

NBER WORKING PAPER SERIES

THE CAPITALIZATION OF CONSUMER FINANCING INTO DURABLE GOODS
PRICES

Bronson Argyle
Taylor D. Nadauld
Christopher Palmer
Ryan D. Pratt

Working Paper 24699
<http://www.nber.org/papers/w24699>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2018

We thank the Editor (Amit Seru), the Associate Editor, and two anonymous referees. We also thank our discussants Tom Chang, Marco Di Maggio, Paul Goldsmith-Pinkham, Christopher Hansman, and John Mondragon, and workshop and conference participants at BYU, Cornell, Federal Reserve Board, Minnesota, MIT, Notre Dame, NYU, Philadelphia Fed, Princeton, Stanford SITE, UT Austin, University of Washington, Washington University in St. Louis, FIRS, SFS Cavalcade, and Natalie Bachas, E Benmelech, Shai Bernstein, Giovanni Favara, Vincent Glode, Brad Larsen, Greg Leiserson, Brigitte Madrian, Jonathan Parker, Antoinette Schoar, David Sraer, Jeremy Stein, Stijn Van Nieuwerburgh, and Emil Verner for helpful comments. We appreciate the research assistance of Lei Ma and Alex Tuft. The data were provided by an anonymous information-technology firm. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w24699.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Bronson Argyle, Taylor D. Nadauld, Christopher Palmer, and Ryan D. Pratt. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Capitalization of Consumer Financing into Durable Goods Prices
Bronson Argyle, Taylor D. Nadauld, Christopher Palmer, and Ryan D. Pratt
NBER Working Paper No. 24699
June 2018, Revised in September 2020
JEL No. E31,E43,E51,G21,H22,L11,L62

ABSTRACT

Using loan-level data on millions of used-car transactions across hundreds of lenders, we study the consumer response to exogenous variation in credit terms. Borrowers offered shorter maturity decrease expenditures enough to offset 60-90% of the monthly payment increase. Most of this is driven by shifting toward lower quality cars, but affected borrowers are able to offset 20-30% of a monthly payment shock by negotiating lower prices for equivalent cars. Our results suggest that durable goods prices adjust to reflect credit terms even at the individual level, with one year of additional loan maturity increasing a given car's price by 2.8%.

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

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

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

Ryan D. Pratt
Department of Finance
Brigham Young University
Provo, Utah 84602
ryan.pratt@byu.edu

The relationship between credit and asset prices has been an important area of study in post-crisis finance. Recent work, focused largely on the housing market, has established a causal link running from aggregate credit supply to average prices. However, studies of aggregate credit supply necessarily obscure differences in access to credit in the cross-section of borrowers. In this paper, we exploit disaggregated data on auto prices and loans to investigate how durable goods markets capitalize credit terms at the individual borrower level. We document significant heterogeneity in transaction prices across borrowers driven by differences in their access to credit, isolating credit terms as a key channel through which credit impacts prices. Our results indicate that prices respond to credit terms at a more granular level than has been documented previously.

In a world with credit-constrained borrowers, we would expect a decrease in the aggregate supply of credit to result in a decrease in aggregate demand for the final good, lowering both the good's unit price and quantity purchased. This conceptually corresponds to the standard interpretation of papers like Favara and Imbs (2015) and Di Maggio and Kermani (2017), who quantify the effect of aggregate credit-supply shocks on average house prices. Zooming in to the individual borrower in the auto market, a classical framework of supply and demand suggests that affected borrowers would respond to a decrease in their available credit by substituting toward lower quality cars, an effect which we find in the data. If enough borrowers are affected to influence aggregate demand, then prices would fall for everyone, but two people buying the same car would pay the same price whether they had selected that car because of a credit shock or not. In contrast, we find borrowers with decreased access to credit pay lower prices holding quality fixed.

Methodologically, we isolate plausibly exogenous changes in the monthly cost of debt service arising from lender maturity policies in the auto-loan market.¹ While average maturity decreases smoothly with car age, the policies in place at any given lender often feature discontinuous drops in offered maturity at particular car ages. For example, a given lender may offer a 72-month loan to a borrower to purchase a car up to three years old but may only be willing to offer a 60-month loan to the same borrower to purchase a four-year-old car. We combine this fact with the observation that car age is measured based on manufacture year such that cars are effectively treated as if they fully

¹As we discuss below, a large body of evidence documents the importance borrowers attach to payment size, for which maturity has first-order importance.

age by one year on January 1.² Importantly, we find step-function maturity schedules with breaks at different car ages for different lenders in the data. Thus, on January 1 cars of a given manufacture year may age across a discontinuity in a particular lender’s maturity schedule while experiencing constant offered maturity at another lender. From an empirical-design standpoint, such maturity policies that feature discontinuities in allowable maturity around the new year map neatly into “pre” and “post” event dates. Similarly, variation across lenders’ maturity policies—even for cars of the same year, make, model, and trim—defines treatment and control samples. Taken together, the pre- and post-event dates and treatment and control samples allow for causal inference in a difference-in-differences framework. We emphasize that the variation we exploit in offered maturity arises from interacting the passage of time with predetermined lender maturity policies, as opposed to any potentially endogenous decision taken by a lender to change its existing maturity schedule.

Armed with plausibly exogenous variation in offered maturity, we first estimate its effect on consumer expenditure. We find that affected borrowers spend about 1.1% less on their car purchase per month of decreased maturity. For a borrower receiving average loan terms, the decreased expenditure is sufficient to offset roughly 90% of the increase in monthly payment that would otherwise result from shorter maturity. Of course, this decrease in expenditure may be driven by consumers substituting toward lower quality vehicles. To account for the role of substitution in explaining our results, we analyze car prices holding quality fixed with a battery of manufacture year-make-model-trim (“*YMMT*”) \times purchase-month fixed effects (e.g., 2014 Honda Accord LXs purchased in February of 2017). Our results suggest that treated and untreated borrowers pay materially different prices for observationally identical cars purchased at the same point in time, controlling for any time-invariant differences across geographic region. We find that one month of exogenously lower maturity is associated with treated borrowers paying 0.3% lower prices. As offered maturity most frequently changes by 12 months, our estimates imply that the modal reduction in offered maturity reduces car prices in our sample by 3.6%, or roughly \$720 on a \$20,000 car.³ These estimates imply that affected borrowers are able to offset around 25% of the potential increase in

²Englmaier et al. (2017) present evidence that in the European used-car market, cars are treated as aging relative to the year they were first registered.

³To assess the magnitude of this estimate, note that estimates of gross margins on used cars are 5–20% (Gavazza et al. (2014), Huang et al. (2015), and Larsen (2018)).

monthly payment through lower negotiated prices on equivalent vehicles.

We focus on maturity primarily because we expect borrowers to significantly value long maturities. Previous work has shown the importance of monthly payments for household decision-making (see, for example, Argyle, Nadauld, and Palmer (2020)), and the modal variation in maturity (12 months) has as large an impact on the monthly payment of an average vehicle as would a seven percentage-point increase in interest rates. Still, we are interested in the impact of credit terms broadly, and we exploit the fact that lenders' policies may feature discontinuities in interest rates that coincide with maturity discontinuities to both refine our maturity estimates and to separately estimate the effect of interest-rate variation. Using a two-stage least-squares procedure in which we simultaneously instrument for both maturity and interest rates, we find that maturity accounts for 70–80% of the impact of our credit-supply shocks. Accounting for contemporaneous shifts in interest rates, borrowers spend 0.74% less per month of shorter maturity, offsetting 60% of the increase in monthly payment. Controlling for substitution, borrowers who receive 12 months shorter maturity pay about 2.8% less when purchasing a car of a given *YMMT* in a given month, corresponding to an elasticity of price with respect to monthly payment size of -0.19. This implies a lower bound on the consumer discount rates needed to rationalize these results of around 11.6%.⁴

Meanwhile, we estimate that, for a one percentage-point increase in interest rates, borrowers spend 1.95% less on their car purchase, enough to offset around 70% of the increase in monthly payment. The fact that borrowers respond similarly to a change in monthly payment whether that change came from interest rate or maturity variation is consistent with monthly payments playing a central role in their decision-making. Finally, for the same *YMMT* in the same month, borrowers with one percentage-point higher interest rate would pay a 0.9% lower price.

We conduct several robustness exercises to address the most plausible alternative explanations for our pricing results, including the possibility that *YMMT* by month fixed effects do not adequately capture the variation in vehicle quality. A finance-induced shock to demand could lead treated borrowers to shift toward unobservably lower quality, less expensive cars. For example, a borrower could opt for a car with the same *YMMT* that is in worse condition or has higher mileage and thus a lower price. We address this possibility in several ways in section B; in one test, we examine

⁴By contrast, Busse, Knittel, and Zettelmeyer (2013) estimate that in the context of fuel economy, the average buyer acts as if having a discount rate roughly equal to her loan interest rate.

a subset of our data comprised of repeat sales—transactions involving the exact same vehicle. If treated-borrower purchases occur at lower values because of differences in unobserved quality, these differences should persist in the second transaction. In a subsample of roughly 8,700 vehicles for which we observe a subsequent transaction, we find no difference in the *second* transaction price of cars originally treated with low maturity in the *first* transaction (relative to other cars of the same *YMMT* sold in the same month).⁵ Instead, treated cars’ prices appear to rebound when sold at a later time, inconsistent with an interpretation in which treated borrowers buy lower quality cars. We further address concerns about unobserved heterogeneity in car quality through an analysis of pricing effects within subsamples of older and newer cars. Unobserved variation in car quality increases with car age, yet our estimates are similar in samples of relatively young and old cars. Using a subsample of cars for which we have mileage information, we also show that treated cars do not have higher mileage at the time of sale. Finally, we show that borrower characteristics do not change in economically meaningful ways with our treatment and that our results are robust to the Oster (2017) adjustment for potential unobserved correlated heterogeneity.

Why would one consumer with different financing terms pay a different price than another consumer purchasing the same car at the same point in time? Our inclusion of commuting zone fixed effects rules out the effect being driven by variation in average prices of cars across geographic markets, while lender fixed effects cast doubt on clientele selection effects across lenders. Instead, we argue that the lower private values caused by poorer financing terms may influence the search or bargaining games inherent in the car market. Car buyers have access to an array of sellers from whom they might purchase a car. However, because it is costly to search across sellers, a buyer visiting a particular seller stands to gain from consummating the transaction with that seller. Similarly, the seller, not knowing when an alternative buyer might show up, minimizes his cost of carry by selling to the current customer. The transaction price is the result of bargaining over the corresponding bilateral surplus. In this setting, affected borrowers have two potential margins of adjustment: i) they can absorb their finance-induced demand shock by searching more for a better deal, or ii) they can offload some of the cost of poor financing terms onto the seller, effectively leveraging their lower reservation price to negotiate a better transaction price.

⁵We also test for endogenous selection into observable resale and find none.

To distinguish between these bargaining and search explanations for our capitalization results, it would be useful to have information on sellers. If borrowers respond by increasing their search intensity, they are exploiting variation in prices *across* sellers, while borrowers who negotiate better prices take advantage of price dispersion *within* seller. Not observing anything about sellers, we examine the subset of lenders for which we have loan application data. If buyers respond to poor financing terms by searching more intensively for a vehicle, we might expect more time to elapse between loan application and origination. However, we find no evidence that treated borrowers spend more time shopping for a vehicle, suggesting that the lower transaction prices are likely the result of more successful negotiations by treated borrowers.

While our dataset is focused on the auto market, similar search and bargaining dynamics are at play in the markets for many big-ticket items in which consumers transact infrequently (real estate, machines, furniture, etc.). Based on our results, it seems reasonable to expect individual financing terms to influence the bargained outcomes across a broad range of markets.⁶

The balance of the paper is organized as follows. Section I discusses the literature and conceptual motivation. Section II discusses the dataset. Section III outlines the empirical strategy that we employ, and section IV discusses results and mechanism. Section V concludes.

I. Contribution to the Literature and Conceptual Motivation

A. Related Literature

The empirical literature studying the causal link between credit and prices has been concentrated in the housing market. See, for example, Mian and Sufi (2009), Glaeser, Gottlieb, and Gyourko (2012), Adelino, Schoar, and Severino (2013), Favara and Imbs (2015), Landvoigt, Piazzesi, and Schneider (2015), Zevelev (2016), Di Maggio and Kermani (2017), Verner and Gyöngyösi (2018), and Davis et al. (2017), none of which has studied loan maturity.⁷ Favara and Imbs (2015) exploit

⁶In section I, we also discuss related literature on the strategic use of debt in bargaining games.

⁷A recent empirical macro literature also studies the causes and consequences of credit shocks for house prices (Jorda, Schularick, and Taylor (2015)), price-discrimination markups (Cornia, Gerardi, and Shapiro (2011)), business cycles (Borio and Lowe (2002), Mian, Sufi, and Verner (2017), and Krishnamurthy and Muir (2017)), and stock markets (Hansman et al. (2018)). See Mian and Sufi (2018) for a survey of recent work on credit-driven business cycles.

state-level exposure to bank branching deregulation as an instrument for credit-supply shocks to demonstrate a causal link between credit expansion and house prices. Di Maggio and Kermani (2017) use state-level variation in anti-predatory lending laws' impact to trace a boom and bust in house prices resulting from credit-supply shocks. These papers feature geographic variation in credit supply shocks that may affect local credit markets in complex ways as opposed to quantifiable, individual-level payment shocks.⁸ In addition, they do not examine the cross-sectional implications of credit capitalization on individual borrowers. Closer in spirit to our work is Adelino, Schoar, and Severino's (2013) analysis of conforming loan limits (CLL) and housing prices. While an increase in the CLL impacts house prices in the cross-section (with prices near the CLL more affected), the differentiated nature of real estate does not permit disentangling whether two borrowers with different access to financing terms would pay different prices for the same house.

We differ from previous work along a second dimension. The set of frictions that are the source of credit-supply shocks in the literature are often macro in nature. Aside from examples cited above, these include credit shocks driven by regulation (Rice and Strahan (2010)), financial innovation (Mian and Sufi (2009), and Nadauld and Sherlund (2013)), government credit subsidies (Lucca, Nadauld, and Shen (2018)), and funding market failures (Benmelech, Meisenzahl, and Ramcharan (2017)). In each of these papers, macroeconomic fluctuations influence the aggregate supply of credit. In contrast, our setting demonstrates the existence of a different class of relevant credit-market frictions. Our results underscore that firm-level institutional idiosyncrasies play an important role in determining the borrower-level supply of credit and that such policies have material effects on consumer outcomes. Moreover, we are the first paper to study the effect of loan maturity on prices.⁹

Our paper also contributes to a literature that explores the importance of finance in the market for used cars. Hortaçsu et al. (2013) document that dealers who purchase used cars at auto auctions incorporate expectations of manufacturers' financial distress into prices, where pricing discounts reflect financing-related uncertainty around manufacturers' ability to deliver on future obligations

⁸For example, reduced-form credit shocks resulting from local regulatory changes could affect credit terms, lending standards, expectations, local aggregate demand, and incomes.

⁹Hertzberg et al. (2018) argue a shared sentiment that the role of maturity has been understudied relative to interest rates in this literature.

such as warranties. Adams, Einav, and Levin (2009) explore the sensitivity of car demand to financing conditions and show that out-of-pocket liquidity plays a large role in purchasing decisions. In a related paper, Einav, Jenkins, and Levin (2012) document how down payment requirements effectively mitigate adverse selection by screening risky borrowers. Our results shed further light on the critical role that financing conditions, and more specifically loan maturity, play in the market for used cars.

Our results are related to the corporate finance literature highlighting the strategic role that debt plays in determining bargaining outcomes.¹⁰ Israel (1991) and Muller and Panunzi (2004) argue that debt can be used to influence bargaining outcomes in the market for corporate control. Spiegel and Spulber (1994) show that debt burdens influence the prices charged by regulated firms such as utilities. Hennessy and Livdan (2009) demonstrate the strategic role of debt in the allocation of surplus between firms and their suppliers, and Matsa (2010) documents the influence of debt on the outcomes of negotiations between firms and organized labor. In each case, debt limits the financial flexibility of firms, which strengthens a firm's bargaining position. We show a similar dynamic in a retail setting. Borrowers who are offered shorter maximum maturity have limited financial flexibility and appear to be able to use this to influence the outcome of the bargaining game with car sellers.¹¹

Within the vast public-finance literature on economic incidence, several papers have looked specifically at the market for new cars and the incidence of taxes and manufacturer subsidies. Although these papers do not examine the incidence of financing shocks per se or the distributional implications of individual-level changes in access to credit, they document capitalization effects of cost shocks into vehicle prices. For example, Busse et al. (2006) examine the effects of manufacturer cash rebates for new cars, documenting that incidence depends on statutory incidence, that is, whether the rebate is issued to buyers or sellers. Consistent with our findings that prices capitalize changes in credit terms, they find that prices rise by 10-30% of the amount of a customer rebate. Sallee (2011) finds that new Toyota Prius prices did not capitalize hybrid vehicle tax incentives at all, attributing the lack of pass-through to Toyota's concerns about future demand given the dynamics of buyer price beliefs. Busse et al. (2012) find that resale prices capitalize exposure

¹⁰See also related work in psychology, for example, Lee and Ames (2017).

