated using the Calonico et al. (2014) robust RD estimator with fourth-degree polynomials for institutions with pricing discontinuities.

Table 1: Summary Statistics

	Count	Mean	Std. Dev.	Percentile 25th	Percentile 50th	Percentile 75th
			I. Loan Applications			
Loan Term (months)	1,119,153	67.25	24.43	60	72	72
Loan Amount ($)	1,320,109	21,927.3	11,660.7	13,296.0	20,000	28,932.1
Loan Rate	1,131,240	0.05	0.05	0.02	0.04	0.06
FICO	898,339	647.9	118.2	605	661	720
Debt-to-Income	833,854	0.26	0.3	0.13	0.27	0.39
Take Up	588,231	0.65	0.48	0	1	1
			II. Originated Loans			
Loan Rate	2,434,049	0.05	0.03	0.03	0.04	0.06
Loan Term (months)	2,434,049	62.73	22.08	48	60	72
Loan Amount ($)	2,434,049	18,136.52	10,808.97	10,094	16,034	23,892
FICO	2,165,173	710.55	74.89	661	714	770
Debt-to-Income (%)	1,276,585	0.25	0.32	0.05	0.26	0.37
Collateral Value ($)	2,434,049	19,895.13	10,929.1	12,046.81	17,850	25,562.28
Monthly Payment ($)	2,434,049	324.4	159.21	210.93	297.02	405.56
			III. Ex-Post Loan Performance Measures			
Days Delinquent	1,589,843	23.41	221.99	0	0	0
Charged-off Indicator	2,434,049	0.02	0.13	0	0	0
Default Indicator	2,434,049	0.02	0.14	0	0	0
Current FICO	1,719,848	705.5	83.28	654	714	772
%ΔFICO	1,697,700	-0.01	0.09	-0.04	0	0.03

Note: Panels I-III, respectively, report summary statistics for loan applications, originated loans, and ex-post loan performance. Loan Rate is the annual interest rate of the loan. Loan Term is the term (in months) of the loan. Debt-to-Income is the ratio of all debt service payments to income. Collateral Value is the value of the car at origination. Days Delinquent is the number of days since 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. 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.

Table 2: Effects of FICO Discontinuities on Loan Rate

	(1)	(2)	(3)	(4)	(5)
Discontinuity Indicator	-0.0076**	-0.0120**	-0.0131***	-0.0137***	-0.0128***
	(0.0038)	(0.0048)	(0.0042)	(0.0043)	(0.0043)
Discontinuity Indicator ×					-0.0018
High Search Cost					(0.0015)
RD Controls	✓	✓	✓	✓	✓
Discontinuity × Lender FEs		✓	✓	✓	✓
CZ × Quarter FEs			✓		
Zip Code × Quarter FEs				✓	✓
Number of Observations	514,834	514,834	514,834	514,834	514,834
R^2	0.05	0.05	0.22	0.46	0.22

Notes: Table reports regression-discontinuity estimates of equation (8) with the indicated fixed effects, normalizing FICO scores around each threshold using a uniform kernel and a bandwidth of 19 FICO points. Discontinuity Indicator is a dummy for whether a borrower's normalized FICO score is positive. High Search Cost is a dummy equal to 1 for borrowers with at most 10 financial institutions within a 20-minute drive from their home. RD controls consist of a quadratic spline in normalized FICO score that is allowed to change at the discontinuity as specified in (9). Robust standard errors in parentheses are double clustered by lender and FICO score.

Table 3: Exclusion Restrictions Tests of High Financing Search Cost Measures

	(1)	(2)
Dependent Variable	Search Cost Measure	
	High Search Cost	High $\widehat{\text{Search}}$ Cost
	I. Borrower Characteristics	
Credit Score	-0.267	1.220
	(0.636)	(1.624)
Debt-to-Income Ratio	-0.026	-0.009
	(0.031)	(0.013)
Minority Indicator	-0.020	-0.028
	(0.015)	(0.020)
Age (years)	1.98	-0.15
	(1.50)	(0.65)
Male Indicator	0.004	-0.010
	(0.027)	(0.023)
	II. Census Tract Characteristics	
Median Income	2251.05	1490.58
	(2185.06)	(911.92)
Poverty Share	-0.013	-0.002
	(0.010)	(0.006)
Average Rent	-34.39**	-4.40
	(13.88)	(16.45)
Log Wage Growth	0.115**	0.070
	(0.047)	(0.042)
Job Growth Rate	0.009**	0.006
	(0.003)	(0.006)

Notes: Table reports OLS estimates testing whether measures of high search costs predict borrower characteristics (panel I) and tract characteristics (panel II). Each row reports the coefficient on High Search Cost (column 1) or predicted High Search Cost (column 2) in a regression on the listed dependent variable. High Search Cost is a dummy equal to 1 for borrowers with at most 10 financial institutions within a 20-minute drive from their home at the time of their loan application. Predicted High Search Cost is an indicator for whether the predicted number of financial institutions within a 20-minute drive is at most 10, with the predicted number defined by (11) using 1990 local branch proximity and subsequent national trends in branching. Tract-level characteristics in panel II are compiled from Chetty et al. (2018). Median Income is tract median income from the 2015 ACS. Poverty Share is the share of tract residents below the poverty line in the 2006-2010 ACS. Average Rent is the tract average rent of two-bedroom apartments in the 2015 ACS. Log Wage Growth is the change in log of average hourly wages for high school graduates between the 2005-2009 and 2010-2014 5-year ACS. Job Growth Rate is the average annual growth rate from 2004-2013 of the number of jobs from the Census Bureau's LODES-WAC files. All specifications include zip code by quarter fixed effects. Robust standard errors in parentheses are double clustered by lender and FICO score.

Table 4: Effect of FICO Discontinuities on Loan Offer Take-up Decisions by Search Costs

	(1)	(2)	(3)	(4)	(5)
Discontinuity Indicator	0.114**	0.129***	0.142**	0.127***	0.141**
	(0.038)	(0.041)	(0.061)	(0.040)	(0.063)
High Search Cost		0.002	-0.008		
		(0.035)	(0.058)		
Discontinuity Indicator × High Search Cost		-0.108***	-0.145**		
		(0.030)	(0.048)		
High $\widehat{\text{Search}}$ Cost				-0.013	-0.004
				(0.025)	(0.037)
Discontinuity Indicator × High $\widehat{\text{Search}}$ Cost				-0.071***	-0.105***
				(0.016)	(0.033)
RD Controls	✓	✓	✓	✓	✓
Discontinuity × Lender FEs	✓	✓	✓	✓	✓
CZ × Quarter FEs	✓	✓		✓	
Zip Code × Quarter FEs			✓		✓
Number of Observations	30,743	30,743	30,743	30,743	30,743
R^2	0.31	0.31	0.61	0.31	0.61

Notes: Table reports regression-discontinuity estimates of whether a borrower accepts an approved loan offer regressed on a discontinuity indicator, a high-search-cost indicator, their interaction, and the indicated controls and fixed effects using the specification in (14). Discontinuity Indicator is a dummy for whether a borrower's normalized FICO score is positive. High Search Cost is a dummy equal to 1 for borrowers with at most 10 financial institutions within a 20-minute drive from their home. The predicted high search cost measure is an indicator for whether the predicted number of financial institutions within a 20-minute drive is at most 10, with the predicted number defined by (11) using 1990 local branch proximity and subsequent national trends in branching. RD controls consist of a quadratic spline in normalized FICO score that is allowed to change at the discontinuity as specified in (9). Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 2 for further details.

Table 5: Effects of FICO Discontinuities on Origination Outcomes

	(1) Purchase Price ($)	(2) Loan Amount ($)	(3) Loan-to-Value Ratio	(4) Monthly Payment ($)
Discontinuity Indicator	375.48**	557.37***	0.014***	-0.08
	(169.55)	(165.70)	(0.003)	(1.02)
Discontinuity × Lender FEs	✓	✓	✓	✓
CZ × Quarter FEs	✓	✓	✓	✓
Number of Observations	514,834	514,834	514,834	514,834
R^2	0.08	0.09	0.17	0.09

Notes: Table reports regression-discontinuity estimates of car-purchase and originated-loan characteristics regressed on the discontinuity indicator, our RD controls, and the indicated fixed effects following equation (8). Discontinuity Indicator is a dummy for whether a borrower's normalized FICO score is positive. Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 2 for more details.

Table 6: Effects of FICO Discontinuities on Vehicle Purchase

	(1) Purchase Price	(2) Purchase Price	(3) Car Age (months)
Discontinuity Indicator	344.38***	71.76	-1.81***
	(120.43)	(48.43)	(0.143)
Discontinuity × Lender FEs	✓	✓	✓
CZ × Quarter FEs	✓	✓	✓
Make-Model FEs	✓		✓
Make-Model-Year FEs		✓	
Number of Observations	468,800	468,800	468,800
R^2	0.38	0.78	0.37

Notes: Table reports regression-discontinuity estimates of car purchase prices (columns 1–2) and car age in months (column 3) regressed on the discontinuity indicator, our RD controls, and the indicated fixed effects following equation (8). Columns 1 and 3 include make × model fixed effects and column 2 includes make × model × model-year fixed effects. Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 2 for more details.

Table 7: Balance of Ex-Post Credit Outcomes Across FICO Discontinuities

	(1) Days Delinquent	(2) Charge-off	(3) Default	(4) %ΔFICO
Discontinuity Indicator	3.55	0.004	0.001	0.002
	(2.75)	(0.002)	(0.003)	(0.003)
Discontinuity Indicator ×	4.55	-0.004	-0.004	0.002
High Search Cost	(4.96)	(0.004)	(0.004)	(0.003)
Discontinuity × Lender FEs	✓	✓	✓	✓
CZ × Quarter FEs	✓	✓	✓	✓
Number of Observations	331,590	514,834	514,834	405,236
R^2	0.49	0.21	0.24	0.06

Notes: Table reports reduced-form RD estimates of equation (8) 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. Percentage change in FICO score is the change in FICO score since origination as a fraction of the FICO score at origination for the subsample of institutions that report credit scores after loan origination. Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 2 for estimation details.

Appendix

A Model of Credit Search with Continuous Demand

In this appendix, we extend a sequential search model to illustrate theoretically an example of how credit-market search frictions can reduce demand for credit and goods and lower aggregate welfare.[61] As emphasized above, much of the modern empirical search literature—including recent work on consumer credit—assumes inelastic consumer demand. Instead, we extend the elastic-demand search model of Reinganum (1979) to a credit-market setting.[62] We then demonstrate that in this model, an increase in search costs leads to higher interest-rate markups, smaller loan sizes, and less durable consumption being purchased. Finally, we highlight the welfare loss from search frictions that is obscured by models which assume inelastic demand. Applying this model to our specific empirical setting, we map our identification strategy to objects in the model to generate a set of empirical predictions.

Borrowers A set of borrowers with measure 1 are ex-ante identical with indirect utility U that is quasi-linear in wealth W and depends on the prices of loans r and durable goods p.

$$U(r, p, W) = V(r, p) + W, \tag{15}$$

where $V(\cdot, \cdot)$ is the indirect utility of consuming the consumption bundle of loan and durables chosen when facing prices r and p.[63] However, instead of inelastic demand, we assume that demand for loans and durables each slope downward such that $V(\cdot, \cdot)$ is strictly decreasing in both its arguments. Further distinct from many search models, note that we do not implicitly assume the cross-price elasticities to be zero. In our setting, for example, it is likely that car loans and car services are strong complements. Without loss of generality, we model the demand curves for loans and durables resulting from borrower utility maximization as having constant own- and cross-price elasticities. This means individual demand for loans q can be represented as

$$q(r, p) = a \cdot r^{\eta_r} \cdot p^{\eta_{rp}}, \tag{16}$$

and demand for durables x can be represented as $x(p, r) = b \cdot p^{\eta_p} \cdot r^{\eta_{pr}}$.[64] Assuming loans and durables are normal goods (i.e., with positive income elasticities), η_r and η_p will be negative given that $V(\cdot, \cdot)$ is downward sloping in demand such that as interest rates or prices fall, borrowers will take out larger loans q and consume more durables x. Given that the set of borrowers is measure 1, $q(\cdot, \cdot)$ and $x(\cdot, \cdot)$ also represent market demand curves.

