NBER WORKING PAPER SERIES

REAL EFFECTS OF SEARCH FRICTIONS IN CONSUMER CREDIT MARKETS

Bronson Argyle
Taylor D. Nadauld
Christopher Palmer

Working Paper 26645
http://www.nber.org/papers/w26645

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2020

We thank our discussants Paul Calem, Anthony DeFusco, Paul Goldsmith-Pinkahm, 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, Glenn Ellison, Maryam Farboodi, Amy Finkelstein, Brigham Frandsen, Peter Ganong, Kyle Herkenhoff, Lars Lefgren, Andres Liberman, Brigitte Madrian, Adrien Matray, 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 and Sam Hughes provided excellent research assistance. Palmer thanks the Fisher Center for Real Estate and Urban Economics for support. An anonymous information-technology firm provided the data. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w26645.ack

Real Effects of Search Frictions in Consumer Credit Markets
Bronson Argyle, Taylor D. Nadauld, and Christopher Palmer
NBER Working Paper No. 26645
January 2020
JEL No. D12,D83,E43,G21,L11

ABSTRACT

We establish two underappreciated facts about costly search. First, unless demand is perfectly inelastic, search frictions can result in significant deadweight loss by decreasing consumption. Second, whenever cross-price elasticities are non-zero, costly search in one market also affects quantities in other markets. As predicted by our model of search for credit under elastic demand, we show that search frictions in credit markets contribute to price dispersion, affect loan sizes, and decrease final-goods consumption. Using microdata from millions of auto-loan applications and originations not intermediated by car dealers, we isolate plausibly exogenous variation in interest rates due to institution-specific pricing rules that price risk with step functions. These within-lender discontinuities lead to substantial variation in the benefits of search across lenders and distort extensive- and intensive-margin loan and car choices differentially in high- versus low-search-cost areas. Our results demonstrate real effects of the costliness of shopping for credit and the continued importance of local bank branches for borrower outcomes even in the mobile-banking era. More broadly, we conclude that costly search affects consumption in both primary and complementary markets.

Bronson Argyle
Department of Finance
Brigham Young University
Provo, UT 84602
bsa@byu.edu

Taylor D. Nadauld
Department of Finance
Brigham Young University
Provo, Utah 84602
taylor.nadauld@byu.edu

Christopher Palmer
MIT Sloan School of Management
100 Main Street, E62-639
Cambridge, MA 02142
and NBER
cjpalmer@mit.edu

1 Introduction

Some of the most important questions in household finance center around how credit-market imperfections affect consumption.[1] In this paper, we demonstrate the special role that financing search frictions play in determining the equilibrium consumption of durables. When demand is elastic, search frictions in one market not only affect equilibrium quantities in that market but also the demand for complementary goods, resulting in an efficiency loss that is obscured by popular search models' assumption of inelastic demand. We provide evidence that costly search in consumer credit markets leads to interest-rate dispersion across similar loans and can affect extensive- and intensive-margin loan and durable consumption choices. To our knowledge, we are the first paper to quantify how search frictions affect consumption.

We begin by extending a search model with elastic demand to a credit-market setting to demonstrate theoretically how financing search costs could affect the distribution of markups, equilibrium credit and durables outcomes, and welfare. Next, using administrative data on 2.4 million auto loans extended by 326 different financial institutions in every state and loan application data on 1.3 million potential loans from 41 institutions, we establish four main empirical facts.[2] First, there is significant price dispersion for the same credit product across providers—most borrowers in our data could access significantly dominating loan offers if they queried two additional financial institutions. Second, such search appears costly, and borrowers' propensity to search for loans with better terms is lower in areas likely to have higher search costs. Third, the segment of the auto lending market we study does not feature pure risk-based pricing; we observe large interest-rate discontinuities at various lender-specific credit-score (FICO) thresholds. Fourth, consumer financing and durable-goods purchasing decisions respond to the resulting interest-rate dispersion around these lending thresholds more in high- versus low-search-cost areas.

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 a plausible setting to look for credit-market search frictions affecting final-goods consumption. Auto debt is the third-largest category of consumer debt (ranking ahead of credit cards), with over 114 million outstanding loans (0.89 per U.S. household) comprising \$1.32 trillion in aggregate auto debt (NY Fed, 2019). Most car purchases are financed (Bartlett, 2013), and vehicles represent over 50% of total

[1] See, for example, influential work on the role of adverse selection in consumer credit markets (Adams, Einav, and Levin, 2009), the importance of credit constraints in explaining high marginal propensities to borrow and consume out of credit (Gross and Souleles, 2002), and the inhibition of credit expansions to the household sector (Scharfstein and Sunderam, 2016; Agarwal et al., 2017).

[2] As noted by Allen, Clark, and Houde (2019), the ability to observe rejected price quotes afforded by our application data is novel in the empirical search literature.

assets for low-wealth households (Campbell, 2006). From an empirical-design standpoint, auto loans are a relatively homogeneous credit product and can be described completely by their interest rate, term, and amount. Focusing on the large segment of used-car loans not intermediated by the car seller allows us to test for credit- and product-market linkages in a non-mechanical setting. Finally, auto loan markets are quite local. The median borrower in our sample borrowing directly from a lender (as opposed to indirect loans originated via auto dealers) originates a loan from a branch that is within a 15-minute drive of her home, whereas the median worker in the United States commutes 26 minutes to work. The stylized fact that direct auto loan markets are more local than labor markets motivates our inquiry into the distortions that physical search frictions might cause in consumer debt markets.

Our empirical strategy features a setting where potential gains to search are high and quasi-randomly assigned. We document large discontinuities in offered loan terms around FICO thresholds within lending institutions.[3] On average, borrowers just above an institution's FICO threshold are offered loans with interest rates 1.27 percentage points lower than otherwise similar borrowers just below a FICO threshold. Importantly, the location of the thresholds 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 show 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 within-lender, across-borrower variation in the returns to search provides a laboratory to test how borrowers with a high return to additional search differentially change their behavior in markets where search costs are high.

Figure 1 provides examples of such interest-rate discontinuities for three different credit unions in our data with detected discontinuities using the lending policy rule estimation procedure described in section 6.2. As discussed in section 6.3 below, the observed FICO thresholds isolate supply-side changes in loan characteristics from demand-driven factors under the assumption that demand-side factors (e.g., preferences, income, financial sophistication) are not likely to also change discontinuously at quasi-random FICO thresholds that vary across institutions within the same geography. Moreover, 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 assumption of quasi-randomly assigned markups 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

[3]Lending policies that jump discontinuously at various FICO thresholds appear to exist in 173 of the 326 lending institutions in our sample.

FICO bin) are balanced around FICO thresholds.[4]

We discuss potential aspects of shopping for a loan that may have particular utility costs associated with them in section 7.1 below; the process often entails time, effort, and stress—each of which may be in short supply while simultaneously shopping for a car. We show that borrowers on the expensive side of FICO thresholds reject high-interest-rate loans more often when the number of nearby alternative lenders is high. Using the physical branch locations of every bank and credit union in the United States, we calculate the number of financial institutions within a 20-minute drive from each borrower as a proxy for search costs.[5] We find that differences in loan take-up rates across FICO thresholds are smaller for borrowers in high-search-cost areas. Borrowers who would presumably have to exert more effort to search for a loan with better terms are more likely to accept the loan pricing they are offered even though these terms are strongly dominated by nearby alternatives. Matching borrowers across loan applications to different lenders, we verify that borrowers are more likely to submit multiple loan applications when our search-cost measure is low. These results on price dispersion and search behavior are inconsistent with a simple story of market concentration leading to high markups, and we show in an appendix our results hold even in relatively high-competition markets.

What impact does sharp variation in loan pricing for otherwise identical borrowers have on borrower outcomes? On average, borrowers quasi-randomly offered expensive credit purchase cars that are two months older, spending an average of $345 less, with stronger effects in high than low-search-cost areas. The balance of borrower characteristics across FICO thresholds suggest that borrowers quasi-randomly drawing high loan markups would also like to 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 correlated with borrower-level demand shocks bu further verifying the balance of borrower characteristics and ex-post borrower outcomes. Post-origination changes in credit scores and ex-post loan performance do not change differentially across discontinuities, consistent with our theoretical model and suggesting that borrowers who accept high markups are not likely to have fundamentally different constraints or preferences. Taken together, our evidence suggests that search costs represent a meaningful market friction that enables the persistence of equilibrium price dispersion and ultimately

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

[5]In Appendix A, we consider a series of robustness tests to address potential omitted variables that could be correlated with our measure of physical search costs. In Appendix B, we estimate search costs using the structural search model of Hortaçsu and Syverson (2004) to support our use of nearby lender density as a measure of search costs.

distorts durable consumption.

The remainder of the paper proceeds as follows. After contextualizing our work in several related literatures in section 2, section 3 uses a model of costly search with elastic demand for both loans and durables to make several empirical predictions motivating our empirical analysis. Section 4 details the administrative data we use throughout the paper, including an analysis of its representativeness. Section 5 documents price dispersion in the market for auto loans. Section 6 introduces our regression-discontinuity identification strategy to isolate exogenous variation in the benefits to search using lender pricing rules. In sections 7 and 8, 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 9 concludes.

2 Related Literature

In this section, we motivate our work in the context of literatures on search frictions, auto loans, and FICO-based discontinuities.

Theories of costly search since Stigler (1961) suggest that under many conditions, when agents find it costly to solicit the full menu of offered prices, equilibrium prices will feature price dispersion.[6] In a credit market, for example, 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. Much of the follow-up to the original Stigler (1961) model modeled sequential search by shoppers with full information about the distribution of available prices (e.g., Stahl, 1989). Particularly relevant is the model of Reinganum (1979) we build on below that generates price dispersion without the search-cost heterogeneity required by Stahl (1989) by assuming inelastic demand and seller-cost heterogeneity.

Several recent papers feature alternative formulations of consumer search. De Los Santos, Hortaçsu, and Wildenbeest (2012) examine online shopping behavior and find support for fixed sample size search instead of the optimal stopping rule of sequential search models. Ellison and Ellison (2009) present evidence consistent with sellers obfuscating to increase search costs via the opacity of the product characteristics and prices menu. Zhu (2012) models the penalty buyers or sellers may face when they return to a previously obtained quote and thereby signal the level of surplus in the transaction. Salz (2017) demonstrates the welfare effects of intermediaries in search markets. While our setting does not allow us to cleanly distinguish between competing models of search, our estimated gains from search

[6] For a comprehensive survey of the theoretical and empirical search and price dispersion literature, see Baye, Morgan, and Scholten (2006).

exceed typical measures of the opportunity cost of time, suggesting perhaps that potential borrowers are unaware of the degree of price dispersion in the auto-loan market.

Multiple empirical papers establish the existence of equilibrium price dispersion and connect it to evidence that consumer search is costly in a given domain—see Baye, Morgan, and Scholten (2006) for an overview. Relevant to our use of natural variation in the returns to search, Sorenson (2000) documents dispersion in prices of prescription drugs that are driven by variation in the financial return to greater search intensity. In consumer finance, Hortaçsu and Syverson (2004) find large dispersion in the fees charged by mutual funds tracking the same index that are driven by search frictions. Woodward and Hall (2012) document that mortgage borrowers overpay for mortgage broker services due to a reluctance to shop for mortgages. 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 consumer surplus loss from inelastic borrowers paying higher prices. Alexandrov and Koulayev (2017) also document price dispersion in mortgage rates and provide survey evidence that close to half of consumers do not shop for a mortgage before origination and are generally unaware of price dispersion.[7] 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 the literature on price dispersion and search intensity that often models consumers as having inelastic demand for a final good, our setting allows for measurement of the real effects on consumption quantities that can result from costly search.

Recent work by Agarwal et al. (2019) shows that in the cross section, intensive loan search is correlated with higher interest rates, running counter to the standard prediction that search and selected prices are inversely correlated. Agarwal et al. (2019) 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 effectively allows us to abstract away from cross-sectional variation in private information and rely on the conceptual argument that for a given borrower, the relationship between search and interest rates should be negative.

We are not the first paper to exploit FICO-based discontinuities in treatment variables. Keys et al. (2009 and 2010) find that the probability of securitization (and thus loan screen-

[7]See also related evidence in Adams, Hunt, Palmer, and Zaliauskas (2019) and Bhutta, Fuster, and Hizmo (2019) that many consumers underestimate the benefits and overestimate the costs of shopping for deposit products and mortgages, respectively.

ing) changes discontinuously at a FICO score of 620. Bubb and Kaufman (2014) provide evidence for other discrete FICO thresholds in the mortgage underwriting process, including detailing the likely genesis of threshold-based policies. More recently, Laufer and Paciorek (2016) evaluate the consequences of minimum credit-score thresholds for mortgage lending, Agarwal et al. (2017) use FICO-based discontinuities in credit limits to estimate heterogeneity in marginal propensities to borrow and lend, and Aneja and Avenancio-León (2019) use credit-score discontinuities to document that ex-convicts with higher credit limits are less likely to recidivate. 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 *pricing* rules and to link those discontinuities to price dispersion, costly consumer search, and effects on final-good consumption.

Finally, we 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, Lanning, Low, and Salz (2019). Two companion papers use a similar dataset to study the importance of loan maturity and payment sizes in the market for cars and car loans. Argyle, Nadauld, Pratt, and Palmer (2018) find that payment-size shocks arising from shocks to maturity affect bargaining outcomes and are capitalized into transaction prices. Argyle, Nadauld, and Palmer (2019) document behavioral frictions associated with household debt decision making, showing that even unconstrained households make high-stakes debt decisions based on a monthly category-specific household budget. In contrast, in this paper, we focus on the industrial organization of local credit markets, highlighting the importance of search frictions in sustaining price dispersion and documenting how interest-rate markups affect substitution patterns in the market for cars. In at least one respect, the identification strategies of Argyle et al. (2018) and Argyle et al. (2019) rely on the search frictions documented in this paper; the various nonlinear lender policies exploited for identification by those papers would have no effect in a perfectly competitive market where consumers are fully informed of all prices.

3 Search Model with Elastic Demand

In this section, we illustrate how search frictions in the market for one good can reduce consumption in both that market and the market for complements, lowering aggregate welfare. We study the specific case of costly search in a credit market and consider downstream effects on durable consumption, but our results apply more broadly to the demand for any complementary products. As emphasized above, most 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.[8] 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 that we test below.

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. Indirect utility is given by

$$U(r,p,W) = V(r,p) + W, \tag{1}$$

where $V(\cdot,\cdot)$ is the indirect utility of consuming the consumption bundle of loan and durables chosen when facing prices r and p.[9] However, instead of inelastic demand, we assume that demand for loans and durables slope downwards 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

$$q(r,p) = a \cdot r^{\eta_r} \cdot p^{\eta_{rp}}, \tag{2}$$

and demand for durables x can be represented as $x(p,r) = b \cdot p^{\eta_p} \cdot r^{\eta_{pr}}$.[10] 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

[8] 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.

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

[10] We represent durables as a composite good x mapping all durable goods onto a scalar measure of durable consumption.

lender by drawing uniformly from the distribution of lenders but incur a search cost k for each quote they obtain. When their current quote is r', the expected utility gain of one additional search is

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

As shown by De Groot (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 (2). 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 (3) 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 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 $\overline{c} > \underline{c}$.[11] Without loss of generality, assume $\overline{c} \leq \underline{c}\eta_r/(1+\eta_r)$.[12] Lenders are perfectly informed of k , $m(k)$, and $F(\cdot)$. Their objective is to choose an interest

[11]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 deposit competition, etc.

[12]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.

rate r_j to maximize expected profits π_j, given by

$$\max_{r_j} E(\pi_j) = \begin{cases} (r_j - c_j)q(r_j, p)E(N_j) & \text{for } r_j \leq m(k) \\ 0 & \text{for } r_j > m(k) \end{cases} \tag{4}$$

where N_j is the mass of borrowers that each take out $q(r_j)$ in debt from lender j when facing r_j.

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 arising from the cost heterogeneity of firms and the elasticity of demand.[13] 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

$$\int_{\underline{r}}^{m(k)} [V(r, p) - V(m(k), p)]dF_{m(k)}(r) = k, \tag{5}$$

where $F_{m(k)}(r)$ is the endogenous distribution of interest rates given reservation price $m(k)$. Note that (5) 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 De Groot (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 (5), 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

$$\max_{r_j} E(\pi_j) = \max_{r_j} (r_j - c_j)ar_j^{\eta_r} p^{\eta_{rp}} \tag{6}$$

whenever $r_j \leq m(k)$. Holding durables prices p fixed, the first-order condition for (6) 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

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

[13]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.