¹¹See also Diamond's (1991) model of debt-contract maturity, in particular theoretical results on borrower preference for longer maturity and lender reluctance for the same in the presence of private information.

to gasoline taxes.¹² We complement this literature by studying the transmission of credit-supply shocks with cross-sectional identification, further emphasizing that disaggregate credit shocks can have disaggregate price effects.

Finally, we note the contribution of this paper relative to Argyle, Nadauld, and Palmer (2017) and Argyle, Nadauld, and Palmer (2019). Though each paper uses the same broad set of auto data, the three papers differ in economic focus. Argyle, Nadauld, and Palmer (2017) documents the relevancy of physical search frictions in consumer loan markets, showing that consumers originate loans that do not represent the plausibly lowest available rate to them and that the number of credit providers within a reasonable commute of a potential borrower’s home leads to significant variation in the take up of offered loans. Argyle, Nadauld, and Palmer (2019) documents that even relatively unconstrained borrowers make loan decisions based on monthly payment amounts. In particular, many consumers anchor around hundred-dollar-multiple amounts (e.g., \$200, \$300, \$400, etc.), supporting the conjecture that cognitive frictions like mental accounting influence how consumers make high-stakes debt decisions.

The present paper builds on the other papers in this research agenda by documenting that monthly payment shocks are capitalized into asset purchase prices in a way that, depending on borrower discount rates, offsets much of the value of easier credit. Ultimately, while costly search and a focus on monthly payments are necessary ingredients into this finding, they do not imply the cross-sectional pricing effects of individual credit shocks that are the contribution of this paper. Lastly, our empirical strategy also differs from these related papers; we exploit vehicle-level discontinuities in offered maturity across the car-age space while the companion papers exploit borrower-level discontinuities across the FICO spectrum.

B. Conceptual Motivation

The central question of the paper is how individual borrowers respond to differences in credit terms. To guide our empirical design and the interpretation of our tests, it is useful to briefly sketch the key economic characteristics of our empirical setting, as these shape the way that financing impacts the durable goods market. Two primary frictions are important in the auto-loan market:

¹²Other relevant incidence papers include Goolsbee (1998), who shows that investment tax credits increase capital goods prices, and MaCurdy (2015), who shows that consumer prices increase following minimum wage increases.

financial constraints and search costs. Consider a financially constrained borrower with utility over transportation services where “quantities” of transportation services correspond to the quality of available vehicles. Because the borrower is constrained, the terms of credit influence the borrower’s demand for transportation services. In particular, a decrease in offered maturity will decrease demand, driving down the private value that the borrower attaches to any particular vehicle. Absent additional frictions, affected borrowers would respond by substituting to lower quality vehicles.

However, search costs open another avenue for affected borrowers. Costly search confers local market power on sellers, supporting equilibria with price dispersion across sellers; indeed, one of the motivating examples in Stigler’s (1961) seminal paper on search was the variation in posted prices across Chevrolet dealers in Chicago. Costly search raises the possibility that affected borrowers might pay less for equivalent vehicles by searching longer for a seller willing to sell at a lower price.

Similarly, search costs also impact the supply side of the auto market as sellers are uncertain when willing buyers will arrive. Because auto sellers pay a significant cost of carry on their inventory, they stand to gain from consummating a transaction with a current customer rather than waiting for a new one to arrive. When the two parties have met, both the buyer’s search costs and the seller’s carrying costs are minimized if they can complete a transaction. Given this surplus, sellers cannot commit to posted prices, and bargaining ensues with the negotiated price dividing the bilateral surplus. Thus, costly search gives rise to bargaining, introducing the prospect that affected borrowers might be able to exploit their lower private value to negotiate a lower price.

To summarize, in a world where financially constrained borrowers face search costs, borrowers have three margins of adjustment to absorb a financing shock: i) they can substitute to lower quality goods, ii) they can pay higher search costs to find a lower price on a given good, or iii) they can offload some of the cost onto the seller by negotiating a lower price. We aim to quantify the extent to which borrowers substitute vs. pay lower prices for equivalent goods and to shed light on whether lower prices arise from increased search or more successful bargaining.

It is worth noting the role that search costs play in the loan market as well. Argyle, Nadauld, and Palmer (2017) provide evidence of significant search costs in the loan market, documenting that borrowers accept loan terms that are inferior to others that would be available to them. Given the lender-specific nature of our financing shocks, our empirical design builds on this fact. Without search costs in the loan market, we would be unable to construct control and treatment groups

defined by their differential loan terms.

Financial constraints and search costs are not unique to the auto market. Similar frictions are at play in a wide variety of markets, notably real estate. While the differentiated nature of housing makes it difficult to hold quality constant when measuring prices, it seems reasonable to expect that differences in access to credit may drive individual-level variation in home price transactions as well.

II. Data

Our data on auto loan originations come from a technology firm that provides data warehousing and analytics services to retail-oriented lending institutions nationwide. We begin with a dataset consisting of over four million auto loans originated by 372 unique lenders covering all 50 states. The data include only loans originated directly through the lending institution, as opposed to so-called indirect loan applications processed through auto dealerships. Direct loan applications occur primarily in one of two ways. First, a borrower may identify the exact car she would like to purchase and then apply for a loan. In this circumstance, lenders evaluate the collateral and offer loan terms specific to the collateral. Alternatively, a borrower may apply for an auto loan without having identified a specific car she would like to buy. In this circumstance, lenders evaluate a potential borrower based on her credit characteristics. An approved application then specifies rough financing terms, conditional on a bundle of collateral characteristics, with loan terms finalized after the borrower selects a specific car. In either case, final negotiations for the car purchase typically occur after the borrower learns about the financing terms.

Our sample includes loans originated between 2005 and 2017, though over 80% were originated between 2011 and 2017. The growth in originations over time is mostly driven by growth in our data provider’s client base, though it also partly reflects increased reporting of loan originations over time within lender as our data provider’s products became more integral to the lenders’ businesses. Moreover, aggregate auto loan originations have increased substantially over our sample period, with outstanding auto debt in the U.S. increasing 56% between 2010 and 2017. Similar data are used in Argyle, Nadauld, and Palmer (2017, 2019).

The dataset, anonymized of any personally identifiable information, includes loan contract fea-

tures such as the purchase price, loan amount, maturity, interest rate, and origination date. We also have information on the underlying collateral, including the VIN number in most cases, which allows us to extract the manufacture year, make, model, and trim (*YMMT*) of the vehicle. Borrower information includes FICO scores and self-reported debt-to-income (DTI) ratios.

We begin with a sample of 4,192,502 loans to detect maturity policies in each lender \times car age \times month cell, as described in detail in the following section. After inferring lender maturity policies, we drop loans that were not originated during a stable policy regime, eliminating roughly two thirds of the observations (mostly consisting of those lender \times car age \times month cells with the fewest observations). We lose an additional quarter of the remaining observations that are missing vehicle trim information that we use in *YMMT* fixed effects to hold observable vehicle quality fixed.

These two restrictions leave us with a final dataset of 972,621 loans originated by 112 unique lenders. Table I reports summary statistics broken out by treatment and control groups. The average borrower in our sample has a credit score at loan origination of 714, slightly above the national average of 700. As most of our lenders are credit unions, our data does not have strong coverage of subprime borrowers. Average back-end DTI ratios, which measure the monthly fraction of total debt-service payments to income, are around 35%. Examining collateral and loan characteristics, most of the car purchases we study involve used cars; the average car in our sample is 3.9 years old and sold for \$20,341. The average loan-to-value is 90.7%, with average maturity of 61.3 months and interest rate of 4.1%. We discuss the comparison of treatment and control groups after defining them in section A.

[Insert
Table I
about
here.]

III. Empirical Strategy

We are interested in measuring consumer responses to variation in credit terms. To what extent do borrowers absorb shorter maturity by substituting toward different goods? Can borrowers offset some of the cost of shorter maturity loans by negotiating better prices for the *same* durable goods relative to buyers who were offered longer maturity loans at the same time? In answering these questions, we face immediate identification challenges, as the relationship between credit and prices may be driven by a variety of economic mechanisms, including simple reverse causality. For example, lenders—willing to offer longer maturities for higher quality collateral—may use price as a proxy

for unobservable collateral quality. In this case, buyers who pay higher prices, perhaps because they have higher private value for the good, would also receive longer maturities. Alternatively, any aspects of quality that are observable to the lender but not to the econometrician may jointly drive both higher prices and longer maturities.

To overcome these empirical challenges, the ideal experiment would feature randomly assigned loan maturities. We do our best to approximate this by exploiting maturity rules used by lenders based on the age of cars. Based on conversations with lenders, the maximum maturity borrowers are offered on an auto loan is frequently determined as a function of car age, a practice that is motivated by lenders' risk management concerns. Longer maturity increases the likelihood that the loan balance exceeds the collateral value during the life of a loan, exposing lenders to losses in the case of default. For older cars with shorter remaining expected life, lenders offer shorter maturity on average (Figure 1). However, instead of smoothly mapping car age into offered maturity, many lender policies feature discrete drops in maximum offered maturity at particular car ages, as illustrated in Figure 2. This leads to a discontinuous drop in maturity offered for a given car as it ages across a break in a given institution's maturity schedule. To the extent that all cars of a given manufacture year are considered the same age, these discontinuities should occur as the calendar moves from December to January, when all cars turn one year older. To the best of our knowledge, there are no industry standard rules mapping car age to loan maturity; indeed, we find variation across lenders in the casuch as desiredr ages at which their maturity schedules feature discrete drops in offered maturity. At any point in time, we observe treated buyers (those borrowing from an institution with a discrete drop in maturities in January for a car of the age being purchased) and untreated buyers (those borrowing from an institution without such a discrete drop for the car age being purchased), even for cars of the same *YMMT*. We use this feature of the data to construct a difference-in-differences quasi-experiment, comparing the change in prices paid before and after January 1 for borrowers treated with an exogenous maturity shock to the corresponding price change for untreated borrowers. To be clear, our variation does *not* arise from some lenders changing their policies in January; rather, predetermined maturity schedules interact with the passage of time to treat individual cars with persistent maturity shocks beginning in January.

[Insert
Figure
1 about
here.]

[Insert
Figure
2 about
here.]

In addition to randomly assigned maturities, the ideal experiment would also hold fixed the quality of the goods purchased by treated and untreated borrowers. Absent this, it would be

impossible to disentangle the extent to which borrowers given exogenously shorter maturity respond by spending less on lower quality goods or negotiating better prices on the same goods. One of the strengths of our dataset is that we can hold fixed the quality of the goods to a significant extent by controlling for *YMMT* fixed effects interacted with the month of sale. Thus, the spirit of our tests is to compare the prices of two cars of the same *YMMT* being purchased in the same month, where one buyer receives exogenously different maturity than the other. While these fixed effects soak up a large majority of the variation in car quality, as we discuss the interpretation of our tests, we will be careful to address the possibility that our results are affected by any residual variation in quality within *YMMT* in a given month.

A. *Identifying Loan Maturity Policies by Car Age*

The first step in our analysis is to empirically identify age-based maturity policies for the 372 lenders in our sample so that we can assign vehicles experiencing a discrete change in maturity on January 1 to a treatment group and compare their prices to a control group of vehicles experiencing no change. A key challenge we face is that some lender maturity policies may be based on variables that are correlated with age, for example, mileage or even price. Misidentifying a price-based maturity policy as an age-based one would bake in a price-maturity relationship, a worst-case scenario for our identification. With this in mind, we design our algorithm to detect age-based maturity policies conservatively. The benefit of not contaminating our instrument with endogenous variation in maturity is worth the cost of potentially missing some actual age-based maturity policies.

Since each lender is likely to have its own maturity policies applying to cars of various ages, we look for rules—sustained periods of stable maturity—at the lender \times car-age level, where car age is defined as the calendar year of loan origination minus the year of manufacture. Here we face a second challenge: the amount of unexplained variation in maturity. Table II shows the distribution of maturity within lender \times car age \times month cells, relative to the maximum within each cell. On average, only 16.5% of borrowers receive the maximum observed maturity in their cell, while between 9% and 17% of borrowers receive maturities in each six-month band within 24 months of the maximum. Clearly, not everyone financing a car of a given age at a given lender in a given month receives the same maturity. Considering this variation in maturity, it is not surprising that the average maturity for a lender \times car age varies a lot from month to month—the standard

[Insert
Table II
about
here.]

deviation of the change in average maturity is 3.6 months. This makes it infeasible to detect stable maturity policies based on average maturity.

Instead, we focus on particular percentiles of the maturity distribution to detect maturity policies. Because maturities cluster significantly on multiples of 3, 6, or 12 months (71% are multiples of 12 months), individual percentiles of maturity have the benefit of being very stable, but they come with the cost of potentially being unaffected by actual policy discontinuities. Suppose, for example, that there is a policy discontinuity that results in the proportion of buyers offered the maximum maturity (say, 72 months) decreasing from 50% to 25% in January. At the 80th percentile of maturity, this would look like a stable regime of 72 months offered maturity. To increase the probability that we successfully detect the policy discontinuities, we pool the 70th, 80th, and 90th percentiles. The choice of these upper percentiles is driven by our focus on the maximum offered maturity; however, we show in Appendix A (and Appendix Tables AV–AVII) that we can reliably detect age-based maturity policies at lower percentiles and that our results are not sensitive to the set chosen.

To illustrate our method of categorizing lender policies, consider Figure 3 which plots the 80th percentile of maturity for three-year-old cars (panel A) and four-year-old cars (panel B) in each month for the largest lender in our sample. The x’s represent the individual monthly observations, and we categorize long periods of identical (or nearly identical) maturities as lender policies, as shown in the boxed areas. For each month, we examine the six months before and after; if at least five of the six months both before and after have the same maturity as the month in question, we consider the entire 13-month period to be part of a lender maximum maturity policy.¹³ For three-year-old cars shown in the figure, we identify two separate lender policies: a 66-month policy lasting from December 2007 through December 2012, followed by a 72-month policy lasting from

[Insert
Figure
3 about
here.]

¹³We require that the endpoints of the 13-month window do not deviate from the prevailing maturity policy. This prevents us from including the first month of a new policy in the time window of an old policy. For the purposes of detecting maturity policies, we round each maturity to the nearest three months such that, for example, maturities of 60 months (the most prevalent maturity at 27% of the data) are grouped with maturities of 61 months (2% of the data) and 59 months (0.5%). While a significant majority (84%) of the loans in our data already have maturities that are multiples of three months, some borrowers receive abnormal terms, perhaps motivated by demand-side factors such as a desired monthly payment level (Argyle, Nadauld, and Palmer (2019)).

January 2013 through July 2017 (the end of our sample). For four-year-old cars, we identify four separate lender policies over time, each shown in dashed red boxes.¹⁴

Having identified stable lender maturity policies, the next step is to follow cars as they age over time to observe any discontinuous transitions from one offered maturity level to another. Figure 4 combines the policies for three-year-old and four-year-old cars from Figure 3 into one plot. The policies for three-year-old cars are shown as dashed bars, while the policies for four-year-old cars are shown as solid bars. Note that a three-year-old car in any given December becomes a four-year-old car the following January, and therefore becomes subject to (potentially) different offered maturity. The vertical dotted lines correspond to the set of year-ends for which cars turning four years old would have experienced a discrete drop in maturities at this particular lender.

[Insert
Figure
4 about
here.]