Borrower Search Borrowers know the distribution of interest rates $F(\cdot)$ on a closed interval of support $[\underline{r}, \overline{r}]$ but not the precise location of each rate. Borrowers can access every lender by drawing uniformly from the distribution of lenders but incur a search cost k for each quote they

[61] To focus on the key friction we study, the model abstracts away from search frictions in the car market, which could interact with financing search frictions. Note, however, that some degree of financing search frictions are necessary to rationalize the empirical patterns we document.

[62] In the model of Reinganum (1979), elastic demand combined with cost heterogeneity on the seller side is sufficient to support price dispersion in equilibrium even without Stahl-like (1989) heterogeneity in search costs. Because of elastic demand, sellers still offer dispersed prices even though sellers know buyers' search costs and reservation prices.

[63] As in other models of search for financing, we abstract away from credit risk for clarity.

[64] We represent durables as a composite good x mapping all durable goods onto a scalar measure of durable consumption. This quantity represents durable services and includes both the flow utility of durable consumption and its resale value.

obtain. When their current quote is r', the expected utility gain of one additional search is

$$\int_{\underline{r}}^{r'} [V(r,p) - V(r',p)]dF(r) - k. \tag{17}$$

As shown by DeGroot (1970) and Lippman and McCall (1976) in this setup, optimal search behavior will constitute a reservation price $m(k)$ above which borrowers will be unwilling to take out a loan and below which, demand will be given by (16). Absent any cost of obtaining an additional price quote, all borrowers would be able to find the lowest-cost provider offering $\underline{r}$.

Characterizing the gains from search using the indirect utility function in (17) is important. When there is elastic demand and negative cross-price elasticities, estimators that infer k purely from price dispersion (e.g., Hong and Shum, 2006) will be biased estimates of true search costs. For example, in our setting, replacing the indirect utility terms $V(r,p)-V(r',p)$ in the integrand of (17) with $r - r'$ will implicitly assume that the only consequence of failing to search is paying a higher markup. However, when demand is elastic, the utility loss associated with a given markup will be different from the size of the markup itself. Intuitively, given that the consequence of a higher price paid for one good also includes the utility loss from lower consumption of complementary goods, the search cost need to rationalize equilibrium markups is larger than would be suggested by the size of that markup in isolation. In our specific setting, inferring search costs simply from the distribution of markups would fail to account for the fact that the borrower would be paying both the interest rate markup and the disutility of a cheaper and older car.

Lenders A continuum of lenders indexed $j \in J$ with J having measure 1 are otherwise identical but each have constant marginal costs c_j per \$1 lent, which are continuously distributed according to distribution function $G(\cdot)$ with positive density almost everywhere on a closed interval of support $[\underline{c}, \overline{c}]$ and $\overline{c} > \underline{c}$.[65] Without loss of generality, assume $\overline{c} \leq \underline{c}\eta_r/(1+\eta_r)$.[66] Lenders are perfectly informed of k , $m(k)$, and $F(\cdot)$. Their objective is to choose an interest rate r_j to maximize expected profits π_j, given by

$$\max_{r_j} E(\pi_j) = \begin{cases} (r_j - c_j)q(r_j,p)E(N_j) & \text{for } r_j \leq m(k) \\ 0 & \text{for } r_j > m(k) \end{cases} \tag{18}$$

where N_j is the mass of borrowers that each take out $q(r_j)$ in debt from lender j when facing r_j. We note that this model does not rationalize the discontinuous pricing policies employed by roughly half of the lenders in our sample. See section 4.1 for a discussion of theoretical and empirical work that explain the supply-side cost structures that have resulted in binned pricing functions by many lenders.

Equilibrium This setup lends itself to a pure-strategy Nash Equilibrium defined by a profit-maximizing lender pricing rule r_j^*, an optimal borrower reservation price $m(k)$, and a distribution of interest-rates $F_{m(k)}(r)$ generated by firms' optimal pricing rule with credit-market price dispersion

[65] We assume a continuum of lenders such that perfect competition would prevail in the absence of search costs. In our setting, it is straightforward to imagine lenders with heterogeneous costs of lending based on heterogeneity in deposits, etc.

[66] Below, we show that this condition ensures that all lenders participate in the market in equilibrium. Otherwise, some lenders would exit because their marginal costs would exceed borrowers' reservation price, in which case the statement could be made about the support of remaining lenders' costs.

arising from the cost heterogeneity of firms and the elasticity of demand.[67] Here, we offer a sketch of the proof as it closely follows Reinganum (1979).

Fixing p, the borrower search indifference condition pins down the reservation price $m(k)$, which satisfies

$$\int_{\underline{r}}^{m(k)} [V(r,p) - V(m(k),p)]dF_{m(k)}(r) = k, \tag{19}$$

where $F_{m(k)}(r)$ is the endogenous distribution of interest rates given reservation price $m(k)$. Note that (19) means that in equilibrium, the reservation price will depend not only on search costs but also in how interest rates paid affect the utility received from the corresponding loan sizes and durable consumption through $V(\cdot,\cdot)$.

By DeGroot (1970), there exists an optimal reservation price in this setting given firms' prices corresponding to optimal borrower search behavior. In equilibrium, firms will not charge above this reservation price and, given the borrower search indifference condition (19), markups will adjust such that each borrower will be indifferent between searching again and accepting the first randomly queried rate quote. Accordingly, borrowers will be uniformly distributed across lenders and, given the relative measures of the sets of borrowers and lenders, $E(N_j) = 1$. The lender's problem is then

$$\max_{r_j} E(\pi_j) = \max_{r_j} (r_j - c_j) a r_j^{\eta_r} p^{\eta_{rp}} \tag{20}$$

whenever $r_j \leq m(k)$. Holding durables prices p fixed, the first-order condition for (20) is satisfied when $r_j = c_j \eta_r/(1+\eta_r)$. However, because lenders face zero demand and make zero profits when offering an interest rate $r_j > m(k)$, lenders charge

$$r_j^* = \begin{cases} \frac{c_j \eta_r}{\eta_r + 1} & \text{if } c_j\eta_r/(1+\eta_r) < m(k) \\ m(k) & \text{if } c_j\eta_r/(1+\eta_r) \geq m(k) \end{cases}.$$

Given that $\bar{c} \leq \underline{c}\eta_r/(1+\eta_r)$, all firms will make positive profits and no firms will exit. This pricing rule induces a continuous distribution of equilibrium interest rates in the market almost-everywhere on $[\underline{r}, m(k)]$ with a point mass at $r = m(k)$ and a CDF given by

$$F_{m(k)}(r) = \begin{cases} G[r(1+\eta_r)/\eta_r] & \text{for } \underline{r} \leq r < m(k) \\ 1 & \text{for } r = m(k) \end{cases}. \tag{21}$$

Given the distribution of interest rates (21) resulting from firms' best response to borrowers' reservation price, (19) holds, borrowers will not deviate from their reservation prices, and borrowers will have insufficient incentive to search further given search costs k. Each borrower will originate a loan of size $q(r_j^*, p)$ with j drawn randomly from J. Then for a given search cost k, the reservation price, the pricing rule, and the interest-rate distribution $(m(k), r_j^*, F_{m(k)}(\cdot))$ constitute an equilibrium.

Welfare Consequences of Costly Search Unlike standard search models with inelastic demand, the equilibrium markups described above will be characterized by an aggregate welfare loss relative to the zero-search-frictions case. The deadweight loss has three components in our setting. First, because of each lender's monopoly power, Bertrand pricing fails and lenders other than the lowest-cost lender are able to survive in equilibrium. Second, each lender is able to mark up her cost c_j to charge monopoly prices $r^*(c_j)$. Third, because of elastic demand, each borrower demands

[67] Note that while many of the predictions of our model could be generated by a model of oligopolistic competition, such models do generally not feature price dispersion.

less than she would if she faced $r = \underline{c}$. The deadweight loss is the difference between the surplus under the first best and the surplus under the costly search equilibrium and is given by

$$DWL = \int_{\underline{c}}^{\overline{c}} \int_{q(r^*(c),p)}^{q(\underline{c},p)} (r(q) - \underline{c})\, dqdG(c) + \int_{\underline{c}}^{\overline{c}} \int_{0}^{q(r^*(c),p)} (c - \underline{c})\, dqdG(c) \tag{22}$$

where $r(q)$ is inverse demand, $q^m(c, p)$ is the quantity lent by a monopolistic lender with constant marginal cost c, and $q^*(\underline{c}, p)$ is the quantity that would prevail under perfect competition with interest rates set to the lowest cost provider's marginal cost $\underline{c}$. The outer integral of each term aggregates the deadweight loss over all lenders' marginal costs, distributed according to measure $G(\cdot)$.

Appendix Figure A7 illustrates the division of surplus and welfare loss given a single lender with cost c_j charging the monopolist's price $r_j(c_j)$ using a logarithmic scale for ease of exposition. Under the first best, lenders other than the lowest-cost provider exit and $r = \underline{c}$ such that borrower surplus is the triangle bounded by log inverse demand $\log r(q)$ and the horizontal line $\log r = \log \underline{c}$. Without search costs, given the large number of lenders, perfect competition would prevail and there would be no lender surplus. Under the costly search equilibrium, borrowers arriving at lender j face an interest rate of $r_j(c_j)$, demand $q(r_j(c_j), p)$, and have the borrower surplus labeled BS in Appendix Figure A7. Under costly search, lenders would earn a lender surplus of monopoly rents denoted LS in Appendix Figure A7 equal to $\log q(r_j(c_j), p) \times (\log r_j(c_j) - \log c_j)$ where the first term is the monopolist's markup and the second term is the monopolist's quantity. This non-zero lender surplus under costly search provides a positive rationale for sellers to endogenously differentiate, obfuscate, or otherwise increase switching or search costs (see also Ellison and Ellison, 2009; Allen, Clark, and Houde, 2019; Adams et al., 2021).

The first term in (22) is the usual Harberger's triangle bounded by the demand curve and the marginal cost curve between the monopolistic equilibrium quantity and the efficient quantity. The inside integral of this term is labeled DWL_A in Appendix Figure A7 and represents the deadweight loss from loan sizes that are smaller than the first best given marginal costs c_j because of search-cost–induced markups. The outside integral of this first term adds the region denoted DWL_B in Appendix Figure A7, which is the additional surplus lost from the first-best quantity and price being even higher and lower, respectively, than when the borrowers can find the lowest-cost lender. Note that when demand is inelastic as in many search models, $q(\underline{c}, p) = q(r^*(c), p)$ for all $c \in [\underline{c}, \overline{c}]$ such that there is no aggregate welfare loss from too little quantity being consumed due to search frictions, simply lost consumer surplus from the transfer from buyers to sellers as in Allen, Clark, and Houde (2019).

The second term in (22) is the loss of surplus from the inefficiency of lenders other than the lowest-cost provider lending $q(r^*(c), p)$ and is denoted DWL_C in Appendix Figure A7. Because interest rates would be lower in the first-best equilibrium, even the inframarginal $q(r^*(c), p)$ dollars that are lent out in both the first best and the costly search equilibrium provide lower aggregate surplus under costly search. Some of this would be a transfer from borrowers to lenders $(r_j(c_j) - c_j)$ and is the source of the lender surplus, but some of it would be lost altogether $(c_j - \underline{c})$.

A.1 Comparative Statics and Testable Predictions

We use the context of the model above to establish several comparative statics that will serve as predictions for our empirical work.[68] Notably, several of the predictions below are inconsistent with models of imperfect competition arising from market concentration. In particular, while high market concentration can sustain markups in many equilibrium models (and thereby affect complementary

[68]In what follows, we hold the unit price of durables p fixed.

demand), such models do not generally feature price dispersion across sellers for a single borrower type. Conversely, search costs need not generate price dispersion—see, for example, the Diamond (1971) Paradox, wherein all sellers charge the monopoly price. However, even in models where search does not cause price dispersion, the comparative statics below that show quantity distortions from the combination of markups and elastic demand would still hold. Still, that we do find price dispersion is useful for distinguishing search and market concentration explanations for our findings.

