The Review of Financial Studies

Real Effects of Search Frictions in Consumer Credit Markets

Bronson Argyle
Marriott School of Business, Brigham Young University, USA

Taylor Nadauld
Marriott School of Business, Brigham Young University, USA

Christopher Palmer
Sloan School of Management, Massachusetts Institute of Technology and NBER, USA

We show that search frictions in credit markets affect accepted interest rates and loan sizes and distort consumption. Using data on car loan applications and originations not intermediated by car dealers, we isolate quasi-exogenous variation in both the costs and benefits to searching for credit. After identifying lender-specific policies that price risk discontinuously, we study the differential response to offered interest rates by borrowers who face high and low search costs. High-search-cost borrowers are 10% more likely to accept loan offers with higher markups, consequently originating smaller loans and purchasing older and less expensive cars than lower-search-cost borrowers. (*JEL* D12, D83, E43, G21, L11)

Received January 26, 2022; editorial decision July 29, 2022 by Editor Gregor Matvos. Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

Some of the most important questions in household finance center on how various credit-market imperfections affect consumption. In this paper, we

We thank our discussants Paul Calem, Anthony DeFusco, Paul Goldsmith-Pinkham, Steven Laufer, Neale Mahoney, and Johannes Stroebel; seminar, conference, and workshop participants at Berkeley, Boston Federal Reserve, BYU, CFPB, Chicago Federal Reserve, Duke, FDIC, Housing-Urban-Labor-Macro Conference, Imperial, LSF, MIT, NBER Household Finance, NYU, OSU, Oxford, Philadelphia Federal Reserve, Red Rock Finance Conference, San Francisco Federal Reserve, UNC, and the WFA; and John Campbell, Claire Celerier, Maryam Farboodi, Amy Finkelstein, Brigham Frandsen, Peter Ganong, Kyle Herkenhoff, Brad Larsen, Lars Lefgren, Andres Liberman, Brigitte Madrian, Adrien Matray, Gregor Matvos, Adair Morse, Holger Mueller, Hoai-Luu Nguyen, Andrew Paciorek, Jonathan Parker, Brennan Platt, Tobias Salz, Amit Seru, David Sraer, Bryce Stephens, Stijn Van Nieuwerburgh, Stephen Zeldes, and Jonathan Zinman for helpful conversations. Tommy Brown, Sam Hughes, Tammy Lee, William Liu, and Caroline Palmer provided excellent research assistance. An anonymous information-technology firm provided the data. This work was supported by the National Science Foundation [grant CAREER 1944138] to C. P. Data and code can be found on the RFS Dataverse at https://dataverse.harvard.edu/dataverse/rfs. Send correspondence to Bronson Argyle, bsa@byu.edu.

The Review of Financial Studies 36 (2023) 2685–2720

https://doi.org/10.1093/rfs/hhac088 Advance Access publication November 24, 2022

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

Our empirical strategy features a setting in which search costs vary across space and the potential gains to search are high and quasi-randomly assigned in the cross-section. We leverage administrative data on 2.4 million used-car loans extended by 326 different financial institutions and 1.3 million loan applications to 41 lenders. Unconditionally, many borrowers in our data could access significantly dominating loan offers if they queried an additional financial institution. The data allow us to exploit large interest-rate discontinuities at various lender-specific credit score (FICO) thresholds to isolate quasi-exogenous variation in the incentives to search. On average, borrowers just above a FICO discontinuity at their lender are offered loans with interest rates 1.3-percentage-points lower than otherwise similar borrowers just below a FICO discontinuity. Because there is no consensus set of thresholds used by a plurality of lenders, we demonstrate that borrowers on the expensive 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 above-threshold borrowers from the same institution.

These rate discontinuities isolate supply-side interest-rate variation under the assumption that demand-side factors do not also change discontinuously at FICO thresholds that vary across institutions within the same geography and time. Intuitively, this is likely to be satisfied given that borrowers are unlikely to know their precise FICO score that will be used in pricing or the location of pricing discontinuities across lenders. We support this identifying assumption of quasi-randomly assigned markup offers with evidence that ex ante borrower characteristics (including age, gender, ethnicity, application debt-to-income ratio (DTI), application loan size, and the number of loan applications per FICO bin) are balanced around FICO thresholds.

Many aspects of loan shopping may be particularly costly, including the time involved, the hassle, and effort, each of which may be in short supply while shopping for a car. To demonstrate the existence of one such distortionary dimension of search costs, we calculate the number of financial institutions within a 20-minute drive from each borrower address using the physical branch locations of every bank and credit union in the United States. We hypothesize that obtaining multiple loan quotes will be less costly for borrowers with more nearby lenders, and, indeed, the strength of this proxy is borne out in the data in multiple ways. For example, we observe potential borrowers without many nearby lenders applying for fewer car loans in our data and applying for fewer mortgages in national Home Mortgage Disclosure Act data.

Combining a regression discontinuity (RD) strategy with cross-sectional variation in predicted lender proximity, we demonstrate real effects of credit-market search frictions on loan take-up and consumption. This RD laboratory allows us to test how borrowers with a high return to additional search (from being randomly assigned an expensive loan) differentially change their behavior when facing high versus low search costs. While borrowers with varying lender proximity are presumably different on multiple dimensions, one of the virtues of combining our RD strategy with geographic variation is that our results cannot be driven by fixed differences across high- and low-search cost areas. To account for unobserved factors that could both affect borrower sensitivity to rate markups and be correlated with our physical search cost proxy, we further construct an instrument that isolates exogenous changes to the local branch network using historical spatial variation in bank branching.

Using our proxies for distance-dependent search costs, we show that borrowers on the expensive side of FICO thresholds reject high-interest-rate loans more often when we measure search costs to be relatively low (i.e., when the number of actual or predicted nearby alternative lenders is high). By contrast, borrowers who we predict would have to exert more effort to search for a loan with better terms are more likely to accept the loan pricing they are offered even though these terms are most often strongly dominated by nearby alternatives. Finally, we show results on price dispersion and search behavior that are inconsistent with a simple market concentration explanation, and we show that our results hold even in relatively low-concentration markets.[1]

Given that borrowers with higher costs of physically shopping for credit are more likely to accept expensive loans, we investigate the real effect of this friction on loan and car-purchase outcomes to consider how the costliness of shopping for credit subsequently affects consumption. We find that both financing and durable-goods purchasing decisions respond to interest rates such that borrowers with high search costs facing markups are more likely to substitute toward older and cheaper cars. On average, borrowers that accept quasi-randomly offered more expensive credit take out \$550 smaller loans and purchase cars that are 2 months older, spending an average of \$375 less. The balance of borrower characteristics across FICO thresholds suggests that borrowers quasi-randomly drawing high loan markups would actually also take

[1] The number of sellers in a market directly affects equilibrium pricing in many imperfect competition models. However, our search-cost measure appears to capture variation in the cost of search instead of other variation in market structure given our finding that higher search-cost borrowers indeed search less despite facing, if anything, larger local price dispersion. Moreover, the substantial price dispersion we document is not a feature of simple market concentration models. See Internet Appendix F.1 for further discussion.

out larger loans and purchase a more expensive and newer car had they not been offered higher interest rates. To attribute these conditional-on-origination consumption outcomes to search frictions, we rule out selection in loan take-up being correlated with borrower-level demand shocks by further verifying the balance of borrower characteristics and outcomes conditional on origination. Postorigination changes in credit scores and ex post loan performance do not change differentially across discontinuities for either high- or low-search-cost borrowers. Taken together, the evidence suggests that the costly search for credit represents an important market friction that ultimately distorts financing and durable consumption.

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, Meisenzahl, and Ramcharan 2017) gives the car-loan market aggregate importance and makes it a plausible setting to look for credit-market search frictions affecting consumption. Auto debt is the third-largest category of consumer debt in the United States, with over 114 million outstanding loans (0.89 per U.S. household) totaling $1.5 trillion (Federal Reserve Bank of New York 2021). Most U.S. car purchases are financed (Bartlett 2013; Zabritski 2021), and vehicles represent over 50% of total assets for U.S. low-wealth households (Campbell 2006). From the standpoint of an empirical design, auto loans are a relatively homogeneous installment credit product and can be simply described by their interest rate, term, and size. Our focus on the large segment of used-car loans *not* intermediated by car dealers allows us to test for credit- and product-market linkages in a nonmechanical setting. Finally, auto loan markets are quite local, motivating our analysis of the distortions that distance-related search frictions might cause in consumer debt markets. The median borrower in our sample originates a loan from a branch that is within a 15-minute drive of her home, contrasted with the median U.S. worker's commute of 28 minutes to work (Burd, Burrows, and McKenzie 2021).

In consumer credit markets, Woodward and Hall (2012), Stango and Zinman (2015), Alexandrov and Koulayev (2018), and Bhutta, Fuster, and Hizmo (2020) establish the role of low borrower search intensity in explaining consumer interest rate dispersion. Allen, Clark, and Houde (2014a,b, 2019) document price dispersion in the Canadian mortgage market, demonstrate its response to market concentration, and quantify the lost consumer surplus from higher markups. Agarwal et al. (2020) show that in the cross-section, intensive loan search is correlated with higher interest rates because low credit-worthy borrowers search until they find a lender who offers them an advantageous interest rate. Relative to this literature on price dispersion and search in consumer credit markets, our setting allows for measurement of the real effects on subsequent consumption quantities that can result from costly search for credit. Moreover, the quasi-random assignment of our RD design allows us to abstract away from unobservable private information and contrast financing and consumption outcomes for borrowers with high and low search costs but similar benefits to search. Our work also contributes to a literature that documents how features of various credit markets drive prices and quantities in other markets (e.g., Zeldes 1989; Gross and Souleles 2002; Melzer 2011; Zaki 2016; Delavande and Zafar 2019; Benetton 2021; Aydin 2022). Relative to studies of the connection between credit markets and related markets, we are the first to highlight how search frictions in credit markets can distort durables-goods consumption.