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

$$F_{m(k)}(r) = \begin{cases} G[r(1+\eta_r)/\eta_r] & \text{for } \underline{r} \leq r < m(k) \\ 1 & \text{for } r = m(k) \end{cases}. \tag{7}$$

Given the distribution of interest rates (7) resulting from firms' best response to borrowers' reservation price, (5) 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 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

$$DWL = \int_{\underline{c}}^{\bar{c}} \int_{q(r^*(c),p)}^{q(\underline{c},p)} (r(q) - \underline{c})\, dqdG(c) + \int_{\underline{c}}^{\bar{c}} \int_{0}^{q(r^*(c),p)} (c - \underline{c})\, dqdG(c) \tag{8}$$

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 $G(\cdot)$.

Figure 2 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 Figure 2. Under costly search, lenders would earn a lender surplus of monopoly rents denoted LS in Figure 2 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., 2019).

The first term in (8) 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 Figure 2 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 Figure 2, 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 (8) 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 Figure 2. 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 $(c_j - \underline{c})$.

3.1 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.[14] 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 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

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

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

$$\int_{\underline{r}}^{m(k')} [V(r,p) - V(m(k'),p)] dF_{m(k')}(r) = k'. \tag{9}$$

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 should therefore be 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 $m(k)$.[15]

In the empirical section below, 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.

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.

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 Using the expression for deadweight loss in (8), we can demonstrate that the aggregate welfare loss from

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

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 (8) 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 (8) 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 (8) 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 2.

3.2 Identification

A challenge with evaluating the above predictions empirically is unobserved heterogeneity across high- and low-search-cost areas. The ideal experiment would randomly assign search costs across markets for identification. Instead, if we observe that loan sizes and car purchase prices are smaller in low-search-cost markets, this could be due to other unobserved demand factors potentially correlated with search cost proxies, such as credit limits, income, financial literacy, and preferences.

To address this challenge, we will focus on quasi-experimental variation that randomly assigns the first rate quote borrowers receive to be high or low, as we detail in section 6. This approach essentially exploits variation in r' in equation (3) to induce exogenous variation among otherwise identical borrowers' returns to search. When search costs are low, the first quote should not matter: given the equilibrium search conditions above, borrowers will have a relatively low reservation price and markups will be lower. In other words, 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 high search costs will prevent them from searching further. Given their downward sloping demand functions for loans and cars, these search-cost-supported higher markups should lead consumers to have lower loan sizes and purchase prices.

By comparing ex-ante identical borrowers in the same market, we can rule out demand shifters correlated with search costs that affect the level of loan sizes and consumption. Moreover, our regression-discontinuity tests in section 6.3 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. If search costs are the driving factor, then we would expect that being assigned a high r' versus a low r' should be of less consequence 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.

4 Data

We analyze the loan contract terms and auto 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 (98.5%) were originated by credit unions ranging between \$100 million and \$4 billion in asset size, with the remainder originated by non-bank finance companies.[16] Borrowers from all 50 states are represented in the data, but the five largest states in the data are Washington, Texas, California, Minnesota, and Tennessee.

The dataset contains information capturing all three stages of a loan's life: application,

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

origination, and ex-post performance, although we have loan application data for only approximately 1.3 million loans from 41 different institutions. The available loan application data report borrower characteristics (ethnicity, age, gender, FICO scores, and debt-to-income (DTI) ratios at the time of application), whether a loan application was approved or denied, and whether it was subsequently withdrawn or originated. For originated loans, the data additionally include information on loan amounts, loan terms, car purchase prices, and collateral characteristics. Using Vehicle Identification Numbers (VINs), we learn about the make, model and model year of the purchased car. We restrict our sample to direct loans (not intermediated by a dealer) in an effort to address concerns that indirect loans are potentially endogenously steered to specific financial institutions (perhaps because car dealers become aware of lenders' pricing rules over time).[17] 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 A, B, and C of Table 1 present summary statistics on loan applications, loan originations, and measures of ex-post performance, respectively. As reported in Panel A of Table 1, the median loan application in our data seeks approval for a six-year $20,000 loan at a median interest rate of 4.00%.[18] 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. Throughout the paper we refer to the number of loans originated divided by the number of applications approved for a particular group as the loan take-up rate. Panel B 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 interquartile range of $195.

Panel C tabulates measures of ex-post loan performance. While the average loan is 23.4 days delinquent, 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%. In untabulated results, we find that the default rate for sub-600 FICO borrowers is 6.8%,

[17]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. See Appendix Table A1 for summary statistics for the excluded indirect loans.

[18]Interest rates in the loan application data refer to approved loans, whether they were subsequently originated or not.

compared to a default rate of 2.6% for borrowers with FICOs between 600 and 700 and 1.6% for over-700 FICO borrowers. Lending institutions periodically check the credit score of their borrowers subsequent to loan origination, creating a novel feature of our data. Summary statistics for ΔFICO represent percent changes in borrowers' FICO scores from the time of origination to the lender's most recent (soft) pull of their FICO score.[19] 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.

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. In the auto loan data made available to our data provider by its clients, roughly half 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 section A.3.

Credit-union borrowers do differ slightly 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 Community Survey.[20] 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 draw inference given that we rely on a regression discontinuity (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 constitutes a very large set of auto-loan borrowers.

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

[20]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.

A final external-validity issue involves the distribution of loan originations through time. Over 70% of loan originations in our sample occurred between 2012 and 2015, despite a sample period that runs from 2005 to 2016. The large increase in loans through time reflects the increase in the client base of our data provider through time rather than auto credit origination in general. CCP data shows that auto loan originations in the general population have increased through time as well, from an aggregate outstanding balance of $725 billion in Q1 2005 to just over $1.15 trillion in Q4 2016, but not at the rate reflected in our dataset. We view the non-representative time series of our data as less relevant given our cross-sectional identification strategy.

5 Documenting Price Dispersion

The model in section 3 describes sufficient conditions for equilibrium price dispersion to persist in a credit market: heterogeneity in lender costs, positive search costs for borrowers, and elastic borrow demand for credit. Under these assumptions, the distribution of interest rates $F_{m(k)}(\cdot)$ defined in (7) will be nondegenerate and the law of one price will not hold. Empirically diagnosing a market with dispersed prices requires ruling out any product differentiation, i.e., establishing that differences in prices truly represent identical goods being sold for different prices in the same market. While our main results rely on quasi-experimental variation for identification, in this section, we provide evidence that significant dispersion persists in the market for car loans even after nonparametrically controlling for many dimensions of product heterogeneity.

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. We consider loans originated to borrowers within the same CZ $\times$ quarter $\times$ price $\times$ FICO $\times$ maturity $\times$ DTI cell to be observably identical.[21] 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. 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

[21] 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.

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 as having the universe of interest rates offered to a given cell could only weakly decrease the best available rate for each type.

Figure 3 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, 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 1.3 percentage points more 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 Figure 3, we find that the average borrower would need to obtain three price quotes to find the best available rate for that borrower's type.

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.

6 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 subsequently to estimate the impact of costly search on equilibrium outcomes in both credit and durables markets. We first isolate exogenous variation in markups *within* lender and then demonstrate that such variation is predictive of the returns to search *across* lenders. 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 product characteristics. To address this while estimating the effect of search-cost–induced markups on outcomes, we exploit quasi-experimental 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. According to equation (3), a prospective borrower will continue to search given initial quote r' if the expected utility gain from searching exceeds her search cost k. Facing a high initial quote r' should be relatively inconsequential for borrowers in markets with low search costs. Borrowers with high and low initial quotes r' should have similar loan and durable-consumption outcomes in low-search-cost markets both because such borrowers treated with high-markup initial quotes should be willing to search, and because, as a result, equilibrium markups in low-search-cost markets should be smaller. In high-search-cost markets, however, markups should be larger and borrowers should be more reticent to undo them by searching more such that rate discontinuities should have stronger effects on search behavior and borrowing and consumption outcomes in high versus low-search-cost markets.

6.1 Detecting Discontinuities

Lenders make underwriting and loan pricing decisions based on information observable ("hard") and unobservable ("soft") to the econometrician (Liberti and Petersen, 2019). Our ability to draw inference 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 for mortgages, no industry standard set of thresholds exist in auto lending.[22] Still, 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

[22]Also in contrast to Keys et al. (2010), FICO thresholds observed in our data have little to do with secondary markets given that most auto loans are retained by the lending institutions in our dataset.

than half of the lenders in our data. See Bubb and Kaufman (2014), Livshits, Mac Gee, and Tertilt (2016), and Agarwal et al. (2017) for models of credit risk processing costs and Al-Najjar and Pai (2014) for a model of overfitting that could 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.[23]

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

$$r_{il} = \sum_k \psi_{kl} 1(FICO_i \in \mathcal{F}_k) + \varepsilon_{il} \tag{10}$$

where ε_{il} includes all other factors that influence loan pricing for a given loan. The five-point FICO bins begin at a FICO score of 500 where the first bin includes FICO scores in the 500-504 range, the second bin includes 505-509, etc., up through FICO scores of 800. The estimated coefficients on each FICO bin represent the average interest rate for loans originated to borrowers with FICO scores in that bin relative to the estimated constant (the omitted category is loans outside this range—we focus on relative magnitudes for this exercise).

Figure 1 presents interest-rate plots for three different financial institutions. The point estimates $\hat{\psi}$ represent how that lender's pricing rules appear to vary with borrower FICO score, and the accompanying 95% confidence intervals provide a sense of how reliant on FICO scores each lender's pricing rule was. Panel A of Figure 1, 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 bin, and the difference in average interest rates between the two bins is statistically significant at the 0.001 level. Panels

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

B and C of Figure 1 illustrate similar rule-of-thumb FICO breaks for unique institutions with approximately 6,000 and 25,000 loans, respectively. 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 standardize our analysis to include every institution that employs discontinuous pricing rules, we empirically identify the existence of discontinuities at each institution (if they exist at all) in our sample through the following criteria. We first estimate the interest-rate FICO bin regressions following equation (10) 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 further refine the set of discontinuities by requiring that differences between leading and following FICO bin coefficients ψ_{ck} have a p-value of at least 0.1.[24] We further examine each potential threshold visually to ensure that the identified discontinuities are well behaved around the candidate thresholds. Following Angrist and Lavy (1999), Table 2 reports summary statistics for the 514,834 loans in the discontinuity sample (the set of loans within 19 points of one of our detected thresholds) used in estimation.

6.2 First-Stage Results

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. 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 any borrower attribute (observed or unobserved) that could influence loan outcomes changes only continuously at interest-rate discontinuities. This smoothness condition requires that borrowers on either side of a FICO threshold are otherwise similar, such that borrowing outcomes on either side of a threshold would be continuous absent the difference in treatment induced by policy differences at the threshold.

In our remaining specifications, we normalize FICO scores to create a running variable $\widetilde{FICO}$ that measures distance from an interest-rate discontinuity. For example, for loans near the 600 FICO score threshold, $\widetilde{FICO}_i = FICO_i - 600$. Panel A of Figure 4 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

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

38-point FICO ranges for the 640 and 700 FICO thresholds in panels B and C of Figure 4. These plots confirm the existence of large interest-rate discontinuities at these thresholds. The narrow confidence intervals in Figures 1 and 4 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. To intuitively introduce our RD design, we first explain the RD estimation equation in the context of a single threshold and with linear controls for the running variable. The resulting specification would take the form

$$y_{iglt} = \pi_1 \widetilde{FICO}_i + \delta \cdot 1(\widetilde{FICO}_i \geq 0) + \pi_2 \widetilde{FICO}_i \cdot 1(\widetilde{FICO}_i \geq 0) + \alpha_{gt} + \gamma_l + \varepsilon_{iglt} \quad (11)$$

where y_{iglt} is the outcome for 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 $\times$ quarter and lender fixed effects, respectively. In this specification, δ is the key RD coefficient and estimates how loan terms y 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 (11) 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 a quadratic. Second, to deal with loans that may be in the discontinuity samples corresponding to multiple discontinuities, we sum across discontinuities d from the set of discontinuities $\mathcal{D}$ to estimate

$$y_{iglt} = \sum_{d \in \mathcal{D}} 1(il \in \mathcal{D}_d) \left(\delta \cdot 1(\widetilde{FICO}_{id} \geq 0) + f(\widetilde{FICO}_{id}; \pi) + \varphi_{dl} \right) + \alpha_{gt} + \varepsilon_{iglt} \quad (12)$$

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

$$f(x; \pi) = \pi_1 x + \pi_2 x^2 + 1(x \geq 0) \left(\pi_3 x + \pi_4 x^2 \right) \quad (13)$$

to allow for a smooth nonlinear effect of the running variable that potentially changes shape discontinuously at the threshold. Standard errors are double clustered by lender and FICO score, and the sample used to estimate (12) is the discontinuity sample described in Table (2).[25]

[25] 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.

Table 3 presents results of this exercise. Interest rates for borrowers with FICO scores immediately above a detected threshold are an average of 1.27 percentage points lower than borrowers just below. Column 2 reports that loan maturities for borrowers just above a FICO threshold are 0.8 months longer than otherwise similar borrowers below the threshold. Given an average interest rate in our estimation sample of 6.0% (Panel B of Table 2), the magnitude of these effects is economically meaningful, amounting to a $13 higher monthly payment and $440 higher present value for otherwise identical loans taken out by borrowers on the expensive side of a FICO discontinuity. In the context of our theoretical model, 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 search is costly, could have material consequences on the ultimate cost of credit for such borrowers.

6.3 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 Figure 5, we plot these values and associated 95% confidence intervals along with a fourth-order polynomial robust RD function estimated following Calonico et al. (2014).[26] Importantly, 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 A–E plot borrower DTI ratios, loan amounts, borrower age in years, borrower gender (an indicator for male), and borrower ethnicity (an indicator for white), respectively. These plots indicate smoothness in ex-ante borrower characteristics around FICO thresholds. Borrowers on either side of FICO thresholds do not appear meaningfully different in terms of their debt capacity, their willingness to borrow, or demographics. Panel F plots the number of applicants within each normalized FICO bin, along with the McCrary (2008) test for bunching in the running variable. While such manipulation of the running variable—a discontinuity in the propensity to apply for a loan at a FICO threshold—would raise selection concerns, Panel F 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. In any event, such targeting 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.

[26]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.

Table 4 uses the specification in (12) 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. 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.

While loan interest rates change by an average of nearly 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.

6.4 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, in the model, if search *across* lenders were costless (or demand inelastic), a high r' quote would have no effect on transaction outcomes because borrowers would find the lowest-price lender. 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.[27]

Figure 6 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 section 5. Dotted and solid lines in each plot are for borrowers just below and above a given threshold, respectively. For each threshold, those offered exogenously higher rates, realize substantially higher average spreads from the lowest available rate (and larger variance in those spreads). 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 A2 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.

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

7 Evidence on Loan Search and Search Costs

The persistence of high and dispersed markups in equilibrium—particularly given their detection using quasi-exogenous pricing discontinuities—is prima facie evidence that borrowers find soliciting the full menu of prices costly. Otherwise, as in our model, we would expect borrowers to search until finding the lowest available price and that the resulting competition would drive dispersion to zero. However, while the results of section 6 suggest that the initial assignment of markups to borrowers on either side of a FICO discontinuity is as good as random, attributing variation in subsequent outcomes conditional on origination to search frictions raises selection concerns. In other words, the possibility that borrowers who accept high markups are different on other dimensions from borrowers accepting lower markups prevents simple comparisons of outcomes conditional on markup size, even when these markups have been randomly assigned. In this section, we introduce a simple geographic-based measure of search costs and use our RD apparatus to demonstrate that search behavior and several key predictions of our model line up with this search-cost proxy. Appendix A contains several robustness checks addressing concerns that the response of borrowers across FICO discontinuities to our search-cost measure could be correlated with other characteristics of borrowers in high-search-cost areas.

7.1 Measures of Loan Search and Search Costs

Can costly search explain why many borrowers randomly assigned expensive rates do not avail themselves of better credit terms available elsewhere? In this section, we evaluate whether a proxy for search costs can explain borrowers' apparent reluctance to shop, necessitating measures of both search behavior and search costs.

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.[28] Such borrowers may seek loan preapproval before negotiating with the seller over purchase price to refine their own budget or avoid double marginalization (Busse and Silva-Risso, 2010; Grunewald et al., 2019). Alternatively, buyers may finalize a purchase price and then shop around for car loans before completing the purchase.