Consider the example of a 2006 Honda Civic LX illustrated in the figure. In December 2009, as a three-year-old-car, this car would be subject to a 66-month maximum maturity policy. Yet the same car sold in January 2010 would be subject to the 60-month maturity policy in effect for four-year-old cars. We group cars experiencing this kind of discrete maturity shock in January into the set of treated observations. In contrast, consider the 2012 Honda Civic LX example shown in the upper right of the figure. By late 2015, this lender's policy allows 72-month loans for both three- and four-year-old cars. Thus, a four-year-old car sold in January 2016 would be subject to the same offered maturity as the three-year-old car sold in December 2015. We group all occurrences in which a given car experiences the same offered maturity from December to January into the control group. The treated subset of our final sample consists of any cars subject to a policy with a discrete drop in offered maturity at any of the 70th, 80th, or 90th percentiles from December to January, while control observations are those subject to a policy with continuous offered maturity at any of these percentiles.¹⁵

Our requirement that we observe a stable maturity percentile for at least a year to infer the

¹⁴The 80th percentile of loan maturities is more volatile in the early part of our sample because there are fewer loans during that time period. The coverage of our data provider improves over the early part of our sample, even within lender.

¹⁵Of course, it is possible for a lender to have a discrete drop in maturity policy at one percentile (say, the 70th) but to have a continuous policy at another (say, the 90th). We consider any such cases as treated since they display a maturity shock.

existence of an age-based maturity policy is important to distinguish age-based policies from other policies that indirectly link car age to maturity. To see this, suppose that a lender has a policy that ties maturity directly to the price rather than the age of a vehicle. As we move through the calendar year, the prices of, say, three-year-old cars would decrease, causing the maturity distribution to shift to the left from one month to the next. At some point during the year, the 80th percentile (or another percentile) of the maturity distribution would drop to a lower level. If we allow for age-based maturity policies that were in place for less than a year, we would risk misidentifying these price-based maturity policies as age-based ones that were updated during the middle of the year. Age-based policies differ from price- or mileage-based policies in that the only month boundary that an age-based policy respects is December to January. To ensure that the policies we identify are truly driven by the age of the car, we need to see stable maturity across the entire calendar year.

This observation suggests a way to assess how effective our approach is at isolating exogenous variation in the supply of maturity arising from age-based maturity policies. To the extent that we have correctly identified age-based maturity policies, any discontinuities in lenders' offered maturity should be concentrated in January, as cars nominally become one year older. In contrast, if we are mistakenly picking up price- or mileage-based policies, we would expect discontinuities to be uniformly distributed throughout the calendar year. Another possibility is that our categorization is contaminated with false positives arising from variation in the unexplained portion of maturity (perhaps periods with high demand for maturity at a particular lender followed by periods with lower demand). We would also expect the timing of such false positives to be roughly uniformly distributed.

Figure 5 shows the timing of maturity discontinuities. We detect 118 lender \times car age \times month combinations for which there is a discrete drop in maturity from one month to the next, as shown in panel A of the figure. Of these, 106 (90%) occur in January, with no other month having more than three. While we cannot know how many of the 12 non-January negative maturity shocks we identify represent actual policy updates vs. false positives, in the worst case Figure 5 suggests that no more than two or three of the 106 January maturity discontinuities that define our treatment are false positives.

While static age-based policies will lead to a concentration of maturity discontinuities in January, it is also possible that lenders re-evaluate and update their maturity policies disproportionately

[Insert
Figure
5 about
here.]

in January. It would be hard to argue that variation arising from proactive lender decisions is exogenous to the price of vehicles, as deteriorating local economic conditions could jointly drive shorter maturities and lower demand for cars. To assess this possibility, in panel B of Figure 5 we show the 71 instances of discrete increases in maturity policies for cars as they age one month. Given that these positive maturity shocks run against the natural age-maturity relationship, they almost certainly arise from proactive policy updates. It is not surprising to find such updates, given that the bulk of the loans in our data were originated during a period of lengthening maturities. The 71 occurrences are distributed roughly evenly across months, with no single month accounting for more than 13 positive shocks. In particular, there is no evidence that lenders are inclined to update maturity policies at the new year, as January accounts for only three of the 71 positive discontinuities.

Because maturity policies are very persistent, we include all months July through December in the pre-period and all months January through June in the post period, although monthly event studies will allow us to focus on the months around the end of the year. This leaves us with a total sample of 972,621 cars, of which 54,757 (5.6%) are treated observations defined by our 106 January discontinuities. Table I tabulates summary statistics for the treatment and control samples. The groups are very similar on observables, including FICO scores at origination, debt-to-income ratios (DTI), and loan-to-value (LTV) ratios. Although the differences in means are mostly statistically significant owing to the precision afforded by our large sample size, the typical difference is a tenth of a standard deviation, suggesting that treatment- and control-group borrowers are balanced for practical purposes. Consistent with being slightly older (4.29 vs. 3.86 years), treated cars have slightly lower prices (\$18,821 vs. \$20,432), shorter maturities (59.3 months vs. 61.4 months), and higher interest rates (4.31% vs. 4.09%). While much of this price difference is absorbed by our rich controls for vehicle heterogeneity, our empirical results below show that some of the price differential is a causal effect of treatment-group borrowers being offered shorter maximum maturities in the post period.

B. First Stage Results

We now turn to estimating the reduced-form impact of our detected maturity shocks on the average borrower’s maturity. Recall that a maturity “shock” in our data does not arise from lenders

changing policies but from borrowers buying a car that has recently crossed a discontinuity in a lender’s maturity policy. In Table III, we estimate

$$Maturity_{iglt} = \beta_1 Treatment_i + \beta_2 Post_t + \beta_3 Treatment_i \times Post_t + X'_{it}\gamma + \varphi_g + \psi_l + \varepsilon_{iglt} \quad (1)$$

where $Maturity_{iglt}$ is the loan maturity of transaction i , in geography g , financed by lender l in month t . Event time runs from July through the following June with $Post$ equal to zero for transactions occurring July through December and one for January through June. Treatment is a dummy equaling one for observations within a $Lender \times EventAge \times EventYear$ group with an identified shock to offered maturity occurring in January and zero for any observations in groups for which maximum allowable maturity is not changing. We use $EventAge$ to refer to the age that cars turn during January of the event year and define $EventYear$ as the calendar year of that January. Controls X_{it} consist of borrower characteristics (DTI and FICO score) and various fixed effects that control for the quality of collateral, such as $YMMT$ by month fixed effects. In some specifications, we also control for commuting zone fixed effects φ_g and lender fixed effects ψ_l . We double cluster our standard errors by month and commuting zone.

Table III reports the results. Column 1, without any fixed effects, shows a first-stage effect on average maturity of -2.4 months, meaning that the maturity for treatment-group borrowers decreased by an average of 2.4 months after their cars aged across a maturity discontinuity on January 1 relative to any change in maturity for control-group borrowers. As shown in Table I, cars in the treated group are slightly older and have slightly shorter maturities than cars in the control group. In column 2 we add car-age fixed effects, which predictably narrow the gap between treatment- and control-group maturities but leave the $Treatment \times Post$ coefficient of interest unchanged, suggesting that our difference-in-differences specification accounts for heterogeneity across car age. In column 3 we add finer collateral fixed effects, controlling for the car age interacted with make (e.g., Honda), model (e.g., Accord), and trim (e.g., LX). Column 4 adds a time dimension, interacting $YMMT$ fixed effects with year-month fixed effects. In this case, the coefficient measures the difference in maturity offered to buyers of the same $YMMT$ during the same month but with different lender maturity policies. In column 5 we add commuting zone fixed effects to account for potential differences in maturity norms across geography. Anecdotally, prices of cars differ by geography, and

[Insert
Table
III
about
here.]

column 5 allows for the same to be true of maturities. Finally, column 6 adds lender fixed effects. The estimated magnitude of our detected maturity shock is robust across all specifications, showing a stable effect on originated maturities of slightly more than two months.

[Insert
Figure
6 about
here.]

As indicated above, auto loan maturities cluster significantly on multiples of three, six, or twelve months. While the most common change in *maximum* allowable maturity for treated borrowers is -12 months (Figure 6), not everyone receives the maximum maturity. Table III shows that average originated maturity decreases by around two months, meaning that many borrowers either do not qualify for the maximum maturity or endogenously choose shorter maturity than the maximum allowable. Borrowers that demand loan maturities lower than the maximum allowable could be unaffected by any changes in maturity policy. Our instrumental-variables strategy below is designed precisely to deal with any such sorting behavior. The key takeaway from Table III is that the members of the treated group are consistently more likely to be treated with shorter maturities than members of the control group.

[Insert
Figure
7 about
here.]

To test for whether the difference-in-differences coefficients in Table III are affected by pretrends, Figure 7 plots the conditional average maturity for each month from July through June for the treatment and control groups.¹⁶ The figure shows stable maturities for the control group throughout the entire event year. The treatment group, in contrast, has stable maturities that are slightly higher than the control group from July through December, followed by a sharp drop in January, continuing to February. Maturities in February through June are stable and significantly lower than those in the control group. It is difficult to say exactly why the drop in maturities spans January and February, though it is possible that this is driven by lenders and borrowers agreeing to terms in December before the car purchase is finalized in January in some cases. This event-study approach supports our difference-in-differences parallel trends identifying assumption and bolsters our interpretation of the $Treatment \times Post$ coefficients in Table III as causal effects of year-end discontinuous maturity policies.

¹⁶Specifically, we control for the expected decrease in maturity as a car ages and any differences across geography by regressing maturity on car age \times month fixed effects and commuting zone fixed effects. We then plot the average residuals within each month for treatment and control groups.

IV. Results

Having identified plausibly exogenous variation in the supply of maturity, we now estimate the effect of maturity on consumer expenditure and car prices in the cross section of borrowers. We run the same specification as in equation (1), replacing the dependent variable with the log of the car price. For consistency, we include the same borrower controls (DTI and FICO). Similarly, we include the same sets of fixed effects in each column as we did in Table III.

We report these reduced-form results in Table IV. In column 1, where we don't include any fixed effects, we find a statistically insignificant effect of -2.6%. Of course, one way in which borrowers are likely to respond to lower maturity is by shifting toward cheaper cars, either older cars or lower-end models. Controlling for car age fixed effects (column 2) sharpens our estimation significantly (as evident in smaller standard errors and the increase in R^2 from 0.06 to 0.37) with little effect on the coefficient magnitude. Holding fixed the age of the car, affected borrowers spend 2.7% less on their car purchase, significant at the 1% level. This includes both substitution toward cheaper make/model/trims as well as any effect of financing terms on the negotiated price of a given car. To separate these two effects, we exploit the rich information on collateral quality available in our dataset.

In column 3 we interact the car age fixed effects with make-model-trim fixed effects, effectively comparing the prices paid by two borrowers with different financing terms who each bought, say, a four-year-old Honda Accord LX. The fixed effects soak up a large proportion of the variation, with the R^2 jumping to 0.87. The coefficient drops to 0.9%, indicating that a significant portion of the effect on expenditure from column 2 is being driven by a shift of affected borrowers toward lower-quality vehicles. While it may not be surprising that decreased access to maturity affects demand, recall that our financing treatment varies by lender. Evidently, many borrowers prefer to buy a lower-quality vehicle rather than to shop for better financing.¹⁷ This specification also highlights the importance of holding the quality of the good fixed when measuring the impact of credit terms on durable goods prices, one of the virtues of our setting and dataset. In column 4 we interact $YMMT$ fixed effects with year-month fixed effects such that the coefficient tells us

[Insert
Table
IV
about
here.]

¹⁷In Section B we test observable borrower characteristics for selection into treatment and find no economically meaningful results.

the difference in price paid by an affected borrower for the same *YMMT* purchased in the same month. These *YMMT*×*month* fixed effects absorb any time-varying shocks to *YMMT* values, for example, because of the introduction of a new model. Because prices may differ systematically across geographic regions, we supplement these fixed effects with commuting zone fixed effects in column 5. Finally, we add lender fixed effects in column 6 to rule out our results being driven by any lender-specific clientele-selection effects. Across all of the more stringent specifications, the estimated effect of a shock to maturity on the price paid for the same *YMMT* in the same month is significant and stable around 0.6–0.7%. Recall that the magnitude of our estimate for the first-stage effect on average maturities was about 2.3 months. The estimates in Table IV, then, indicate that a borrower who is shocked with 12 months shorter maturity would pay about 3.6% less for an observationally identical car. Directly estimating the value of an extra month of offered maturity by two-stage least squares, Table V (with columns corresponding to those in Tables III and IV) shows a price effect around 0.3% per month of maturity.¹⁸

[Insert
Table V
about
here.]

The estimates in Table IV compare average prices during the post-period (January through June) to average prices throughout the pre-period (July through December) for treatment and control groups. As shown in Figure 7, the trends in maturity moved roughly in parallel across treatment and control groups except around year end. To the extent that the difference in prices is being driven by shifts in lender maturity policies, we would expect the time series pattern of prices to match that of maturities. In Figure 8, we plot the average residuals from a regression of log price on car age × month fixed effects and commuting zone fixed effects, as we did in Figure 7 for maturities. Control vehicles show a flat pattern over the entire event year, while treated vehicles’ prices are largely flat but for a large drop in January and February, matching the pattern of maturities shown in Figure 7. One of the virtues of our difference-in-differences setup is that the presence of a control group helps rule out the possibility that our results are driven by seasonality in the car market. For example, Adams, Einav and Levin (2009) show that demand for used cars increases significantly among subprime borrowers from late January through early March, when

[Insert
Figure
8 about
here.]

¹⁸See De Chaisemartin and D’Haultfoeuille (2017) and Hudson, Hull, and Liebersohn (2017) for a precise discussion of the identification conditions needed for the consistency of difference-in-differences instrumental-variables estimators. In particular, the necessary assumptions around parallel trends, treatment exogeneity, monotonicity, and stability of treatment effects across time and subgroups are each quite plausible in our setting.

many buyers would receive EITC rebates. Our month fixed effects, the parallel price patterns in Figure 8, and the persistence of our price effects throughout the year confirm that our control group effectively addresses any such seasonal patterns.

One potential concern with our empirical approach is that our results could be driven by the fact that we use the same sample to determine discontinuities in offered maturity—i.e., the assignment of treatment and control—as we do to estimate potential pricing effects. Though we have tried to provide evidence that we are capturing actual shocks to the supply of maturity, we attempt to further mitigate these concerns with the following validation exercise. Within a given lender \times car age \times month cell, we randomly assign half the loans to a training sample and the other half to a hold-out sample. We use the training sample to identify lender maturity policies in order to define treatment and control observations, following the procedure outlined in Section A. We then use the hold-out sample to estimate our reduced form pricing regressions. In this way, we estimate the effect of maturity on prices out of sample relative to the data we use to detect offered maturity discontinuities. The results, shown in Appendix Table AI, closely mirror those in Table IV in terms of magnitudes and significance.

To evaluate the magnitude of these estimates, consider a borrower buying an average car priced at \$20,000 financed by a 72-month loan at an interest rate of 4.1% and an LTV of 90%. Under these parameters, the borrower would put \$2,000 down and have a monthly payment of \$282.43 for 72 months. If we ignore our estimated consumer response to shorter financing, a counterfactual treated borrower receiving a 60-month loan (the modal maturity shock in our data is 12 months) would have a monthly payment of \$332.31, or \$49.88 higher. In response to the shorter financing, column 2 of Table V indicates that a treated borrower would spend about 13.4% less on a vehicle (1.12% per month of maturity \times 12 months), or \$17,312. The resulting monthly payment would be \$287.65, only \$5.22 higher than that paid by untreated borrowers. In other words, borrowers affected with shorter maturity spend less on a car so as to offset 90% of the increase in their monthly payment. The majority of this comes from substitution to lower quality vehicles. Controlling for substitution, columns 4–6 indicate that a treated borrower would pay about 3.6% less (0.30% per month of maturity), or \$19,280, for an observationally equivalent car. The resulting monthly payment of \$320.35 indicates that borrowers are able to offset 24% of the monthly payment increase from shorter maturity through lower negotiated prices.

We can also quantify our estimates in terms of the implied borrower discount rates. Affected borrowers pay monthly payments of \$320.35 for 60 months, while untreated borrowers pay \$282.43 for 72 months. The lower price paid by treated borrowers also results in a lower down payment (\$1,928 vs. \$2,000). The internal rate of return on the marginal cash flows for treated borrowers is about 14.2%—the annual discount rate that would make borrowers indifferent between shorter and longer maturity. To the extent that borrowers are not able to fully offload the cost of decreased access to finance onto sellers, this represents a lower bound on borrower discount rates.