Traditionally speaking, most search models feature inelastic and discrete demand, inhibiting the model's ability to speak to quantity or welfare effects. By contrast, our focus on costly search for credit highlights the importance of continuous and elastic demand in driving equilibrium outcomes.[2] Modern work on search and differentiated products emphasizes how buyer valuations could change because of the structure of search costs in various settings (Zhou 2014; Moraga-González, Sándor, and Wildenbeest 2017, 2021). Our emphasis on how costly search for credit affects the choice of goods purchased with that credit highlights a new search-cost source of demand elasticity distinct from other characterizations of search with elastic demand.

We also contribute to a growing literature studying the automobile-loan market and its frictions (Attanasio et al. 2008; Adams, Einav, and Levin 2009; Busse and Silva-Risso 2010; Einav, Jenkins, and Levin 2012, 2013; Melzer and Schroeder 2017; Grunewald et al. 2020). Two other contemporaneous papers use a similar data set to this paper to answer distinct questions on the importance of loan maturity and budgeting heuristics in the markets for cars and car loans. Argyle et al. (2021) (ANPP hereafter) use *vehicle*-level variation to document that payment-size shocks arising from loan-maturity shocks affect bargaining outcomes and are capitalized into transaction prices. While ANPP hold the purchased durable fixed to ask whether financing affects prices paid for the *same* car, this paper examines how the cost of searching for better interest rates changes the borrower's decision of which car to ultimately purchase. By illustrating the distortions induced by credit-market search frictions, this paper demonstrates a new dimension of credit-market specialness and extends the search literature that often neglects any welfare loss from search by assuming inelastic demand.

Argyle, Nadauld, and Palmer (2020) (ANP hereafter) use FICO-score-based pricing discontinuities to identify behavioral frictions associated with household debt decision making, showing that even financially unconstrained households bunch at round-number payment sizes and smooth monthly payments when facing payment shocks. While we also use FICO-score-based pricing discontinuities in this paper, here we focus on the consequences of

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

search costs in credit markets and document how interest-rate markups affect final-goods substitution patterns. The identification strategies of both ANPP and ANP are ultimately rationalized by the search frictions documented in this paper. The various nonlinear lender policies exploited for identification by ANPP and ANP and many other papers using lender-specific rules would have no effect in a perfectly competitive market in which consumers are fully informed of all prices because consumers would simply reject any loan offer featuring an arbitrary markup in rate or monthly payment.

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

1. Data

We analyze the loan contract terms and used-car purchasing decisions of 2.4 million individual borrowers in the United States from 326 retail lending institutions between 2005 and 2016. The loan data are provided by a technology firm that provides administrative data warehousing and analytics services to retail-oriented lending institutions nationwide. The majority of the loans in our data (99%) were originated by credit unions ranging between \$100 million and \$4 billion in asset size, with the remainder originated by non-bank-finance companies.[3] Given that the bulk of our data come from credit unions, we discuss data representativeness in Internet Appendix B.

Unlike the data sets used by most studies on secured credit, our data set contains information capturing all three stages of a loan's life: application, origination, and ex post performance, although our loan application data are for a subset of our origination data, consisting of approximately 1.3 million loans from 41 different institutions (owing to the smaller number of lenders that share applications data with our provider). The loan application data report borrower characteristics (age, gender, imputed minority status, FICO scores, and debt-to-income (DTI) ratios at the time of application), whether a loan application was approved or denied, and whether it was subsequently withdrawn or originated. For originated loans, the data additionally include information on loan amounts, loan terms, car purchase prices, loan performance, and collateral characteristics. Using Vehicle Identification Numbers (VINs), we observe the make, model, and model year of the purchased car. We restrict our sample

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

Table 1
Summary statistics

	Count	Mean	SD	25th	Percentile 50th	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

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.

to direct loans (those not intermediated by a dealer) to address concerns that indirect loans are potentially endogenously steered by sellers to specific financial institutions (perhaps because car dealers become aware of lenders' pricing rules over time).[4] Finally, to measure ex post loan performance, we observe 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

[4] The terms direct and indirect loans refer, respectively, to whether the borrower applied for a loan directly to the lending institution or through an auto dealership that then sent the loan application to lending institutions on the buyer's behalf. While indirect loans are more common overall, direct loans are more common among borrowers financing used cars through credit unions. Private transactions, a large share of the used-car market, are necessarily financed by direct loans. See Table A1 in the Internet Appendix for summary statistics for the indirect loans excluded from our estimation sample. The comparability of the original and discontinuity samples mitigates concerns that the average treatment effect we estimate is representative only of the population of borrowers local to one of our many detected discontinuities (Bertanha 2020).

seeks approval for a six-year \$20,000 loan at a median interest rate of 4.0%.[5] Borrowers applying for loans in our data have an average credit score of 648. The percentage of loans approved is 50%, with 65% of the approved borrowers subsequently originating a loan (which we refer to as the 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.

Panel C tabulates measures of ex post loan performance. Defining default as a loan that is at least 90 days delinquent, default rates average 2.2%.[6] Lending institutions periodically check the credit score of their borrowers subsequent to loan origination, allowing us to monitor borrowers' financial performance across their liabilities and to assess the extent of adverse selection at the origination stage. Summary statistics for ΔFICO represent percentage changes in borrowers' FICO scores from the time of origination to the lender's most recent (soft) pull of their FICO score. Updated FICO scores indicate that borrowers on average experience a 1% reduction in FICO score since origination, although borrowers with FICO scores below 600 on average realize a 5.7% increase in FICO score.

2. Conceptual Framework

In this section, we discuss the variation in the costs and benefits of searching for credit and how they interact to affect real outcomes (see Internet Appendix A for a theoretical model of costly search for credit with complementary goods). When borrowers face credit constraints that depend on their debt installment payments, cheaper credit allows them to borrow more, potentially changing their consumption. Borrowers will search when their expected net benefit of searching further for credit is positive. The expected benefit of more search is the increase in indirect utility consumers receive from the interest-rate reduction they expect to realize. The process of searching, in turn, entails incurring some search cost k, measured in utility terms. Key components of the cost of searching for credit include costs of one's time and the hassle, such as researching potential lenders and submitting signed loan applications. Other frictions captured by the concept of search costs include concerns about perceived credit-score impacts of loan applications (Liberman, Paravisini, and Pathania 2021) and the need to find a loan quickly to finalize the purchase of a car.

Many search models assume discrete demand where consumers choose whether to buy a single unit of a final good, allowing researchers to characterize

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

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

search benefits in price terms (e.g., Hong and Shum 2006). However, when demand is elastic, the utility loss associated with a given markup will be different from the size of the markup itself. Under continuous and elastic demand, the return to search will include the change in utility from elastic borrowers increasing both their loan amounts and good expenditure when finding a lower interest rate offer. In our specific setting, inferring financing search costs simply from the distribution of markups would fail to account for the fact that the borrower would be paying both the interest rate markup and the disutility of a cheaper and older car.

Combining search costs and expected benefits, borrowers search if the expected utility gain from another price draw given the current interest rate quote r' exceeds the cost of obtaining another quote. Fixing an outside option interest rate r', more search is likely to be optimal when search costs are lower. Comparing two otherwise identical borrowers with different search costs, the borrower with lower k will be more likely to search more and find a lower interest rate $r < r'$. An important implication of costly search for credit then is that borrowers facing high search costs will be more likely to accept higher interest rates. It follows given downward-sloping demand and the complementarity between credit and durables that weakly higher accepted interest rates from higher search costs are predicted to lead to smaller loan sizes, lower purchased car services, and less car expenditure.

2.1 Benefits to search

In assessing the expected utility gain from search, a borrower must evaluate the distribution of potential rates, including the likelihood that additional search would yield a better draw from the available distribution. In conceptualizing expected gains to search, consider a simple empirical measure of the interest-rate improvement a borrower would expect were she to engage in one additional round of search. We group applicants and borrowers that are in the same commuting zone and in the same five-point FICO bin, 10-point DTI bin, 6-month time period, \$1,000 collateral-value bin, and have the same maturity. We then compare the rate a given borrower has to the rates of otherwise similar borrowers, adjusting for the likelihood that comparable rates are an improvement. See Internet Appendix C for further details on the measurement of expected interest-rate improvement. This calculation suggests that the average borrower would expect to improve their interest rate by 81–113 bps were they to engage in additional search.

While this simple calculation illustrates the concept of expected search benefits, the endogeneity of interest rates and the data limitation of only observing approximately 5% of the used car market prevent us from being able to use this variation in the incentives to search directly in our estimation. Instead, we isolate measurable and quasi-exogenous variation in the expected benefits to search using an RD framework, as we describe in detail in Section 3.1.

2.2 Costs of search

Potential borrowers face a variety of nonmonetary 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. Such borrowers may purchase used cars from a seller that does not have a financing arm (e.g., private transactions), seek loan preapproval before negotiating with the seller over purchase price to refine their own budget, or seek to avoid double marginalization (Busse and Silva-Risso 2010; Grunewald et al. 2020). As car-loan pricing is specific to the credit risk of each individual, obtaining rate quotes in the direct market most often entails filling out a loan application, undergoing a credit check, and potentially verifying assets and income. For measurability and the potential to isolate exogenous variation, 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.[7] However, we note 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 affect credit scores (Liberman, Paravisini, and Pathania 2021).[8] Given the many contributors to the reduced-form concept of search costs, we view our results as providing a lower bound on the role of search and information frictions in affecting consumer borrowing and consumption.

To proxy for distance-based search costs, we use FDIC and NCUA data to identify the precise physical location of every bank branch and credit-union branch in the United States for each year in our application data. We then create a measure of the number of nearby financial institutions by calculating the driving-time lender density for each borrower. To do so, we geocode and count the number of physical branch locations within a 20-minute drive of the borrower's address.[9] This density measure is designed to capture the effort (proxied by time and distance) for each borrower to shop for an additional interest-rate quote from a lending institution that is within a reasonable distance from their

---

7 We discuss the option borrowers have to search for loans online in Internet Appendix F.2.

8 Note that we do not consider several other plausible proxies for search costs in our data because of their ambiguous 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.