As car-loan pricing is specific to the credit risk of each individual, obtaining price quotes in this market most often entails filling out a loan application, undergoing a credit check,

[28] Recall that our sample consists of so-called 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., 2019).

and potentially verifying assets and income. Here, 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.[29] 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, 2016).[30]

To proxy for distance-based search costs, we use FDIC and NCUA data to identify the physical location of every bank branch and credit union branch in the United States for each year in our application data. We then create a measure of proximity to financial institutions (PFI) 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, although our results are robust to the precise time cutoff used (see section 7.1.1).[31] This driving-time density measure PFI 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.[32] Supporting this search-cost proxy, Degryse and Ongena (2005) find evidence of the important role of transportation costs in local credit markets.

The median borrower in our application data lives within a 20-minute drive of 72 lending institutions. Borrowers in the 25th percentile of driving distance live less than a 20-minute drive from 23 institutions, as compared to 168 institutions for borrowers in the 75th percentile. Just under 4% of applicants in our sample live in an area with one or fewer lending institutions within a 20-minute drive. In contrast, 45% of applicants live within a 20-minute drive of at least 100 different lending institutions. Our baseline results categorize borrowers with fewer than 10 lending institutions within a 20-minute drive as facing high search costs (lowest 15% in the density distribution) to capture nonlinearity in the effect of additional

[29]We discuss the option borrowers have to search for loans online in Appendix Section A.3.

[30]Note that we do not consider several other plausible correlates of search costs in our data because of their likely 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.

[31]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.

[32]While distance can also proxy for soft-information producing relationships (see Nguyen, 2019; Granja, Leuz, and Rajan, 2018), we do not believe auto-loan lending to be 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.

nearby lenders on search costs.[33] In section 7.1.1 below, we verify robustness of our results to the definition of high-search-cost area. In Appendix B, we demonstrate that the structural search model of Hortaçsu and Syverson (2004) also estimates low-PFI areas to have higher search costs.

Our first measure of search propensity is whether an applicant accepts a given loan offer. Loan take-up proxies for search insofar as prospective borrowers who applied for but then decline an offered loan do so in favor of accepting a different loan. While applicants could decline offered loans for other reasons, including deciding not to originate any loan, we view differences in the decision to accept an offered loan around lending thresholds as a reasonable measure of differences in borrowers' shopping behavior.

Using application FICO scores, we estimate differences in take-up rates around FICO thresholds. If PFI is correlated with search costs and thus influences the propensity to search, differences in loan take-up-rates around FICO thresholds should be larger in areas with higher PFI (i.e., more lenders live close by, thus lowering search costs). We predict that applicants in low-PFI areas, facing higher search costs, will be less likely to reject unfavorable loan terms and search for better terms elsewhere. Importantly, in a difference-in-differences spirit, our empirical specification measures *differences* in loan take-up rates around FICO thresholds. This allows us to exploit the 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 space.[34] We then compare differences in loan take-up rates for borrowers in high versus low-PFI areas.

Using our standard RD framework described by equation (12), we specify the dependent variable as an indicator equal to one for accepted loans and estimate differences in the take-up rate around FICO thresholds. In Table 5, we estimate effects for the full sample and separately for borrowers with more and less than 10 lending institutions within a 20-minute drive (low and high search costs, respectively). For the full sample, column 1 documents differences in take-up rates of 12.1 pp around the threshold; that is, borrowers randomly drawing high markups from the distribution of interest rates are 12.1 percentage points less likely to accept the offered loan. We estimate the same regression in high- and low-

[33]In unreported results, we split the sample in other ways based on quantiles of the driving-density distribution. Top quartile versus bottom quartile driving densities produce estimates qualitatively similar to our reported results. Above- and below-median splits produce similarly signed (but smaller) differences in the difference of take-up rates across thresholds.

[34]While our RD empirical strategy does not require that applicants in high and low search cost areas are otherwise identical, we note that they do look similar on observables. For example, 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 lower FICO scores (6 points) but also lower DTI ratios (1.5 points). See Appendix A for further results addressing endogeneity concerns.

search-cost areas and report results in columns 2 and 3, respectively. In high-search-cost areas, take-up rates differ by only two percentage points between borrowers on either side of the discontinuity. In contrast, the difference in take-up rates around FICO thresholds is 13.7 pp in low-search-cost areas (column 3). In other words, borrowers in low-search-cost areas are much less likely to accept a loan if they are assigned a high interest-rate markup, whereas high-search-cost borrowers are likely to accept loan terms regardless of where they fall in the distribution of prices.[35] Finally, column 4 shows that the difference between the discontinuity coefficients in the high- and low-search-cost samples is statistically significant. Given the mean take-up rate of 0.51 (Table 2), being in a high-search-cost has over a 20% effect on the likelihood a borrower will accept an expensive loan.

7.1.1 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 augment our RD specification (12) with two additional controls: a High Search Cost indicator and an interaction between this dummy and the discontinuity indicator $1(\widetilde{FICO}_{id} \geq 0)$. We then vary the cutoff used for the definition of the High Search Cost dummy and plot the results in Figure 7. Each point plots the coefficient on the interaction between High Search $\text{Cost}_i \times 1(\widetilde{FICO}_{id} \geq 0)$ from a separate regression. Along the x-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 PFIs on the x-axis is 10, the point estimate corresponds to column 4 of Table 5.

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 127 bp lower) than just-below applicants. Being in a high-search-cost area nearly 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 first interest rate r' because they would find

[35] One threat to our interpretation of Table 5 is that the magnitude of the interest-rate discontinuity 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. Appendix Table A3 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. See Appendix Section A.1 for further discussion.

shopping for a better rate particularly costly relative to low-search-cost borrowers. Second, while borrowers with five or fewer PFIs are rare and so the resulting take-up discontinuity is estimated imprecisely, the other estimates for borrowers who have less than 25 PFIs are significantly less than zero. Third, estimates for high-search-cost definitions near our baseline cutoff of 10 PFIs 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 PFIs 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 PFI diminishes (i.e., each lender matters a lot 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.

7.1.2 Explaining Variation in Price Dispersion with Search Costs

Our first comparative static in section 3.1 is that markets with lower search costs should result have less price dispersion and smaller markups, as both of these objects should respond to consumers facing lower search costs being more likely to be informed about the complete distribution of available prices. We test this prediction using our PFI measure of search costs. In Figure 8, 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 A4 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.

7.2 Inferring Loan Shopping from Multiple Applications

To construct a more direct measure of loan shopping, we exploit the ability of our data to identify the same borrower filing a loan application at multiple financial institutions. 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. As before, we define high-search-cost areas to be places where borrowers live within a 20-minute drive of at most 10 lending institutions.

Every prospective borrower in our application data, by construction, applies for at least one loan. Table 6 documents that we observe applicants in high-search-cost areas (column 1) applying for an average of 0.34 additional loans loans. In contrast, we observe applicants in low-search-cost areas (column 2) applying for an average of 0.41 loans per vehicle purchase. Column 3 shows that this difference is statistically significant at a 1% significance level. Although the difference (0.067 loans) is small in magnitude, these estimates suggest that one in every 15 borrowers applies for an extra loan only if in a low-search-cost area. Importantly, these average differences are a lower bound as we do not observe applications to any lender not in our data. Note, too, that equilibrium applications need not change to be affected by search—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 7.1 that borrowers facing high search costs search less and accept worse rates than borrowers facing relatively lower search costs.

7.3 Identification

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. However, there could be several time-varying or time-invariant factors correlated with our search cost measure that could affect equilibrium price dispersion and take-up behavior. 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. Instead, such factors bias our estimates 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.

To address this possibility, in Appendix A, we present several robustness exercises and a discussion of identification issues related to our measure of search costs. First, a Bartik strategy based on the 1990 network of lender branches in the United States allows us to address potential time-varying endogeneity in the number of proximate lenders (e.g., that banks close branches in response to hyperlocal economic conditions that also determine interest-rate elasticities). Because shift-share instruments essentially rely on the exogeneity of preexisting conditions, we also present a difference-in-differences strategy and a hyperlocal fixed effects strategy that allow us to rule out time-invariant explanations for our results (for example, the concern that low financial sophistication or brand loyalty covaries across space with interest-rate sensitivity and the density of financial providers). The results are consistent with those presented above and support our use of nearby lender density combined with our RD strategy to measure the consequences of search frictions in consumer credit markets.

8 Effects on Loan and Consumption Quantities

To test the second and third predictions of our model—that search frictions in credit markets can have real effects on consumption—we next establish that being treated with a higher interest rate affects loan and purchase decisions. Whether a given credit-market imperfection constrains consumption is difficult to ascertain 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 result that borrowers are ex-ante similar around FICO thresholds, our operating assumption is 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 that each borrower borrowers and spends 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.

Figure 9 plots car purchase amounts 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 first-stage analysis above, we formally test for statistical differences in purchase amounts. As before, we estimate equation (12) 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.

Table 7 presents these reduced-form results. Borrowers quasi-randomly offered more expensive loans spend an average of \$377 less on the cars they purchase (a 2.4% effect). Column 2 presents results with loan amounts as the dependent variable. Originated loan sizes are an average of \$566 lower (4%) on the expensive side of a detected FICO discontinuity. The fact that loan sizes increase by larger amounts around the threshold than purchase amounts indicates that, ex-post, borrowers on the right side of the cutoff are approved for and take up higher loan-to-value (LTV) ratios. Column 3 of Table 7 indicates that ex-post LTV ratios are an average of 1.3 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 (Table 4), we interpret these results as further evidence of the easing of credit terms for above-threshold borrowers. That is, ex-post, borrowers with FICO scores just above a pricing discontinuity are offered lower rates, longer terms, and apparently allowed higher ex-post LTV and DTI ratios.

Microdata on loan amounts and loan terms allow us to calculate the implied monthly payment of borrowers on either side of the thresholds. In column 4 of Table 7, we test whether ex-post monthly payments are different around the thresholds. On average, monthly payments increase by a statistically and economically insignificant \$0.17 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.[36]

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, an essential component of the welfare effects predicted by our model. Absent search frictions, we would expect all borrowers to find the lowest-price provider, and given the balance of borrowers across discontinuities, we would also expect to observe similar purchasing decisions. 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.[37]

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 8 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 \$345 less on cars, suggesting that the bulk of the purchasing behavior we observe in Table 7 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 \$80.[38] Because fixing the model

[36]See Argyle et al. (2018) for related evidence on borrowers' monthly payment targeting.

[37]See Argyle et al. (2019) for evidence that individual-level used-car prices capitalize financing terms.

[38]We hesitate to overinterpret column 2 as strong evidence of a null effect on purchase prices conditional

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.

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 (low) interest-rate markup by buying a car that is one model year older (newer), keeping their monthly payments roughly constant. These car-age effects are likely to represent some borrowers preferring to purchase older cars over searching for better financing. Appendix Figure A1 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. (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 age are consistent with search frictions changing the quantity of car services purchased by borrowers with high search costs, affirming real (as opposed to purely financial) effects of search frictions. Under the assumption that borrowers on either side of pricing discontinuities are otherwise identical, we would expect them to find similar financing opportunities and purchase similar cars absent search frictions. Our model also predicts that any quantity distortions should be stronger in markets with higher search costs. To reinforce our interpretation that these effects are driven by costly search per se, we again divide the sample into high- and low-search-cost areas and test whether our effects are stronger for borrowers in higher-search-cost areas. Table 9 shows that borrowers in both relatively high- and low-search-cost areas are affected by the interest-rate markup offered by a given firm; FICO pricing discontinuities affect loan sizes and purchase prices in both types of markets. However, comparing discontinuity effects for high-search-cost areas (columns 1 and 3) to discontinuity effects in low-search-cost areas (columns 2 and 4), a borrower's draw of a given markup r' is much more consequential when search costs are high, as predicted by

on the car chosen given the loss of power from including so many fixed effects. Borrowers with more affordable credit could be paying slightly more for the same make-model-year either because their marginal willingness to pay increased and was extracted by the dealer or by choosing a nicer car within a make-model-year (extra add-on features, lower mileage, etc.). Still, we can reject the column 1 effect size using the column 2 estimates.

the model. While markups would also distort quantities in an oligopolistic market, too, we show in Appendix A.4 that the contrast between high- and low-search-cost borrowers holds in both concentrated and less concentrated markets.

How do borrowers respond to being arbitrarily offered high financing markups? The evidence presented in Tables 7 and 8 indicates that borrowers offered 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, potentially bargaining harder on purchase price and originating smaller loans at higher loan-to-value ratios. Moreover, these effects are economically and statistically significantly larger in areas we measure as having high search costs. We view this as evidence that borrowers' inability to costlessly identify the best available loan terms distorts consumption away from the efficient quantities that would prevail absent search frictions.

8.1 Assessing Selection into Origination

Limiting our sample to direct loans eliminates the possibility of borrower selection driven by car-seller steering of borrowers to lenders. 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, as mentioned previously, 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. For example, perhaps borrowers who accept high loan markups are particularly inelastic, and the larger car expenditures for below-threshold borrowers are the result of conditional-on-origination effects drawing disproportionately from higher-demand, inelastic borrowers. 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.

The above evidence demonstrates that interest-rate markups seem quasi-randomly assigned (borrowers are similar on observable dimensions at the ex-ante application stage). 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 follow through with the origination of 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.

We test for the possibility that such selection drives the observed equilibrium outcomes in our data by comparing the balance around the FICO discontinuities of borrower characteristics and ex-post borrower performance conditional on origination. 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 A2 repeats the exercise of Figure 5 by checking for the smoothness of borrower and loan characteristics around FICO discontinuities. We merge our application data to our origination data to demonstrate that the same smoothness in borrower attributes that we saw at the application stage persists conditional on origination. 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 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 B and F show that loan sizes and the fraction of borrowers accepting loans do respond to interest-rate markups, consistent with the extensive- and intensive-margin effects of the discontinuities we estimate in Tables 5 and 7, respectively. The results of the Appendix Figure A2 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.

To test for selection on (unobserved) creditworthiness, we specify as a dependent variable in our RD setting various ex-post credit outcomes as any selection on credit quality should eventually manifest itself in the average ex-post credit performance of selected borrowers. The coefficient in column 1 of Table 10 estimates that above-threshold borrowers are an average of 4.2 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.[39] 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 0.2 percentage points more likely to be in default (over 90 days past due, column 3), both of which we view as relatively precise zeroes.

[39]The sample size differs across columns in Table 10 because of inconsistent data coverage of all monitoring variables across lenders.

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.[40] Ex-post credit scores allow us to calculate changes in credit scores over time, capturing broad changes in borrower distress and financial responsibility that extend 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 test incorporating credit performance for other products, too, 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 10 show no meaningful differences (0.1 percentage points) in credit score changes for borrowers around the threshold.[41]

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 high- and low-markup borrowers 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. We address the robustness of our results to alternative sources of variation in local branch networks in Appendix A.

[40]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.

[41]While our effects on monthly payments suggest that there is unlikely to be a causal effect of higher interest rates themselves on default outcomes, note that this would work against our finding of null effects on credit outcomes for high-markup borrowers.

9 Conclusion

Mounting evidence indicates that 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. Motivated by our model of search for credit when credit demand is both elastic and values complementary durable consumption, we present empirical evidence connecting these two literatures.

First, we present evidence that substantial price dispersion exists in retail auto lending markets. About 54% of borrowers in our sample for whom a comparable borrower exists did not originate a loan with the lowest available interest rate. The average borrower in our data pays 130 bp more ($440 in present-value terms for the average loan) than the best rate available to observationally similar borrowers. In this setting where the gains to search are high, we show that borrowers' acceptance of dominated loan terms is related to measures of search costs. In particular, because arbitrary 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 should be unlikely to accept dominated loan terms. Proxying for the costliness of loan shopping with the density of nearby lenders, we show that borrowers in higher-search-cost areas face more dispersed prices, are more likely to accept quasi-randomly offered dominated loan terms, and apply for fewer loans.

Second, as illustrated by our model, we demonstrate that search costs have material effects on loan quantities and durable consumption. On average, borrowers that accept quasi-randomly offered higher interest rates borrow $566 less, spend $377 less on their car purchases, and buy 1.8 months older cars than otherwise similar borrowers that accept quasi-randomly offered lower rates. These results highlight the importance of well-functioning consumer credit markets in determining durable goods consumption patterns. Moreover, relative to popular search models with inelastic demand where dispersed prices have no associated deadweight loss and just represent a transfer from buyers to sellers, there appear to be aggregate welfare consequences of costly search in the real world. When search costs are non-negligible, consumers facing firm-specific markups for a given good 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 credit market imperfections such as costly search. 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 lowest available prices. Even with the possibility of shopping for interest rates online, searching for consumer credit products remains an opaque, local, and costly process for many borrowers. This relationship between costly search and 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 D, Stefan Hunt, Christopher Palmer, and Redis Zaliauskas (2019). "Testing the Effectiveness of Consumer Financial Disclosure: Experimental Evidence from Savings Accounts." NBER Working Paper No. 25718.