A. *Isolating Maturity Effects from Interest Rate Effects*

Our results presented thus far have focused on the maturity dimension of the financing contract, motivated by evidence in Argyle, Nadauld, and Palmer (2020) that constrained consumers have stronger preferences over maturity than over interest rates. However, given that maturities and interest rates frequently move together in a contract bundle, the empirical strategy discussed in Section III is subject to the concern that identified breaks in maturity policies may be coincident with breaks in lenders’ interest rate policies. While any interest-rate effect contained in our estimates would not invalidate a claim that we are estimating a causal effect of credit on prices generally, it would compromise the interpretation that estimated price effects are driven by changes in maturity. In this section, we turn our attention to disentangling the effects due to changes in maturity from those due to changes in interest rates.

We begin by examining whether the maturity shocks we detect coincide with changes in interest rates by re-estimating the difference-in-differences specifications from Table III, with interest rates for each loan replacing maturity as the dependent variable. We report the results in Table VI. With no fixed effects, the coefficient of interest on $Treatment \times Post$ is six basis points and statistically insignificant. As we control for increasingly fine collateral fixed effects in columns 2–4, the estimate remains insignificant, ranging in magnitude from four to nine basis points. In column 5 we add commuting zone fixed effects to the $YMMT \times month$ fixed effects of column 4, which increases the coefficient to 12 basis points, marginally significant at the 10% level. Finally, with the addition of lender fixed effects in column 6, the coefficient is 16 basis points, significant at the 1% level. Figure 9 plots the time series pattern of interest rates during the event year for the treatment and control groups. Consistent with the estimates in the table, rates appear to be somewhat higher in the post-

[Insert
Table
VI
about
here.]

[Insert
Figure
9 about
here.]

period months for treated cars, though the pattern appears much less stark than the corresponding pattern for maturities in Figure 7. This is not surprising given that our empirical design is built to detect maturity breaks rather than interest rate breaks. Still, though the relationship is not strong, the consistent message from Table VI and Figure 9 is that interest rates appear to increase somewhat for treated cars in the post-period, which could potentially be driving some of the lower prices that we observe for treated borrowers in the post period.¹⁹

In an effort to pin down a causal estimate of maturity accounting for any interest-rate impact, we estimate a two-stage least squares regression, instrumenting for both maturity and interest rates. We estimate separate first stages for maturity and rates as follows:

$$Maturity_{igt} = \sum_k \pi_k^{mat} \mathbb{I}_{k,ilt} \times Post_t + \sum_k \alpha_k^{mat} \mathbb{I}_{k,ilt} + X'_{igt} \gamma^{mat} + v_{igt}^{mat} \quad (2)$$

$$Rate_{igt} = \sum_k \pi_k^{rate} \mathbb{I}_{k,ilt} \times Post_t + \sum_k \alpha_k^{rate} \mathbb{I}_{k,ilt} + X'_{igt} \gamma^{rate} + v_{igt}^{rate} \quad (3)$$

As before, X_{igt} contains borrower controls, lender fixed effects, commuting zone fixed effects, and rich collateral fixed effects. The key innovation with respect to the reduced-form formulations is the instrument set, which is a full set of treatment-cell indicators interacted with $Post$. In our notation, k indexes the individual $Lender \times EventAge \times EventYear$ cells that make up our treatment group, with $k = 0$ corresponding to the control group. The $\mathbb{I}_{k,ilt}$ indicator variables identify whether a given borrower i financing their purchase with lender l at time t was in treatment cell k , as defined in section A. The key feature of this specification is to allow unique magnitudes of the difference-in-differences coefficients π_k for each treated cell. Identifying the unique impact of maturity policy breaks separately from interest rate breaks relies on the magnitudes of π_k^{mat} and π_k^{rate} not being perfectly correlated across each of the 106 identified policy breaks. The exclusion restriction is satisfied if variation in π_k^{mat} and π_k^{rate} is exogenous to pricing outcomes, only affecting prices through loan maturities and interest rates.

Consider the following illustrative example. A particular institution has a policy in place that

¹⁹Note that the evidence that any interest-rate movements coincident with our shifts in maturity are positive further supports the claim that we have identified true shifts in the supply of maturity. If the changes were driven by demand for maturity, we would expect to see lower interest rates associated with the lower maturities, as borrowers often have a menu of maturity-interest rate pairs from which to choose, with an upward-sloping term structure.

decreases allowable maturity by six months as a vehicle manufactured in 2006 rolls from three years old in December 2009 to four years old in January 2010. The same institution’s policy calls for a 20 basis point increase in interest rates in this scenario. Meanwhile, a different lender’s policy results in a 12-month decrease in maturity and a 10 basis-point interest-rate increase as three-year-old vehicles age by a year. The variation in the magnitudes of the maturity and interest-rate breaks across $Lender \times EventAge \times EventYear$ combinations allows us to simultaneously identify the causal impact of maturity and interest rate changes on prices.

With 106 coefficients on policy discontinuities, we do not report a full set of first stage results. In each column and for every endogenous variable, though, the F -statistic for joint significance of the instruments has a p -value less than 0.1%. Meanwhile, the correlation between the estimated coefficients on maturity and interest rates is small, varying between -0.19 and -0.13 across the six columns. The takeaways from the first stage regressions are that the instrument set strongly predicts each endogenous variable and that there is sufficient independent variation to separately identify the price impact driven by each variable.

Instruments for maturity and interest rates allow for a second-stage specification that utilizes equations (2) and (9) and is given by

$$\log Price_{igt} = \sum_k \alpha_k \mathbb{I}_{k,igt} + \eta^{mat} Maturity_i + \eta^{rate} Rate_i + \eta^{LTV} LTV_i + X'_{igt} \mu + \varepsilon_{igt} \quad (4)$$

such that the η coefficients are semi-elasticities of price with respect to maturity and interest rates and represent the local average treatment effects on prices for complier borrowers affected by the instruments. We report estimates in Table VII, where each column corresponds to the same column of Table IV. In general, the estimated impact of a one-month change in maturity is 70–80% as large once we account for concurrent changes in interest rates. The estimate in column 6, for example, indicates that for a car of the same $YMMT$ bought in the same month, holding fixed average differences in prices across commuting zones and lenders, an additional month of supplied maturity translates into 23 basis points higher price, significant at the 1% level. This compares to an estimate of 29 basis points per month of maturity in the two-stage least squares specification without interest rates in Table V, indicating that roughly 80% of the effect on prices previously estimated is coming through the maturity channel. The coefficient on $Rate$ indicates that for a one percentage-point

[Insert
Table
VII
about
here.]

increase in interest rates, prices fall by 90 basis points. Given our first-stage estimate of a change in interest rates of 16 basis points, these estimates imply that roughly 14 basis points of price impact is driven by changes in interest rates, compared to the total price impact of 70 basis points reported in Table IV.

The estimates in Table VII allow us to refine our interpretation of the impact of maturity on durable goods markets, accounting for changes in interest rates. Using the same average borrower/average car calibration as above, the estimates in column 2 imply that a borrower offered 12 months shorter maturity spends 8.9% less on a car, enough to offset 60% of the increase in monthly payment. The estimates in columns 4–6 indicate that around 20% comes from a lower negotiated price, while the remaining 40% is due to a shift to lower quality vehicles.

We can also interpret the price discount that affected borrowers receive in terms of implied discount rates. For buyers purchasing an average car as described above, these estimates imply a break-even discount rate of 11.6%, meaning that at that borrower discount rate sellers would bear the entire cost of shorter maturity. Buyers with discount rates below this IRR of 11.6% would happily accept lower maturity if it meant attaining the discounts we estimate; however we note that maturity shocks could affect the reservation prices of such buyers sufficiently to alter bargaining outcomes. Buyers and sellers share the cost of lower maturity equally at a borrower discount rate of 30%.

Table VII also allows us to independently assess a borrower's response to changes in interest rates. Column 2 indicates that borrowers respond to a one percentage-point higher interest rate by spending 1.95% less on a car. While the impact on monthly payments coming from such an interest rate change is comparatively small (only \$8.29 in this case), the lower expenditure offsets 68% of the increase in monthly payment, similar to the consumer response to shorter maturity. The fact that borrowers respond similarly to differences in monthly payments regardless of the source of the change is consistent with monthly payments being a focal point of consumer decision-making. In the case of interest rates, the 68% lower expenditure is split roughly equally between substitution to lower quality cars (36%) and lower negotiated prices for equivalent cars (32%), as columns 4–6 indicate that borrowers pay about 0.90% less controlling for vehicle quality.

In addition to interest rates, it is natural to ask whether lenders have LTV policies that feature discontinuities that correspond to our inferred maturity discontinuities. If so, part of the result

that we are attributing to differences in maturity might be driven by differences in required down payments. This would complement Adams, Einav, and Levin (2009), who show that vehicle demand is very sensitive to available liquidity for a population of subprime borrowers. To assess this possibility, we might again re-run our main difference-in-difference specification with LTV as the dependent variable. However, Argyle, Nadauld, and Palmer (2019) show that borrowers demand smaller loans when offered less favorable maturity, so it is difficult to know whether a significant effect on LTV should be interpreted as a discontinuity in lender LTV policies or an endogenous response by borrowers to decreased access to maturity. Moreover, car age explains only 0.2% of the variation in LTV in a simple univariate regression, compared to 17% of maturity and even 7% of interest rates. Given the likelihood that LTV effects may be driven by endogenous responses in borrower demand, we interpret any regressions involving LTV cautiously. In Appendix B, we present results that affected borrowers have lower LTVs by around two percentage points (relative to a baseline of 90%), although this has very little effect on the estimated price response to a change in offered maturity or interest rates (see Appendix Tables AVIII and AIX).

B. Unobserved Heterogeneity

One novel aspect of our empirical strategy is that our ability to control for $YMMT \times month$ fixed effects substantially reduces the scope for our estimates of price effects to be driven by substitution toward lower quality goods. Indeed, the R^2 of our pricing model in Table IV is over 0.9. Still, the fixed effects cannot conclusively rule out the possibility that unobserved vehicle or borrower heterogeneity plays some role in our results. Vehicles of a given $YMMT$ in a given month may still exhibit meaningful differences in vehicle condition, including mileage, accident history, and whether they were owned by smokers or driven by aggressive, pizza-delivering teenagers. Similarly, borrowers who take up loans with lower maximum maturities may also be different in some way correlated with their demand for cars. We address these concerns in several ways: by analyzing repeat-sales prices, testing for heterogenous effects in subsamples with relatively less scope for unobserved heterogeneity, examining the mileage of treated car purchases for a subsample of cars for which we have odometer data, and checking for changes in borrower composition in our difference-in-differences framework.

We first attempt to address unobserved quality concerns by evaluating prices for cars that sold multiple times in our sample. If our pricing results are driven by consumers that shift demand to

cars with unobservably lower quality in response to being offered lower maturity loans, the lower quality would presumably be manifest in a lower relative price when the car is sold again. Relaxing our sample selection criteria for power considerations, our entire dataset features 8,697 cars with at least two transactions. We require the two transactions to occur at least 18 months apart to avoid contamination resulting from aggressive purchasers looking to quickly flip cars. Our repeat-sales pricing analysis begins with the calculation of pricing residuals for every transaction in our data, conditioning on $YMMT \times month$ fixed effects $\delta_{YMMT(i),t}$, lender fixed effects φ_l , and commuting zone fixed effects α_g in the following regression

$$\log Price_{igt} = \alpha_g + \varphi_l + \delta_{YMMT(i),t} + u_{igt}. \quad (5)$$

We then evaluate whether the fitted pricing residuals for *second* sale transactions are unusually low if the *first* sale for that car was a transaction with $Treatment \times Post = 1$ by running a difference-in-differences regression, with $Treatment$ and $Post$ taking their values as of the first sale at t_0 :

$$\hat{u}_{igt} = \beta_1 Treatment_{it_0} + \beta_2 Treatment_{it_0} \times Post_{t_0} + \varepsilon_{igt}. \quad (6)$$

If our price results are being driven by unobserved differences in quality, these differences would likely be persistent, resulting in lower prices for those same cars when sold a second time.

Before estimating equation (6), we assess the scope for our sample of repeat sales to be selected in important ways. Specifically, a concern for this exercise is the possibility for cars to only be observed selling twice in our data if they have not decreased in value significantly. Such endogenous resale behavior would bias our estimates of price effects at the second sale upwards if correlated with $Treatment \times Post$. In Appendix Table AII, we estimate a linear-probability model to see whether cars that faced financing with exogenously lower maturity due to a lender’s maturity discontinuity ($Treatment \times Post = 1$) are less likely to be sold again. We find no evidence of differential selection into resale.

Table VIII presents results estimating equation (6). Column 1 reports that cars previously treated with low maturity sell for a statistically insignificant higher price (60 basis points) when sold a second time. Of course, this sample is different from our main sample, so in column 2 we

[Insert
Table
VIII
about
here.]

estimate the difference-in-differences regression for the first sale of the same 8,697 cars, essentially the same specification as in column 6 of Table IV.²⁰ The estimate shows that our main price result—a pricing discount for cars treated with low maturity relative to otherwise comparable cars—holds in this subsample, though the statistical significance is muted due to a substantially smaller sample size. The difference between the first- and second-sale estimates is significant at the 10% level, indicating that financing-related discounts appear to rebound when the same car is sold subsequently. While we acknowledge that we only observe a small subsample of cars with repeat transactions, this price rebound at second sale is inconsistent with many forms of unobserved vehicle quality (accident history, high mileage, etc.) driving our results.

A complementary approach to testing whether our results are driven by unobserved vehicle heterogeneity is to examine subsamples of our data where the scope for unobserved heterogeneity is reduced. Younger cars, for example, have less time to accumulate quality differences such as the beneficial effects of fastidious maintenance or the negative impacts of heavy use or accidents. As one measure of this, we show using National Household Travel Survey microdata that the standard deviation of vehicle mileage is strongly increasing in vehicle age (see Appendix Figure A1). To the extent that dealers may specialize in older and younger cars, this analysis also helps us understand whether dealer heterogeneity could be driving some of our results. Dividing our sample by the sample median age (five years), we re-estimate our 2SLS specification in equation (4) for young (average age three years old) and old cars (average 8.5 years old) and report these results in Appendix Table AIII. Though the R^2 indicates greater scope for substitution to lower priced cars of a given $YMMT$ among older vehicles, we find very similar effects of a month of offered maturity on prices in both samples, and a formal test fails to reject the equality of the maturity coefficients in the two samples.

For a subsample of our cars, we were able to collect mileage information from the California Bureau of Automotive Repair and merge by VIN. For these cars we can directly observe mileage to test whether cars purchased by treated borrowers have higher mileage. In Appendix Figure A2 we plot an event study of mileage in this subsample. The figure shows no apparent difference in mileage for treated cars sold in the first half of the calendar year.

Next, we follow the exercise of Oster (2017) and adjust our estimates for omitted variables

²⁰The only difference is that the fixed effects are accounted for in the creation of the pricing residuals rather than being estimated directly in the regression.

bias. In Appendix Table AIV, we adjust the $Treatment \times Post$ coefficients in the six reduced-form price regressions from Table IV. Unsurprisingly, the potential bias is meaningfully large without adequate fixed effects (columns 1–2), though it still does not change the sign of our estimates. With the addition of Age \times MMT fixed effects, however, the adjusted estimate is within our original 95% confidence interval (column 3). This is true of the more aggressive fixed effect structures as well (columns 4–6), where we see estimates very similar to the unadjusted coefficients reported earlier. The results of Appendix Table AIV suggest that the scope for our price results to be driven by unobserved product or borrower heterogeneity correlated with maturity policies is quite limited thanks to the richness of our controls.

Finally, we examine difference-in-differences estimates using the borrower controls (FICO and DTI) as dependent variables to explore potential changes in borrower composition. As expected given implied changes in monthly payments, we find a slight increase in reported DTIs in some specifications (point estimates range from $\frac{1}{25}$ to $\frac{1}{30}$ of a standard deviation) but no significant result when lender fixed effects are included. We find no significant change in FICO scores, regardless of the fixed effect structure. Figure 10 plots event studies of FICO and DTI by month of the year for treatment and control groups separately. The magnitudes of any differences are economically small with no consistent nor statistically significant pattern that would suggest the composition of treated borrowers changing in January relative to control-group borrowers.

[Insert
Figure
10
about
here.]

Taken together, these results suggest that unobservable heterogeneity in borrowers or car quality is not a likely source of bias in our estimates of the causal impact of maturity on prices.