9 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.

home.[10] Supporting this search-cost proxy, Degryse and Ongena (2005) find evidence of the important role of transportation costs in local credit markets.[11]

Borrowers in the 25th percentile of driving distance live less than a 20-minute drive from 23 lending institutions compared to 168 institutions for borrowers in the 75th percentile. Our baseline results categorize borrowers as having high search costs if their home address is within a 20-minute drive of at most 10 lending institutions, although our results are robust to alternative definitions (see Internet Appendix F.5 and Figure A1 in the Internet Appendix). This definition classifies roughly 15% of borrowers as living in high search-cost areas and is designed to capture the diminishing effect of an additional nearby lender on search costs. Figure A2 in the appendix geocodes the borrowers in our data in the Jacksonville, Florida metropolitan area to illustrate the spatial distribution of our search-cost classification. Relative to the blue dots for high-search-cost borrower locations, the red dots for borrowers with more than ten nearby lenders are concentrated closer to downtown Jacksonville.

To provide a more direct test of whether search behavior varies with our search cost measure, Internet Appendix D shows that borrowers facing low search costs indeed submit a higher number of loan applications to lenders covered by our data than borrowers facing high search costs. Panel A of Table A2 in the Internet Appendix uses HMDA data on the near-universe of mortgage applications to show that borrowers in tracts in which our search-cost measure is on average high submit 0.16 fewer applications than borrowers in tracts where our average search-cost measure is low. Panel B of Table A2 uses our car-loan application data to show a similar result; applicants facing higher search costs submit 0.07 fewer applications, on average, than applicants facing lower search costs. The lower coefficient is unsurprising given that we do not observe applications to institutions outside of our data. Overall, Table A2 supports our interpretation of the number of nearby lenders affecting the costliness of shopping for credit.

Because our measure of search costs could conflate differences in financial sophistication, income, creditworthiness, or market power, we employ a measure of predicted search costs that isolates exogenous variation in the number of nearby lenders. We discuss this measure in detail in Section 3.2.

[10] While distance can also proxy for soft-information-producing relationships (see Nguyen 2019; Granja, Leuz, and Rajan 2022), auto loans are not a particularly relationship-intensive credit product. Consistent with this, we find a lack of adverse selection around discontinuities and a high R^2 in our interest-rate regressions based on lender pricing rules. While brand loyalty effects could still affect take-up (Allen, Clark, and Houde 2019), our RD design allows within-lender analysis. Internet Appendix F.1 also presents findings inconsistent with the number of nearby lenders directly affecting outcomes through market concentration.

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

3. Identification Design

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

$$takeup_{it} = \beta_0 + \beta_1 r_i + \beta_2 k_{it} + \beta_3 r_i \times k_{it} + \varepsilon_{it}. \tag{1}$$

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

The ideal experiment would randomly assign both the benefits and costs of search to break the correlation between r, k, and ε. To randomize the benefits of search, we exploit quasi-exogenous within-lender markup variation in our data that serves as a laboratory in which the potential gains to search are quasi-randomly assigned across borrowers (detailed in Section 3.1). To isolate exogenous variation in the costs of search (introduced in Section 2.2), we develop an instrument that predicts high search costs using preperiod differences in branch density interacted with national changes to bank branches (detailed in Section 3.2).

3.1 Identifying exogenous variation in the benefits of search

As noted, individual-level heterogeneity in transacted prices could be driven by market mismeasurement or could be correlated with unobserved heterogeneity in search costs or taste shocks that could plausibly be correlated with other outcomes or durable-good product characteristics.

To address this, our RD design exploits discontinuous lender-level pricing schedules to assign otherwise nearly identical borrowers to high or low offered interest rates. Facing a high initial quote r' should be relatively inconsequential for borrowers with low search costs. Among borrowers with low search costs, those borrowers with high initial quotes r' should have similar loan and durable consumption outcomes as borrowers with low initial quotes because borrowers drawing high initial quotes should be willing to search further. As such, rate discontinuities should have stronger effects on search behavior and borrowing and consumption outcomes for high versus low-search-cost borrowers.

Unlike in mortgage lending, in auto lending no industry standard exists for FICO thresholds, and credit unions were prohibited from securitizing auto loans until 2017 (after our sample ends). This heterogeneity in FICO pricing discontinuities across lenders plausibly makes each given lender's FICO bin locations unknown to most potential borrowers. While auto-loan lending institutions do not adhere to a common set of FICO cutoffs, the use of discontinuous pricing at some point across the FICO spectrum is prevalent for more than half of the lenders in our data.[12] See Bubb and Kaufman (2014), Al-Najjar and Pai (2014), Matějka and McKay (2015), Livshits, MacGee, and Tertilt (2016), and Agarwal et al. (2017) for models of credit risk processing that each rationalize binning risk types in pricing decisions. FICO discontinuities may have been incorporated into software systems as a holdover from a time when pricing was done via rate sheets instead of automated algorithms (Hutto and Lederman 2003) and could persist in part because costly consumer search prevents more accurately risk-based pricers from gaining market share. See Internet Appendix E for a description of the procedure we follow to detect pricing discontinuities in our data.

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

First, the RD design assumes that the probability of borrower treatment (i.e., offered interest rates) with respect to loan terms is discontinuous at detected FICO thresholds. Second, valid RD requires that borrower attributes (observed or unobserved) that could influence loan outcomes change only continuously at interest-rate discontinuities. This smoothness condition requires that borrowers on either side of a FICO threshold are otherwise similar, such that borrowing outcomes on either side of a threshold would be continuous absent the difference in treatment induced by policy differences at the threshold. We provide evidence that the smoothness condition is satisfied in Internet Appendix F.3 and Figure A3 in the Internet Appendix. Across a large set of observables, we do not see any statistically or visually significant changes at a FICO discontinuity, which we take to be strong evidence that unobservables are also not changing across FICO discontinuities.

Panel A of Figure 1 plots average interest rates against normalized borrower FICO scores for a sample restricted to loans with borrower FICO scores between 581 and 619. The plots demonstrate smoothness in the conditional expectation function, except for the points corresponding to FICO scores of 599 and 600, where interest rates jump discontinuously. We repeat the plot

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

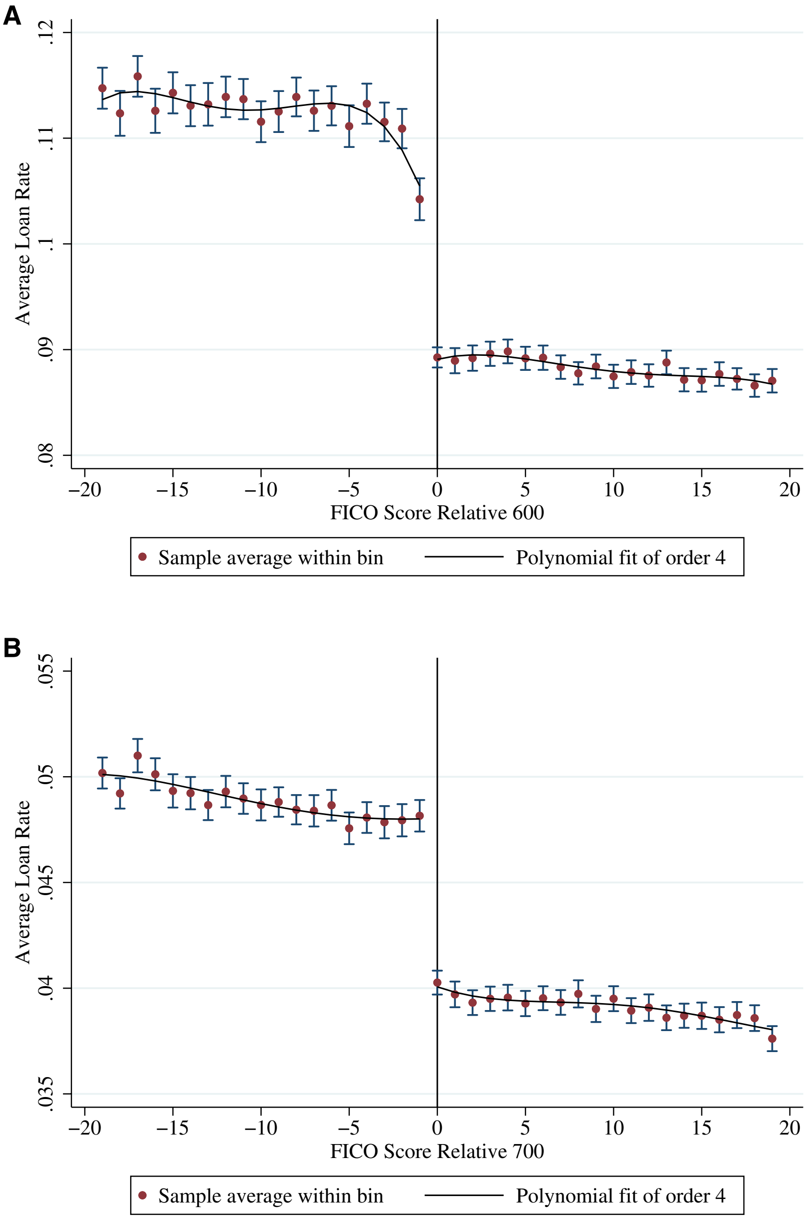

Figure 1
Interest-rate FICO regression discontinuity plots
The panels plot average interest rates against borrower FICO scores normalized to pricing discontinuities. Ninety-five percent confidence intervals are double clustered by lender and FICO score. Plotted RD functions are estimated using the Calonico, Cattaneo, and Titiunik (2014) robust RD estimator with fourth-degree polynomials for institutions with pricing discontinuities detected at FICO scores of 600 and 700 in panels A and B, respectively.

using similar 38-point FICO ranges for the 700 FICO thresholds in panel B of Figure 1. These plots confirm the existence of large and precise interest-rate discontinuities at these thresholds.