Price dispersion and loan markups increasing in search costs If search costs are exogenously higher for all borrowers in a market, then the amount of price dispersion rises because $m(k)$ is increasing in k. To see this note that for $k' > k$, the corresponding reservation price $m(k')$ must still satisfy

$$\int_{\underline{r}}^{m(k')} [V(r,p) - V(m(k'),p)] dF_{m(k')}(r) = k'. \tag{23}$$

Given that $V(r,p) - V(m(k),p) > 0$ since V is decreasing in its first argument and $r < m(k)$, this equation will only be satisfied for $k' > k$ for $m(k') > m(k)$. Price dispersion in this model is therefore higher in markets where search costs are higher. Given that a smaller range of pure monopolistic markups r_j^* are censored by the reservation price, this means average markups are weakly increasing in search costs and strictly increasing for a strictly positive mass of lenders with costs satisfying $m(k) < c_j\eta_r/(1+\eta_r)$, i.e., costs sufficiently high to be censored by reservation price $m(k)$.[69]

In the empirics above, we proxy for the search costs k_g of consumers in market g using the density of nearby lenders. For markets characterized by a low number of nearby lenders for the typical consumer, we hypothesize search costs will be higher. Our model therefore predicts that price dispersion will be greater and average loan markups will be higher in such areas relative to areas with a large number of nearby lenders. We test this prediction using our measure of search costs. In Appendix Figure A8, we plot kernel density functions of the spread to lowest available rate for each borrower in our data separately for borrowers in high- and low-search-cost areas. The solid line is the estimated probability density function for borrowers facing higher search costs and the dotted line is for lower-search-cost borrowers. The plots indicate that a smaller fraction of high-search-cost borrowers accept loans with lower spreads to the lowest available rate and a higher fraction of lower-search-cost borrowers accept larger spreads. Appendix Table A7 shows that the mean spread for higher-search-cost borrowers is 27 bp higher than the mean spread for lower-search-cost borrowers. The standard deviation of spreads is also 22 bp higher in high-search-cost areas. Although these magnitudes are likely attenuated by the limited coverage of our data relative to the entire local auto-loan market, Kolmogorov-Smirnov tests confirm that the distributions are statistically different at the 1% confidence level. An additional implication of Appendix Figure A8 is that variation in the benefits to search cannot explain less search behavior in high-search-cost areas; if anything, borrowers in such areas have *higher* returns to search.

Loan sizes decreasing in search costs Given that credit demand $q(r,p)$ is decreasing in r, the increase in markups generated by an increase in search costs decreases loan sizes. Accordingly, in markets with fewer and geographically dispersed lenders, our model predicts that loan sizes should be smaller.

[69] One caveat in this setting is that if the reservation price is never binding, i.e., $m(k) > \bar{r}\eta_r/(1+\eta_r)$, then markups will be invariant to k and will depend only on demand elasticities.

Durables consumption decreasing in search costs Given that the demand for durables $x(p, r)$ is also decreasing in r, the increase in markups generated by an increase in search costs decreases durable consumption. Thus, in higher search-cost places, we predict that purchased car quality should be lower.

Welfare loss increasing in search costs and the elasticity of demand The expression for deadweight loss in (22) predicts that the aggregate welfare loss from search frictions will be increasing in search costs. Given that demand is strictly downward sloping in interest rates, we have that inverse demand $r(q) > \underline{c}$ whenever $q < q(\underline{c}, p)$ such that the inner integrand in the first term of (22) is positive over the limits of the integral. Because each lender's markups are weakly increasing in search costs, $q(r^*(c), p)$ weakly decreases as k increases, and the first term in the deadweight loss expression grows. Intuitively, as search costs rise, lenders are able to charge larger markups, and because borrowers dislike both high interest rates and the lower levels of durable consumption they end up with in equilibrium when interest rates are high, utility falls. By the same argument, the second term of the (22) is decreasing in k. However, because $r(q) > c$ for all q, the first term increases by more than the second term decreases.

The amount of welfare loss is also increasing in the elasticity of demand. As η_r increases, demand becomes more sensitive to changes in prices and the gap between the limits of the inner integral in the first term (22) grows. Moreover, the stronger the complementarity between credit and durables, the larger the welfare loss from search. Given that demand for cars falls when interest rates are higher, the cross-price elasticity of demand for cars with respect to interest rates η_{pr} is negative. This means that indirect utility decreases with markups both because borrowers dislike high interest rates and because they like durables. Accordingly, as η_{pr} increases in magnitude, the costliness of higher interest rates is a larger drop in durable consumption resulting in a larger deadweight loss.

Market shares invariant to markups when search costs are high When search costs are sufficiently high, the Nash equilibrium described above will entail myopic shopping with consumers borrowing from the first lender they query. Market shares will thus be similar across lenders and invariant to markups. When search costs are sufficiently low, lenders with higher markups will be punished with lower market shares. In the limit of perfect competition, lenders with positive markups will have zero market share. We note, however, that in practice other dimensions of product differentiation may prevent the exit of higher-markup lenders. For example, our model provides a positive rationale for firms to endogenously undertake obfuscation efforts as in Ellison and Ellison (2009) or invest in brand loyalty effects as in Allen, Clark, and Houde (2019) to inhibit search across products and capture the lender surplus depicted in Figure A7.

B Documenting Price Dispersion

While our main results rely on quasi-exogenous variation for identification, in this section, we provide suggestive evidence that significant interest-rate dispersion persists in the market for car loans even after nonparametrically controlling for many dimensions of product heterogeneity. Although such price dispersion is a feature of many search model formulations, empirically diagnosing price dispersion requires ruling out any product differentiation to establish that differences in prices represent identical goods being sold for different prices in the same market.

For any given borrower with an observable set of attributes, we estimate the spread between that borrower's interest rate and the lowest available interest rate at another lender in our data for another borrower with very similar attributes. To calculate this spread, we group borrowers in the same Commuting Zone (CZ), six-month transaction date window, five-point FICO bin, \$1,000 purchase-price bin, loan maturity, and 10 percentage-point DTI bin.[70] We consider used-car loans originated to borrowers within the same CZ × quarter × price × FICO × maturity × DTI cell to be observably identical.[71] Although there may be some degree of residual heterogeneity within a cell, the magnitude of the variation we find is sufficiently large that it would be difficult to explain solely with remaining borrower-level heterogeneity within these borrower types.[72] In particular, in roughly half of borrower-type cells, the best rate in the cell is achieved by a borrower with a lower FICO and higher DTI than other borrowers in the cell, suggesting that any coarseness in our borrower typology cannot explain the residual rate variation. Moreover, our RD design below also establishes the existence of large pricing disparities for arbitrarily similar credit risks. Note, too, that 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 lowest available rate for each borrower type.

Appendix Figure A9 plots the density of the spread to the best available rate in percentage points for the 54% of borrowers who did not attain the best rate in their cell. The mean and median of this distribution are 234 and 125 basis points (bps), respectively. Including the 46% of borrowers who are getting the best available rate given their borrower type, the average borrower in our data is thus paying a 1.3 percentage point higher interest rate than an observationally equivalent borrower at the same time in the same place. Simulating random markup draws from the distribution implied by the density in Appendix Figure A9, we find that the average borrower would need to obtain three price quotes to find an offer within 10 bp of the best available rate for that borrower's type.

C Inferring Loan Shopping from Multiple Applications

To perform a more direct test for whether our search cost measure actually predicts loan shopping behavior, we construct two measures of search for financing: one from national mortgage applications data and one using our own data on auto loan applications. First, we take Home Mortgage Disclosure (HMDA) data from 2012-2016 on mortgage applications per origination at the census tract level. We collapse our borrower-level count of nearby lenders data to define tract-level measures of the number of institutions within a 20-minute drive for the average borrower in the tract. For the second measure of search behavior, we exploit the ability of our data to identify the same borrower filing a loan application at multiple financial institutions. In both exercises, we define high-search-cost areas to be places where borrowers live within a 20-minute drive of at most 10 lending institutions, and we then evaluate whether we are more likely to observe borrowers in low-search-cost areas submitting more loan applications than borrowers in high-search-cost areas.

Every prospective borrower in our application data, by construction, applies for at least one loan. Panel I of Appendix Table A8 documents that using the HMDA data, we observe that applicants complete an average of 1.08 additional applications per origination in high-search-cost tracts (column 1) in contrast to 1.24 additional applications, on average, in low-search-cost tracts. Turning to our individual-level auto data in panel II, borrowers in high-search-cost areas (column 1) apply for an average of 0.34 additional loans. By contrast, we observe applicants in low-search-cost areas (column 2) applying for an average of 0.41 additional loans per vehicle purchase. Column 3 shows that these differences are statistically significant at a 1% significance level for both panels I and II. These estimates suggest that one in every 5-15 borrowers applies for an extra loan only if

[70]While we condition on purchase prices to homogenize loans as much as possible, this price dispersion exercise is robust to not conditioning on purchase prices, and our main empirical strategy does not condition on prices.

[71]While many borrowers in our data are in singleton cells (for whom we cannot calculate price dispersion) because of the strictness of our matching criteria, the richness of our data coverage across hundreds of providers provides us with thousands of cells with multiple borrowers.

[72]We address the issue of some credit unions restricting membership to narrowly defined groups in Appendix D.4.

in a low-search-cost area. For panel II, these average differences are a lower bound as we do not observe applications to any lender not in our data. In light of our data covering only a small fraction of all lenders, we view any detectable difference in additional applications to be additional positive evidence of differential search behavior. Note, too, that equilibrium applications need not change for outcomes to nevertheless be affected by search frictions—in a sequential search equilibrium (including our model above), consumers always buy from the first seller in equilibrium. Regressions of loan applications per applicant on the count of lending institutions within a 20-minute drive also confirm a positive and significant relationship. These results inferring loan shopping behavior from submitting multiple loan applications are consistent with the evidence presented in section 6 that borrowers facing high search costs search less and accept worse rates than borrowers facing relatively lower search costs.

D Identification Appendix

In this appendix, we discuss additional discussion and robustness checks that support our interpretation of the results in the main body of the paper.

D.1 Distinguishing Search Costs from Market Concentration

Proxying 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 market concentration that determine the degree of price competition among lenders. Moreover, some of the predictions of the model in Appendix A would also be consistent with more imperfect competition arising from market concentration. In particular, markups being sustainable in equilibrium and these markups affecting own-good quantities and complementary consumption would also be a feature of a concentrated market.

We first reiterate that in a standard oligopolistic-competition setup, there would not be price dispersion or differential effects on take-up or search behavior across markets. Appendix Figure A10 supports this argument by plotting the relationship in our data between markups and market shares in high- and low-search-cost markets. Markups are estimated as lender fixed effects in a regression of loan-level interest rates on controls for a FICO cubic, LTV, loan size, loan term, and quarter fixed effects. Market shares are calculated at the lender $\times$ Commuting Zone $\times$ quarter level. In low-search-cost markets (where the density of nearby lenders is high), charging a higher markup corresponds to a lower market share. In high-search-cost areas, market shares are statistically unrelated to markups.

Appendix Table A9 confirms that in a multivariate regression of log market shares on average markups with commuting zone fixed effects and lender fixed effects, we cannot reject a zero effect of higher markups on market shares for high-search-cost markets. For low-search-cost markets, the semi-elasticity of market share with respect to interest-rate markups in column 3 is -7, meaning that for every 100 bp higher markup a lender charges, it can expect to lose 7% market share. These results confirm the last prediction of our theoretical model in section A.1 that market shares should be unrelated to markups when search costs are high.

To further differentiate between search costs and a market concentration story, we construct empirical measures of lending competition within CZs. We calculate the share of originated mortgage loans by each HMDA lender within a given CZ and use the origination shares to construct a CZ-level Herfindahl index to capture the idea that two locations with identical branching networks could face differing degrees of competition based on the distribution of market shares across branches.[73] Dividing loans into high and low (above and below median) competition areas based

[73]See Scharfstein and Sunderam (2016), Drechsler, Savov, and Schnabl (2017), and Liebersohn (2019) for recent

on our constructed Herfindahl index, we reestimate the take-up RD specification of Table 4 for all four combinations of high and low search cost and market concentration combinations.

The results of this exercise in Appendix Table A10 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. Even for markets with a highly competitive banking sector, borrowers in high-search-cost areas are much more likely to accept dominated loan terms. Within low-competition CZs in low-search-cost areas, the difference in take-up rates around lending thresholds is 14 percentage points; borrowers quasi-randomly assigned high RD rate markups are 14 pp more likely to walk away from a loan offer. In comparison, borrowers in high-search-cost areas in the same low competition bin do not appear to respond to markups at the extensive margin—the point estimate is only marginally significant but suggests that such borrowers are more likely to accept a likely dominated loan. Within high-competition markets, we find similar results: borrowers with lower search costs are 10 pp less likely to accept a firm specific markup than borrowers in high-search-cost areas. These results suggest that regardless of the overall level of market concentration, borrowers in areas we expect to have high search costs are much less sensitive to interest rates in their extensive-margin loan take-up decisions and thus more affected by firm-specific pricing.