C. Mechanisms

What is the underlying mechanism that would result in two buyers of observationally identical goods paying different prices? Our discussion of the economic framework, outlined in Section B, described the set of possible borrower reactions to shorter maturity loans. Borrowers could satisfy a binding budget constraint by substituting into lower quality cars, searching for a better price on their car of first choice, or bargaining over the surplus defined by the difference between buyer and seller private valuations (i.e., marginal willingness to pay and marginal willingness to accept). We separate price effects from substitution by comparing estimates without tight collateral controls (a 2.7% expenditure impact in columns 1–2 of Table IV) to estimates with tight collateral controls (a

0.7% impact in columns 4–6). Though collateral controls facilitate the estimation of “same car” price effects, they do not convincingly distinguish between search or bargaining as mechanisms driving price effects.

In an effort to disentangle these mechanisms, we exploit a subsample of lenders in our data that provide details on loan applications. Merging application data to subsequently originated loans yields a sample of 54,929 loans, allowing for the construction of a variable that measures the number of days elapsed between a loan application and origination. The average (interquartile range) of the time between loan application and origination is 6.4 days (0 to 8 days). We again employ our difference-in-differences framework to estimate whether borrowers treated with exogenously shorter maturity loans wait longer to originate a loan after filling out an application, under the presumption that borrowers are searching in the interim. We report the results in Table IX. The limited loan application sample size does not allow for the *YMMT* collateral controls featured in previous tables, so we control for collateral quality with car age fixed effects. The negative and insignificant estimates indicate that *Treatment* × *Post* observations are not associated with a meaningfully larger number of days between application and origination, suggesting that treated borrowers are not spending more days shopping for cars in response to lower maturity loans. These results cast some doubt on the possibility that borrowers respond to maturity shocks by increasing their search effort, though our test is admittedly limited in addressing the possibility that treated borrowers search more intensely over the same number of days as non-treated borrowers.

[Insert
Table
IX
about
here.]

Unfortunately, our data do not facilitate direct measurement of bargaining behavior. Demonstrating bargaining as the exact mechanism would require, at the very least, differences between posted prices and realized prices. The ideal data to find positive evidence of bargaining would involve a detailed set of offers and counteroffers leading to equilibrium prices as in Larsen (2018). Our data do provide compelling evidence that substitution cannot fully explain the decreased expenditure of treated borrowers, and suggestive evidence from application data indicates that additional search by treated borrowers is not playing a substantial role. The totality of evidence indicates that bargaining likely plays a meaningful role in the way that finance is capitalized into the price of equivalent-quality used cars.

V. Conclusion

We study the consumer response to individual-level variation in credit terms. We find that borrowers that are treated with shorter maturity spend less on a car so as to offset 60% of the increase in monthly payment. Two-thirds of this decrease in expenditure comes from substitution toward lower quality vehicles, while the remaining third arises as a result of negotiating better prices on equivalent vehicles. Borrowers treated with 12 months shorter maturity pay roughly 2.8% less than unaffected borrowers for cars of the same manufacture year-make-model-trim (*YMMT*) at the same point in time. These results are not driven by changes in the interest rates of the accompanying loans. Moreover, the prices of cars bought by affected borrowers, if anything, rebound when sold in later transactions, indicating that initial price differentials were unlikely to be driven by unobservable quality differences within *YMMT*. Our interpretation is that constrained buyers, pinched by lower maturity and the associated higher monthly payments, have a lower private valuation for cars in their choice set. This lower private valuation affects their incentives in the search and bargaining processes inherent in the auto market, resulting in lower realized prices for observationally equivalent vehicles. Frictions in the auto market likely play a significant role in facilitating the pass-through of finance terms to prices at the individual level—including search and bargaining or sticky demand driven by consumer preferences for a certain car type, brand, or dealership. While it is an open question whether our results would generalize to other markets, many markets for big-ticket items are characterized by similar frictions (real estate, machines, furniture, higher education, etc.).

Our focus on the cross section of prices raises interesting questions regarding the incidence of credit-supply shocks. Our results suggest that the price impact of changes in credit terms is concentrated among affected borrowers, rather than being spread across all borrowers through an aggregate demand channel. This serves to decrease any wedge in surplus between treated and untreated borrowers caused by differential access to credit. Meanwhile, sellers are sorted into winners and losers based on the financing of their buyers.²¹

Our analysis also speaks to the transmission of policy actions through to final-goods prices. For

²¹Note that our findings of the impact of consumer financing disruptions on sellers provides a positive economic rationale for commonplace vertical integration between lenders and dealers. See Murfin and Pratt (2019) for theoretical and empirical evidence on vendor financing in heavy construction equipment.

example, one goal of monetary policy is to influence consumer demand through the interest-rate channel. Our results demonstrate how capitalization effects can blunt monetary policy's ability to affect demand by changing monthly payments. Moreover, demand, and ultimately prices, can be influenced through dimensions of the credit surface besides rates, such as maturity. Given the importance of monthly debt service capacity to consumer demand and equilibrium prices, a parameter of interest is an estimate of the sensitivity of durable goods prices to changes in monthly debt service payments. The elasticity of price to changes in monthly payments can be calculated by dividing estimated price changes (2.8% from a 12-month maturity shock) by estimated changes in monthly payment amounts (14.4%). This calculation implies an elasticity estimate of -0.19 , suggesting that policy actions that increase monthly payment amounts by 10% would be associated with price declines of 1.9%.

We view our results as a novel contribution to the literature investigating the link between credit and prices. While most studies that link credit and prices evaluate credit shocks in the time series and examine their impact on aggregate price levels, our cross-sectional identification documents the existence of a transmission mechanism between aggregate shocks like monetary policy and aggregate prices. A monetary policy shock plausibly impacts many factors in general equilibrium, including interest rates, lending standards, expectations, investment, and aggregate demand. We use a microeconomic empirical strategy to isolate the credit-terms channel and demonstrate its importance. Additional evidence suggests that financing terms may affect the dynamics of the bargaining game between sellers and retail buyers, consistent with the literature showing the effect of corporate debt on various forms of negotiations. Finally, our results also have implications for the optimal design of macroprudential policy. Given the tight link between payment size, asset prices, and demand, maturity is an important if presently overlooked lever in affecting prices and consumption. Overall, our results call for further examination of the attributes of loan contracts that consumers most value with potential implications for credit product design.

References

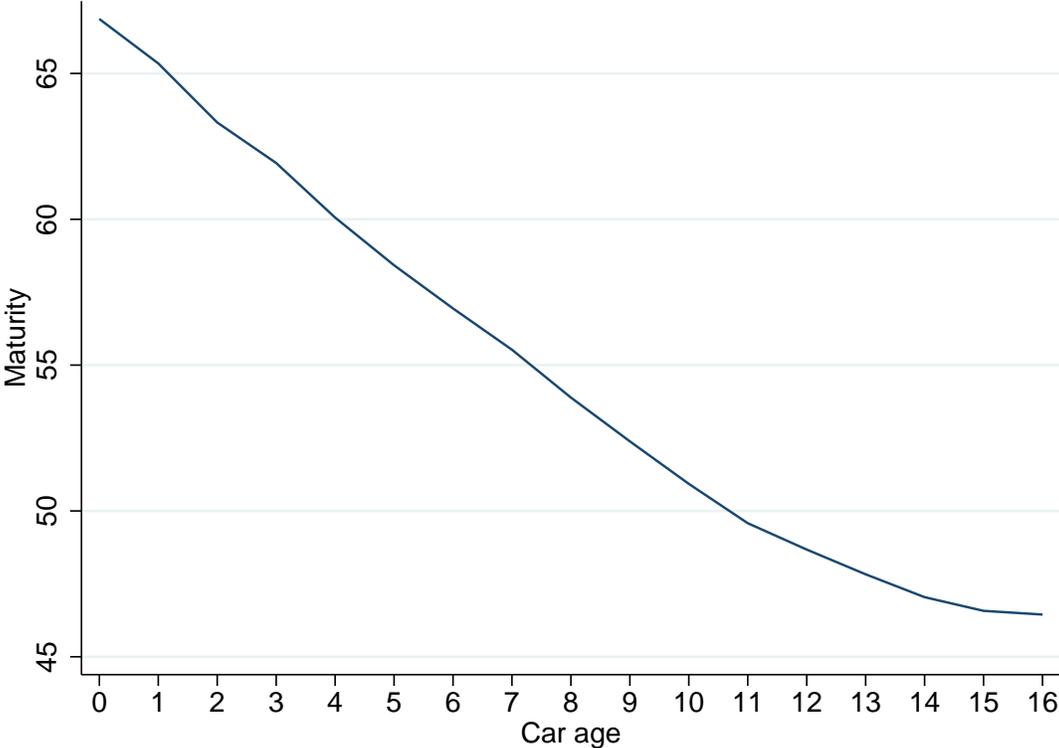
- Adelino, Manuel, Antoinette Schoar, and Felipe Severino (2013), Credit Supply and House Prices: Evidence from Mortgage Market Segmentation, NBER Working Paper No. 17832.
- Argyle, Bronson, Taylor Nadauld, and Christopher Palmer (2017), Real Effects of Search Frictions in Consumer Credit Markets, MIT Sloan Working Paper 5242-17.
- (2020), Monthly Payment Targeting and the Demand for Maturity. *Review of Financial Studies*, Forthcoming.
- Benmelech, Efraim, Ralf Meisenzahl, and Rodney Ramcharan (2017), The real effects of liquidity during the financial crisis: Evidence from automobiles. *Quarterly Journal of Economics* 132(1), 317–365.
- Bhutta, Neil and Daniel Ringo (2017), The Effect of Interest Rates on Home Buying: Evidence from a Discontinuity in Mortgage Insurance Premiums, SSRN Working Paper No. 3085008.
- Borio, Claudio and Philip William Lowe (2002), Asset prices, financial and monetary stability: exploring the nexus, BIS Working Paper No. 114.
- Busse, Meghan, Christopher Knittel, and Florian Zettelmeyer (2012), Stranded Vehicles: How Gasoline Taxes Change the Value of Households' Vehicle Assets, Working Paper.
- Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer (2006), \$1,000 cash back: The pass-through of auto manufacturer promotions. *American Economic Review* 96(4), 1253–1270.
- Cornia, Marco, Kristopher Gerardi, and Adam Shapiro (2011), Price Discrimination and Business-Cycle Risk, Federal Reserve Bank of Atlanta Working Paper 2011-3.
- Correia, Sergio (2016), *Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator*. Tech. rep. Working Paper.
- Davis, Morris, Stephen Oliner, Tobias Peter, and Edward Pinto (2017), The Impact of Interest Rates on House Prices and Housing Demand: Evidence from a Quasi-Experiment, Working Paper.
- De Chaisemartin, Clément and Xavier D'Haultfœuille (2017), Fuzzy differences-in-differences. *The Review of Economic Studies* 85(2), 999–1028.
- Di Maggio, Marco and Amir Kermani (2017), Credit-induced boom and bust. *The Review of Financial Studies* 30(11), 3711–3758.
- Englmaier, Florian, Arno Schmöller, and Till Stowasser (2017), Price Discontinuities in an Online Market for Used Cars. *Management Science*.
- Favara, Giovanni and Jean Imbs (2015), Credit Supply and the Price of Housing. *American Economic Review* 105(3), 958–992.
- Gavazza, Alessandro, Alessandro Lizzeri, and Nikita Roketskiy (2014), A quantitative analysis of the used-car market. *American Economic Review* 104(11), 3668–3700.

- Glaeser, Edward, Joshua Gottlieb, and Joseph Gyourko (2012), Can Cheap Credit Explain the Housing Boom? *Housing and the Financial Crisis*. University of Chicago Press, 301–359.
- Goolsbee, Austan (1998), Investment tax incentives, prices, and the supply of capital goods. *Quarterly Journal of Economics* 113(1), 121–148.
- Hansman, Christopher, Harrison Hong, Wenxi Jiang, Yu-Jane Liu, and Juan-Juan Meng (2018), Riding the Credit Boom, NBER Working Paper No. 24586.
- Hennessy, Christopher and Dmitry Livdan (2009), Debt, bargaining, and credibility in firm–supplier relationships. *Journal of Financial Economics* 93(3), 382–399.
- Hertzberg, Andrew, Andres Liberman, and Daniel Paravisini (2018), Screening on loan terms: evidence from maturity choice in consumer credit. *The Review of Financial Studies*, 3532–3567.
- Huang, Guofang, Hong Luo, and Jing Xia (2015), Invest in information or wing it? A model of dynamic pricing with seller learning, SSRN Working Paper No. 2668838.
- Hudson, Sally, Peter Hull, and Jack Liebersohn (2017), Interpreting Instrumented Difference-in-Differences, MIT Mimeo.
- Israel, Ronen (1991), Capital structure and the market for corporate control: The defensive role of debt financing. *Journal of Finance* 46(4), 1391–1409.
- Jordà, Òscar, Moritz Schularick, and Alan M Taylor (2015), Betting the house. *Journal of International Economics* 96, S2–S18.
- Krishnamurthy, Arvind and Tyler Muir (2017), How credit cycles across a financial crisis, NBER Working Paper No. 23850.
- Landvoigt, Tim, Monika Piazzesi, and Martin Schneider (2015), The Housing Market(s) of San Diego. *American Economic Review* 105(4), 1371–1407.
- Larsen, Bradley (2018), The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions, NBER Working paper No. 20431.
- Lee, Alice J and Daniel R Ames (2017), “I can’t pay more” versus “It’s not worth more”: Divergent effects of constraint and disparagement rationales in negotiations. *Organizational Behavior and Human Decision Processes* 141, 16–28.
- Lucca, David O, Taylor Nadauld, and Karen Chen (2016), Credit supply and the rise in college tuition: Evidence from the expansion in federal student aid programs. FRB of NY Staff Report No. 733.
- Matsa, David (2010), Capital structure as a strategic variable: Evidence from collective bargaining. *Journal of Finance* 65(3), 1197–1232.
- Mian, Atif and Amir Sufi (2009), The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis. *Quarterly Journal of Economics* 124(4), 1449–1496.
- (2011), House prices, home equity-based borrowing, and the US household leverage crisis. *American Economic Review* 101(5), 2132–56.

- Mian, Atif and Amir Sufi (2018), Finance and business cycles: the credit-driven household demand channel, NBER Working Paper No. 24322.
- Mian, Atif, Amir Sufi, and Emil Verner (2017), Household debt and business cycles worldwide. *Quarterly Journal of Economics*.
- Müller, Holger and Fausto Panunzi (2004), Tender offers and leverage. *Quarterly Journal of Economics* 119(4), 1217–1248.
- Nadauld, Taylor and Shane Sherlund (2013), The impact of securitization on the expansion of subprime credit. *Journal of Financial Economics* 107(2), 454–476.
- Oster, Emily (2017), Unobservable Selection and Coefficient Stability: Theory and Evidence. *Journal of Business & Economic Statistics*, 1–18.
- Rice, Tara and Philippe Strahan (2010), Does Credit Competition Affect Small-Firm Finance? *Journal of Finance* 65(3), 861–889.
- Sallee, James (2011), The surprising incidence of tax credits for the Toyota Prius. *American Economic Journal: Economic Policy* 3(2), 189–219.
- Spiegel, Yossef and Daniel Spulber (1994), The capital structure of a regulated firm. *The RAND Journal of Economics*, 424–440.
- Stanton, Richard, Chris Strickland, and Nancy Wallace (2015), A New Dynamic House-Price Index for Mortgage Valuation and Stress Testing, Working Paper.
- Verner, Emil and Gyozo Gyöngyösi (2018), Household Debt Revaluation and the Real Economy: Evidence from a Foreign Currency Debt Crisis, Working Paper.
- Zevelev, Albert (2016), Does Collateral Value Affect Asset Prices? Evidence from a Natural Experiment in Texas, SSRN Working Paper No. 2815609.