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

$$r_{iglt} = \pi_1 \widetilde{FICO}_i + \delta \cdot 1(\widetilde{FICO}_i \geq 0) + \pi_2 \widetilde{FICO}_i \cdot 1(\widetilde{FICO}_i \geq 0) + \alpha_{gt} + \gamma_l + \upsilon_{iglt} \quad (2)$$

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

There are two differences between (2) and our actual estimating equation. First, we allow for the effect of the running variable $\widetilde{FICO}$ above and below the cutoff at $\widetilde{FICO} = 0$ to be quadratic. Second, to deal with loans that may be within 19 FICO points of multiple discontinuities, we sum across discontinuities d from the set of discontinuities $\mathcal{D}$ to estimate

$$r_{iglt} = \sum_{d \in \mathcal{D}} 1(il \in \mathcal{D}_d) \Big(\delta \cdot 1(\widetilde{FICO}_{id} \geq 0) + f(\widetilde{FICO}_{id}; \pi) + \varphi_{dl} \Big) + \alpha_{gt} + \upsilon_{iglt} \quad (3)$$

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 and φ_{dl} are discontinuity $\times$ lender fixed effects to allow for each lender to have a different selection of borrowers around each threshold. The function $f(\cdot;\cdot)$ is defined as

$$f(x;\pi) = \pi_1 x + \pi_2 x^2 + 1(x \geq 0)\left(\pi_3 x + \pi_4 x^2\right) \quad (4)$$

to allow for a smooth but nonlinear effect of the running variable that potentially changes shape discontinuously at the threshold.[13] Standard errors are double clustered by lender and FICO score, and the sample used to estimate (3) is the discontinuity sample described in Table A3 in the Internet Appendix.[14]

Table 2 presents results of this exercise for varying levels of fixed effects. Using the stringent ZIP code $\times$ quarter and discontinuity $\times$ lender fixed effects in column 4, interest rates for borrowers with FICO scores immediately above a detected threshold are an average of 1.37-percentage-points lower than

---

[13] The specification in (3) also allows us to accommodate loans on the left of one threshold and the right of another, similar to Agarwal et al. (2017). Pooling discontinuities this way estimates the average treatment effect for the population local to the discontinuities (Bertanha 2020).

[14] 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 2
Effects of FICO discontinuities on loan rate

	(1)	(2)	(3)	(4)	(5)
Discontinuity indicator	−0.0076**	−0.0120**	−0.0131***	−0.0137***	−0.0128***
	(0.0038)	(0.0048)	(0.0042)	(0.0043)	(0.0043)
Discontinuity indicator × High search cost					−0.0018
					(0.0015)
RD controls	✓	✓	✓	✓	✓
Discontinuity × lender FEs		✓	✓	✓	✓
CZ × quarter FEs			✓		
ZIP code × quarter FEs				✓	✓
Number of observations	514,834	514,834	514,834	514,834	514,834
R^2	.05	.05	.22	.46	.22

This table reports regression-discontinuity estimates of Equation (3) with the indicated fixed effects, normalizing FICO scores around each threshold using a uniform kernel and a bandwidth of 19 FICO points. Discontinuity indicator is a dummy for whether a borrower's normalized FICO score is positive. High search cost is a dummy equal to one for borrowers with at most 10 financial institutions within a 20-minute drive from their home. RD controls consist of a quadratic spline in normalized FICO score that is allowed to change at the discontinuity as specified in (4). Robust standard errors in parentheses are double clustered by lender and FICO score.

borrowers just below. Given an average interest rate in our estimation sample of 6.0% (panel B of Table A3), the magnitude of these effects is economically meaningful, amounting to a $488 higher present value and a $9 higher monthly payment for otherwise identical loans taken out by borrowers on the expensive side of a FICO discontinuity. In the context of our conceptual framework, drawing a rate quote from the expensive side of an interest-rate discontinuity constitutes a significantly higher initial interest rate r', and, depending on whether searching for financing is costly, could have material consequences on the ultimate cost of credit for such borrowers. As we discuss in Section 4 after introducing our measure of financing search costs, column 5 of Table 2 shows that the average FICO discontinuity has similar markup effects for borrowers facing high and low search costs.

Figure A4 in the appendix illustrates the link between pricing discontinuities and the benefits to search. Because below-threshold borrowers on average start out farther to the right in the distribution of local interest rates than above-threshold borrowers, their expected benefits to search are significantly larger. Figure A4 plots the density of the spread to the lowest available rate for left- and right-of-threshold borrowers using the matching strategy described in Internet Appendix C. Dotted and solid lines represent borrowers just below and above a given threshold, respectively. The average spread to the lowest available rate and the variance of those spreads are larger for left-of-threshold borrowers, implying that the incentives to search are higher for those below-threshold borrowers offered exogenously higher rates. The discontinuities thus offer a quasi-random proxy for the incentives to search. Internet Appendix C discusses the robustness of this illustrative exercise to adjustments for the probability of finding a better offer and for expected approval rates.

3.2 Identifying exogenous variation in the costs of search

Borrowers who live within 20 minutes of more or less than 10 lenders may be different from each other on dimensions other than search costs. Thus, without exogenous variation in search costs, a positive correlation between accepted interest rates and search costs could be due to other unobserved demand factors ε potentially correlated with search-cost proxies, such as credit limits, financial literacy, preferences, or market concentration.

Our RD setup mostly accounts for even unobservable differences across high- and low-search-cost borrowers by comparing borrowers within lender, commuting zone, and time who are randomly assigned to above or below pricing discontinuities; that is, differences in income, creditworthiness, or financial literacy are differenced out by our RD specifications before comparing across search-cost areas. In addition, applicants facing high and low search costs also look similar on observables; as we discuss, two measures of creditworthiness tell mixed but muted stories about applicant differences across our search cost measure. On average, applicants facing high search costs have slightly lower FICO scores (0.3 points) but also lower DTI ratios (2.6 points). To address any unobserved heterogeneity correlated with both our search-cost proxy and borrowers' response to rate discontinuities, we employ a measure of predicted search costs that isolates exogenous variation in the number of nearby lenders. In Internet Appendix F.1, we further address the possibility that the number of local lenders may directly affect the level of competition in ways unrelated to consumer search.

To understand how unobserved heterogeneity could affect our contrasting of effects across high- and low-search-cost borrowers, consider the following example. If a given location has few nearby lenders because lenders anticipate future local economic conditions will be poor, then our proxy for high search costs may be correlated with borrowers in such areas being constrained and more likely to accept high-markup loan offers. Such local endogeneity of nearby lender density could generate the results we present in Section 4 even if the number of nearby lenders had remained high.

Such factors bias our estimates of β_3 in (1) only to the extent low-search-cost borrowers respond to pricing discontinuities differently from high-search-cost borrowers for reasons other than differences in search costs. To separately identify β_3 from the effect of any such unobservables, we introduce an instrument for search costs based on historical variation in lender branch density combined with time-series variation in the nationwide total number of lender branches. By exploiting quasi-exogenous time-varying variation in our search-cost proxy, we address possible bias from such time-varying omitted variables that are correlated with current search costs $k_{it} = HighSearchCost_{it}$. Internet Appendix F.6 formalizes this identification argument.

Specifically, we calculate the number of nearby bank and credit union branches for each borrower in our sample using the density of nearby financial institutions as of 1990 using NETS historical data on the location of every financial institution in the United States. We then grow the 1990 lender density measure using the national growth rate in financial institutions from 1990 through the year of each observation in our sample period. The identifying assumption behind this instrument is that conditional on borrower controls and local geographic fixed effects, the local cross-sectional variation interacted with the national time-series variation in the predicted search cost measure is unrelated to time-varying local demand shocks. The predicted high-search-cost indicator is

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

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

$$\#\ \widehat{\text{nearby branches}}_{it} = \#\ \text{nearby branches}_{i,1990} \times \frac{\text{nationwide branches}_t}{\text{nationwide branches}_{1990}}. \tag{6}$$

Variation in this predicted high-search-cost measure is driven by local branching concentration in 1990 and aggregate variation in national branching trends. Table A4 in the Internet Appendix demonstrates that $\widehat{HighSearchCost}_{it}$ is a strong predictor of $HighSearchCost_{it}$, with partial F-statistics ranging from 12 to 83, depending on the controls.

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

Table 3 shows that our search cost measure is generally unrelated to credit characteristics. The statistically and economically insignificant credit score coefficient in column 1 indicates that borrowers in areas predicted to have higher search costs had FICO scores that were 0.3 points lower than borrowers in lower search cost areas. Both search costs and predicted search costs are

Table 3
Exclusion restrictions tests of high financing search cost measures

	(1)	(2)
Dependent variable	Search cost measure	
	High search cost	High $\widehat{\text{search}}$ cost
	A. Borrower characteristics	
Credit score	−0.267	1.220
	(0.636)	(1.624)
Debt-to-Income ratio	−0.026	−0.009
	(0.031)	(0.013)
Minority indicator	−0.020	−0.028
	(0.015)	(0.020)
Age (years)	1.98	−0.15
	(1.50)	(0.65)
Male indicator	0.004	−0.010
	(0.027)	(0.023)
	B. Census tract characteristics	
Median income	2251.05	1490.58
	(2185.06)	(911.92)
Poverty share	−0.013	−0.002
	(0.010)	(0.006)
Average rent	−34.39**	−4.40
	(13.88)	(16.45)
log(Wage growth)	0.115**	0.070
	(0.047)	(0.042)
Job growth rate	0.009**	0.006
	(0.003)	(0.006)

This table reports OLS estimates testing whether measures of high search costs predict borrower characteristics (panel A) and tract characteristics (panel B). Each row reports the coefficient for high search cost (column 1) or predicted high search cost (column 2) in a regression on the listed dependent variable. High search cost is a dummy equal to one for borrowers with at most 10 financial institutions within a 20-minute drive from their home at the time of their loan application. Predicted high search cost is an indicator for whether the predicted number of financial institutions within a 20-minute drive is at most 10, with the predicted number defined by (6) using 1990 local branch proximity and subsequent national trends in branching. Tract-level characteristics in panel B have been compiled from Chetty et al. (2018). Median income is the tract median income from the 2015 ACS. Poverty share is the share of tract residents below the poverty line in the 2006–2010 ACS. Average rent is the tract average rent of two-bedroom apartments in the 2015 ACS. The logarithm of wage growth is the change in the logarithm of average hourly wages for high school graduates between the 2005–2009 and 2010–2014 five-year ACS. Job growth rate is the average annual growth rate from 2004 to 2013 of the number of jobs from the Census Bureau's LODES-WAC files. All specifications include ZIP code by quarter fixed effects. Robust standard errors in parentheses are double clustered by lender and FICO score.