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

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 (2019). "Searching for Approval." Working Paper.

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

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

— (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 (2019). "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 (2018). "Monthly Payment Targeting and the Demand for Maturity." *Review of Financial Studies* Forthcoming.

Argyle, Bronson, Taylor Nadauld, Christopher Palmer, and Ryan Pratt (2019). "The Capitalization of Consumer Financing into Durable Goods Prices." *The Journal of Finance* Forthcoming.

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.

Bartlett, Jeff (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.

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

Bertrand, Marianne and Adair Morse (2011). "Information disclosure, cognitive biases, and payday borrowing." *The Journal of Finance* 66.6, pp. 1865–1893.

Bhutta, Neil, Andreas Fuster, and Aurel Hizmo (2019). "Paying Too Much? Price Dispersion in the US Mortgage Market." SSRN Working Paper No. 3422904.

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

Bulow, Jeremy I and Paul Pfleiderer (1983). "A note on the effect of cost changes on prices." *Journal of Political Economy* 91.1, pp. 182–185.

Busse, Meghan R., Nicola Lacetera, Devin G. Pope, Jorge Silva-Risso, and Justin R. Sydnor (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." *The 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.

De Los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest (2012). "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior." *American Economic Review* 102.6, pp. 2955–80.

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.

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.

Granja, João, Christian Leuz, and Raghuram Rajan (2018). “Going the Extra Mile: Distant Lending and Credit Cycles.” NBER Working Paper No. 25196.

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

Grunewald, Andreas, Jonathan Lanning, David Low, and Tobias Salz (2019). “Auto Dealer Loan Intermediation: Consumer Behavior and Competitive Effects.” Working Paper.

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 Singh Pathania (2016). “High-Cost Debt and Borrower Reputation: Evidence from the UK.” Working Paper.

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

Liebersohn, Jack (2019). “Does Local Bank Competition Still Matter? Quasi-Experimental Evidence from Bank Mergers.” Unpublished 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 Mac Gee, and Michele Tertilt (2016). "The democratization of credit and the rise in consumer bankruptcies." *The Review of Economic Studies* 83.4, pp. 1673–1710.

McCrary, Justin (2008). "Manipulation of the running variable in the regression discontinuity design: A density test." *Journal of Econometrics* 142.2, pp. 698–714.

New York, Federal Reserve Bank of (2019). "Quarterly Report on Household Debt and Credit."

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.

Salz, Tobias (2017). "Intermediation and competition in search markets: An empirical case study." SSRN Working Paper No. 2961795.

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

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

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

Stango, Victor and Jonathan Zinman (2009). "Exponential growth bias and household finance." *The Journal of Finance* 64.6, pp. 2807–2849.

— (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.

U.S. Department of Transportation, Federal Highway Administration (2017). *National Household Transportation Survey*. http://nhts.ornl.gov.

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

Zhu, Haoxiang (2012). "Finding a good price in opaque over-the-counter markets." *The Review of Financial Studies* 25.4, pp. 1255–1285.

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

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

A. Sample Lender #1

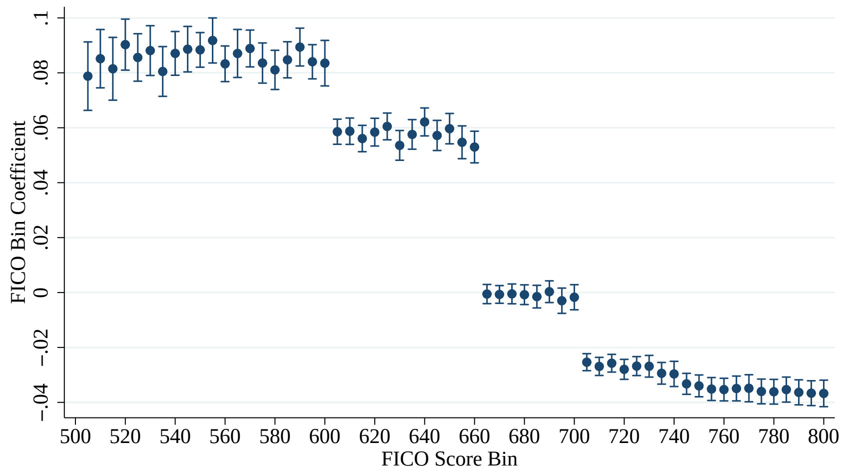

B. Sample Lender #2

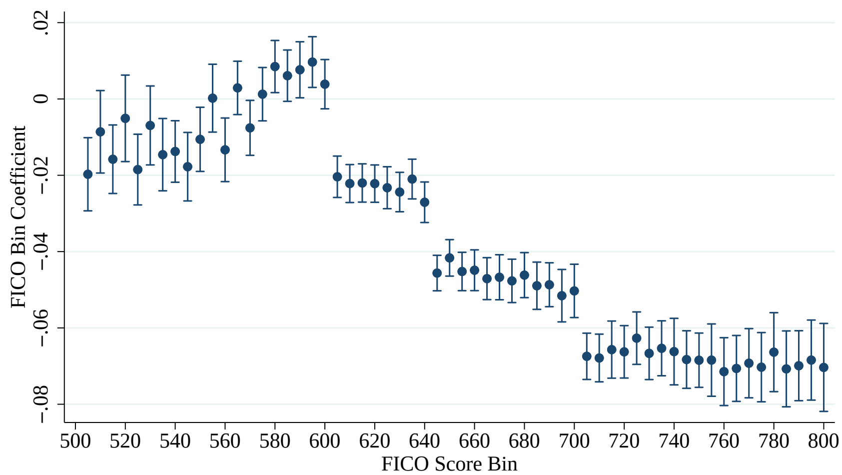

C. Sample Lender #3

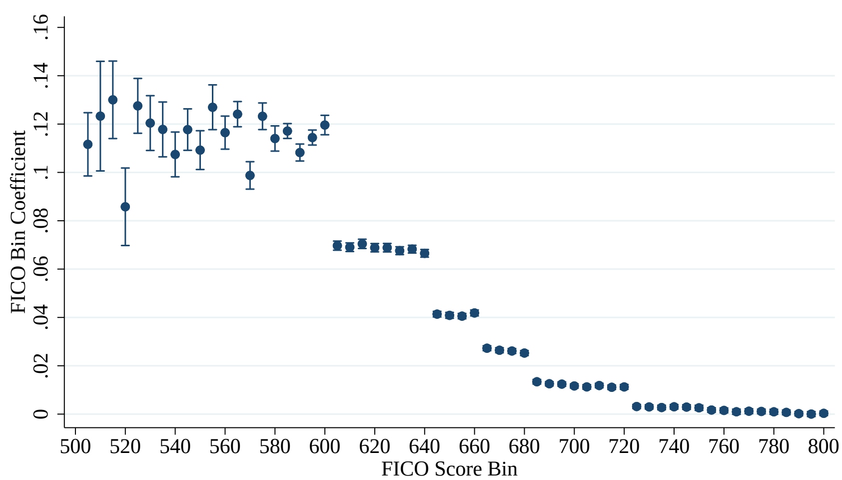

Notes: Each panel plots estimated interest-rate rules (with 95% confidence intervals) for a different lender in our sample. Loan rates are regressed on five-point FICO bin indicators as in equation (10).

Figure 2: Illustration of Surplus and Welfare Loss from Search Frictions

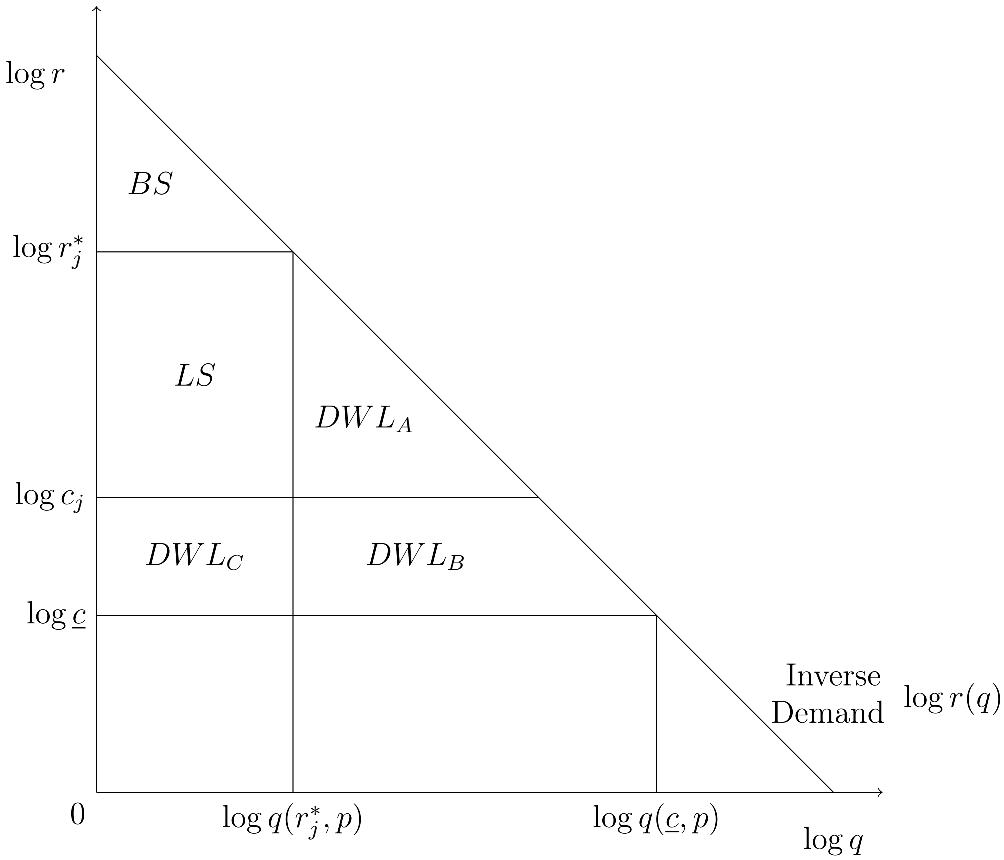

Notes: Figure plots equilibrium outcomes under perfect competition and costly search in log interest rate (y-axis) and log quantity (x-axis) space. Log inverse demand is the diagonal line denoted $\log r(q)$. The perfect-competition equilibrium quantity and price are denoted $q(\underline{c}, p)$ and $\underline{c}$, 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.

Figure 3: Density of Spread to Lowest Available Rate

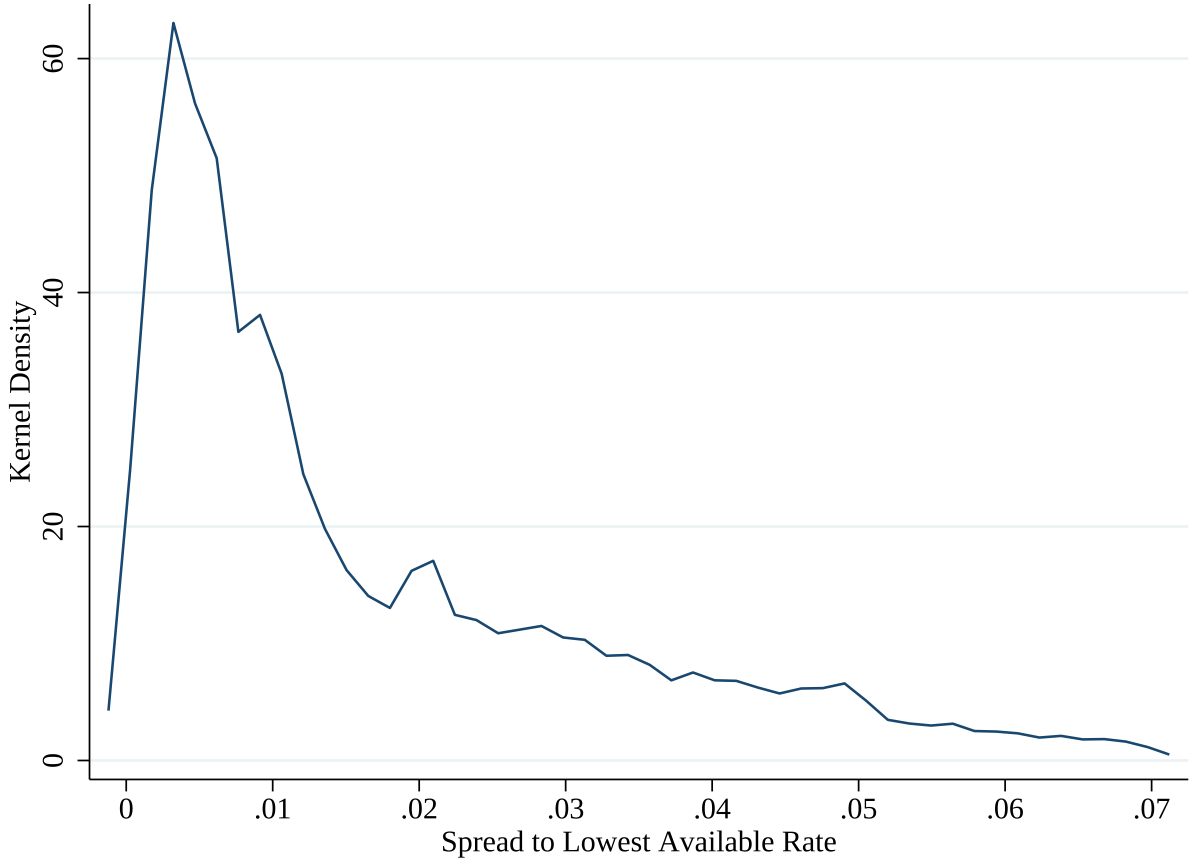

Notes: Figure reports the kernel density of the spread (in percentage points) to the lowest available rate for borrowers not receiving the best available rate in their cell. A cell is defined as all borrowers in the same commuting zone taking out a loan in the same $1,000 collateral-value bin, five-point FICO bin, 10-point DTI bin, six-month time period, and loan maturity. Estimated density uses an Epanechnikov kernel with a bandwidth of 0.0012.

Figure 4: Interest-rate FICO Regression Discontinuity Plots

A. Interest Rates Around FICO = 600 Discontinuities

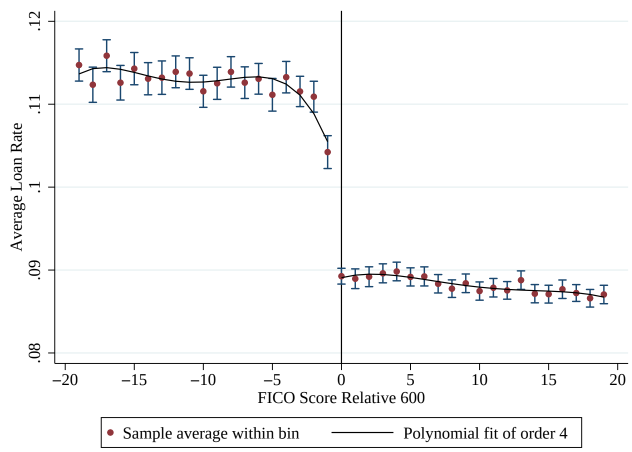

B. Interest Rates Around FICO = 640 Discontinuities

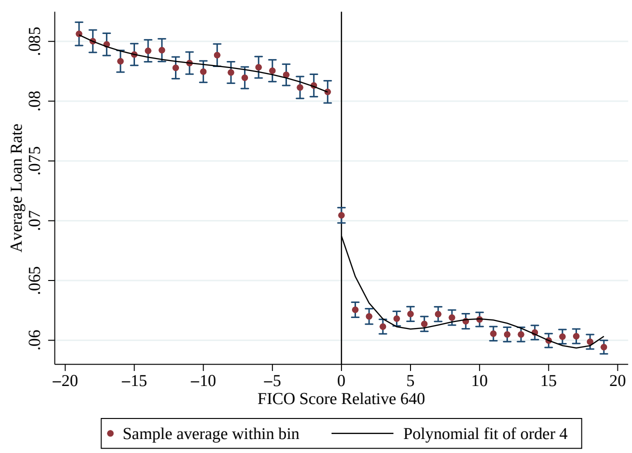

C. Interest Rates Around FICO = 700 Discontinuities

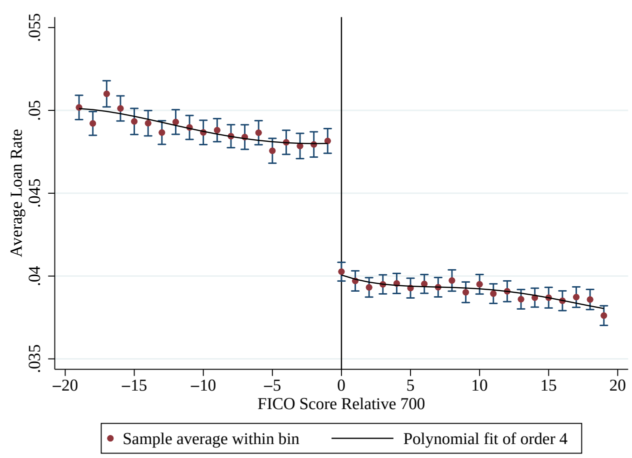

Notes: Figures plot average interest rates against borrower FICO scores normalized to pricing discontinuities. 95% confidence intervals are double clustered by lender and FICO score. Plotted RD functions are estimated using the Calonico et al. (2014) robust RD estimator with fourth-degree polynomials for institutions with pricing discontinuities detected at FICO scores of 600, 640, and 700, respectively.