Finally, these results are also useful in ruling out more nuanced alternative explanations for our results based on a combination of uniformly costly search and market structure differences correlated with our search-cost measure. For example, in many structural search models, the number of sellers directly affects the equilibrium price distribution. Might our results contrasting areas with many or few nearby lenders be driven by variation in the number of sellers in the local market and not by any real differential in borrower search costs? If so, our inference that areas with few nearby lenders can proxy for high search costs would be undermined. However, we view such an explanation as unlikely for two reasons. First, if search were uniformly costly, then we would expect a given sized markup to have a similar market-share consequence in markets with many and few nearby lenders. Instead, we see high markups punished with lower market shares only in markets with fewer nearby lenders and find our main results to hold across levels of market concentration within markets with many or few nearby lenders. Second, we observe differential search behavior across areas with many or few nearby lenders, both directly in application counts and indirectly in take-up rates. Similar to the markup and market-share results, if market structure and not differential search costs were driving our results, we would expect to see similar search behavior (or lack thereof) across both types of local markets.

D.2 Online 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. One large lender in our sample suggests that formal digital search (actually filing out an application to receive an interest-rate offer online) is less common than might be expected: only 8% of their total applications are digital. While credit-union clientele skew slightly older than the general population, another potential explanation for the ability of physical search measures to explain variation in search propensity is that, although borrowers can be easily preapproved on the internet, the actual closing of loans (signing documents, transfer of title, etc.) still most frequently occurs at physical branch or dealer locations. We also note that the literature on online search has not found e-commerce to be a panacea driving search costs and price dispersion to zero (see, e.g., Ellison and Ellison, 2009).

The option to search online could affect the interpretation of our empirical results in a few ways.

alternative measures of retail banking market concentration.

First, our price dispersion results are lower bounds given the possibility of digital lending given that each calculated spread from a given borrower's rate to the best available rate could only be weakly higher if taking into account other lenders' rates. On representativeness, there still may be a large category of borrowers not in our data that are unaffected by nearby lender density (although they may well be affected by other types of search costs). However, our results provide positive evidence of the persistent important of physical search costs for millions of borrowers. Ultimately, given trends in digital banking and the possibility of online search for credit, we view our results with respect to nearby lender density and loan search even more noteworthy.

D.3 Testing Quasi-Random Assignment

To test whether other observables beside loan terms also change discontinuously at our detected FICO thresholds, we plot the average value of other borrower and loan application characteristics by FICO score normalized to each threshold. In Appendix Figure A2, we plot these values and associated 95% confidence intervals along with a fourth-order polynomial robust RD function estimated following Calonico et al. (2014).[74] These graphs are constructed with loan application data in order to ensure that borrowers are similar along observable characteristics around FICO thresholds at the time of application. Panels I–V plot borrower DTI ratios, loan amounts, borrower age in years, borrower gender (an indicator for male), and imputed 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 demographics. Although statistically insignificant, panel II seems to show that borrowers with credit scores immediately at a FICO threshold apply for larger loans. That this effect seems to hold only for $\widetilde{FICO} = 0$ and not at higher credit scores is inconsistent with a selection story, but nevertheless, we test for differences conditional on our fixed effects below and find much smaller and still statistically insignificant differences in application loan size.

Panel VI plots the number of applicants within each normalized FICO bin, along with the McCrary (2008) estimated density to test for manipulation of the running variable. While bunching in the running variable—a discontinuity in the propensity to apply for a loan at a FICO threshold—would raise selection concerns, panel VI shows that borrowers do not appear to select into applying for a loan based on where their FICO score falls relative to a lender's cutoff. Such targeting also seems a priori unlikely given the uncertainty applicants face about their own credit scores (owing to the volatility of FICO scores, uncertainty about which credit bureau(s) a lender will query, and general unawareness) and the low likelihood that prospective borrowers are aware of the precise thresholds used by a given lender.

Appendix Table A5 uses the specification in (8) allowing for overlapping discontinuity samples to report the magnitude and significance of the discontinuity coefficients on the loan characteristics variables using the loan-application data. Reassuringly, the estimates in column 1 indicate no statistical difference in requested loan amounts for borrowers on either side of the threshold. Column 2 shows that ex-ante debt-to-income ratios of borrowers on either side of the thresholds are statistically indistinguishable. Column 3 is equivalent to a McCrary test, counting the number of applications received from borrowers of each normalized FICO score and examining these counts at the normalized-FICO-score level using our RD estimator. The point estimate suggests that the number of borrowers applying for loans is also not statistically different on either side of a discontinuity.

[74]To facilitate this graphing exercise, we make the mapping between a loan and its normalized FICO score unique by keeping only loans within 10 points of a discontinuity. No loans in our data are less than 10 FICO points of more than one discontinuity at a given lender.

While loan interest rates change by an average of about 130 basis points at institution-specific FICO discontinuities, borrowers on either side of these thresholds are demographically similar and apply for similarly sized loans at similar frequencies, supporting the validity of our RD design.

D.4 Exclusivity of Credit Unions

A credit union is a member-owned cooperative financial institution that provides members with financial services. To satisfy the membership requirement, credit unions typically incorporate a brief membership application into their loan applications and require eventual borrowers to open a deposit account with a nominal minimum balance. Because credit-union membership eligibility is often restricted to well-defined groups, one concern is whether a given borrower could have joined the credit union providing the corresponding best available rate in our data. For example, if the lowest available interest rate for a particular cell was offered by a firefighters' credit union, then the borrower's search cost to obtain an offer from such a credit union would not only involve finding the low rate but also the costs associated with becoming a firefighter. To address this concern, we recalculate the spread-to-lowest-available rate measures using a sample comprised entirely of credit unions whose primary membership requirement is residence in a specified geographic area. In other words, all borrowers in our CZ-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 CZ as others in their cell. Our results are nearly identical after making this restriction. We also note that the finance companies in our sample have no membership requirements.

Figure A1: Examples of FICO-Based Discontinuities in Lender-Specific Interest-Rate Policies

I. Sample Lender #1

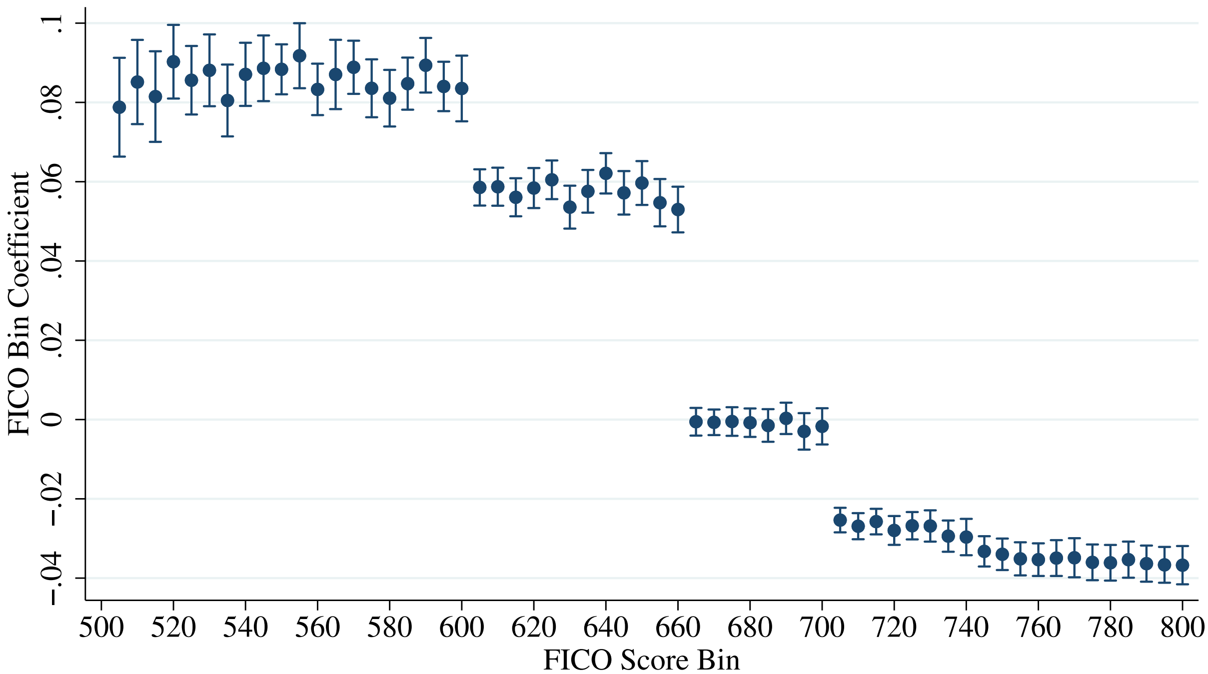

II. Sample Lender #2

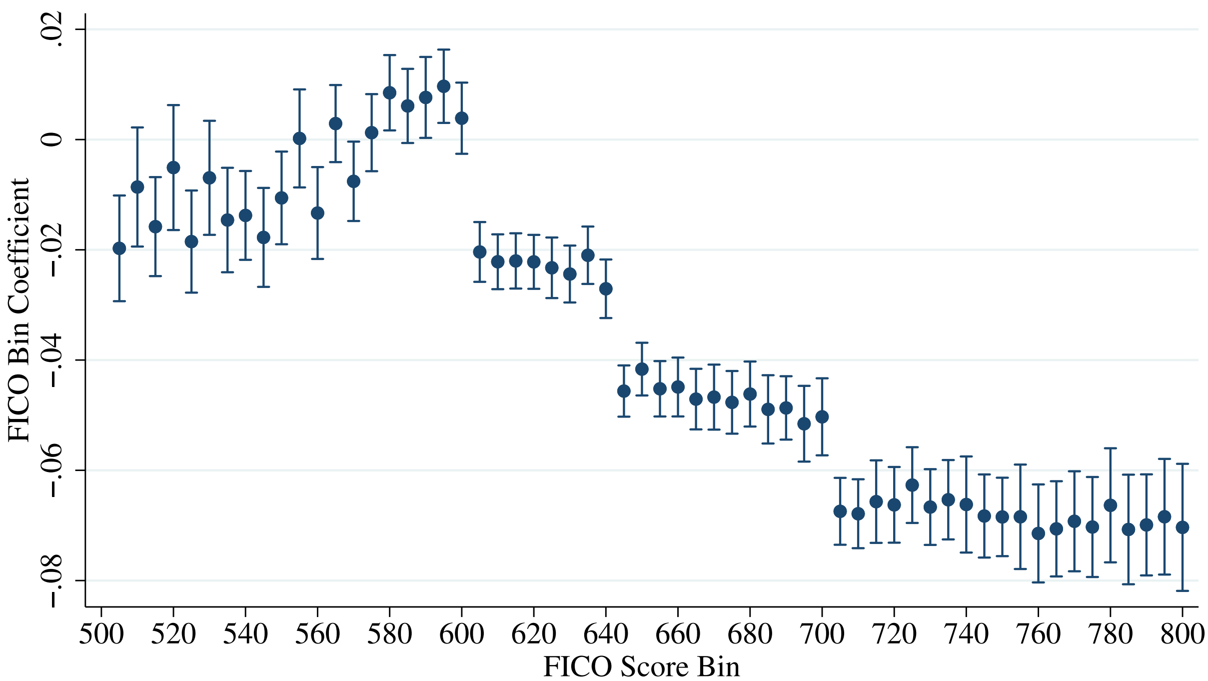

Notes: Each panel plots estimated interest-rate rules (with 95% confidence intervals) for a different lender in our sample. Loan rates are regressed on five-point FICO bin indicators as in equation (6).