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

Overall, the results of Table 3 support our use of predicted search costs to ensure that our results on loan demand are driven by costly search. Combining these results with those in Table A2 in the Internet Appendix rules out most forms of unobserved heterogeneity that could confound our interpretation of the effect of nearby lenders on outcomes. For an unobserved correlate of our search-cost measure to drive our results instead of search costs, it would have to affect borrowers on one side of a pricing discontinuity differently than borrowers on the other, be correlated with the number of mortgage and car loan applications submitted by the average borrower, and yet be uncorrelated with average borrower and borrower tract characteristics.

Finally, while markups would also distort quantities in an oligopolistic market, too, simple market concentration explanations for our results fail to predict the price dispersion we observe. In addition, we show in Table A5 in the Internet Appendix that the contrast between high- and low-search-cost borrowers holds in both concentrated and less concentrated markets. Market concentration is not random, but to explain our results, concentration would need to vary discontinuously within the discontinuity-lender and ZIP code-quarter pairs. Finally, Figure A5 in the Internet Appendix shows that market shares seem unrelated to markups for high-search-cost applicants and negatively related for low-search-cost applicants (see Internet Appendix F.1 for a further discussion). Taken together, these plots confirm that independent of market concentration, 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.

4. Effects of Search Frictions on Loan Take-Up

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

In the spirit of a difference-in-differences specification, our empirical specification measures how differences in loan take-up rates around FICO thresholds vary with search costs.[15] This allows us to exploit the quasi-random assignment of borrowers in the neighborhood of a FICO discontinuity to

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

high and low markups to control for any unobserved differences in borrowers across search-cost categories. Even if borrowers facing high and low search costs do vary meaningfully on some unobservable dimension, the conditional variation in the FICO-discontinuity indicator variable allows us to difference out any demand-side variation common to borrowers at the same lender, FICO threshold, market, and time period and then compare the resultant responses across search-cost areas.

To implement the difference-in-discontinuities strategy described in Internet Appendix F.6, we augment our RD specification (3) with two additional controls: a high search cost indicator that proxies for k_{high} and an interaction between this dummy and the discontinuity indicator $1(\widetilde{FICO}_{id} \geq 0)$ that proxies for $r_{high} = 0$:

$$takeup_{iglt} = \sum_{d \in \mathcal{D}} 1(il \in \mathcal{D}_d)\Big(\beta_1 \cdot 1(\widetilde{FICO}_{id} > 0) + \beta_2 \cdot HighSearchCost_{it} + \beta_3 \cdot 1(\widetilde{FICO}_{id} \geq 0) \times HighSearchCost_{it} + f(\widetilde{FICO}_{id}; \theta) + \lambda_{dl}\Big) + \eta_{gt} + \varepsilon_{iglt}. \quad (7)$$

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

Table 4 reports estimates of (7), beginning with η_{gt} defined as commuting zone × quarter fixed effects. Column 1 reports an 11.4 pp difference in take-up rates across FICO discontinuities pooling borrowers with high and low search costs; that is, borrowers quasi-randomly drawing high markups from the distribution of interest rates are 11.4-percentage-points less likely to accept the offered loan. Column 2 adds controls for high search costs and the interaction term of high interest rates and high search costs, defined as borrowers with at most 10 lending institutions within a 20-minute drive. The coefficient β_2 on the high-search-cost indicator is statistically and economically insignificant. That lower-interest-rate borrowers have quite similar take-up rates in our

Table 4
Effect of FICO discontinuities on loan offer take-up decisions by search costs

	(1)	(2)	(3)	(4)	(5)
Discontinuity indicator	0.114**	0.129***	0.142**	0.127***	0.141**
	(0.038)	(0.041)	(0.061)	(0.040)	(0.063)
High search cost		0.002	−0.008		
		(0.035)	(0.058)		
Discontinuity indicator × high search cost		−0.108***	−0.145**		
		(0.030)	(0.048)		
$\widehat{\text{High search cost}}$				−0.013	−0.004
				(0.025)	(0.037)
Discontinuity indicator × $\widehat{\text{high search cost}}$				−0.071***	−0.105***
				(0.016)	(0.033)
RD controls	✓	✓	✓	✓	✓
Discontinuity × lender FEs	✓	✓	✓	✓	✓
CZ × quarter FEs	✓	✓		✓	
ZIP code × quarter FEs			✓		✓
Number of observations	30,743	30,743	30,743	30,743	30,743
R^2	.31	.31	.61	.31	.61

This table reports regression-discontinuity estimates of whether a borrower accepts an approved loan offer regressed on a discontinuity indicator, a high-search-cost indicator, their interaction, and the indicated controls and fixed effects using the specification in (7). Discontinuity indicator is a dummy for whether a borrower's normalized FICO score is positive. High search cost is a dummy equal to one for borrowers with at most 10 financial institutions within a 20-minute drive from their home. The predicted high search cost measure is an indicator for whether the predicted number of financial institutions within a 20-minute drive is at most 10, with the predicted number defined by (6) using 1990 local branch proximity and subsequent national trends in branching. RD controls consist of a quadratic spline in normalized FICO score that is allowed to change at the discontinuity as specified in (4). Robust standard errors in parentheses are double clustered by lender and FICO score. See the legend to Table 2 for further details.

specification supports our identifying assumption of the comparability of borrowers in these two areas conditional on our RD controls and fixed effects. While low-search-cost borrowers are $\hat{\beta}_1 = 12.9$ pp more likely to accept a loan offer when they are just above a pricing discontinuity than below, high-search-cost borrowers are equally likely to accept loans on either side of a discontinuity—we cannot reject that $\beta_1 + \beta_3 = 0$ in column 2. Coupled with FICO pricing discontinuities being indistinguishable for high- and low-search-cost borrowers, this suggests that borrowers who have higher search costs are much more willing to accept high interest rates than borrowers with lower search costs.[16] Column 3 shows that this result is robust to using within ZIP code × quarter variation.[17]

16 A threat to our interpretation of Table 4 is that the magnitude of the interest-rate discontinuity ($r_{high} - r_{low}$ in Equation (18) in Internet Appendix F.6) may be different for high- and low-search-cost borrowers, naturally leading to differences in responses to the discontinuities. Note, however, that for this to explain our results, low-search-cost borrowers would have to face *larger* discontinuities in rates. Column 5 of Table 2 finds a small and insignificant difference in the size of the average interest-rate change at a FICO threshold across our measure of search costs, confirmed by graphical evidence in panel A of Figure 2. Similarly, we demonstrate in Internet Appendix C that differences in the benefits to search across areas cannot explain our results given high-search-cost borrowers if anything can expect higher benefits to search from higher price dispersion.

17 See panel B of Figure 2 for graphical evidence parallel to that in column 3 that the discontinuity in take-up rates is only present for low-search-cost borrowers.

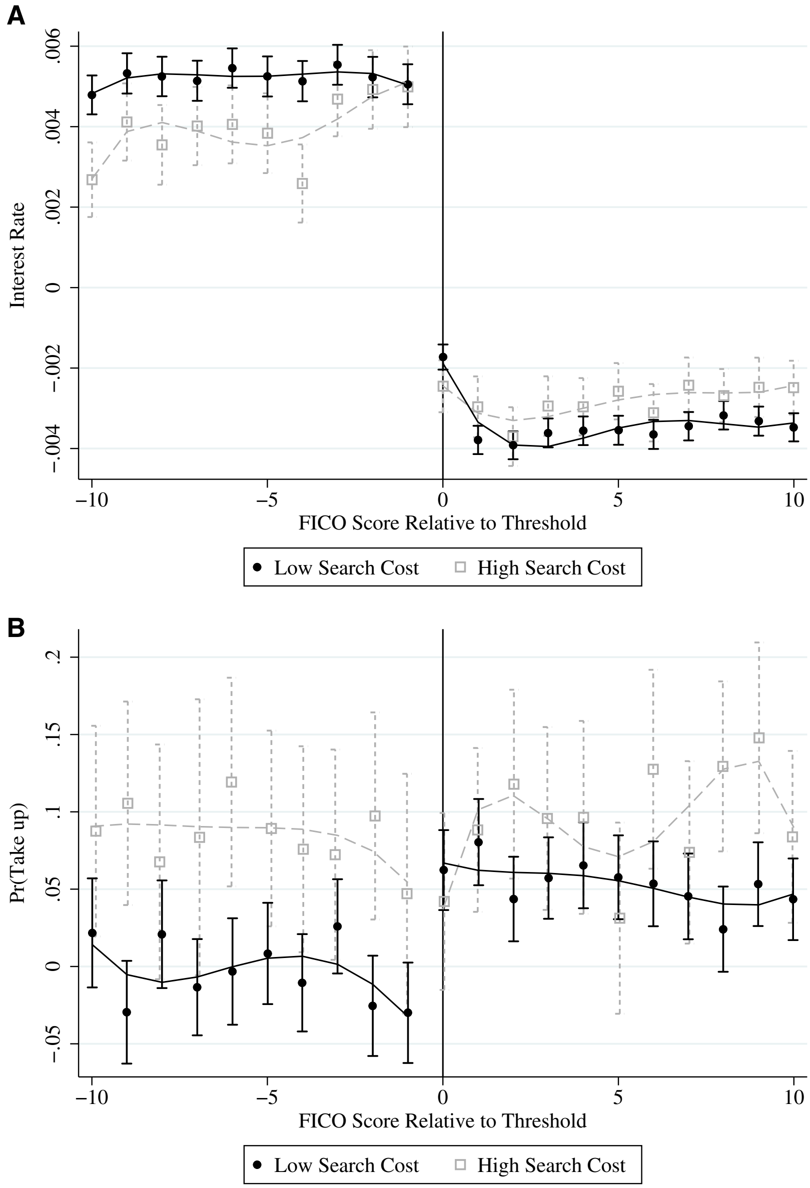

Figure 2
Interest rates and take-up around discontinuities by search costs
The panels plot average interest rates (panel A) and loan take-up rates (panel B) by search costs and FICO scores normalized to detected pricing discontinuities. Hollow gray squares and solid black circles represent the sample of borrowers with high and low search costs, respectively, defined as having at most or more than 10 financial institutions within a 20-minute drive of their home address. Outcomes are residualized by origination-quarter × ZIP code. Ninety-five percent confidence intervals are double clustered by lender and FICO score. Plotted RD functions are for institutions with pricing discontinuities and estimated using the Calonico, Cattaneo, and Titiunik (2014) robust RD estimator with fourth-degree polynomials.