Figure 5: Balance of Borrower Characteristics Across FICO Discontinuities

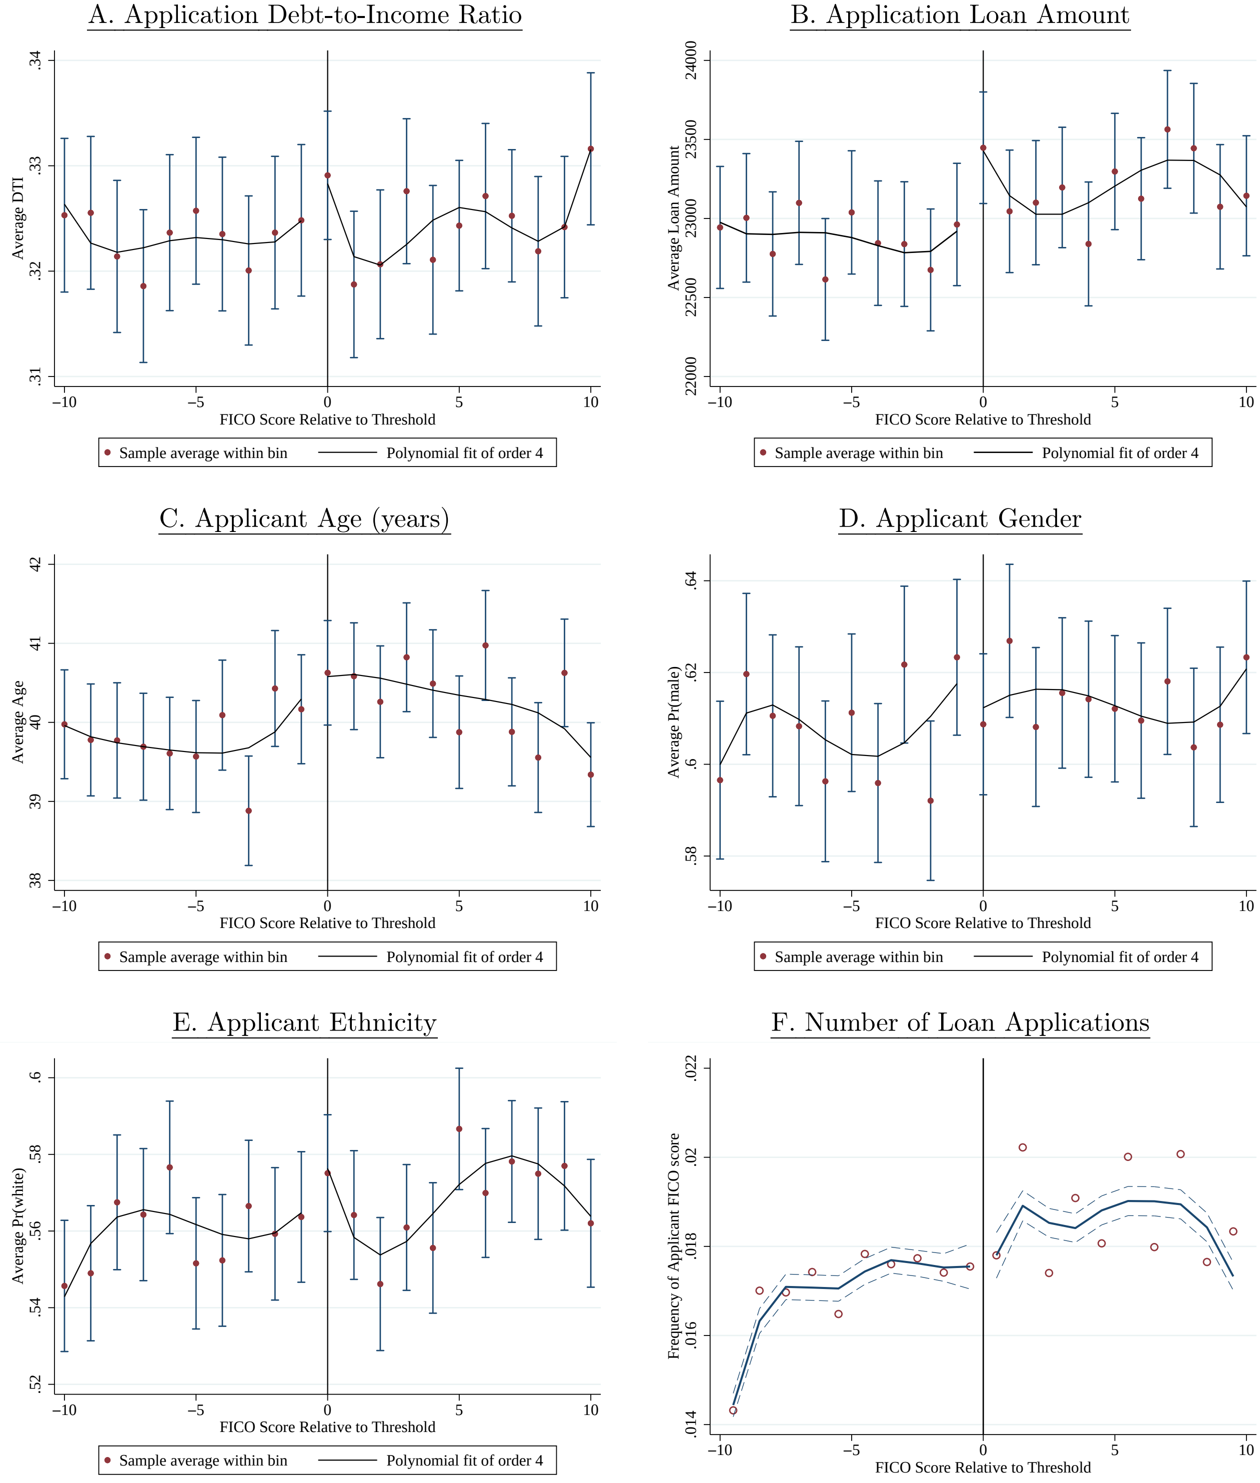

Notes: Figures plot average values of ex-ante borrower characteristics around FICO thresholds for institutions with detected discontinuities. 95% confidence intervals are double clustered by lender and FICO score. Plotted RD functions are estimated using the Calonico et al. (2014) robust RD estimator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

Figure 6: Variation in the Returns to Search Around FICO Discontinuities

A. Borrowers Around a FICO = 600 Threshold

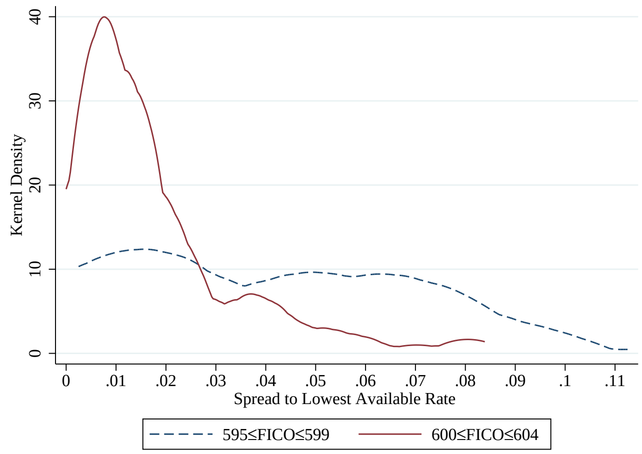

B. Borrowers Around a FICO = 640 Threshold

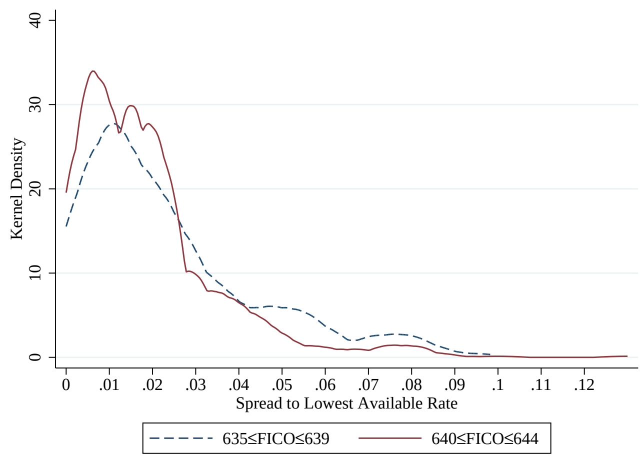

C. Borrowers Around a FICO = 700 Threshold

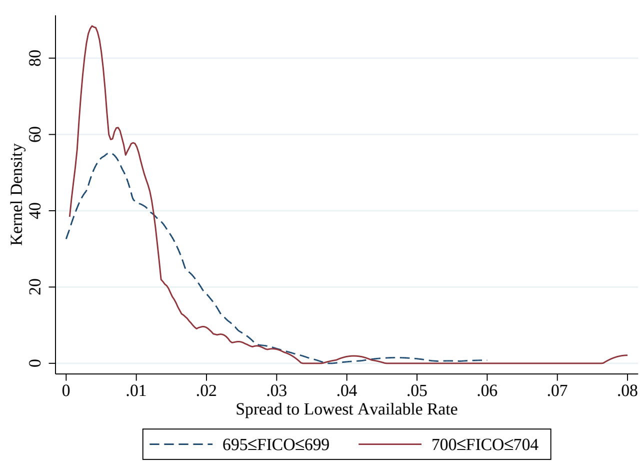

Notes: Figure plots the kernel densities of the spread from the originated interest rate for borrowers at lenders with FICO discontinuities at 600, 640, and 700, respectively, to the lowest rate available to similar borrowers in the same market. Dashed blue lines and solid red lines, respectively, plot the densities for borrowers with FICO scores just below and above each FICO threshold.

Figure 7: Robustness of Take-up Rate Differentials to High Search Cost Definition

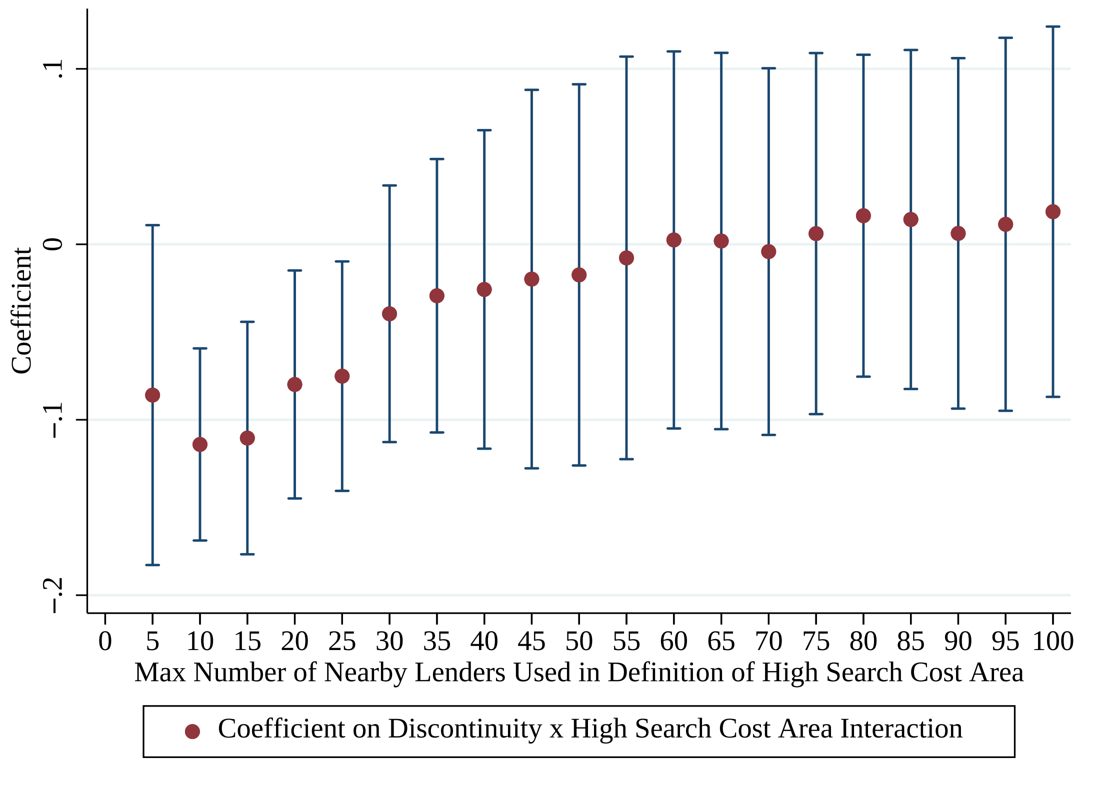

Notes: Figure reports estimates of the coefficient on an interaction term between the Discontinuity indicator ($\widetilde{FICO} \geq 0$) and a High Search Cost indicator in an RD regression of loan take-up on the specification in (12) along with a High Search Cost main effect. High Search Cost is defined 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. 95% confidence intervals are double clustered by lender and FICO score.

Figure 8: CDFs of the Spread to Lowest Available Rate by Search Costs

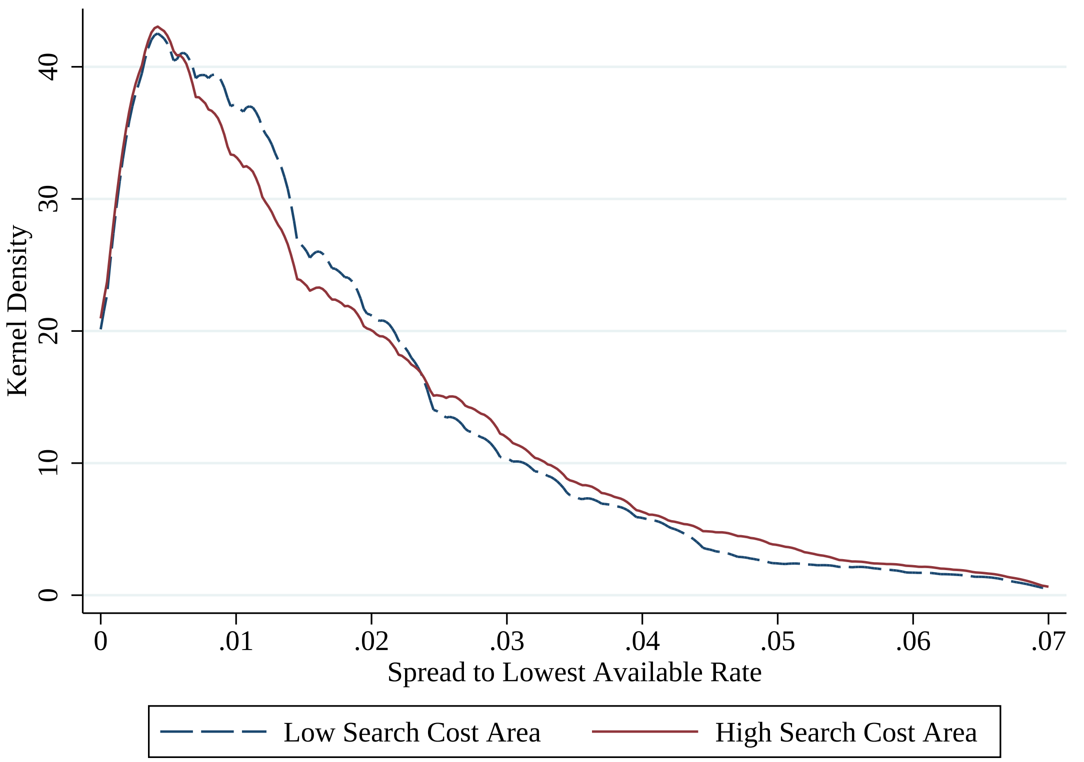

Notes: Figure plots kernel density of the spread to the lowest available rate for each loan (the difference between a loan's interest rate and the best rate among similar borrowers in its cell) for borrowers in high- and low-search-cost areas. The estimated distributions for high- and low-search-cost areas are plotted in solid red and dashed blue lines, respectively. High-search-cost areas are defined as locations where borrowers have less than 10 lenders within a 20-minute drive of their home. See notes to Figure 3.