Figure A2: Balance of Borrower Application Characteristics Across FICO Discontinuities

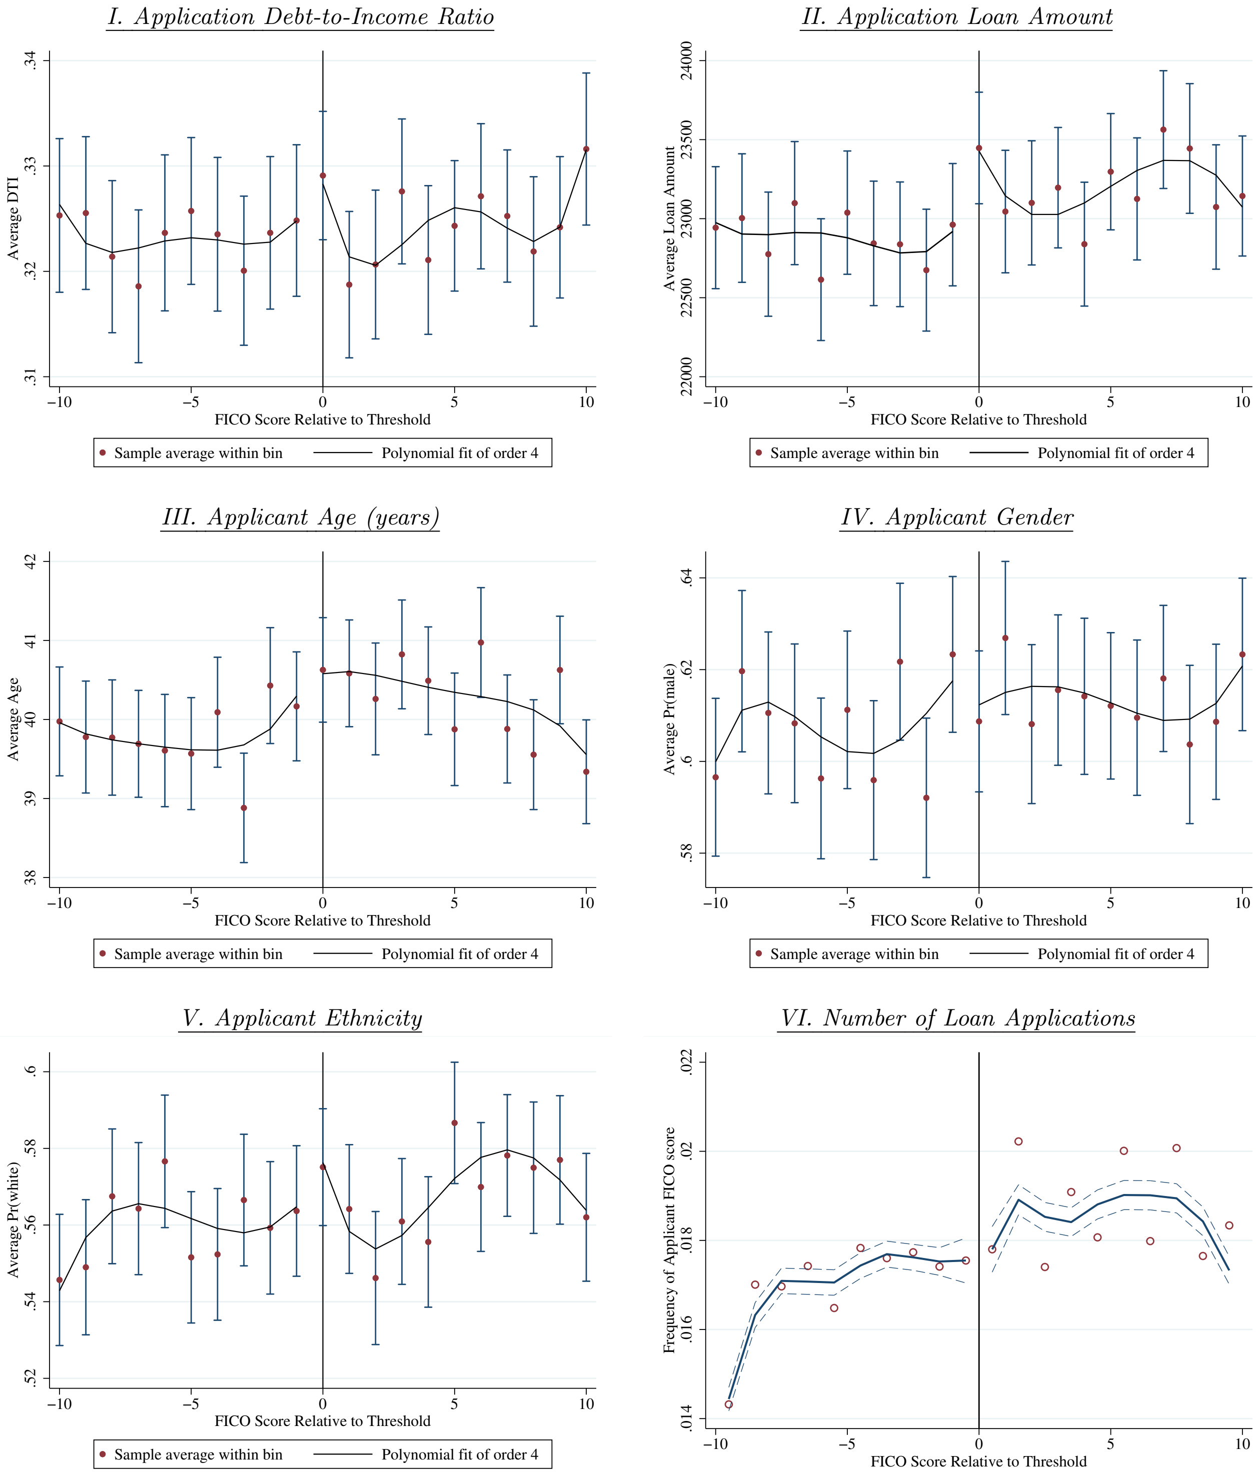

Notes: Figures plot average values of ex-ante borrower characteristics around FICO thresholds for institutions with detected discontinuities. 95% confidence intervals are double clustered by lender and FICO score. Plotted RD functions are estimated using the Calonico et al. (2014) robust RD estimator with fourth-degree polynomials. Applicant gender in panel IV is an indicator for male, and ethnicity in panel V is an indicator for whether the applicant is estimated as white by the lender. Panel VI plots the application count within each normalized FICO bin along with the estimated McCrary (2008) density test.

Figure A3: Loan Characteristics Around Discontinuities by Search Costs

I. Interest Rate

II. Loan Take-up

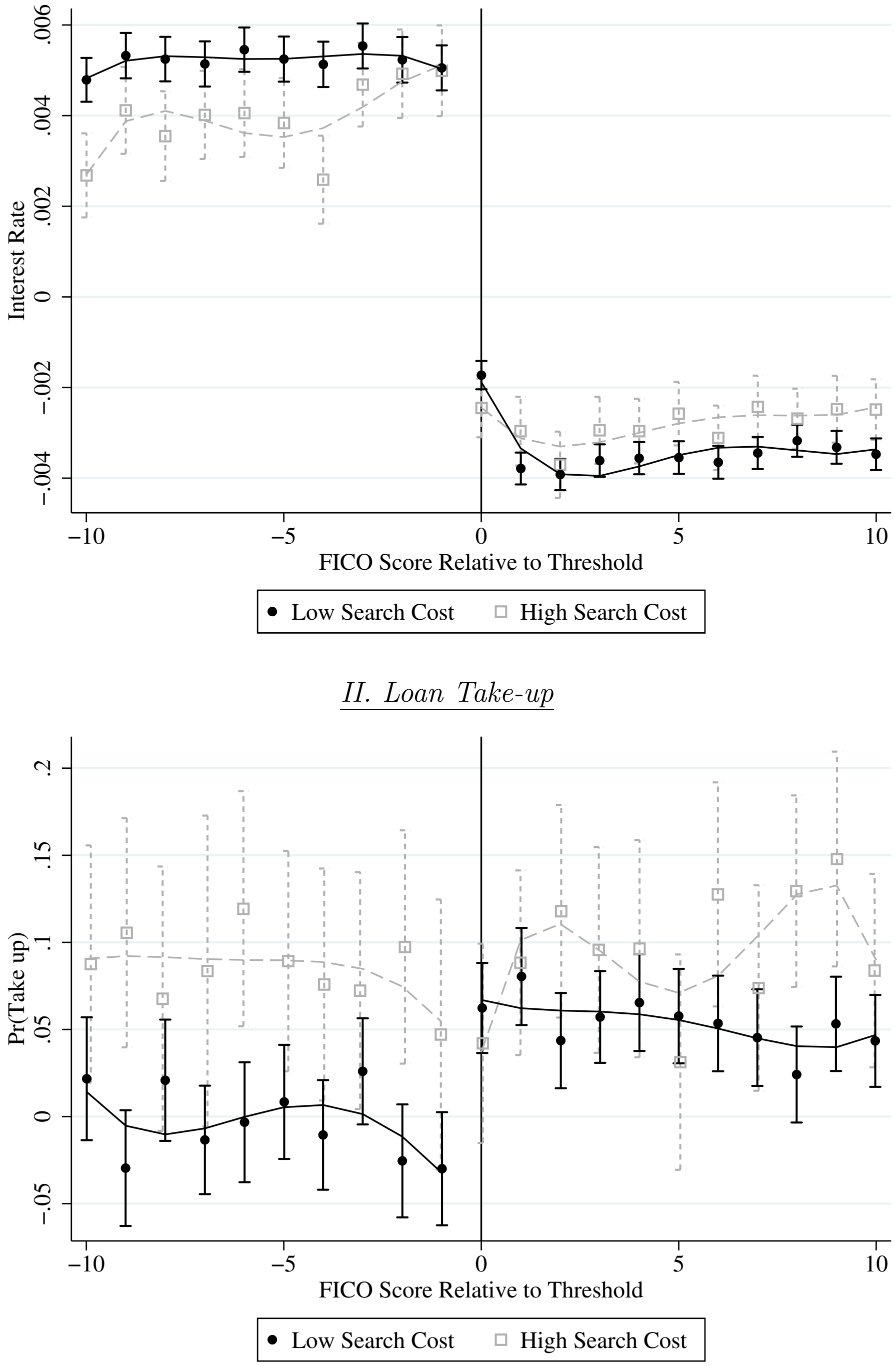

Notes: Figures plot average interest rates (panel I) and loan take-up rates (panel II) by search costs and FICO scores normalized to detected pricing discontinuities. Hollow gray squares and solid black circles represent the sample of borrowers with high and low search costs, respectively, defined as having at most or more than 10 financial institutions within a 20-minute drive of their home address. Outcomes are residualized by origination-quarter × zip code. 95% confidence intervals are double clustered by lender and FICO score. Plotted RD functions are for institutions with pricing discontinuities and estimated using the Calonico et al. (2014) robust RD estimator with fourth-degree polynomials.

Figure A4: Average Vehicle Mileage by Car Age

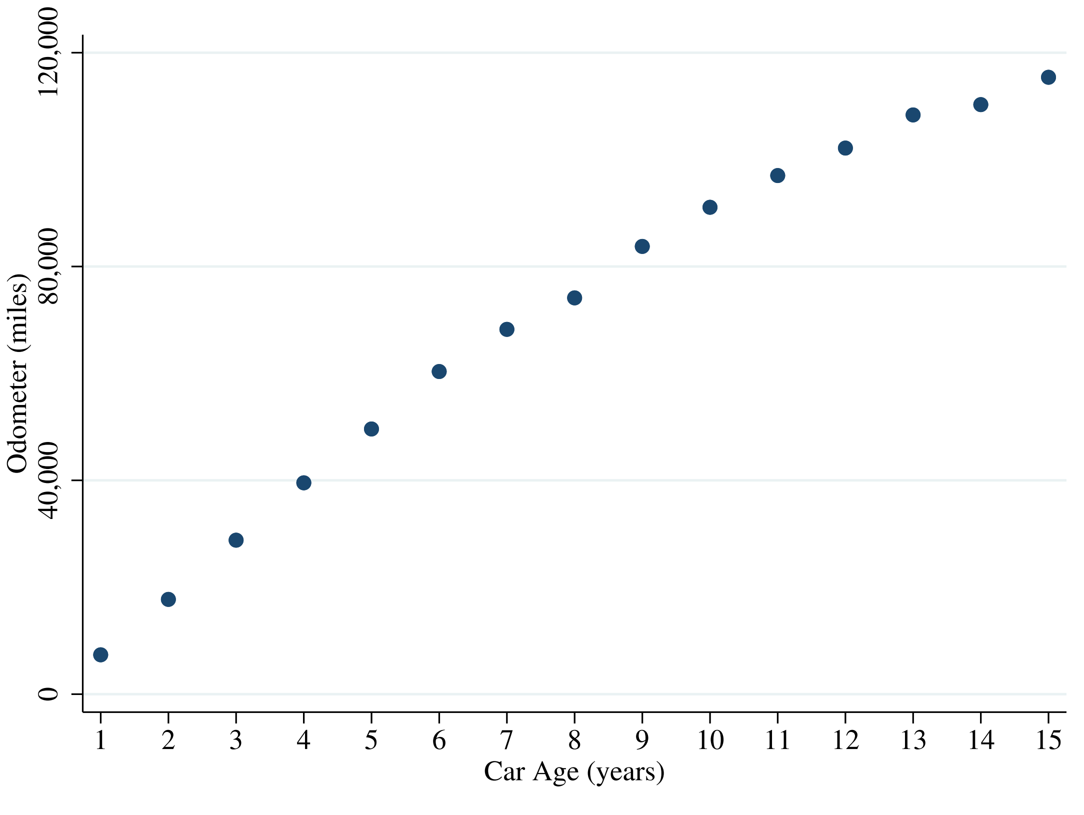

Notes: Figures plots average used-car odometer readings as a function of car age using microdata from the National Household Travel Survey (U.S. Department of Transportation, 2017).