Given the mean take-up rate of 0.51 (Table A3 in the appendix), being a high-search-cost borrower has a 20-30% effect on the likelihood a borrower will accept an expensive loan. These differences in take-up rates are consistent with our conjectured mechanism where borrowers with high search costs are more likely to accept high interest-rate markups rather than search, potentially adjusting their consumption on other margins as discussed below.

Columns 4 and 5 of Table 4 repeat the specifications of columns 2 and 3 but replace our high-search-cost indicator with a predicted measure of search costs to address the possibility that omitted variables correlated with $r \times k$ could bias our estimates of β_3. To do so, we reestimate (7) instead conditioning on predicted search costs $\hat{k}_{it} = \widehat{HighSearchCost}_{it}$ as defined in (5). While an unobservable like $low\ financial\ literacy \times r_{high}$ may be correlated with $k_{high} \times r_{high}$, such unobservables are conceptually unlikely to be correlated with $\hat{k}_{high} \times r_{high}$, supported empirically by Table 3. Again, the rationale behind this identifying assumption used in columns 4 and 5 is that the current unobserved demand factors that might affect high- and low-markup borrowers differently for borrowers facing low versus high search costs are unlikely to be correlated with the 1990 local lender density used to define $\hat{k}$. As before, borrowers predicted to have higher search costs are significantly less sensitive to interest-rate discontinuities, with the point estimates for β_3 in columns 4 and 5, suggesting they are 7- to 10-pp less likely to accept a higher-markup offer. Borrowers with low search costs exhibit strong and statistically significant reactions to loan offers. The estimated main effect $\hat{\beta}_2$ of higher search costs is again statistically and economically insignificant, consistent with our specification rendering low-markup borrowers in these areas roughly comparable. These results suggest that differences in loan take-up rates across borrowers predicted to have high and low search costs are not driven by the endogeneity of the local branching network to time-varying economic shocks. Finally, combining these results with Table 3 also rules out some sort of persistent spatial variation in credit-demand elasticities being correlated with our search-cost measure.

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

5. Real Effects on Loan and Consumption Outcomes

To directly demonstrate how search frictions in credit markets can have real effects on consumption, we next establish that being treated with a higher offered interest rate affects subsequent loan and purchase decisions. Whether a given credit-market imperfection constrains consumption is usually challenging to identify empirically because it requires estimating counterfactual consumption in the absence of the alleged friction. However, our RD setup

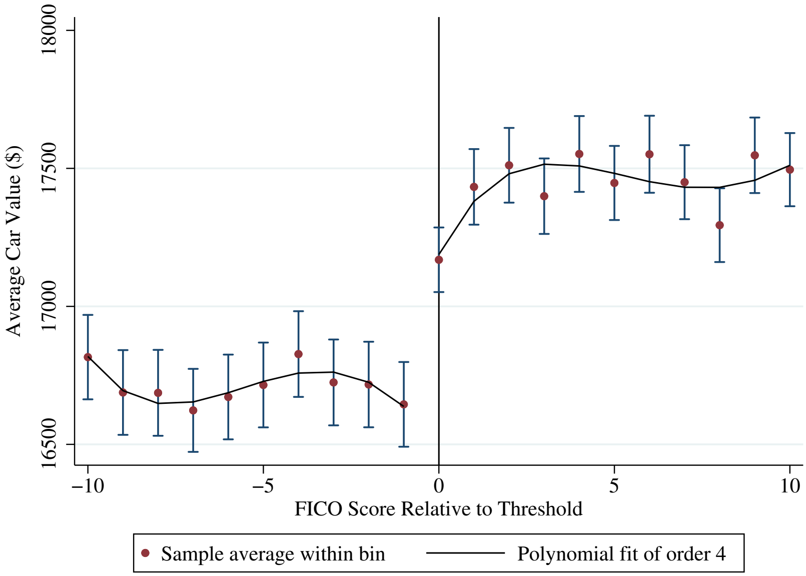

Figure 3
Effect of FICO discontinuities on value of car purchased
This figure plots average car transaction prices by FICO scores normalized to detected pricing discontinuities. Ninety-five percent confidence intervals are double clustered by lender and FICO score. Plotted RD functions are estimated using the Calonico, Cattaneo, and Titiunik (2014) robust RD estimator with fourth-degree polynomials for institutions with pricing discontinuities.

allows us to test for quantity effects by using the borrowing and purchasing decisions of borrowers on one side of a FICO threshold as a counterfactual for borrowers on the other. Given the empirical results in Internet Appendix F.3 that borrowers are ex ante similar around FICO thresholds, we assume that borrowers around FICO thresholds would have similar demand for loans and cars if quoted the same set of financing terms. Exploiting our ability to observe the exact amount borrowed and spent on a car, we test whether borrowers spend differently around the observed FICO thresholds and whether the composition of borrowers accepting loans changes across thresholds. Combined with higher interest rates on the expensive side of FICO thresholds being much more likely to be accepted by high-search-cost borrowers, these results demonstrate that search costs for credit affect subsequent consumption.

Summarizing our main result graphically, Figure 3 plots car purchase expenditures around the normalized FICO threshold. Purchase amounts are smooth leading up to the FICO threshold and then jump up discontinuously at the threshold. Using the same RD design used in our pricing discontinuity analysis above, we formally test for statistical differences in purchase amounts. As before, we estimate Equation (3) by controlling for commuting zone × quarter-of-origination fixed effects and discontinuity × lender fixed effects and

Table 5
Effects of FICO discontinuities on origination outcomes

	(1) Purchase price ($)	(2) Loan amount ($)	(3) Loan-to-value ratio	(4) Monthly payment ($)
Discontinuity indicator	375.48**	557.37***	0.014***	−0.08
	(169.55)	(165.70)	(0.003)	(1.02)
Discontinuity × lender FEs	✓	✓	✓	✓
CZ × quarter FEs	✓	✓	✓	✓
Number of observations	514,834	514,834	514,834	514,834
R^2	.08	.09	.17	.09

This table reports regression-discontinuity estimates of car-purchase and originated-loan characteristics regressed on the discontinuity indicator, our RD controls, and the indicated fixed effects following Equation (3). Discontinuity indicator is a dummy for whether a borrower's normalized FICO score is positive. Robust standard errors in parentheses are double clustered by lender and FICO score. See the legend to Table 2 for more details.

allow for a quadratic function of the running variable using a bandwidth of 19 around the normalized FICO threshold with a uniform kernel.[18]

Table 5 presents these reduced-form results. Borrowers quasi-randomly offered more expensive loans spend an average of $375 less on the cars they purchase (a 2.4% effect evaluated at the mean purchase price). Column 2 presents results with loan amounts as the dependent variable. Originated loan sizes are an average of $557 (4%) lower on the expensive side of a detected FICO discontinuity.[19] The fact that loan sizes increase by larger amounts around the threshold than purchase amounts indicates that, ex post, borrowers on the right side of the cutoff are approved for and take up higher loan-to-value (LTV) ratios. Column 3 of Table 5 indicates that ex post LTV ratios are an average of 1.4-percentage-points higher for borrowers to the right of FICO thresholds. Given that ex ante DTI ratios in the loan application data are continuous around the thresholds (Table A6 in the Internet Appendix), we interpret these results as further evidence of the easing of credit terms for above-threshold borrowers. Borrowers with FICO scores just above a pricing discontinuity are offered lower rates, longer terms, and allowed higher ex post LTV ratios.

Microdata on loan amounts and loan terms allow us to calculate the implied monthly payment of each borrower on either side of the thresholds. In column 4 of Table 5, we test whether ex post monthly payments are different around the thresholds. On average, monthly payments decrease by a statistically and economically insignificant $0.08 for above-threshold borrowers. Shorter terms and higher interest rates lead below-threshold borrowers to purchase less expensive cars and use less financing in their purchase than above-threshold

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

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

borrowers, essentially purchasing less car and using less credit to keep the same monthly payment.[20] We decompose how high-search-cost borrowers accommodate high interest rates instead of searching: roughly 50% of accepted higher loan costs is purely offset by lower LTVs, 20% is offset by buying cheaper cars, and the remainder is offset by the combination of the two.

This evidence of otherwise similar borrowers spending different amounts on the cars they purchase as a result of the financing terms they are offered is consistent with a quantity effect of search frictions. Absent search frictions, we would expect all borrowers to find the lowest-price provider leading to infinitely elastic residual demand curves for a given lender. Instead, these results suggest each lender faces a finite residual demand elasticity, consistent with search frictions giving direct auto-loan lenders market power. Given the balance of borrowers across discontinuities, we would also expect to observe similar purchasing decisions for above- and below-threshold borrowers absent their differential interest-rate draws.

One concern with this interpretation is that purchase-price effects may not represent quantity effects if borrowers only *pay* different amounts but actually purchase the same cars they would have otherwise, deriving the same flow utility from their purchase. For example, if dealers can use financing terms to price discriminate, they may exploit above-threshold borrowers' increased marginal willingness to pay by charging more for the exact same car than otherwise similar borrowers with more expensive financing.[21] We test for this possibility by controlling for year-make-model (e.g., 2013 Honda Accord) fixed effects in our RD regressions. Column 1 of Table 6 reports results when controlling for make-model fixed effects. Even within a make and model category, borrowers quasi-randomly assigned expensive credit continue to spend $344 less on cars, suggesting that the bulk of the purchasing behavior we observe in Table 5 is not driven by people choosing to purchase different model cars as a result of their assigned credit. Contrasting the coefficients in columns 1 and 2 provides indirect evidence on the nature of the substitution patterns in this market. When we include *year*-make-model fixed effects in column 2, we find a much smaller change in purchase price at the discontinuity of $72. Because fixing the model year of a car has such large explanatory power on the effect of an interest-rate markup, we conclude that much of the effect in column 1 is explained by substitution within a model and across model years.[22]

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

---

[20] See Argyle, Nadauld, and Palmer (2020) for related evidence on borrowers' monthly payment targeting.