Figure 9: Effect of FICO Discontinuity on Value of Car Purchased

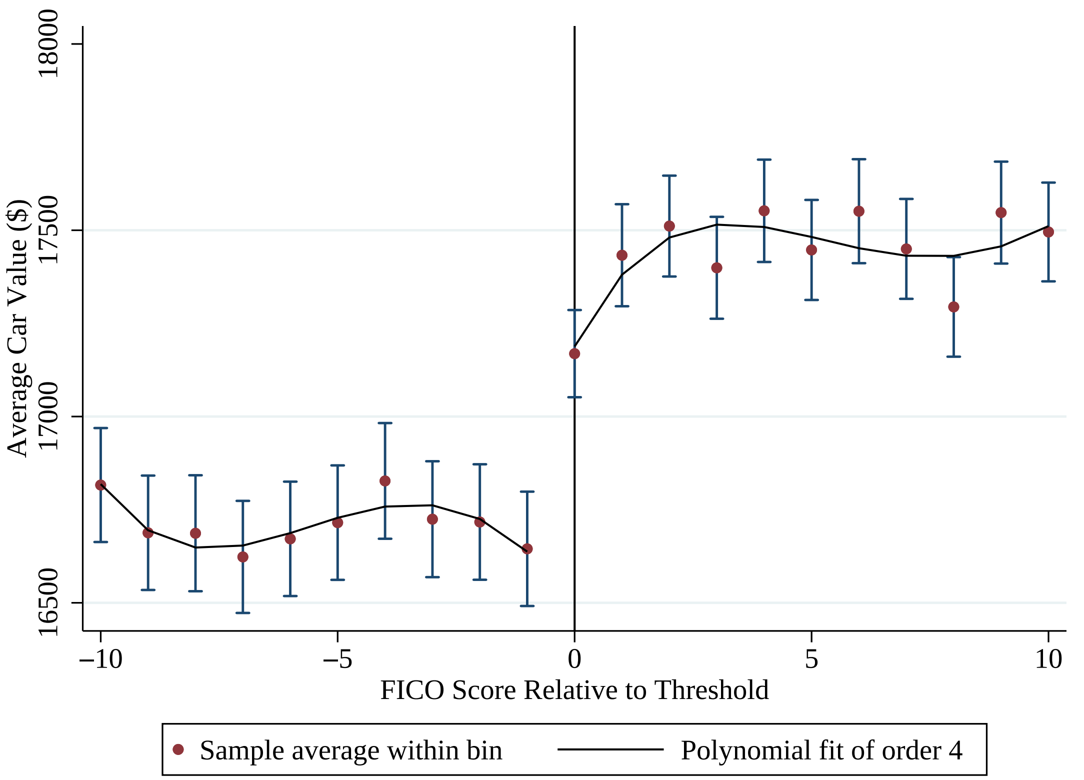

Notes: Figures plot average car transaction prices by FICO scores normalized to detected pricing discontinuities. 95% confidence intervals are double clustered by lender and FICO score. Plotted RD functions are estimated using the Calonico et al. (2014) robust RD estimator with fourth-degree polynomials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

Table 1: Summary Statistics

	Count	Mean	Std. Dev.	Percentile 25th	Percentile 50th	Percentile 75th
	A. Loan Applications					
Loan Term (months)	1,119,153	67.25	24.43	60	72	72
Loan Amount ($)	1,320,109	21,927.3	11,660.7	13,296.0	20,000	28,932.1
Loan Rate	1,131,240	0.05	0.05	0.02	0.04	0.06
FICO	898,339	647.9	118.2	605	661	720
Debt-to-Income	833,854	0.26	0.3	0.13	0.27	0.39
Take Up	588,231	0.65	0.48	0	1	1
	B. Originated Loans					
Loan Rate	2,434,049	0.05	0.03	0.03	0.04	0.06
Loan Term (months)	2,434,049	62.73	22.08	48	60	72
Loan Amount ($)	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
	C. 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 A–C, 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: Summary Statistics for Estimation Sample with Identified FICO Discontinuities

	Count	Mean	Std. Dev.	Percentile 25th	Percentile 50th	Percentile 75th
	A. Approved Loan Applications					
Loan Term (months)	30,743	63.75	13.02	60	66	72
Loan Amount ($)	30,743	21,310.6	11,106.2	13,188.5	19,626	27,606
Loan Rate	26,938	0.056	0.082	0.037	0.05	0.07
FICO	30,743	678.2	31.4	657	682	703
Debt-to-Income	27,197	0.292	0.133	0.197	0.296	0.388
Take-up	30,743	0.511	0.5	0	1	1
	B. Originated Loans					
Loan Rate	514,834	0.061	0.031	0.037	0.055	0.078
Loan Term (months)	514,834	59.616	19.813	48	60	71
Loan Amount ($)	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
	C. 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 A, B, and C describe loan applications, loan originations, and ex-post loan performance, respectively. See notes to Table 1 for further details.

Table 3: Effects of FICO Discontinuities on Loan Rate and Loan Term

	(1) Loan Rate	(2) Loan Term
Discontinuity Coefficient	-0.0127***	0.822***
	(0.004)	(0.187)
Discontinuity × Lender FEs	✓	✓
Commuting Zone × Quarter FEs	✓	✓
Number of Observations	514,834	514,834
R^2	0.169	0.083

Notes: Table reports regression discontinuity estimates of equation (12), normalizing FICO scores around each threshold using a uniform kernel and a bandwidth of 19 FICO points. Both specifications include commuting-zone-by-quarter fixed effects and discontinuity-by-lender fixed effects. Robust standard errors in parentheses are double clustered by lender and FICO score.

Table 4: Loan Application Covariate Balance Regressions

	(1) Loan Amount	(2) Debt-to-Income	(3) Number of Loan Applications
Discontinuity Coefficient	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	0.0584	0.009	0.466

Notes: Table reports reduced-form RD regressions for the subset of institutions for which we have detailed loan application data. See notes to Table 3 for details. 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.

Table 5: Effect of FICO Discontinuities on Loan Offer Take-up Decisions by Search Costs

Search Costs Sample	Full	High	Low	Difference
	(1)	(2)	(3)	(4)
Discontinuity Coefficient	0.121***	0.020***	0.137***	-0.116***
	(0.015)	(0.005)	(0.016)	(0.006)
Discontinuity × Lender FEs	✓	✓	✓	
Commuting Zone × Quarter FEs	✓	✓	✓	
Number of Observations	30,743	4,436	26,307	
R^2	0.27	0.45	0.25	

Notes: Table reports results for reduced-form RD regressions of loan take-up (conditional on being approved for a loan offer) separately for the full sample (column 1) and for borrowers in areas with high and low search costs in columns 2 and 3, respectively, using the specification in equation (12). High search costs are defined as borrowers living at locations with less than 10 lenders within a 20-minute drive. Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 3 for estimation details.

Table 6: Number of Additional Observed Loan Applications per Borrower by Search Costs

Search Costs	High	Low	Difference
	(1)	(2)	(1) - (2)
Mean	0.342	0.409	-0.067***
Standard Deviation	(0.009)	(0.004)	(0.011)
Number of Observations	6,042	44,655	

Notes: Table reports average number of additional applications observed for each loan application in our data for applications with non-missing birthdates and nine-digit zip codes 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 less than 10 lenders within a 20-minute drive.

Table 7: Effects of FICO Discontinuities on Origination Outcomes

	(1) Purchase Price ($)	(2) Loan Amount ($)	(3) Loan-to-Value Ratio	(4) Monthly Payment ($)
Discontinuity Coefficient	376.58**	566.21***	0.0130**	0.17
	(175.72)	(167.93)	(0.005)	(1.02)
Discontinuity $\times$ Lender FEs	✓	✓	✓	✓
Commuting Zone FEs	✓	✓	✓	✓
Quarter FEs	✓	✓	✓	✓
Number of Observations	514,834	514,834	514,834	514,834
R^2	0.052	0.058	0.029	0.056

Notes: Table reports reduced-form RD estimates of equation (12). Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 3 for more details.

Table 8: Effects of FICO Discontinuities on Vehicle Purchase

	(1) Purchase Price	(2) Purchase Price	(3) Car Age (months)
Discontinuity Coefficient	344.69***	79.71	-1.76***
	(123.78)	(49.25)	(0.043)
Discontinuity × Lender FEs	✓	✓	✓
Commuting Zone FEs	✓	✓	✓
Quarter FEs	✓	✓	✓
Make-Model FEs	✓		✓
Make-Model-Year FEs		✓	
Number of Observations	468,800	468,800	468,800
R^2	0.353	0.767	0.352

Notes: Table reports reduced-form RD estimates of equation (12) on car purchase prices (columns 1–2) and car age in months (column 3). Columns 1 and 3 include make × model fixed effects and column 2 includes make × model × model-year fixed effects. Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 3 for more details.

Table 9: 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 Coefficient	754.41***	535.41***	634.48**	335.96*
	(226.18)	(171.10)	(260.62)	(171.75)
Discontinuity × Lender FEs	✓	✓	✓	✓
Commuting Zone FEs	✓	✓	✓	✓
Quarter FEs	✓	✓	✓	✓
Number of Observations	75,206	439,628	75,206	439,628
R^2	0.064	0.059	0.054	0.054

Notes: Table reports reduced-form RD estimates of equation (12) for borrowers with high and low search costs. High search costs are defined as borrowers living at locations with less than 10 lenders within a 20-minute drive. Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 3 for more details.

Table 10: Balance of Ex-Post Credit Outcomes Across FICO Discontinuities

	(1) Days Delinquent	(2) Charge-off	(3) Default	(4) %ΔFICO
Discontinuity Coefficient	4.185	0.004	0.002	0.001
	(3.101)	(0.003)	(0.004)	(0.003)
Discontinuity × Lender FEs	✓	✓	✓	✓
Commuting Zone FEs	✓	✓	✓	✓
Quarter FEs	✓	✓	✓	✓
Number of Observations	331,590	514,834	514,834	405,236
R^2	0.162	0.073	0.089	0.015

Notes: Table reports reduced-form RD estimates of equation (12) on ex-post loan and borrower outcomes. Days delinquent is the number of days a borrower is delinquent as of our data extract. Charge-off is an indicator for whether a loan has been written off the books of the lending institution. Default is an indicator for whether a borrower has been delinquent for at least 90 days. ΔFICO is the change in FICO score since origination as a fraction of the FICO score at origination for the 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 3 for estimation details.

A Identification Appendix

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

A.1 First-stage Heterogeneity

Our take-up estimates represent unbiased estimates of borrowing elasticities at the extensive margin to the extent that the discontinuities satisfy the identifying assumptions required of valid RD estimation. However, when comparing these second-stage RD estimates across subsamples, an important caveat is that variation in the proximity of financial institutions (PFI) could influence first-stage estimates of the magnitude of FICO pricing discontinuities. If so, it would not be surprising to find that below-threshold borrowers search more in high-PFI geographies if differences in offered rates across thresholds are correlated with PFI. Appendix Table A3 reports estimates of differences in offered interest rates around thresholds for high- and low-search-cost geographies. High-search-cost areas (low PFI) have estimated average differences of 1.4 percentage points across thresholds, compared to differences of 1.2 percentage points in low-search-cost areas (high PFI). The difference in the first-stage estimates between high and low-PFI areas is not statistically significant (as reported in column 3), suggesting that our measure of search costs does create bias through differences in first-stage effects.

A.2 Potential Omitted Variables Correlated with High Search Costs

Another challenge to our inference is the interpretation of differences in borrowing elasticities across high- and low-search-cost geographies. While our RD estimation allows us to estimate internally valid responses to interest-rate markups, differences in RD coefficients across areas could be related to other factors besides search costs. Our search-cost proxy (the number of PFIs) is plausibly correlated with several unobserved variables, including the financial sophistication of borrowers, the advertising of credit (which plausibly influences search costs), local (within-CZ) preferences, and economic or demographic conditions. Candidate omitted variables could be time varying (e.g., hyper-local economic conditions) or relatively time invariant (e.g., financial sophistication, brand loyalty preferences).

A.2.1 Addressing Time-varying Omitted Variables

If a given location has few nearby lenders because, for example, lenders anticipate forthcoming local economic shocks, then our proxy for high search costs may be correlated with a growing inelasticity of local demand that could generate our results even if the number of PFIs had remained high. We address possible bias from such time-varying omitted variables in the following way. We create a Bartik measure of PFI, calculated as the number of PFIs for each borrower in our sample using the density of nearby financial institutions as of 1990 using NETS data on the location of every financial institution in the U.S. We then grow the 1990 PFI measure using the national growth rate in financial institutions from 1990 through the year of each observation in our sample period, calculating the national growth rate for each location excluding the contribution of that location to the aggregate changes. Variation in the Bartik PFI measure is thus driven by local branching concentration as of 1990 and aggregate variation in national branching trends, neither of which is likely to be correlated with time-series variation in, e.g., local economic conditions during our sample period of 2005–2015. We sort our sample based on high and low Bartik PFIs and estimate take-up elasticities as in equation (12). Results are reported in Appendix Table A5. These results indicate that borrowing elasticities in high-Bartik-PFI geographies (low search costs) are higher than in low-Bartik-PFI geographies (high search costs), with high-search-cost areas being relatively insensitive to interest rates and low-search-cost areas exhibiting strong and statistically significant reactions. These results suggest that time-varying economic shocks that occur during our sample period do not appear to be a sufficient explanation for differences in loan take-up rates across high- and low-PFI geographies.

A.2.2 Addressing Time-invariant Omitted Variables

A second challenge to our inference is correlated time-invariant omitted variables, such as financial sophistication or brand loyalty. Ultimately, the strength of the Bartik instrumentation strategy relies on the exogeneity of the 1990 branch network, which may itself be correlated with local characteristics. For example, borrowers with low financial sophistication could exhibit lower borrowing elasticities if they do not appreciate the costs associated with differences in loan interest rates (see, e.g., Stango and Zinman, 2009, and Bertrand and Morse, 2011). Related, borrowers may have a preference for borrowing from their current depository (although this, too, could be driven by unawareness of price dispersion arising from high search costs). While our regression-discontinuity design uses random assignment to one side of a FICO threshold and thus conceptually holds such unobservables fixed, when we compare the effect of discontinuities in high and low-PFI areas, we need to address possible time-invariant correlates of the number of nearby lenders that could affect spatial variation in RD coefficients. In general, our specifications above attempted to address bias caused by time-invariant omitted variables with commuting-zone fixed effects. However, while commuting-zone fixed effects address time-invariant omitted variables at the metropolitan-area level, they do not address likely variation in sophistication correlated with high or low PFI within a commuting zone. We address this concern in two ways: using eight-digit zip code fixed effects in our RD specification and with a difference-in-differences design.

First, we augment our RD specification (12) to include fixed effects $\eta_{zip8(i)}$ for the eight-digit zip code of borrower i. We then estimate take-up RD regressions separately for high- and low-search-cost areas, as before. Results in Appendix Table A6 are identified off of regression-discontinuity variation within a hyperlocal geography, allowing us to compare borrowers on the opposite sides of a FICO threshold who live in an area smaller than a census tract. While our High Search Cost indicator may be correlated with other unobservables that affect take-up decisions (e.g., stronger brand preferences or a weaker understanding of personal finance), to the extent that there is spatial correlation in such unobservables, this specification allows us to hone our RD apparatus on more comparable borrowers. Again, column 1 of Appendix Table A6 shows that even within an eight-digit zip code, there is essentially zero contrast between the take-up decisions of borrowers on either side of a discontinuity, whereas within an eight-digit zip code in low-search-cost areas, there is a large and significant difference in take-up rates across a FICO threshold (column 2). The difference between high- and low-search-cost area take-up rates shown in column 3 is even larger than the effects estimated using CZ fixed effects in column 3 of Appendix Table 5. Supporting our interpretation of the effects of the discontinuities being driven by search costs, firm-specific pricing has significant effects on prospective borrowers only when our search-cost proxy is high even when comparing borrowers that live in the same hyperlocal neighborhood.