Figure A5: Balance of Borrower and Loan Characteristics Conditional on Origination

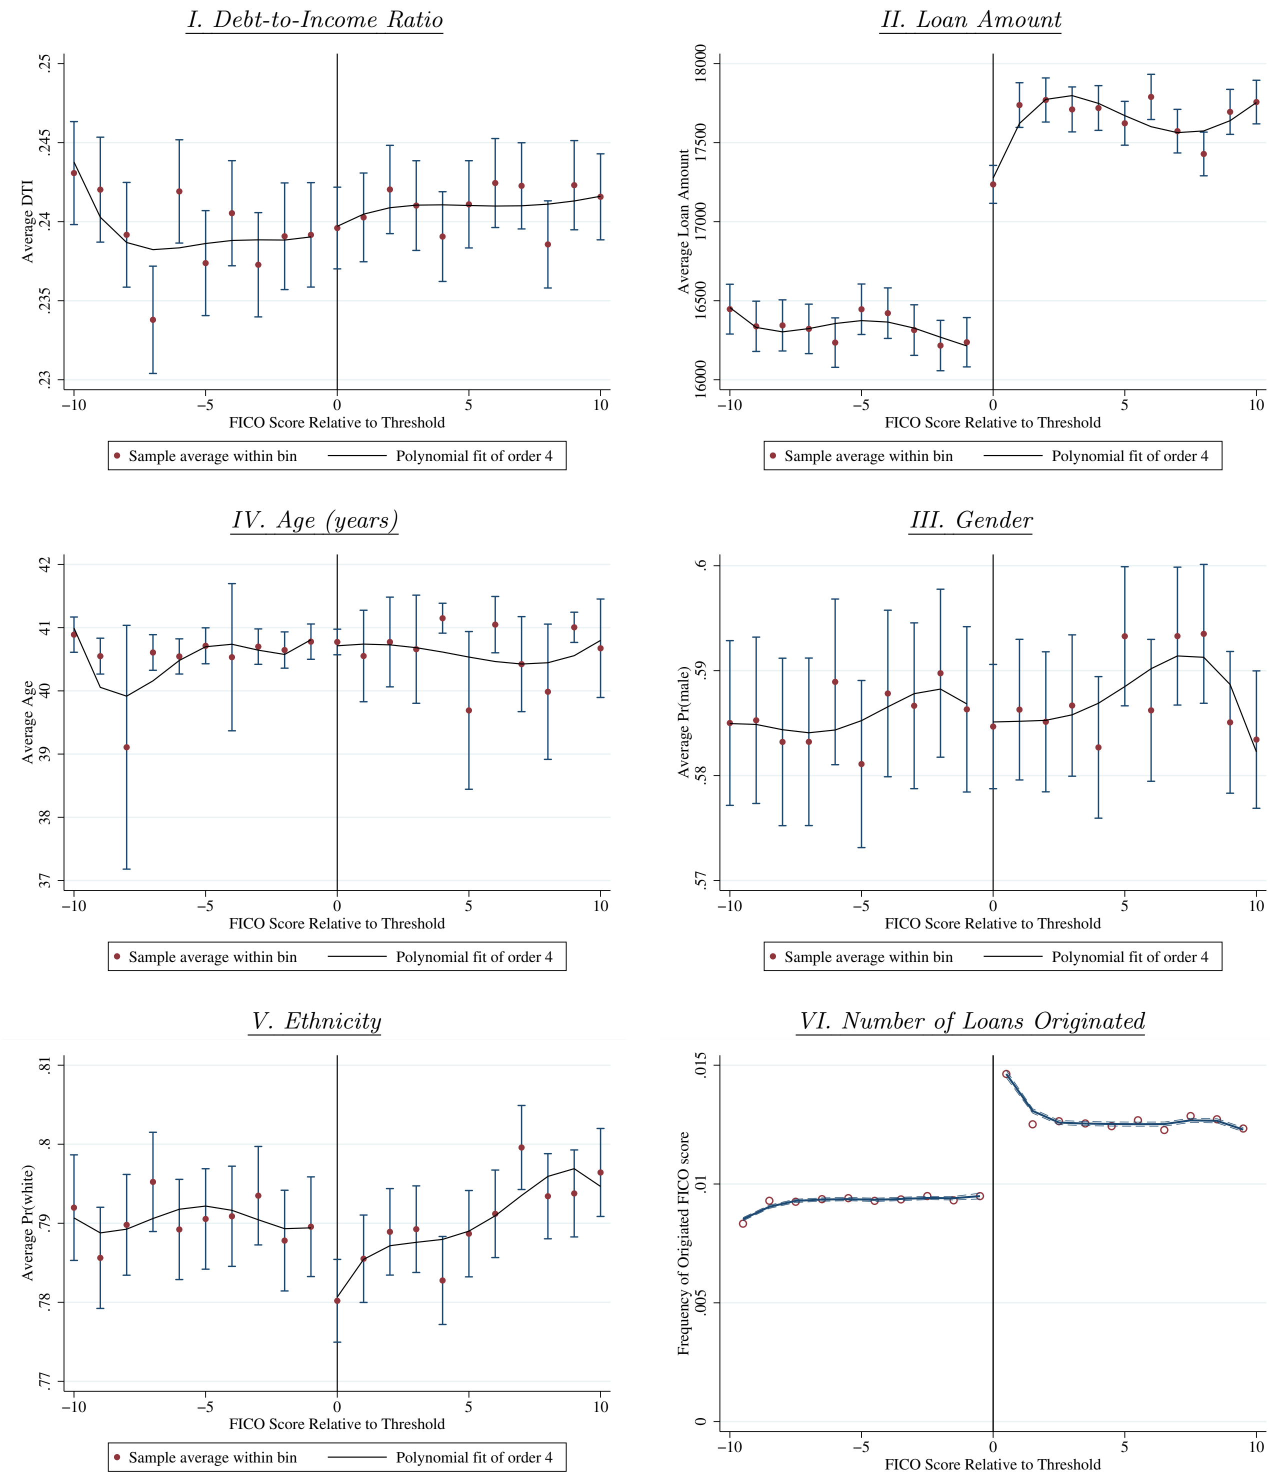

Notes: Figures plot average values of borrower and loan characteristics around FICO thresholds conditional on observing the application of an originated loan in our data. See notes to Appendix Figure A2 for further details.

Figure A6: Loan Performance Around Discontinuities by Search Cost

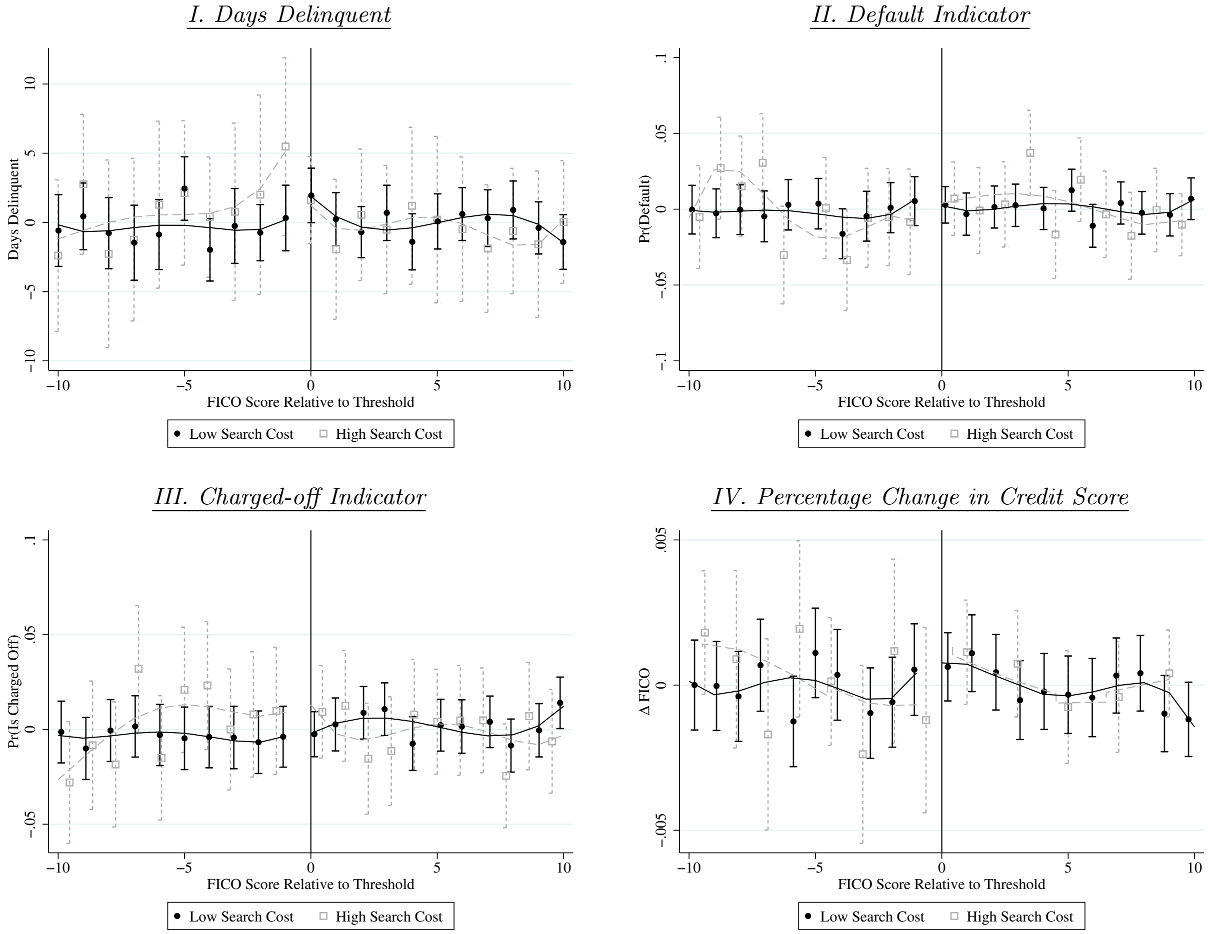

Notes: Figures plot average loan performance measures by search costs and FICO scores normalized to detected pricing discontinuities. Hollow gray squares and solid black circles represent the sample of borrowers with high and low search costs, respectively, defined as having at most and more than 10 financial institutions within a 20-minute drive of their home address. All outcomes are residualized by quarter $\times$ zip code. 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. Percentage change in FICO score is the change in FICO score since origination as a fraction of the FICO score at origination for the subsample of institutions that report credit scores after loan origination. 95% confidence intervals are double clustered by lender and FICO score. Plotted RD functions are for institutions with pricing discontinuities and estimated using the Calonico et al. (2014) robust RD estimator with fourth-degree polynomials.