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

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

Table 6
Effects of FICO discontinuities on vehicle purchase

	(1) Purchase price	(2) Purchase price	(3) Car age (months)
Discontinuity indicator	344.38***	71.76	−1.81***
	(120.43)	(48.43)	(0.143)
Discontinuity × lender FEs	✓	✓	✓
CZ × quarter FEs	✓	✓	✓
Make-model FEs	✓		✓
Make-model-year FEs		✓	
Number of observations	468,800	468,800	468,800
R^2	.38	.78	.37

This table reports regression-discontinuity estimates of car purchase prices (columns 1 and 2) and car age in months (column 3) regressed on the discontinuity indicator, our RD controls, and the indicated fixed effects following Equation (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 the legend to Table 2 for more details.

would be collinear with year-make-model and time fixed effects). Borrowers with access to easier credit purchase cars that are on average 1.8 months newer, suggesting that roughly one in seven borrowers respond to a high interest-rate markup by buying a car that is one model year older, keeping their monthly payments roughly constant. These car-age effects are consistent with some borrowers preferring to purchase older cars over searching for better financing. Figure A6 in the Internet Appendix uses data from the National Highway Transportation Survey (Federal Highway Administration 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 expenditures and car age are consistent with search frictions distorting the quantity of car services purchased. These changes in durable purchases affect consumers' utility in two ways, by lowering flow utility of their purchased car and the vehicle's durability and resale value. The estimated differences in purchases across interest-rate discontinuities—combined with the higher propensity of higher-search-cost borrowers to accept quasi-randomly higher interest rates—are sufficient to imply real (as opposed to purely financial) effects of financing search frictions. Overall, high financing search costs induce borrowers facing high rate markups to take out higher-interest-rate and smaller loans and purchase older and cheaper cars.

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

5.1 Assessing selection into origination

Limiting our sample to direct loans eliminates the possibility of borrower selection driven by car dealers steering of borrowers to lenders via so-called "indirect loans." Moreover, borrowers are unlikely to be aware of their precise FICO score as calculated by the credit bureau queried by the lender pricing their loan application and even less aware of that lender's or alternative lenders' FICO discontinuities. However, another possibility is that borrowers *who accept loans* on either side of FICO thresholds might differ systematically (even if borrowers are balanced at the application stage), violating the smoothness condition required for valid RD inference of second-stage effects.[23] 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.[24]

Theories of adverse selection in credit markets, such as Stiglitz and Weiss (1981), postulate that the population of borrowers who accept high interest-rate markups is disproportionately made up of borrowers with private information that they are more likely to default. Such a selection story could be at play in our setting, too, in which conditional on origination, the sample of borrowers on the expensive side of FICO discontinuities is somehow selected relative to the sample of borrowers with similar FICOs who apply but then withdraw after observing the interest rate. However, the premise of the Stiglitz and Weiss (1981) unraveling argument is that although such borrower heterogeneity is unobserved by the credit market at the origination stage, it reveals itself in higher defaults later. Our setting enables us to follow borrowers after origination and observe outcomes for borrowers on both sides of discontinuities, similar to the Chiappori and Salanié (2000) test for private information in insurance markets. Using several ex post creditworthiness measures, we find no evidence that borrowers taking up loans on either side of discontinuities are systematically different ex ante.

The evidence in Internet Appendix F.3 demonstrates that interest-rate markups seem quasi-randomly assigned (i.e., at the ex ante application stage, borrowers are similar on observable dimensions). An alternative explanation

---

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

[24] Consistent with the results below, Jansen et al. (2022) also find no evidence of selection on price in the car loan market.

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

To test for the possibility that some form of selection drives the observed equilibrium outcomes in our data, we compare the balance around the FICO discontinuities of borrower characteristics and ex post borrower performance for the accepted loan applications in our applications data. If some correlated selection process guides differences in who accepts expensive loan offers, this should be revealed by an imbalance of offer-accepting borrower characteristics or ex post credit outcomes. Figure A7 in the Internet Appendix repeats the exercise of Figure A3 by checking for the smoothness of borrower and loan characteristics around FICO discontinuities.[25] We find smoothness for the four borrower attributes that should not respond to the thresholds if there is not selection into origination based on borrower demand. Although the estimates are noisier in this applications-originations merged sample due to smaller sample sizes, age, DTI, gender, and ethnicity have economically small and statistically insignificant differences across FICO discontinuities, indicating that borrowers who accept high markups seem similar to those above FICO discontinuities that accept low markups. Panels 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 4 and 5, respectively. The results of the Figure A7 exercise suggest that the observed borrower responses to FICO discontinuities are causal effects of the search-cost-induced interest-rate markups and not explainable by selection into which borrowers accept higher markups.

As any selection on credit quality should eventually manifest itself in the average ex post credit performance of selected borrowers, we further test for selection by specifying as a dependent variable in our RD setting various ex post credit outcomes. We also interact our discontinuity indicator with our high-search-cost indicator to verify that the selection into origination is not different for high-search-cost borrowers. The coefficient in column 1 of Table 7

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

Table 7
Balance of ex post credit outcomes across FICO discontinuities

	(1) Days delinquent	(2) Charge-off	(3) Default	(4) %ΔFICO
Discontinuity indicator	3.55	0.004	0.001	0.002
	(2.75)	(0.002)	(0.003)	(0.003)
Discontinuity indicator × high search cost	4.55	−0.004	−0.004	0.002
	(4.96)	(0.004)	(0.004)	(0.003)
Discontinuity × lender FEs	✓	✓	✓	✓
CZ × quarter FEs	✓	✓	✓	✓
Number of observations	331,590	514,834	514,834	405,236
R^2	.49	.21	.24	.06

This table reports reduced-form RD estimates of Equation (3) on ex post loan and borrower outcomes. Days delinquent is the number of days a borrower is delinquent as of our data extract. Charge-off is an indicator for whether a loan has been written off the books of the lending institution. Default is an indicator for whether a borrower has been delinquent for at least 90 days. Percentage change in FICO score is the change in FICO score since origination as a fraction of the FICO score at origination for the subsample of institutions that report credit scores after loan origination. Robust standard errors in parentheses are double clustered by lender and FICO score. See the legend to Table 2 for estimation details.

estimates that above-threshold borrowers are an average of 3.5 more days delinquent than below-threshold borrowers, indicating that borrowers on either side of the threshold do not exhibit economically meaningful or statistically significant differences in delinquency.[26] Similarly, above-threshold borrowers are 0.4-percentage-points more likely to have their loan charged off (written off as a loss by the lender, column 2) and 0.1-percentage-points more likely to be in default (over 90 days past due, column 3), which we consider precise zeroes. For each of these outcomes in columns 1–3, the magnitude of the discontinuity in ex post credit outcomes is statistically indistinguishable for borrowers with high financing search costs.[27] At least one reason we do not observe any evidence of ex post differences in default is likely the monthly payment smoothing behavior of borrowers documented in Table 5. Given that borrowers adjust loan sizes and expenditure amounts to keep monthly payments roughly constant, there is less scope for differential defaulting by borrowers with higher markups.

A novel feature of our data set 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 data set pull credit scores on borrowers after loan origination roughly once a year. Ex post credit scores allow us to calculate changes in credit scores over time, capturing broad changes in borrower distress and financial responsibility incorporating other credit products beyond the given auto loan in question. Any unobserved heterogeneity driving selection into loan take-up should affect 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

[26] The sample size varies across columns because of inconsistent data coverage across lenders.

[27] Figure A8 in the Internet Appendix presents these results graphically.

originate such loans. This omnibus test is particularly valuable given that the effect of markups on equilibrium monthly payments is close to zero, possibly resulting in negligible effects on the performance of the auto loan itself. Using the subsample of institutions that collect updated FICO scores after origination, we use the percentage change between credit scores at origination and the most recently observable credit score as the dependent variable in our RD framework. Results presented in column 4 of Table 7 show no meaningful differences (0.2 percentage points) in credit score changes for borrowers around the threshold, with small and statistically insignificant differences for borrowers with high financing search costs.

While adverse selection motivates many features of retail car-loan markets (Adams, Einav, and Levin 2009; Jansen et al. 2022), 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 creditworthiness is not the only alternative explanation for our observed results around thresholds, and we are implicitly assuming that variation in take-up is driven by idiosyncratic factors uncorrelated with car and loan preferences that we interpret as heterogeneous individual search costs. For example, FICO thresholds could promote the steering of financially unsophisticated borrowers into higher-rate loans. However, any explanation such as borrower naïveté would also have to fail to result in differences across thresholds at loan application or origination, significantly higher prices paid for the same make-model-year, differential ex post default rates, or differences in ex post credit scores. Given this set of outcomes showing borrowers offered high and low interest rate markups to be otherwise quite similar, we find it unlikely our results are driven by a missing factor that drives the differential response of above- and below-discontinuity borrowers with high and low search costs.

6. Conclusion

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

Considering observationally similar borrowers, the average borrower in our data would pay 81–113 bps less if they applied to one more loan. In this setting where the gains to search are high, we show that borrowers' acceptance of dominated loan terms is related to measures of the cost of searching for retail credit. Because discontinuous pricing schedules vary across lenders within the same commuting zone, borrowers on the expensive side of FICO discontinuities in loan pricing at one institution would be more likely to find favorable pricing at a different institution. Absent search frictions, borrowers are unlikely to accept seemingly dominated loan terms. Moreover, under the assumption that borrowers on either side of pricing discontinuities are otherwise identical, we would expect them to ultimately find similar financing opportunities and purchase similar cars absent search frictions. Proxying for the costliness of loan shopping with the density of nearby lenders, we show that borrowers with higher search costs face weakly more dispersed prices, are more likely to accept quasi-randomly offered dominated loan terms, and apply for fewer loans. Robustness exercises show this result is distinct from a simple market-concentration explanation and holds using alternative identifying variation. Reinforcing our interpretation, we show that to bias our estimates, any omitted variables correlated with our search-cost proxy would have to also have differential effects for high- and low-markup borrowers and be uncorrelated with the characteristics of both borrower and their Census tracts.