An alternate approach is to explicitly use temporal variation in the branch network in a difference-in-differences setting that looks at changes in take-up rates as a function of changes in the nearby lender density within a very narrow location. While the identifying assumption here is the exogeneity of time-series changes to the branch network, using this variation is supported by the results of section A.2.1 that potentially endogenous changes to the branch network cannot explain our results. For this exercise, we set aside the discontinuity sample and instead focus on ascertaining the reaction of shopping behavior of locations g that become high-search-cost areas after not being high-search-cost areas previously. We first restrict our attention to only those nine-digit zip codes g in our data that eventually transition from having 10 or more PFIs to fewer than 10. This results in a small sample of 608 observations from locations that we observe before and after becoming high-search-cost areas (note that nine-digit zip codes can often be a single address). We estimate the effect γ of becoming a high-search-cost areas in a panel setting as follows

$$\text{Take-up}_{igt} = \eta_g + \delta_t + \gamma \text{High Search Cost}_{gt} + \beta FICO_{igt} + \varepsilon_{igt}, \tag{14}$$

where $HighSearchCost$ is a dummy for whether location g was a high-search-cost location in quarter t. The virtue of this specification is that it absorbs fixed differences across extremely local areas (e.g., financial sophistication or credit-union loyalty) and thereby identifies the effect of search costs solely off of the timing of changes in search costs, which we argued in Appendix Section A.2.1 seem unrelated to unobservables driving variation in demand elasticities across space. We control for the FICO score of the given borrower to control for changes in the composition of borrowers at a nine-digit address over time. Column 1 of Appendix Table A7 reports estimates of equation (14), showing that when borrowers face a reduction in nearby potential lenders, they are 11.0 percentage points more likely to accept a given loan offer.

To include locations which did not change search-cost status in the control group to identify the coefficients on the control variables, we also estimate a version of equation (14) in first-differences. For each location g in our sample that we observe more than once, we calculate the change in take-up rates $\Delta\text{Take-up}_{gt}$, which we then regress on commuting-zone and quarter-pair fixed effects (a fixed effect for the pair of quarters in which location g was observed to calculate $\Delta\text{Take-up}_{gt}$); the change in the location's average FICO score; and an indicator $\Delta\text{High Search Cost}_{gt}$ equal to one if location g was a high-search-cost area in quarter t but not in the last period location g was observed before period t.

$$\Delta\text{Take-up}_{gt} = \eta_{cz(g)} + \delta_{t,\Delta t} + \gamma\Delta\text{High Search Cost}_{gt} + \beta\Delta FICO_{gt} + \varepsilon_{gt}. \tag{15}$$

Again, the identifying assumption behind this specification is that the timing of decreases in the number of PFIs is unrelated to counterfactual trends in take-up rates, supported by the Bartik results above. Including commuting-zone fixed effects in the differences specification allows for each metropolitan area to change its propensity to shop for a loan differentially. The prediction of our search-cost explanation for high take-up rates in high-search-cost areas for this setting is that cross-sectional *changes* in take-up rates, having differenced out time-invariant factors, will be positively related to changes in our search-cost measure. Column 2 of Appendix Table A7 shows that the estimated $\hat{\gamma}$ is positive and significant such that take-up rates increase as borrowers shop around less when nearby lenders exit. The results of the levels difference-in-differences specification and the first-differences specification further confirm that borrowers seem to be more affected by firm-specific markups when search costs are high.

A.3 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, even for direct loans. 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. 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.

A.4 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 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 A3 supports this argument by plotting the relationship in our data between markups and market shares in high- and low-PFI 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 × Commuting Zone × quarter level. In high-PFI (low-search-cost) markets, where the density of nearby lenders is high, charging a higher markup corresponds to a lower market share. In low-PFI (high-search-cost) areas, market shares are statistically unrelated to markups.

Appendix Table A8 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 3.1 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.[42] Dividing loans into high and low (above and below median) competition areas based on our constructed Herfindahl index, we reestimate the take-up RD specification of Table 5 for all four combinations of high and low search cost and market concentration combinations.

The results of this exercise in Appendix Table A9 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.

B Estimates from Structural Search Model

In this appendix, we demonstrate that search-cost estimates from the structural search model of Hortaçsu and Syverson (2004) (HS) are consistent with our characterization of high- and low-search-cost areas. Although HS models inelastic demand for a single final good (mutual funds), it provides a useful framework to relate data on market shares and

[42]See Scharfstein and Sunderam (2016), Dreschler, Savov, and Schnabl (2017), and Liebersohn (2019) for recent alternative measures of retail banking market concentration.

prices to the distribution of search costs in a sequential search equilibrium. We first introduce the version of the HS model that we apply to our setting and then discuss our estimation procedure. We then reestimate the model separately using only market-share and markup data from borrowers in high- or low-PFI areas (our proxy for low and high search costs, respectively) and discuss our results.

In the version of the HS model we consider, buyers are homogenous except for heterogeneous search costs k_i distributed according to $H(\cdot)$, which we take to be a log-normal distribution with parameters $E(\log k) = \mu$ and $Var(\log k) = \sigma^2$. When searching, buyers draw a price quote from seller j with probability

$$\rho_j = \frac{Z_j^\alpha}{\sum_{k=1}^N Z_k^\alpha}, \tag{16}$$

where Z_j is a vector of characteristics about seller j that determine its sampling probability. For a shifter Z_j of sampling probability, we use NUCA data to calculate the number of years since a credit union first opened in a market.[43]

Analogous to (3), borrowers in the HS model continue searching insofar as the expected gain from search given their current quote r' exceeds their search cost k_i

$$k_i \leq \int_{\underline{r}}^{r'} \left[V(r) - V(r')\right] dF(r) \tag{17}$$

where $F(\cdot)$ is the distribution of interest rates and $V(\cdot)$ is the indirect utility function. This yields a set of search-cost critical values k_j^*, where borrowers that have drawn a quote from firm j and who have individual search costs $k_i \geq k_j^*$ will accept the offer from firm j, with

$$k_j^* = \sum_{l=j}^{N} \rho_l \left[V(r_l) - V(r_j)\right] \tag{18}$$

where N is the number of (discrete) lenders and lenders are indexed in descending order according to the interest rate r_l that they charge such that $V(r_l) \geq V(r_j)$ whenever $l > j$. Given (17), borrower search will again feature an optimal stopping rule, where borrowers will search until they find a lender whose search cost cutoff $k_j^* < k_i$.

The cutoff values established by (18) connect a lender's market share given price r_j to the search cost distribution. Borrowers with search cost $k_i \in (k_j^*, k_{j-1}^*)$ will make up some of the market share of every lender $l \geq j$, as some of these borrowers will have arrived at lender l first and will stop because their search cost is above the critical value for lender l (because likelihood of finding a sufficiently lower-cost lender than l is too small). The contribution of borrowers satisfying $k_i \in (k_j^*, k_{j-1}^*)$ to the market share of lender $l \geq j$ will be

$$\frac{\rho_l}{1 - \rho_1 - \rho_2 - \ldots - \rho_{j-1}} \left[H(k_{j-1}^*) - H(k_j^*)\right] \tag{19}$$

[43] We find that that estimates of α are positive, suggesting that borrowers are indeed more likely to sample from lenders with longer tenure in a market.

and the total market share s_l for lender l will then be the sum of these values for all segments in the search cost distribution

$$
\begin{aligned}
s_l = \rho_l \Bigg[1 + \frac{\rho_1 H(k_1^*)}{1-\rho_1} + \frac{\rho_2 H(k_2^*)}{(1-\rho_1)(1-\rho_1-\rho_2)} \\
+ \sum_{j=3}^{l-1} \frac{\rho_j H(k_j^*)}{(1-\rho_1-\ldots-\rho_{l-1})(1-\rho_1-\ldots-\rho_l)} - \frac{H(k_l^*)}{(1-\rho_1-\ldots-\rho_{l-1})} \Bigg].
\end{aligned}
\tag{20}
$$

Intuitively, borrowers from firm l are borrowers that drew from l first with search costs sufficiently high as to fail to justify additional search, or else happened to sample from only firms with sufficiently higher prices than l before sampling from, and being satisfied with, firm l.

Lenders choose rates r_j to maximize current profits, given by

$$\pi_j = S \cdot s_j(r_j)(r_j - c_j) \tag{21}$$

where S is the total size of the market, $s_j(r_j)$ is firm market share given prices, and c_j is firm's marginal cost of lending. Note that HS abstracts away from continuously elastic demand or complements, which is sufficient for our objective here of illustrating whether a search model estimates different search costs for high- and low-PFI areas. The FOC is then

$$s_j(r_j) + (r_j - c_j)\frac{\partial s_j(r_j)}{\partial r_j} = 0. \tag{22}$$

The relevant derivative of market shares with respect to prices in (22) is knowable from (20):

$$
\begin{aligned}
\frac{\partial s_j}{\partial r_j} = -\frac{\rho_1 \rho_j^2 h(k_1^*)}{1-\rho_1} - \frac{\rho_2 \rho_j^2 h(k_2^*)}{(1-\rho_1)(1-\rho_1-\rho_2)} \\
- \sum_{l=3}^{j-1} \left[\frac{\rho_l \rho_j^2 h(k_l^*)}{(1-\rho_1-\rho_2-...-\rho_{l-1})(1-\rho_1-\rho_2-...-\rho_l)} - \frac{\rho_j(\sum_{l=j+1}^{N} \rho_l) h(k_j^*)}{1-\rho_1-\rho_2...-\rho_{j-1}} \right]
\end{aligned}
$$

where the dependence of k^* on r comes from (18) and $h(\cdot) = dH(\cdot)$ is the PDF of the search-cost distribution.

Assuming constant marginal cost c as in HS, we can then estimate the parameters (μ, σ, α, c) by non-linear least squares, taking (22) as our estimating equation and an objective function

$$(\hat{\mu}, \hat{\sigma}, \hat{\alpha}, \hat{c}) = \arg\min \sum_{j=1}^{N} \left(s_j(r_j) + (r_j - c)\frac{\partial s_j(r_j)}{\partial r_j} \right)^2. \tag{23}$$

Because this model assumes homogenous products and borrowers, we restrict the data to loans with the modal maturity (60 months) and residualize interest rates of credit risk and other product characteristics using lender fixed effects as described in section A.4. Appendix Figure A4 plots the CDF of search costs for high- and low-PFI areas implied by our log-normal distributional assumption and the estimated parameters μ and σ. The model estimates that given the joint distribution of market shares and prices, low-PFI borrowers have higher search costs, with the CDF of search costs for borrowers in low-PFI areas always less than the search-cost CDF for borrowers in high-PFI areas. When the y-axis is 0.5, the horizontal difference between the two CDFs corresponds to the difference in median search costs, which we show below is 70 basis points.

Appendix Table A10 reports key moments of the search-cost distribution in each area. The units are the same as interest-rate units such that 0.01 corresponds to one percentage point. We reject the equality of the two distributions. Low-PFI areas have much higher average search costs, as well as more dispersed search costs. Because the search-cost distribution is significantly skewed, HS emphasize differences in median search costs, which in our case is 70 basis points. Discounting at 5%, for a typical, \$20,000 loan with a 60-month maturity, a 70 basis point change in interest rates from 5% to 5.70% corresponds to a present value difference of total loan payments of \$342. These estimates therefore imply that the median borrower in a high-PFI area has a search cost \$342 less than the median borrower in a low-PFI area.

The Hortaçsu and Syverson (2004) search model differs from our model in several key ways. However, its framework for estimating search-cost distributions given market shares and prices provides an independent characterization of the relative costliness of searching for credit in areas with high and low nearby lender densities. The results in Figure A4 and Appendix Table A10 confirm our use of driving-time lender density above and below a cutoff as a proxy for search costs that is indeed predictive of model-implied estimates of relative search costs.

Figure A1: Average Vehicle Mileage by Car Age

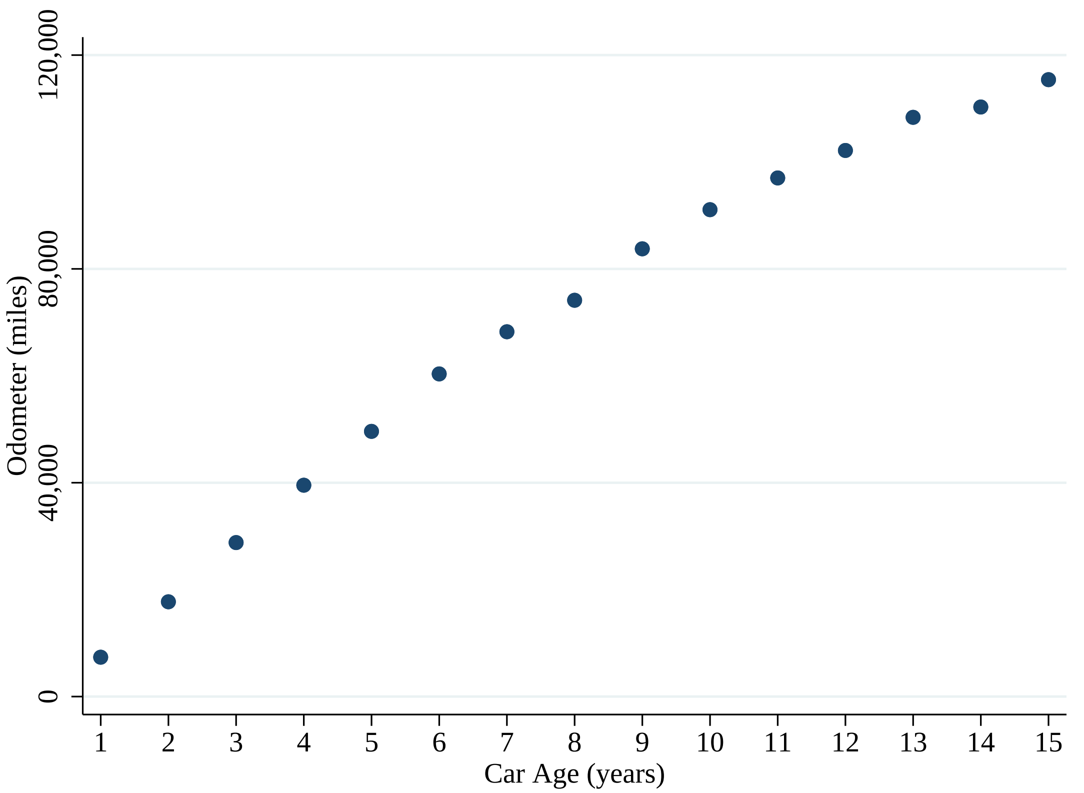

Notes: Figures plots average used-car odometer readings as a function of car age using microdata from the National Household Transportation Survey (U.S. Department of Transportation, 2017).