Figure A7: Illustration of Surplus and Welfare Loss from Search Frictions

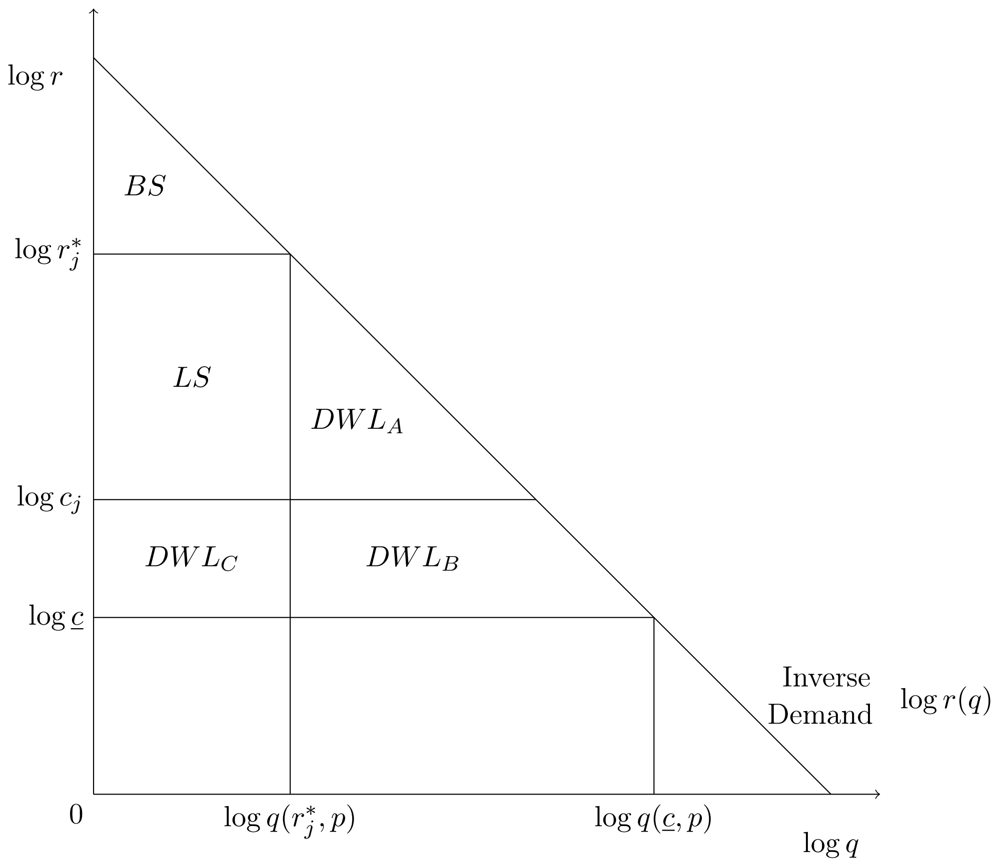

Notes: Figure plots equilibrium outcomes under perfect competition and costly search in log interest rate (y-axis) and log quantity (x-axis) space. Log inverse demand is the diagonal line denoted $\log r(q)$. The perfect-competition equilibrium quantity and price are denoted $q(\underline{c}, p)$ and $\underline{c}$, respectively, where $\underline{c}$ is the minimum support of the marginal cost distribution. The costly search equilibrium quantity and price are $q(r_j^*, p)$ and r_j^*, respectively, where r_j^* is the monopolist price charged by a lender with marginal cost c_j. BS and LS denote borrower surplus and lender surplus, respectively. The three deadweight loss components are denoted DWL. See Appendix A for more details.

Figure A8: Cumulative Density of the Spread to Lowest Available Rate by Search Costs

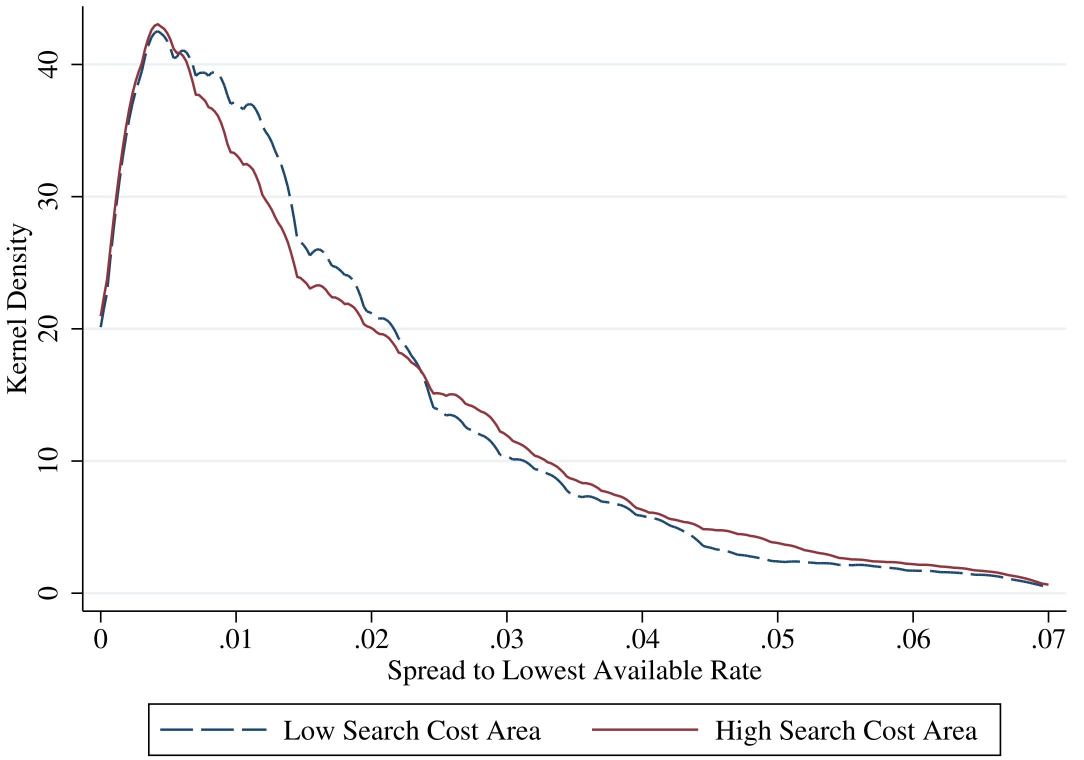

Notes: Figure plots kernel density of the spread to the lowest available interest rate for each loan (the difference between a loan's interest rate and the best rate among similar borrowers in its cell) for borrowers in high- and low-search-cost areas. The estimated distributions for high- and low-search-cost areas are plotted in solid red and dashed blue lines, respectively. High-search-cost areas are defined as locations where borrowers have at most 10 lenders within a 20-minute drive of their home. A cell is defined as all borrowers in the same commuting zone taking out a loan in the same $1,000 collateral-value bin, five-point FICO bin, 10-point DTI bin, six-month time period, and loan maturity.

Figure A9: Density of Spread to Lowest Available Rate

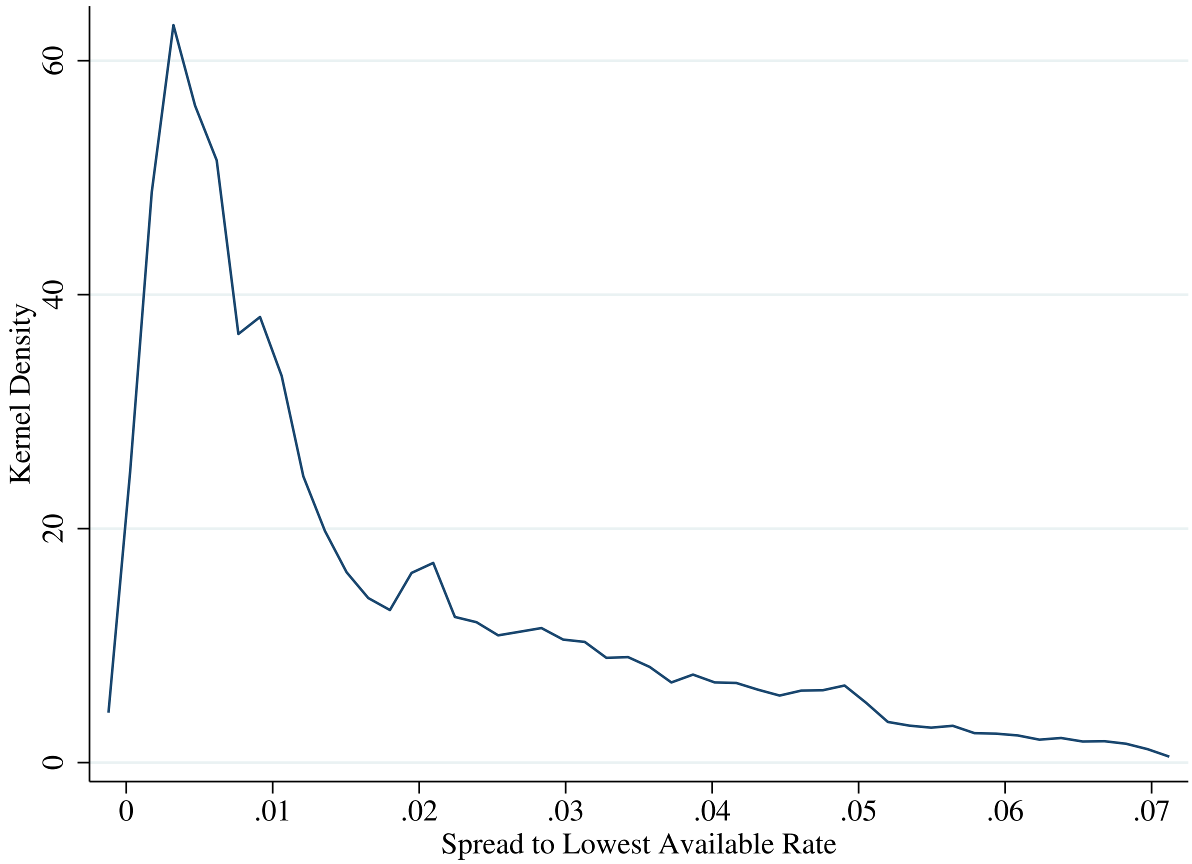

Notes: Figure reports the kernel density of the spread to the lowest available interest rate for borrowers not receiving the best available rate in their cell. Estimated density uses an Epanechnikov kernel with a bandwidth of 0.0012. See notes to Appendix Figure A8.

Figure A10: Relationship between Market Shares and Average Mark-ups by Search Costs

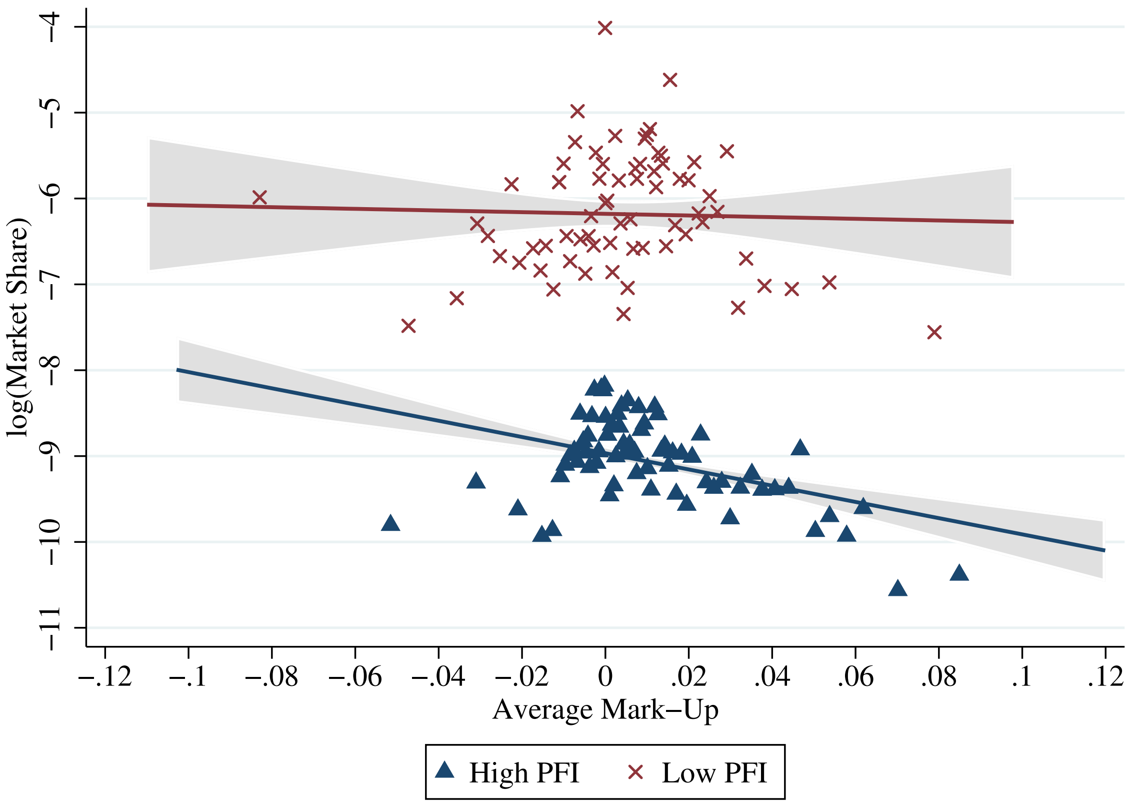

Notes: Figure shows a bin-scatter plot of average log market shares against average interest-rate markups for low-search-cost and high-search-cost borrowers (denoted with ▲ and × symbols, respectively) along with bivariate regression lines and 95% confidence intervals. Low (high) search costs are defined as borrowers living at locations with more than (at most) 10 lenders within a 20-minute drive. Market shares are calculated at the lender × Commuting Zone level using the origination data of lenders in our sample. Markups are calculated as lender fixed effects in a loan-level regression of interest rates on controls for a cubic in FICO, loan size, loan term, LTV, and quarter fixed effects.