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

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

References

Adams, P., S. Hunt, C. Palmer, and R. Zaliauskas. 2021. Testing the effectiveness of consumer financial disclosure: Experimental evidence from savings accounts. *Journal of Financial Economics* 141:122–47.

Adams, W., L. Einav, and J. Levin. 2009. Liquidity constraints and imperfect information in subprime lending. *American Economic Review* 99:49–84.

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel. 2017. Do banks pass through credit expansions to consumers who want to borrow? *Quarterly Journal of Economics* 133:129–90.

Agarwal, S., J. Grigsby, A. Hortaçu, G. Matvos, A. Seru, and V. Yao. 2020. Searching for approval. Working Paper, National University of Singapore.

Al-Najjar, N. I., and M. M. Pai. 2014. Coarse decision making and overfitting. *Journal of Economic Theory* 150:467–86.

Alexandrov, A., and S. Koulayev. 2018. No shopping in the u.s. mortgage market: Direct and strategic effects of providing information. Working Paper, Consumer Financial Protection Bureau.

Allen, J., R. Clark, and J.-F. Houde. 2014a. The effect of mergers in search markets: Evidence from the canadian mortgage industry. *American Economic Review* 104:3365–96.

———. 2014b. Price dispersion in mortgage markets. *Journal of Industrial Economics* 62:377–416.

———. 2019. Search frictions and market power in negotiated-price markets. *Journal of Political Economy* 127:1550–98.

Aneja, A. P., and C. F. Avenancio-León. 2020. No credit for time served? Incarceration and credit-driven crime cycles. Working Paper, Berkeley Law.

Argyle, B., T. Nadauld, and C. Palmer. 2020. Monthly payment targeting and the demand for maturity. *Review of Financial Studies* 33:5416–62.

Argyle, B., T. Nadauld, C. Palmer, and R. Pratt. 2021. The capitalization of consumer financing into durable goods prices. *Journal of Finance* 76:169–210.

Attanasio, O. P., P. Koujianou Goldberg, and E. Kyriazidou. 2008. Credit constraints in the market for consumer durables: Evidence from micro data on car loans. *International Economic Review* 49:401–36.

Aydin, D. 2022. Consumption response to credit expansions: Evidence from experimental assignment of 45,307 credit lines. *American Economic Review* 112:1–40.

Bartlett, J. 2013. Consumers rely on car financing more than ever. *Consumer Reports*, September 6. https://www.yahoo.com/news/consumers-rely-car-financing-more-ever-203000128.html.

Benetton, M. 2021. Leverage regulation and market structure: A structural model of the UK mortgage market. *Journal of Finance* 76:2997–3053.

Benmelech, E., R. R. Meisenzahl, and R. Ramcharan. 2017. The real effects of liquidity during the financial crisis: Evidence from automobiles. *Quarterly Journal of Economics* 132:317–65.

Bertanha, M. 2020. Regression discontinuity design with many thresholds. *Journal of Econometrics* 218:216–41.

Bhutta, N., A. Fuster, and A. Hizmo. 2020. Paying too much? Price dispersion in the us mortgage market. Working Paper, Federal Reserve Board of Governors.

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

Burd, C., M. Burrows, and B. McKenzie. 2021. Travel time to work in the united states: 2019. Report, U.S. Census Bureau. https://www.census.gov/content/dam/Census/library/publications/2021/acs/acs-47.pdf.

Busse, M. R., N. Lacetera, D. G. Pope, J. Silva-Risso, and J. R. Sydnor. 2013. Estimating the effect of salience in wholesale and retail car markets. *American Economic Review* 103:575–79.

Busse, M. R., and J. M. Silva-Risso. 2010. ‘one discriminatory rent’ or ‘double jeopardy’: Multicomponent negotiation for new car purchases. *American Economic Review* 100:470–4.

Calonico, S., M. D. Cattaneo, and R. Titiunik. 2014. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica* 82:2295–326.

Campbell, J. Y. 2006. Household finance. *Journal of Finance* 61:1553–604.

Chetty, R., J. N. Friedman, N. Hendren, M. R. Jones, and S. R. Porter. 2018. The opportunity atlas: Mapping the childhood roots of social mobility. Working Paper, Harvard University.

Chiappori, P.-A., and B. Salanié. 2000. Testing for asymmetric information in insurance markets. *Journal of Political Economy* 108:56–78.

DeGroot, M. 1970. *Optimal statistical decisions*. New York: McGraw-Hill.

Degryse, H., and S. Ongena. 2005. Distance, lending relationships, and competition. *Journal of Finance* 60:231–66.

Delavande, A., and B. Zafar. 2019. University choice: The role of expected earnings, nonpecuniary outcomes, and financial constraints. *Journal of Political Economy* 127:2343–93.

Diamond, P. A. 1971. A model of price adjustment. *Journal of Economic Theory* 3:156–68.

Einav, L., M. Jenkins, and J. Levin. 2012. Contract pricing in consumer credit markets. *Econometrica* 80:1387–432.

———. 2013. The impact of credit scoring on consumer lending. *RAND Journal of Economics* 44:249–74.

Ellison, G., and S. F. Ellison. 2009. Search, obfuscation, and price elasticities on the internet. *Econometrica* 77:427–52.

Federal Highway Administration. 2017. National household travel survey. Washington, DC.

Federal Reserve Bank of New York. 2021. Quarterly report on household debt and credit.

Finkelstein, A., and J. Poterba. 2004. Adverse selection in insurance markets: Policyholder evidence from the u.k. annuity market. *Journal of Political Economy* 112:183–208.

Granja, J., C. Leuz, and R. Rajan. 2022. Going the extra mile: Distant lending and credit cycles. *Journal of Finance* 77:1259–324.

Grembi, V., T. Nannicini, and U. Troiano. 2016. Do fiscal rules matter? *American Economic Journal: Applied Economics* 8:1–30.

Gross, D. B., and N. S. Souleles. 2002. Do liquidity constraints and interest rates matter for consumer behavior? evidence from credit card data. *Quarterly Journal of Economics* 117:149–85.

Grunewald, A., J. Lanning, D. Low, and T. Salz. 2020. Auto dealer loan intermediation: Consumer behavior and competitive effects. Working Paper, Frankfurt School of Finance and Management.

Hong, H., and M. Shum. 2006. Using price distributions to estimate search costs. *RAND Journal of Economics* 37:257–75.

Hutto, G. W., and J. Lederman. 2003. *Handbook of mortgage lending*. Washington, DC: Mortgage Bankers Association of America.

Jansen, M., F. Nagel, C. Yannelis, and A. L. Zhang. 2022. Data and welfare in credit markets. Working Paper, University of Utah.

Keys, B. J., T. Mukherjee, A. Seru, and V. Vig. 2009. Financial regulation and securitization: Evidence from subprime loans. *Journal of Monetary Economics* 56:700–20.

———. 2010. Did securitization lead to lax screening? evidence from subprime loans. *Quarterly Journal of Economics* 125:307–62.

Laufer, S., and A. Paciorek. 2016. The effects of mortgage credit availability: Evidence from minimum credit score lending rules. Working Paper, Federal Reserve Board of Governors.

Liberman, A., D. Paravisini, and V. Pathania. 2021. High-cost debt and perceived creditworthiness: Evidence from the UK. *Journal of Financial Economics* 142:719–36.

Lippman, S. A., and J. J. McCall. 1976. Job search in a dynamic economy. *Journal of Economic Theory* 12:365–90.

Livshits, I., J. C. MacGee, and M. Tertilt. 2016. Democratization of credit and the rise in consumer bankruptcies. *Review of Economic Studies* 83:1673–710.

Matějka, F., and A. McKay. 2015. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review* 105:272–98.

McCrary, J. 2008. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics* 142:698–714.

Melzer, B. T. 2011. The real costs of credit access: Evidence from the payday lending market. *Quarterly Journal of Economics* 126:517–55.

Melzer, B. T., and A. Schroeder. 2017. Loan contracting in the presence of usury limits: Evidence from automobile lending. Working Paper, Consumer Financial Protection Bureau.

Moraga-González, J. L., Z. Sándor, and M. R. Wildenbeest. 2017. Prices and heterogeneous search costs. *RAND Journal of Economics* 48:125–46.

———. 2021. Consumer search and prices in the automobile market. Working Paper, VU University Amsterdam.

Nguyen, H.-L. Q. 2019. Are credit markets still local? evidence from bank branch closings. *American Economic Journal: Applied Economics* 11:1–32.

Reinganum, J. F. 1979. A simple model of equilibrium price dispersion. *Journal of Political Economy* 87:851–8.

Stahl, D. O. 1989. Oligopolistic pricing with sequential consumer search. *American Economic Review* 79:700–12.

Stango, V., and J. Zinman. 2015. Borrowing high versus borrowing higher: Price dispersion and shopping behavior in the us credit card market. *Review of Financial Studies* 29:979–1006.

Stiglitz, J. E., and A. Weiss. 1981. Credit rationing in markets with imperfect information. *American Economic Review* 71:393–410.

Woodward, S. E., and R. E. Hall. 2012. Diagnosing consumer confusion and sub-optimal shopping effort: Theory and mortgage-market evidence. *American Economic Review* 102:3249–76.

Zabritski, M. 2021. State of the automotive finance market. Working Paper, Experian.

Zaki, M. 2016. Access to short-term credit and consumption smoothing within the paycycle. Working Paper, University of Maryland.

Zeldes, S. P. 1989. Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy* 97:305–46.

Zhou, J. 2014. Multiproduct search and the joint search effect. *American Economic Review* 104:2918–39.

Zinman, J. 2014. Consumer credit: Too much or too little (or just right)? *Journal of Legal Studies* 43:S209–37.