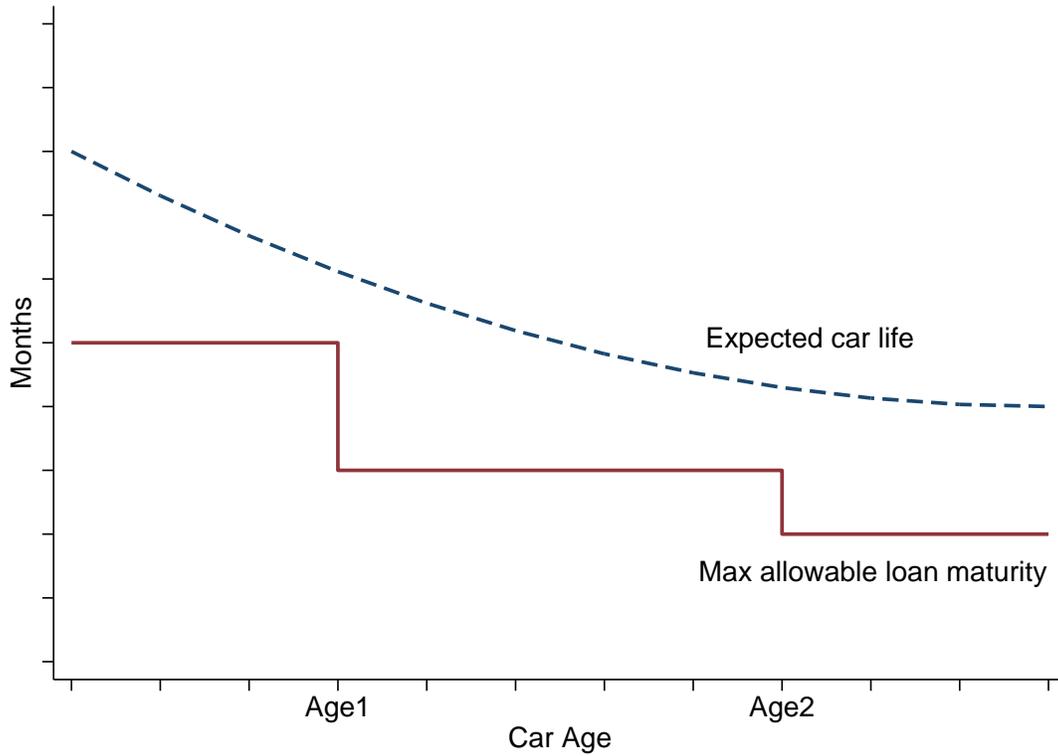
Figures and Tables

Figure 1: Average Maturity By Car Age



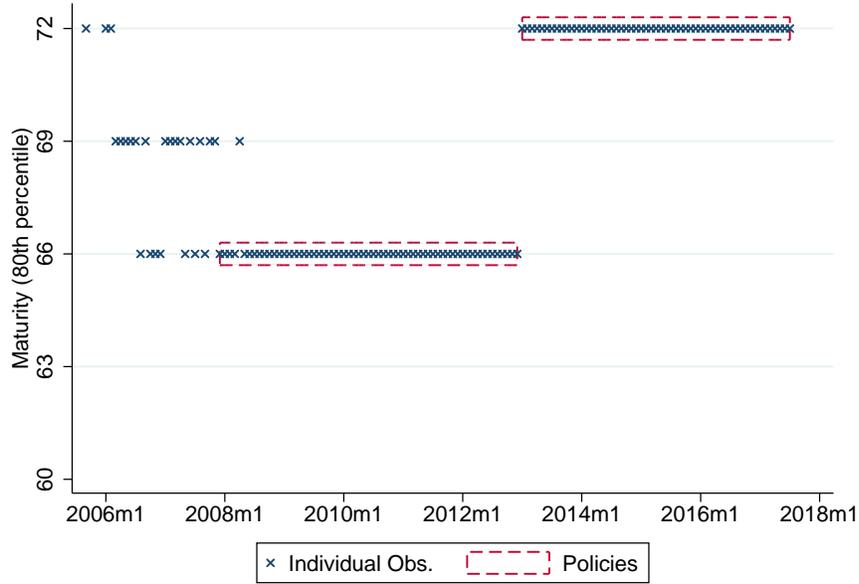
Note: Figure plots the average maturity of an auto loan in our data as a function of the age of the underlying collateral.

Figure 2: Hypothetical Lender Maximum Maturity Rule

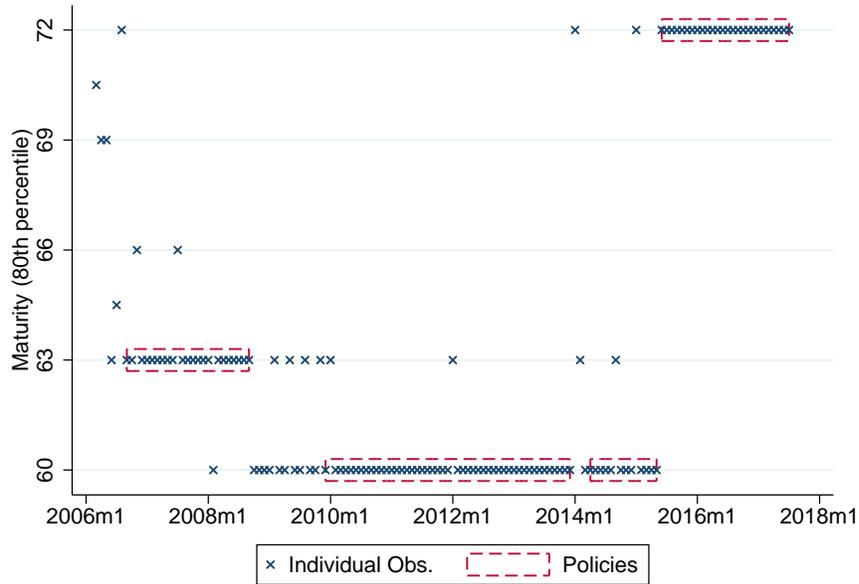


Notes: Figure depicts a hypothetical discontinuous lender rule. While the conditional expected remaining life of a car decreases smoothly with car age, many lenders have discrete maximum allowable loan maturity rules that discontinuously decrease maximum allowable maturity once a car age reaches certain cutoffs.

Figure 3: Example-Lender Maturity Policies
 Panel A: Example Maturity Policy for 3-year-old Cars



Panel B: Example Maturity Policy for 4-year-old Cars



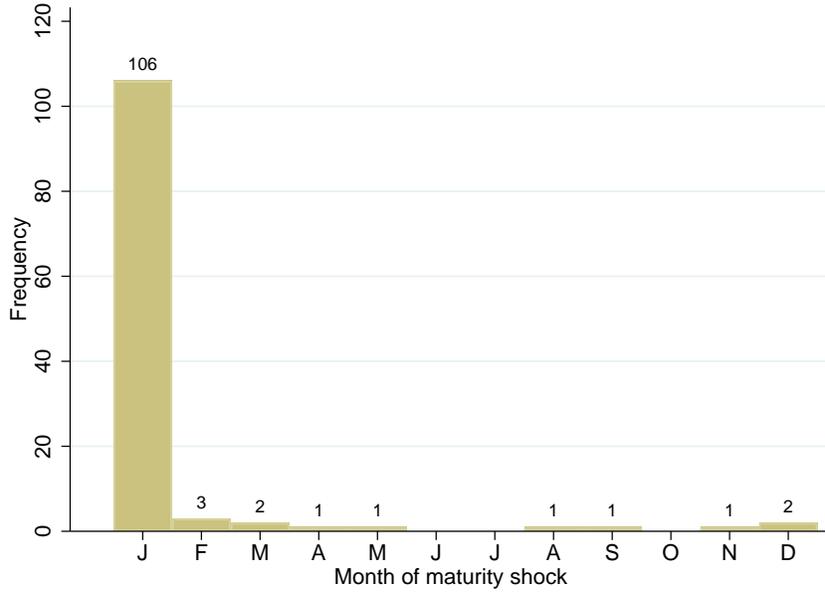
Notes: Figure plots the maturity policies for the largest lender in our data. Individual observations x capture the 80th percentile of maturity for a given car age within a given month at this lender. Boxed areas correspond to the maturity policies identified by our algorithm (see section A for details). Panels A and B show the policies for three- and four-year-old cars, respectively.

Figure 4: Example-Lender Maturity Shocks

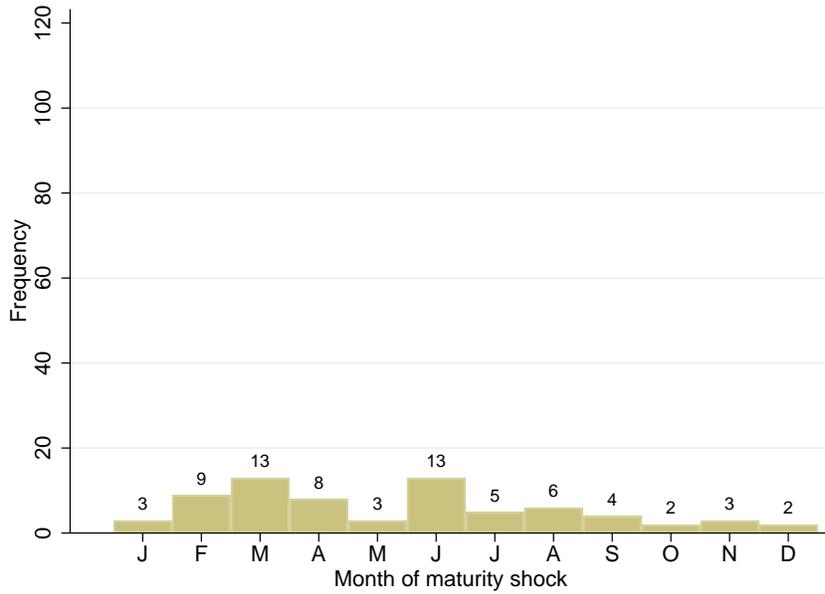


Notes: Figure illustrates our empirical design by plotting treatment- and control-group observations against identified maturity policies for the example lender used in Figure 3. Dashed rectangles correspond to the lender’s maturity policy for three-year-old cars, as identified in Figure 3. Similarly, the solid rectangles represent the maturity policy for four-year-old cars for the same lender. Dotted vertical lines represent instances of maturity shocks in which a given vehicle would receive discontinuously lower maturity in January relative to December. As an example of a treated transaction in our sample, a 2006 Honda Civic LX bought in December 2009 would have had a maturity policy of 66 months, whereas the same vehicle purchased in January 2010 would have a 60 month maximum allowable maturity. In contrast, by 2016 the policies for three and four year old cars are both 72 months. Thus, a 2012 Honda Civic LX bought in December 2015 or January 2016 has the same allowable maturity and is assigned to the control group.

Figure 5: Distribution of Maturity Shock Timing
Panel A: Negative Maturity Shocks

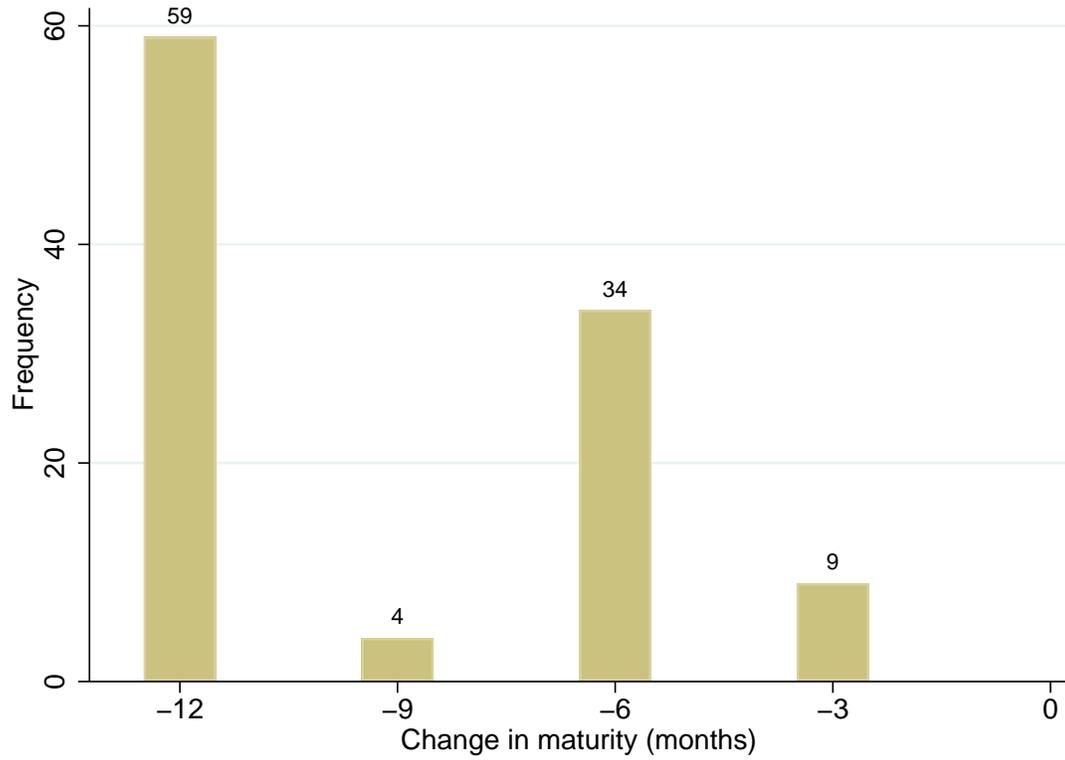


Panel B: Positive Maturity Shocks



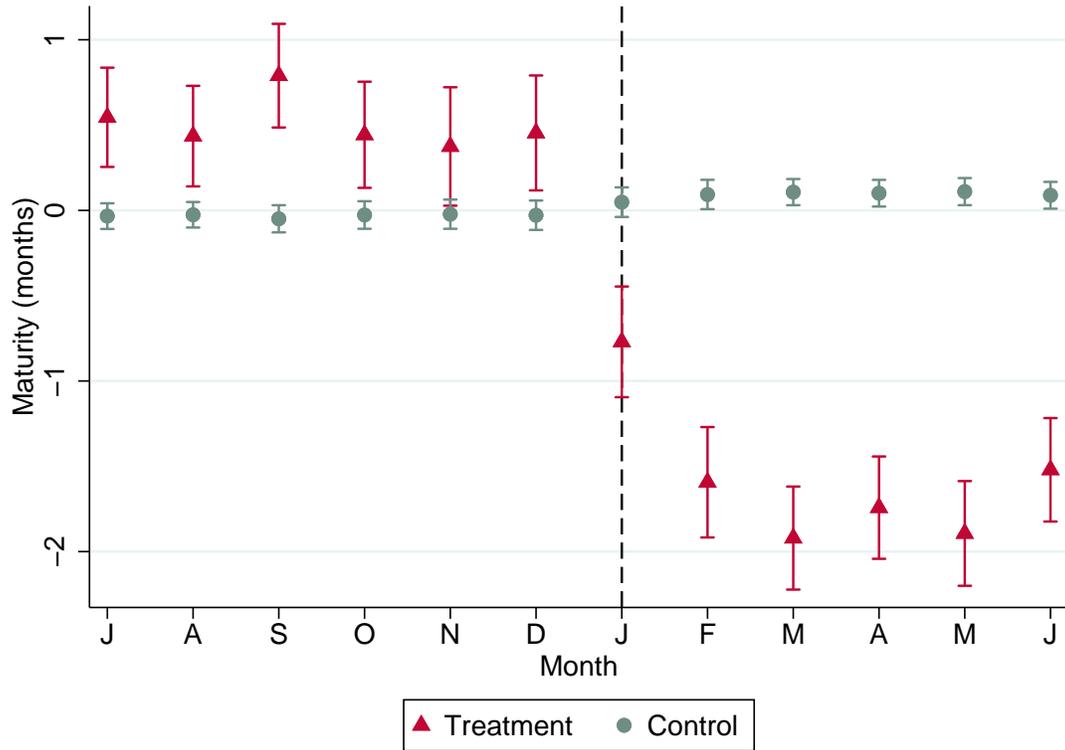
Notes: Figure plots the number of discrete changes in maturity identified in our data that occur in each month for both negative (panel A) and positive (panel B) shocks to maturity.