Figure A2: Balance of Borrower and Loan Characteristics Conditional on Origination

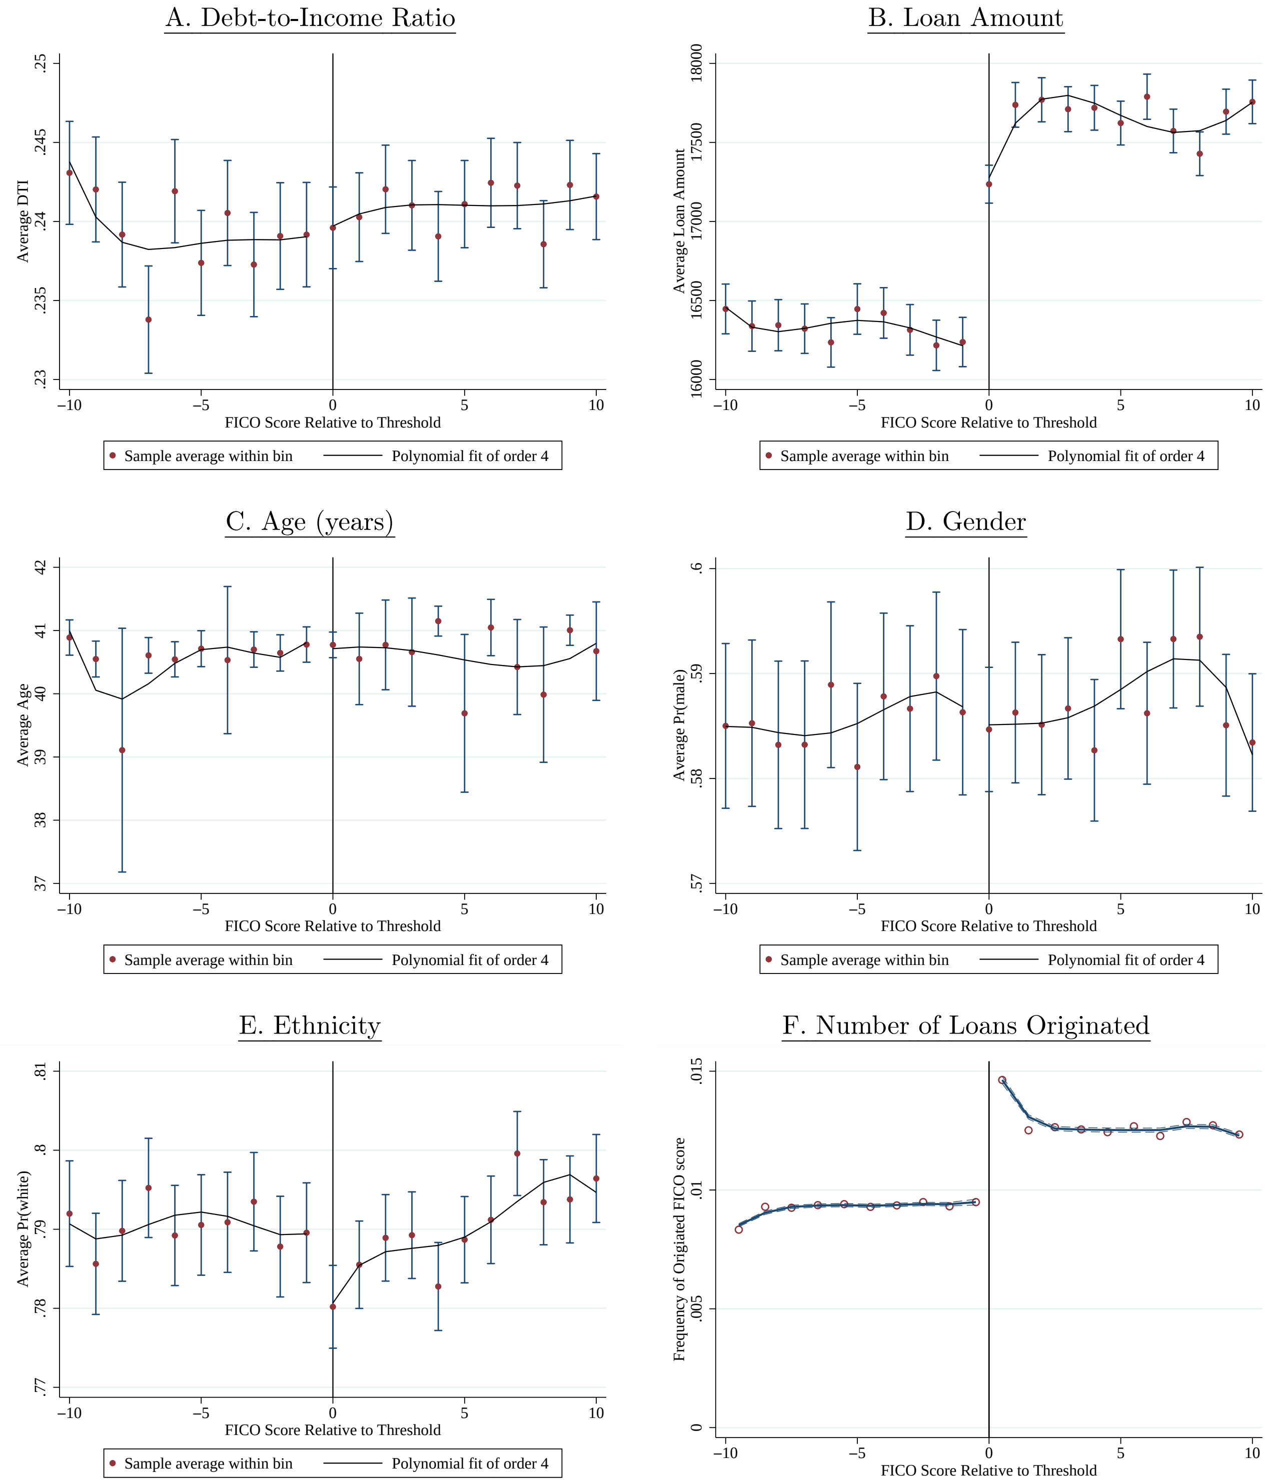

Notes: Figures plot average values of borrower and originated loan characteristics around FICO thresholds conditional on observing the application of an originated loan in our data. See notes to Figure 5 for further details.

Figure A3: Relationship between Market Shares and Average Mark-ups by Search Costs

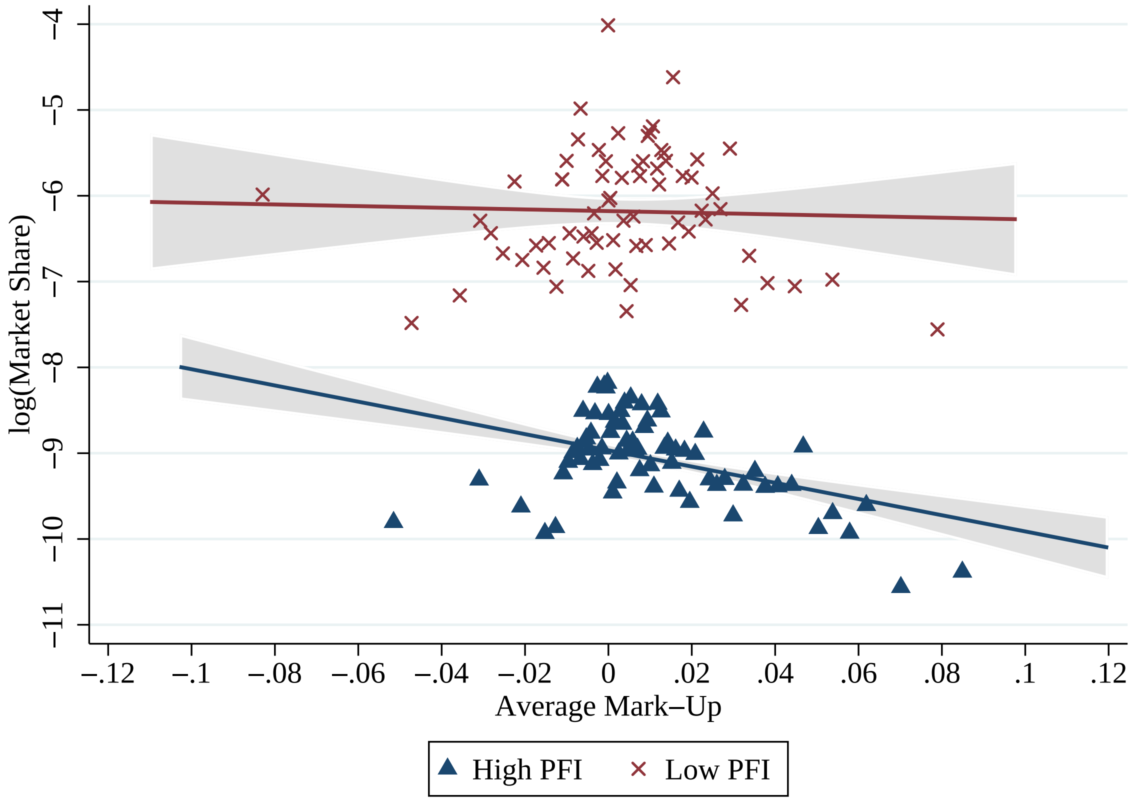

Notes: Figure shows a bin-scatter plot of average log market shares against average interest-rate markups for high-PFI (low-search-cost) and low-PFI (high-search-cost) markets (denoted with ▲ and × symbols, respectively) along with bivariate regression lines and 95% confidence intervals. High (low) PFI 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.

Figure A4: Model-Estimated Search Cost Distributions

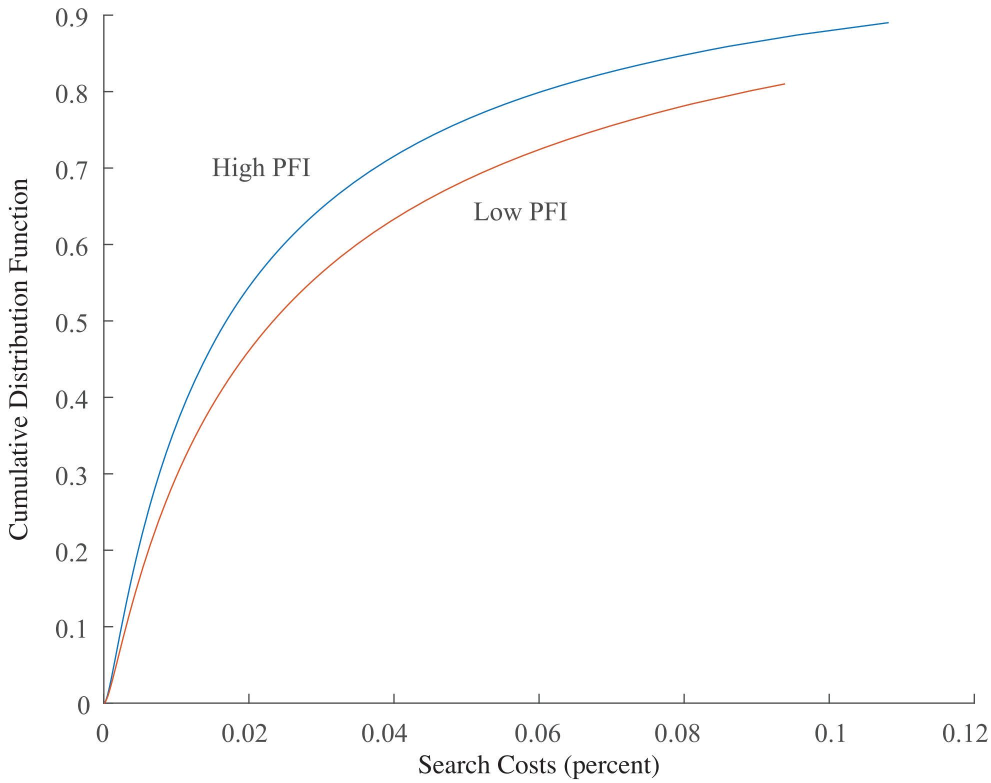

Notes: Figure plots search cost CDFs estimated using the sequential search model of Hortaçsu and Syverson (2004). The blue line on the left is the estimated distribution of search costs for borrowers in high-PFI areas that proxy for low search costs, and the orange line on the right is estimated using data from borrowers in low-PFI areas that proxy for for high search costs. The x axis is measured in the same units as loan interest rates.

Table A1: Summary Statistics for Excluded Sample of Indirect Loans

				Percentile		
	Count	Mean	Std. Dev.	25th	50th	75th
	A. 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
	B. 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 for the balance of the analysis of the paper. See notes to Table 1 for further details.

Table A2: 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 A3: Effects of FICO Discontinuities on Origination Interest Rates by Search Costs

Search Costs	High	Low	Difference
	(1)	(2)	(3)
Discontinuity Coefficient	-0.014***	-0.012***	-0.002
	(0.004)	(0.005)	(0.002)
Discontinuity $\times$ Lender FEs	✓	✓	
Commuting Zone $\times$ Quarter FEs	✓	✓	
Number of Observations	75,206	433,825	

Notes: Table repeats RD estimates of first-stage equation (12) reported in Table 3, splitting the sample into high- and low-search-cost areas. High search costs are defined as borrowers living at locations with less than 10 lenders within a 20-minute drive. Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 3 for more details.

Table A4: Spread to Lowest Available Rate Summary Statistics by Search Costs

	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 less than 10 lenders within a 20-minute drive. Summary statistics are reported for only those cells that contain at least two borrowers.

Table A5: Effect of FICO Discontinuities on Loan Offer Take-up Decisions by Bartik Search Costs

Bartik Search Costs	High	Low	Difference
	(1)	(2)	(3)
Discontinuity Coefficient	0.050	0.135***	-0.085***
	(0.045)	(0.037)	(0.006)
Discontinuity × Lender FEs	✓	✓	
Commuting Zone × Quarter FEs	✓	✓	
Number of Observations	5,591	25,152	
R^2	0.40	0.25	

Notes: Table reports results for reduced-form RD regressions of loan take-up (conditional on being offered a loan) separately for borrowers in areas with low and high search costs in columns 2 and 3, respectively. High search costs are defined as borrowers living at locations predicted to have less than 10 lenders within a 20-minute drive using the Bartik instrument discussed in section A.2.1. Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 3 for estimation details.

Table A6: Robustness of FICO Discontinuities Effects on Take-up to Zip-8 Fixed Effects

Search Costs Sample	High	Low	Difference
	(1)	(2)	(3)
Discontinuity Coefficient	0.066	0.190***	-0.125***
	(0.057)	(0.035)	(0.009)
8-digit Zip Code FEs	✓	✓	
Quarter FEs	✓	✓	
Number of Observations	4,436	26,307	
R^2	0.91	0.87	

Notes: Table reports results for reduced-form RD regressions of loan take-up (conditional on being offered a loan) separately for borrowers in areas with low and high search costs in columns 1 and 2, respectively, including eight-digit zip code fixed effects. Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 3 for estimation details.

Table A7: Effects of Search Costs and Competition on Take-up Decisions

Dependent Variable	Take-up (1)	ΔTake-up (2)
High Search Cost Area	0.11***	
	(0.04)	
Δ High Search Cost Area		0.03*
		(0.017)
FICO	-0.0004	
	(0.0003)	
ΔFICO		-0.0002***
		(0.00003)
Geographic FEs	Zip-8	CZ
Time FEs	Quarter	Quarter Pair
Number of Observations	608	29,321
R-squared	0.60	0.05

Notes: Table reports difference-in-differences regression results relating local take-up rates to a given location's high-search-cost status. Column 1 estimates this relationship in levels at the individual $\times$ quarter level, restricting the sample to only those locations (9-digit zip codes) that transition from low to high search cost. Column 2 estimates this relationship in changes, specifying the dependent variable as the change in the take-up rate at the location $\times$ quarter level, and includes all locations for which we observe originated loans in multiple quarters. Quarter Pair fixed effects are time fixed effects for each pair of quarters over which the differences in column 2 are calculated. ΔFICO is the change in the average FICO score of a given location between observations for that location. Robust standard errors in parentheses are double clustered by FICO and quarter (column 1) or quarter pair (column 2).

Table A8: 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 Cost	2.788***	2.856***	2.474***
	(0.180)	(0.199)	(0.222)
Avg. Mark-up × High Search Cost	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 less than 10 lenders within a 20-minute drive. Robust standard errors in parentheses are clustered by Commuting Zone.

Table A9: Effects of Search Costs and Market Concentration on Take-up Decisions

		Market Concentration: LOW	Market Concentration: HIGH
Search Costs	LOW	0.14***	0.10***
		(0.01)	(0.01)
	HIGH	-0.07*	-0.01
		(0.04)	(0.01)

Notes: Table reports results for reduced-form RD regressions of loan take-up for borrowers in each combination of markets with low and high search costs and high and competition. High search costs are defined as borrowers living at locations with less than 10 lenders within a 20-minute drive. 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. All regressions include discontinuity by lender fixed effects and commuting zone by quarter fixed effects. Robust standard errors in parentheses are double clustered by lender and FICO score. See notes to Table 3 for estimation details.

Table A10: Search Cost Distribution Estimation Results

PFI Sample	Mean	Std. Dev.	Median	Observations
	(1)	(2)	(3)	(5)
High	0.068	1.512	0.017	2,525
	(0.040)	(0.010)	(0.013)	
Low	0.102	1.585	0.023	614
	(0.066)	(0.025)	(0.020)	

Notes: Table reports key moments of the search cost distribution (measured in percents) corresponding to estimates of (22) by Nonlinear Least Squares as in (23) separately for loans in high-PFI and low-PFI areas. Low-PFI areas proxy for high-search-cost markets and are defined as areas where borrowers live within a 20-minute drive of less than 10 lenders. Means, standard deviations, and medians of the search cost distribution for each type of area are calculated from estimated log-normal parameters μ and σ with standard errors computed by the delta method.