Table A1: Summary Statistics for Excluded Sample of Indirect Loans

	Count	Mean	Std. Dev.	Percentile 25th	50th	75th
			I. Originated Loans			
Loan Rate	1,166,822	0.05	0.03	0.03	0.04	0.06
Loan Term (months)	1,166,822	69.97	17.97	60	72	75
Loan Amount ($)	1,166,822	22,051.64	11,318.28	13,790	20,146	28,324
FICO	1,013,915	718.77	68.06	672	719	770
Debt-to-Income (%)	462,116	0.25	0.52	0	0.22	0.35
Collateral Value ($)	1,166,822	21,997.7	11,176.38	13,983	19,965	27,800
Monthly Payment ($)	1,166,822	360.87	161.71	246.8	334.25	445.04
			II. Ex-Post Loan Performance Measures			
Days Delinquent	799,144	39.16	645.64	0	0	0
Charged-off Indicator	1,166,822	0.03	0.16	0	0	0
Default Indicator	1,166,822	0.03	0.17	0	0	0
Current FICO	705,754	704.76	81.71	656	712	769
%ΔFICO	695,114	-0.02	0.08	-0.05	-0.01	0.02

Note: Table reports summary statistics for the indirect loan portion of the original dataset. This portion is excluded from the analysis of the paper. See notes to Table 1 for further details.

Table A2: Summary Statistics for Estimation Sample with Identified FICO Discontinuities

	Count	Mean	Std. Dev.	25th	Percentile 50th	75th
			I. Approved Loan Applications			
Loan Term (months)	30,743	63.75	13.02	60	66	72
Loan Amount ($)	30,743	21,310.6	11,106.2	13,188.5	19,626	27,606
Loan Rate	26,938	0.056	0.082	0.037	0.05	0.07
FICO	30,743	678.2	31.4	657	682	703
Debt-to-Income	27,197	0.292	0.133	0.197	0.296	0.388
Take-up	30,743	0.511	0.5	0	1	1
			II. Originated Loans			
Loan Rate	514,834	0.061	0.031	0.037	0.055	0.078
Loan Term (months)	514,834	59.616	19.813	48	60	71
Loan Amount ($)	514,834	14,297.87	6,539.55	9,209	13,704	18,805
FICO	514,834	662.386	40.242	636	664	691
Debt-to-Income (%)	296,748	0.243	0.161	0.1	0.271	0.38
Collateral Value ($)	514,834	15,768	6,808.878	10,825	14,906	19,750
Monthly Payment ($)	514,834	279.952	108.67	199.888	269.717	348.356
			III. Ex-Post Loan Performance Measures			
Days Delinquent	331,590	30.905	245.453	0	0	0
Charged-off Indicator	514,834	0.025	0.157	0	0	0
Default Indicator	514,834	0.024	0.154	0	0	0
Current FICO	405,236	659.352	69.749	621	667	706
%ΔFICO	405,236	-0.005	0.09	-0.045	0.003	0.047

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

Table A3: Spread to Lowest Available Rate Summary Statistics

FICO Range	# of Cells	Mean # Borrowers in Cell	Mean	Std. Dev.	Percentile 25th	Percentile 50th	Percentile 75th
$595 \leq FICO \leq 599$	74	2.19	0.038	0.029	0.01	0.03	0.06
$635 \leq FICO \leq 639$	250	2.23	0.023	0.021	0.01	0.02	0.03
$695 \leq FICO \leq 699$	161	2.15	0.011	0.01	0.003	0.01	0.02

Notes: Table reports summary statistics for the spread between a left-of-threshold borrower's interest rate and the best available interest rate for borrowers in the same cell for measured discontinuities at around three FICO scores (600, 640, and 700). Cells are defined as borrowers within the same commuting zone, 5-point FICO bin, $1,000 purchase-price bin, 10 percentage point DTI bin, maturity, and who take out loans in the same six-month window. Within each of the matched bins, we calculate the average difference between the lowest interest rate in the cell and each individual loan in the cell. Summary statistics are reported for only those cells that contain at least two borrowers.

Table A4: First Stage Estimates for High Search Costs Using Historical Variation

	(1)	(2)	(3)	(4)	(5)
High $\widehat{\text{Search}}$ Cost	0.72***	0.72***	0.72***	0.68***	0.19***
	(0.08)	(0.08)	(0.08)	(0.08)	(0.06)
Discontinuity × Lender FEs		✓	✓	✓	✓
RD Controls			✓	✓	✓
CZ × Quarter FEs				✓	
Zip × Quarter FEs					✓
Number of Observations	30,740	30,740	30,740	30,740	30,740
Partial F-statistic	81.66	82.11	82.87	70.83	11.77
R^2	0.62	0.62	0.62	0.69	0.94

Notes: Table reports OLS estimates testing the degree to which predicted high search costs predict our main high search cost measure using the applications data. The dependent variable High Search Cost is a dummy equal to 1 for borrowers with at most 10 financial institutions within a 20-minute drive from their home. Predicted High Search Cost is an indicator for whether the predicted number of financial institutions within a 20-minute drive is at most 10, with the predicted number defined by (11) using 1990 local branch proximity and subsequent national trends in branching. RD controls consist of a quadratic spline in normalized FICO score that is allowed to change at the discontinuity as specified in (9). Robust standard errors in parentheses are double clustered by lender and FICO score.

Table A5: Loan Application Covariate Balance Regressions

	(1) Loan Amount	(2) Debt-to-Income	(3) Number of Loan Applications
Discontinuity Indicator	128.43	-0.084	-270.18
	(187.75)	(0.447)	(760.48)
Discontinuity × Lender FEs	✓	✓	✓
Commuting Zone × Quarter FEs	✓	✓	✓
Number of Observations	117,985	91,923	39
R^2	0.06	0.01	0.47

Notes: Table reports regression-discontinuity estimates of loan amount, debt-to-income ratios, and the number of observed loan applications for the subset of institutions for which we have detailed loan application data. Discontinuity Indicator is a dummy for whether a borrower's normalized FICO score is positive. Each observation in the data used for column 3 represents a normalized FICO score. Robust standard errors in parentheses for columns 1 and 2 are double clustered by lender and FICO score. See notes to Table 2 for details.

Table A6: Effects of FICO Discontinuity on Origination Outcomes by Search Costs

	(1)	(2)	(3)	(4)
Dependent Variable	Loan Amount ($)		Purchase Price ($)	
Search Cost Sample	High	Low	High	Low
Discontinuity Indicator	766.80***	515.70***	552.47***	313.74***
	(194.57)	(129.83)	(213.09)	(157.45)
Discontinuity × Lender FEs	✓	✓	✓	✓
CZ × Quarter FEs	✓	✓	✓	✓
Make-Model FEs	✓	✓	✓	✓
Number of Observations	68,417	400,383	68,417	400,383
R^2	0.36	0.28	0.42	0.38

Notes: Table reports regression-discontinuity estimates of loan amount and purchase price for borrowers with high and low search costs regressed on the discontinuity indicator, our RD controls, and the indicated fixed effects. High search costs are defined as borrowers living at locations with at most 10 lenders within a 20-minute drive. Discontinuity Indicator is a dummy for whether a borrower's normalized FICO score is positive. Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 2 for more details.

Table A7: Spread to Lowest Available Rate Summary Statistics by Search Costs

				Percentile		
	N	Mean	Std. Dev.	25th	50th	75th
High Search Costs	29,589	0.0258	0.0288	0.005	0.014	0.035
Low Search Costs	236,031	0.0231	0.0266	0.005	0.0125	0.031

Notes: Table reports summary statistics for the spread between the rate a borrower accepted and the best available interest rate for borrowers in the same cell. Cells are defined as borrowers within the same commuting zone, 5-point FICO bin, $1,000 purchase-price bin, 10 percentage point DTI bin, maturity, and who take out loans in the same six-month window. Within each of the matched bins, we calculate the average difference between the lowest interest rate in the cell and each individual loan in the cell. High search costs are defined as borrowers living at locations with at most 10 lenders within a 20-minute drive. Summary statistics are reported for only those cells that contain at least two borrowers.

Table A8: Number of Observed Loan Applications per Borrower by Search Costs

Search Costs	High	Low	Difference
	(1)	(2)	(1) - (2)
	I. Mortgage Applications per Origination		
Mean	2.08	2.24	-0.18***
Standard Deviation	(0.51)	(1.00)	(0.01)
Number of Observations	3,584	23,596	
	II. Car Loan Applications per Borrower		
Mean	1.34	1.41	-0.07***
Standard Deviation	(0.01)	(0.004)	(0.01)
Number of Observations	6,183	44,514	

Notes: Table reports average number of applications observed for each loan origination in HMDA data (panel I) and in our data for applications with non-missing birthdates and nine-digit zip codes (panel II) for high- and low-search-cost areas in columns 1 and 2, respectively. Standard deviations are reported in parentheses. Column 3 calculates the difference in means along with the robust standard error for the statistical significance of the difference between columns 1 and 2. High search costs are defined as borrowers living at locations with at most 10 lenders within a 20-minute drive using the predicted number of lenders as described in section 5.1.

Table A9: Effects of Search Costs and Mark-ups on Market Shares

	(1)	(2)	(3)
Average Mark-up	-9.452***	-8.058***	-7.067***
	(2.032)	(1.917)	(1.896)
High Search Costs	2.788***	2.856***	2.474***
	(0.180)	(0.199)	(0.222)
Avg. Mark-up × High Search Costs	8.485*	6.141*	7.929*
	(4.396)	(2.837)	(3.941)
Commuting Zone FEs		✓	✓
Lender FEs			✓
R-squared	0.26	0.42	0.42
Number of Observations	3,133	3,133	3,133

Notes: Table reports estimation results of regressions of log(market share) on average mark-ups at the lender × Commuting-Zone level. Market shares are calculated at the lender × Commuting Zone level using the origination data of lenders in our sample. Markups are calculated as lender fixed effects in a loan-level regression of interest rates on controls for a cubic in FICO, loan size, loan term, LTV, and quarter fixed effects. High search cost is an indicator for borrowers living at locations with at most 10 lenders within a 20-minute drive. Robust standard errors in parentheses are clustered by Commuting Zone.

Table A10: Effects of Search Costs and Market Concentration on Take-up Decisions

Market Concentration	Low	High
	(1)	(2)
Discontinuity Indicator	0.177***	0.086***
	(0.049)	(0.023)
Discontinuity Indicator ×	-0.059***	-0.101**
High Search Cost	(0.019)	(0.044)
Discontinuity × Lender FEs	✓	✓
CZ × Quarter FEs	✓	✓
Number of Observations	12,736	15,586
R^2	0.22	0.37

Notes: Table reports regression-discontinuity estimates of whether a borrower accepts an approved loan offer for borrowers in markets with low and high concentration in columns 1 and 2, respectively. Loan take-up is regressed on a discontinuity indicator, a high-search-cost indicator, their interaction, RD controls, discontinuity by lender fixed effects, and commuting zone by quarter fixed effects using the specification in (14). Market concentration uses CZ-level lender mortgage market shares in HMDA data to construct an Herfindahl index (HHI) of competition. High and low market concentration are defined as above and below median HHI. Discontinuity Indicator is a dummy for whether a borrower's normalized FICO score is positive. High Search Cost is a dummy equal to 1 for borrowers with at most 10 financial institutions within a 20-minute drive from their home. Included RD controls consist of a quadratic spline in normalized FICO score that is allowed to change at the discontinuity as specified in (9). Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 2 for estimation details.