Figure 6: Distribution of Maximum Maturity Shocks by Size in Months



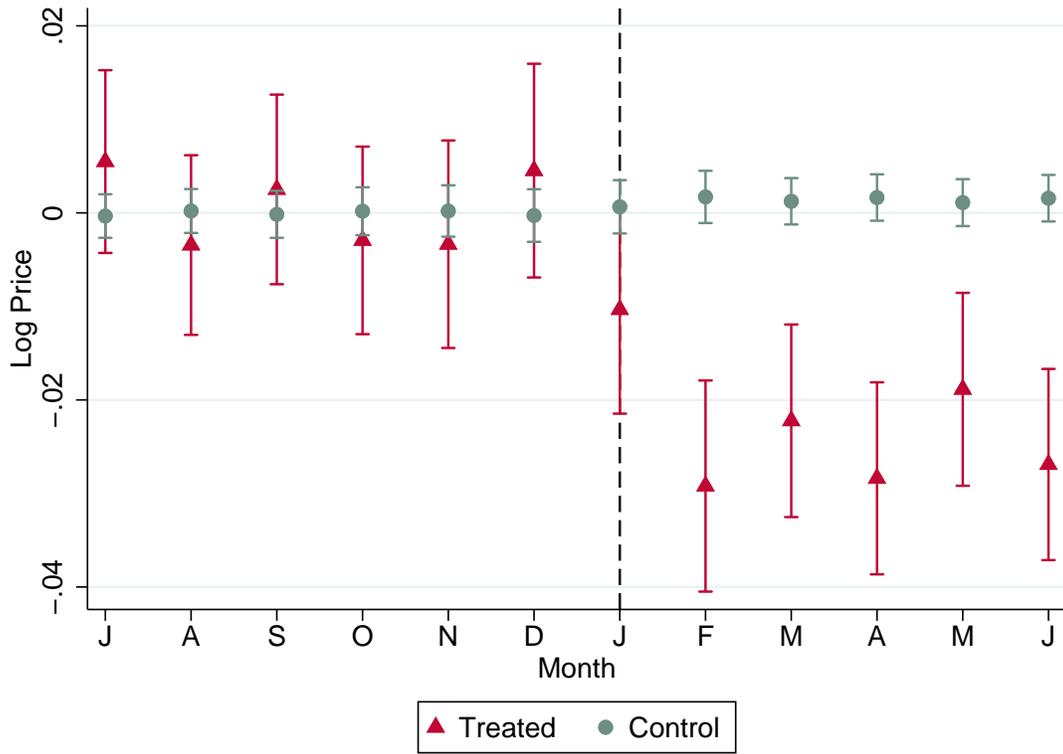
Notes: Figure plots number of occurrences of detected discontinuous maturity drops in January based on the size of the drop in maximum offered maturity.

Figure 7: Average Maturities Around Year End by Treatment



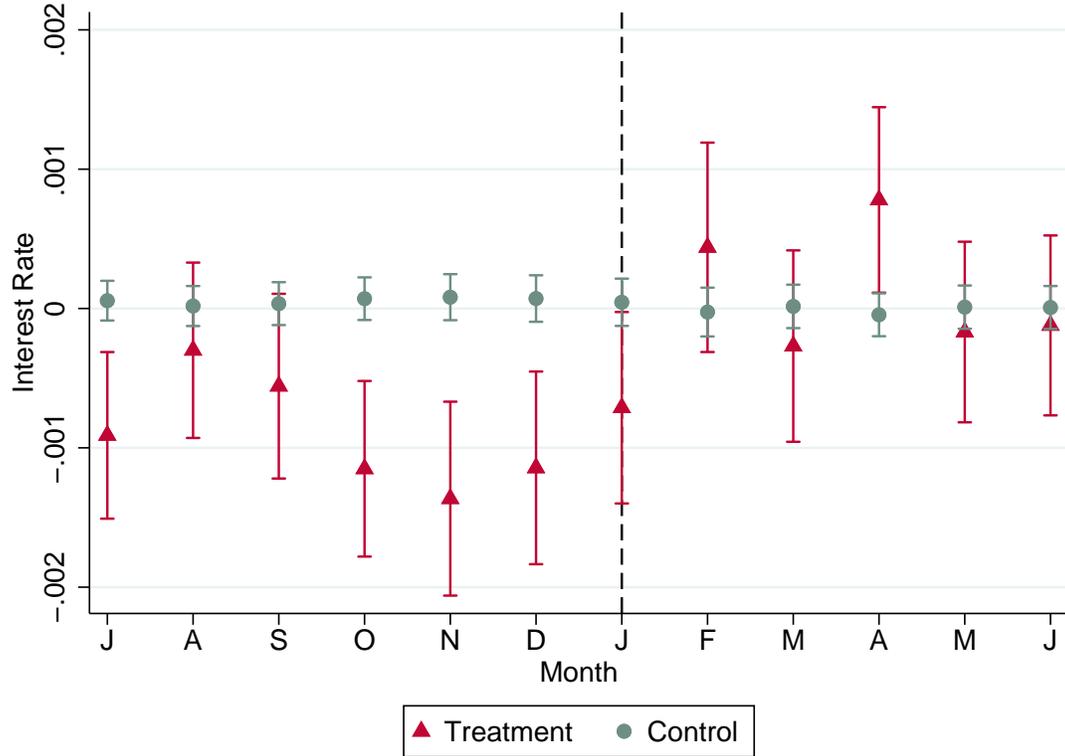
Notes: Figure plots the average conditional maturity by month around the year end for both treatment and control groups. We first regress maturity on car age \times month-of-sale fixed effects and commuting zone fixed effects and then plot average residuals within each month.

Figure 8: Average Price Around Year End by Treatment



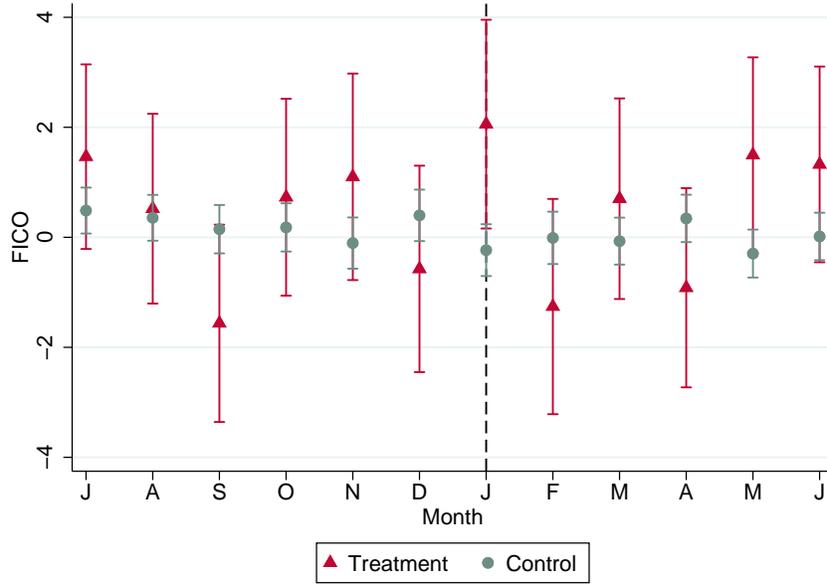
Notes: Figure shows average conditional log price around year end for both treatment and control groups. We first regress the log of car price on car age \times month-of-sale fixed effects and commuting zone fixed effects and then plot the average residuals within each month.

Figure 9: Average Interest Rate Around Year End by Treatment

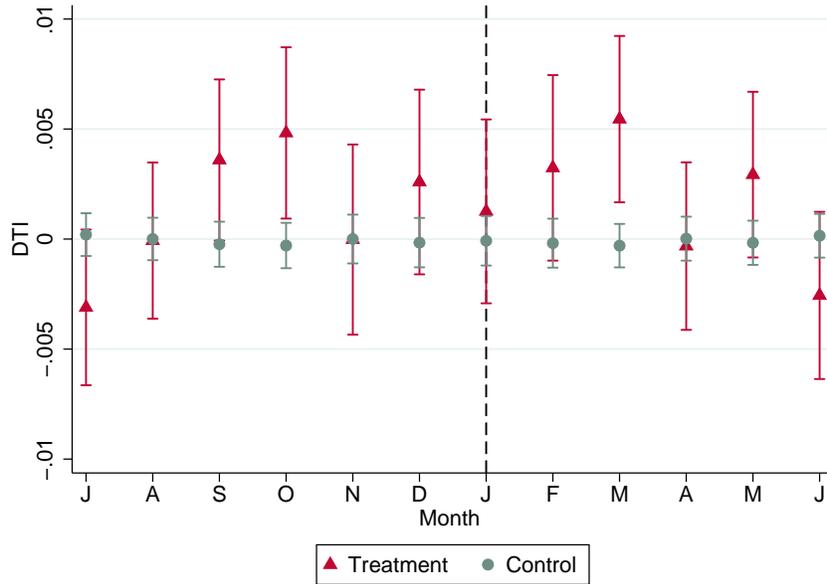


Notes: Figure shows the average conditional interest rate around the new year for both treatment and control groups. We first regress interest rate on car age \times month-of-sale fixed effects and commuting zone fixed effects and then plot the average residuals within each month.

Figure 10: Average Borrower Characteristics Around Year End by Treatment
 Panel A: FICO Scores



Panel B: Debt-to-Income Ratios



Notes: Figure plots an event study of average borrower characteristics by treatment and control group and month of year. Panel A plots average FICO scores, and panel B plots average debt-to-income ratios. We first regress each dependent variable on car age \times month-of-sale fixed effects and commuting zone fixed effects and then plot the average residuals within each month.

Table I: Summary Statistics

	Overall Estimation Sample (1)	Control Sample (2)	Treatment Sample (3)	Difference (2) - (3)
FICO	714.1 (69.0)	714.5 (69.0)	706.9 (68.2)	7.6 [0.3]
DTI	0.346 (0.256)	0.347 (0.259)	0.331 (0.199)	0.016 [0.001]
LTV	0.907 (0.222)	0.907 (0.222)	0.907 (0.217)	0.000 [0.001]
Car Age	3.88 (2.95)	3.86 (2.94)	4.29 (3.09)	-0.43 [0.01]
Car Price	20,341 (9,432)	20,432 (9,460)	18,821 (8,951)	1,611 [39.5]
Maturity	61.3 (12.8)	61.4 (12.8)	59.3 (12.4)	2.1 [0.05]
Interest Rate	0.0410 (0.0244)	0.0409 (0.0244)	0.0431 (0.0246)	-0.0022 [0.0001]
Observations	972,621	917,864	54,757	

Notes: Table shows means and standard deviations in parentheses for the overall estimation sample (1), alongside statistics for the control (2) and treatment (3) samples. The differences in means (2) - (3) are reported in the final column with standard errors in brackets. FICO is the credit score of the borrower as of the origination date of the loan; DTI is the back-end debt-to-income ratio of the borrower; LTV is the loan-to-value ratio for the vehicle being financed.

Table II: Maturity Distribution within Lender \times Car Age \times Month

Maturity (Relative to Max)	Percent of Borrowers	Cumulative Percent
0	16.5	16.5
[-6, 0)	9.2	25.7
[-12, -6)	17.2	43.0
[-18, -12)	10.3	53.2
[-24, -18)	14.8	68.0
< -24	32.0	100.0

Notes: The table shows the maturity distribution relative to the maximum maturity within lender \times car age \times month cells. Within each cell, we report the average percentage of borrowers who receive maturity within each six month band relative to the maximum maturity in that cell.

Table III: First-Stage Difference-in-Differences Results on Maturity

Maturity	(1)	(2)	(3)	(4)	(5)	(6)
Treatment \times Post	-2.404*** (0.664)	-2.390*** (0.303)	-2.021*** (0.277)	-2.157*** (0.304)	-2.284*** (0.271)	-2.290*** (0.265)
Treatment	-0.932 (0.955)	-0.395 (0.406)	-0.325 (0.403)	-0.371 (0.365)	0.561** (0.282)	0.368 (0.263)
Post	-0.872*** (0.197)	0.913*** (0.188)	0.754*** (0.102)			
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes
Car Age FE		Yes				
Car Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
Commuting Zone FE					Yes	Yes
Lender FE						Yes
Observations	972,621	972,621	972,621	972,621	972,621	972,621
R-squared	0.004	0.132	0.207	0.350	0.407	0.447

Notes: Table reports difference-in-differences regressions of loan maturities measured in months over an event year from July to June. Treatment is a dummy equal to one for loans that originate from a lender whose maximum maturity policy changed discontinuously for the transacted car on January 1, as discussed in section III. Post is a dummy equal to one for observations after January 1st. Borrower controls include FICO (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). *MMT* signifies combinations of make-model-trim; *YMMT* signifies combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table IV: Reduced-Form Difference-in-Differences Results on Log Price

log(Price)	(1)	(2)	(3)	(4)	(5)	(6)
Treatment \times Post	-0.026 (0.033)	-0.027*** (0.006)	-0.009*** (0.003)	-0.006** (0.003)	-0.007*** (0.003)	-0.007*** (0.002)
Treatment	-0.059 (0.050)	-0.025*** (0.008)	-0.009* (0.005)	-0.007 (0.006)	0.006 (0.005)	0.006 (0.005)
Post	-0.052*** (0.007)	0.061*** (0.006)	0.060*** (0.006)			
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes
Car Age FE		Yes				
Car Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
Commuting Zone FE					Yes	Yes
Lender FE						Yes
Observations	972,621	972,621	972,621	972,621	972,621	972,621
R-squared	0.060	0.369	0.872	0.909	0.911	0.914

Notes: Table reports difference-in-differences regressions results of log(prices) over an event year from July to June. Treatment is a dummy equal to one for loans that originate from a lender whose maximum maturity policy changed discontinuously for the transacted car on January 1, as discussed in section III. Post is a dummy equal to one for observations after January 1st. Borrower controls include FICO (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). *MMT* signifies combinations of make-model-trim; *YMMT* signifies combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table V: Two-Stage Least Squares Effects of Maturity on Log Price

log(Price)	(1)	(2)	(3)	(4)	(5)	(6)
Maturity	0.0109 (0.0112)	0.0112*** (0.0025)	0.0042*** (0.0014)	0.0027** (0.0012)	0.0032*** (0.0012)	0.0029*** (0.0010)
Post	-0.0423*** (0.0134)	0.0503*** (0.0055)	0.0566*** (0.0054)			
Treatment	-0.0486 (0.0469)	-0.0209** (0.0085)	-0.0079 (0.0049)	-0.0061 (0.0051)	0.0044 (0.0038)	0.0048 (0.0041)
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes
Car Age FE		Yes				
Car Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
Commuting Zone FE					Yes	Yes
Lender FE						Yes
Observations	972,621	972,621	972,621	972,621	972,621	972,621
R-squared	0.210	0.411	0.874	0.911	0.913	0.916

Notes: Table reports difference-in-differences regressions of log prices for an event year running from July to June using two-stage least squares. The excluded instrument is Treatment \times Post, where Post is a dummy equal to one for observations after January 1st. Treatment is a dummy equal to one for loans that originate from a lender that experienced a discontinuous policy change, as discussed in section III. Borrower controls include FICO (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). *MMT* signifies combinations of make-model-trim; *YMMT* signifies combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table VI: Difference-in-Differences Results on Interest Rates

Interest Rate	(1)	(2)	(3)	(4)	(5)	(6)
Treatment \times Post	0.0006 (0.0010)	0.0004 (0.0011)	0.0007 (0.0008)	0.0009 (0.0007)	0.0012* (0.0007)	0.0016*** (0.0005)
Treatment	0.0001 (0.0014)	-0.0002 (0.0017)	-0.0020* (0.0011)	-0.0030*** (0.0008)	-0.0009 (0.0005)	-0.0005 (0.0004)
Post	0.0002 (0.0007)	-0.0006 (0.0006)	-0.0001 (0.0002)			
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes
Car Age FE		Yes				
Car Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
Commuting Zone FE					Yes	Yes
Lender FE						Yes
Observations	972,621	972,621	972,621	972,621	972,621	972,621
R-squared	0.426	0.443	0.499	0.604	0.640	0.664

Notes: Table reports difference-in-differences regressions of loan interest rates over an event year from July to June. Treatment is a dummy equal to one for loans that originate from a lender whose maximum maturity policy changed discontinuously for the transacted car on January 1, as discussed in section III. Post is a dummy equal to one for observations after January 1st. Borrower controls include FICO (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). *MMT* signifies combinations of make-model-trim; *YMMT* signifies combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table VII: Two-Stage Least Squares Effects of Maturity and Rate on Log Price

log(Price)	(1)	(2)	(3)	(4)	(5)	(6)
Maturity	0.0044* (0.0025)	0.0074*** (0.0015)	0.0034*** (0.0007)	0.0022*** (0.0004)	0.0024*** (0.0004)	0.0023*** (0.0004)
Interest Rate	-0.365 (0.846)	-1.951*** (0.628)	-1.467*** (0.418)	-0.920*** (0.323)	-0.863*** (0.328)	-0.905*** (0.333)
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes
Car Age FE		Yes				
Car Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
Commuting Zone FE					Yes	Yes
Lender FE						Yes
Int. Rate 1st Stage p-value	0.000	0.000	0.000	0.000	0.000	0.000
Maturity 1st Stage p-value	0.000	0.000	0.000	0.000	0.000	0.000
Corr(Mat/Rate Coeffs)	-0.134	-0.162	-0.192	-0.160	-0.147	-0.187
Observations	972,621	972,621	972,621	972,621	972,621	972,621
R-squared	0.164	0.421	0.875	0.911	0.914	0.916

Notes: Table reports two-stage least-squares regressions of log transaction prices on loan maturity and interest rate. Excluded instruments are the interactions of Post with Treatment dummies that identify treated lender \times event age \times event year combinations (as discussed in section III). Borrower controls include FICO scores, DTI, and Treatment dummies. Corr(Mat/Rate Coeffs) is the correlation across cells of the first-stage coefficients for the maturity and interest rate instruments. *MMT* signifies combinations of make-model-trim; *YMMT* signifies combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table VIII: Repeat Sales Reduced-form Effects of Treatment on Log Price

	Second-sale log(Price) (1)	Initial-sale log(Price) (2)	Difference (1) - (2)
Treatment \times Post	0.006 (0.007)	-0.012 (0.010)	0.018* (0.011)
Treatment	-0.005 (0.004)	0.009 (0.010)	
YMMT \times Month FE	Yes	Yes	
CZ FE	Yes	Yes	
Lender FE	Yes	Yes	
Observations	8,697	8,697	
R-squared	0.001	0.001	

Notes: Table reports difference-in-differences regressions results of log price residuals for a subsequent second sale (1) and the initial sale (2) for those cars that we observe transacting twice in the entire dataset, at least 18 months apart. These pricing residuals \hat{u} are calculated from equation (5) by controlling for manufacture year-make-model-trim by month fixed effects, commuting zone, and lender fixed effects. Post is a dummy equal to one for observations for which the first sale occurred after January 1st. Treatment is a dummy equal to one for observations for which the first loan was originated by a lender whose maximum maturity policy changed discontinuously for the transacted car on January 1, as discussed in section III. *YMMT* signifies combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table IX: Reduced-form Difference-in-Differences Results on Days Spent Car Shopping

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment \times Post	-1.18 (1.03)	-2.11 (1.61)	-1.33 (1.38)	-2.11 (1.81)	-1.94 (1.75)	0.42 (1.70)
Treatment	-3.69** (1.73)	-4.12** (1.66)	-1.47 (1.85)	-1.94 (1.50)	-1.74 (1.55)	-1.05 (1.35)
Post	0.12 (0.68)	0.30 (0.64)	1.31 (1.07)	0.95 (1.12)	1.03 (1.12)	
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes
Car Age FEs		Yes		Yes	Yes	Yes
Lender FEs			Yes	Yes	Yes	Yes
CZ FEs					Yes	Yes
Loan Month FEs						Yes
Observations	54,929	54,929	54,929	54,929	54,929	54,929
R-squared	0.01	0.01	0.05	0.05	0.05	0.06

Notes: Table reports difference-in-differences regression results where the dependent variable measures the number of days elapsed between a loan application and loan origination. An event year runs from July to June, and the variable Post is a dummy equal to one for observations after January 1st. Treatment is a dummy equal to one for loans that originate from a lender whose maximum maturity policy changed discontinuously for the transacted car on January 1, as discussed in section III. Borrower controls include FICO (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Appendix

[Insert
Figures
A1 to
A2
about
here.]

[Insert
Tables
A1 to
AIV
about
here.]

A. Robustness to Maturity Discontinuities Classification Criteria

In Section A, we describe the tradeoffs that we face in classifying maturity policies and discontinuities. In the main results, we focused on discontinuities defined at any of the 70th, 80th, or 90th percentiles of the maturity distribution for a given lender by car age. However, we find that age-based maturity policies are reliably detectable at lower percentiles than those we focused on in the draft, down to at least the 40th percentile. Appendix Table AV shows the number of negative January discontinuities that we identify at each percentile, as well as the percentage of the negative discontinuities detected using that percentile that occurred in January. While the absolute maximum maturity in any given month is not a particularly reliable statistic (likely due to lenders' willingness to make exceptions to maturity rules in special cases), all percentiles from the 90th to the 40th show more than 30 January discontinuities and have a percentage of negative discontinuities that occur in January around 90%. For reference, Appendix Table AV shows 293 total January discontinuities; no other month had more than six.

[Insert
Table
AV
about
here.]

In Appendix Tables AVI and AVII, we revisit our main first-stage and reduced-form specifications

[Insert
Table
AVI
about
here.]

using discontinuities detected at any percentile from 50 to 90. Because discontinuities can occur in a given cell at more than one percentile, this results in 150 January discontinuities (vs. 106 in the main results). The sample is also slightly larger (just under 1.1 million vs. just under 1.0 million) and has a slightly larger percentage of treated observations (6.7% vs. 5.6%). The results are in all respects very similar to those in our main specifications.

[Insert
Table
AVII
about
here.]

B. Loan-to-value Ratios

Anecdotally, lenders often tie their LTV policies to the blue book value of a car as opposed to the price of each transaction. We do our best to imitate this practice by estimating the value associated with a given transaction as the average transaction price for cars of the same *YMMT* sold in the same state during the same half-year, excluding the vehicle in question. In cases for which there were no other vehicles of the same *YMMT* sold in the same state, we pool all transactions nationally. We pool transactions during a six-month period to maximize the number of transactions from which we estimate the average price without crossing the pre-post boundary of our event year. This will give us correct values on average in each pre- and post-period, though values will be overestimated in January and July and underestimated in June and December, with inverse effects on LTVs. Jackknifing LTVs in this way also helps to avoid any mechanical relationship between LTV and $Treatment \times Post$, as we have seen that prices are lower for affected borrowers. We report the results in Table AVIII. Consistent with tighter credit, we find that LTVs decrease by around two percentage points for treated borrowers across all specifications, from a baseline of around 90%. Note that, in the event that some lenders have LTV policies tied to the actual purchase price of each transaction, these results overstate the difference between treated and control LTVs, as our measurement of loan-to-book-value will mechanically result in lower LTVs among treated vehicles with lower purchase prices. Moreover, as discussed in Section A, LTV differences may also represent an endogenous response of borrowers who demand smaller loans when offered less favorable terms.

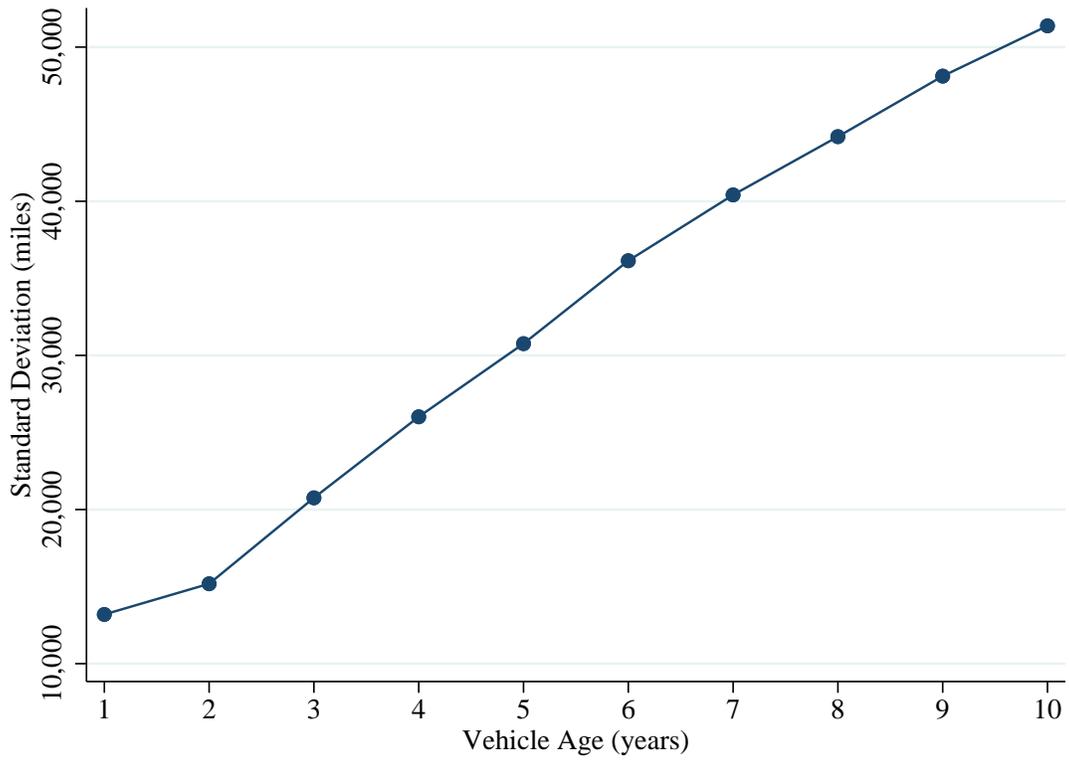
[Insert
Table
AVIII
about
here.]

To evaluate the effect of LTV on price, we estimate a two-stage least squares regression, instrumenting for maturity, interest rates, and LTVs under the assumption that these LTV results represent exogenous policy-determined variation and not an endogenous borrower response to shorter maturity. We report the results in Table AIX. As described in Section A, the instrument set in-

[Insert
Table
AIX
about
here.]

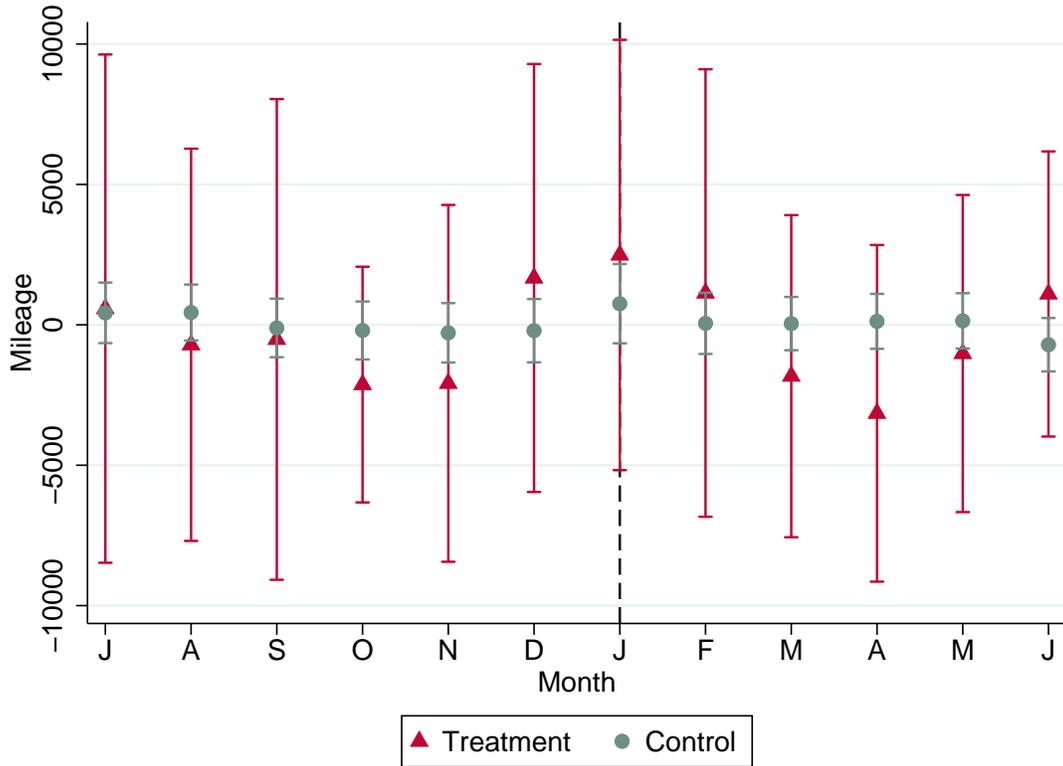
cludes differences in the magnitudes of maturity, interest rate, and LTV policy discontinuities. The results suggest that LTV is negatively related to price for the specifications with weaker fixed effects but becomes positively (but insignificantly) related to price with stronger fixed effects. Relative to Table VII, we cannot reject equivalence of coefficients on maturity nor interest rate for any column and the coefficient on maturity is about 90% as large as reported in Table VII. While we stress that we wish to be cautious in interpreting these results, the table suggests that LTV is not the major driver of the price-financing relationship in our dataset.

Figure A1: Mileage Variation by Car Age



Notes: Figure plots standard deviation of vehicle odometer readings in miles by vehicle age for used cars 1–10 years old observed in the 2017 National Household Travel Survey.

Figure A2: Average Mileage Around Year End by Treatment



Notes: Figure plots an event study of average mileage by treatment and control group and month of year for a subsample of cars sold in California for which we have odometer data. We first regress mileage on car age \times month-of-sale fixed effects and commuting zone fixed effects and then plot the average residuals within each month.

Table AI: Reduced-Form Difference-in-Differences Results for Log Price Using Hold-out Sample

log(Price)	(1)	(2)	(3)	(4)	(5)	(6)
Treatment \times Post	-0.0003 (0.025)	-0.031** (0.012)	-0.012*** (0.004)	-0.009* (0.005)	-0.009** (0.004)	-0.009*** (0.003)
Treatment	-0.101*** (0.037)	-0.026** (0.012)	-0.007 (0.009)	-0.008 (0.009)	0.006** (0.003)	0.008*** (0.002)
Post	-0.062*** (0.008)	0.055*** (0.005)	0.063*** (0.006)			
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes
Car Age FE		Yes				
Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
Commuting Zone FE					Yes	Yes
Lender FE						Yes
Observations	232,984	232,984	232,984	232,984	232,984	232,984
R-squared	0.070	0.325	0.885	0.923	0.925	0.926

Notes: Table reports difference-in-differences regressions results of log prices for an event year running from July to June using a different sample to define treatment and control than the one used in estimation. We randomly divide observations into two groups at the lender \times car age \times month-of-sale level. Using the first group to define treatment and control (see section III), we repeat the reduced-form analysis of prices from Table IV using only the hold-out sample. Borrower controls include FICO (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). Robust standard errors (in parentheses) are double clustered by month and commuting zone. See notes to Table IV for more details. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table AII: Linear-Probability Model of Observed Resale Selection

$\mathbb{I}(\text{Second Sale Observed})$	(1)	(2)	(3)	(4)	(5)	(6)
Treatment \times Post	0.0009 (0.0013)	0.0007 (0.0018)	0.0015 (0.0017)	0.0014 (0.0019)	0.0016 (0.0017)	0.0017 (0.0017)
Treatment	-0.0006 (0.0014)	-0.0012 (0.0019)	-0.0012 (0.0019)	-0.0014 (0.0019)	-0.0021 (0.0013)	-0.0026** (0.0013)
Post	0.0025*** (0.0006)	-0.0004 (0.0006)	-0.0014*** (0.0003)			
Car Age FE		Yes				
Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
Commuting Zone FE					Yes	Yes
Lender FE						Yes
Observations	963,612	963,612	963,612	963,612	963,612	963,612
R-squared	0.001	0.006	0.0256	0.190	0.198	0.199

Notes: Table reports linear-probability model difference-in-differences regressions of a dummy for whether a given vehicle was observed being resold in our sample. Borrower controls include *FICO* (the credit score of the borrower at loan origination) and *DTI* (the back-end debt-to-income ratio of the borrower at origination). Sample includes the first sale per VIN from the main estimation sample subject to observing at least two sales per *YMMT* category. Robust standard errors (in parentheses) are double clustered by month and commuting zone. See notes to Table VIII for more details. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table AIII: Two-Stage Least Squares Effects of Maturity and Rate on Log Price by Vehicle Age

log(Price)	0–5 year old cars		6+ year old cars	
	(1)	(2)	(3)	(4)
Maturity	0.0023** (0.0011)	0.0025** (0.0011)	0.0022*** (0.0006)	0.0019*** (0.0006)
Interest Rate	0.3552 (1.0143)	0.5287 (1.0523)	-1.1577*** (0.4177)	-1.2614*** (0.4129)
YMMT × Month FE	Yes	Yes	Yes	Yes
Commuting Zone FE	Yes	Yes	Yes	Yes
Lender FE		Yes		Yes
Observations	523,648	523,648	448,943	448,943
R-squared	0.907	0.909	0.870	0.876

Notes: Table reports 2SLS regressions of log transaction prices on maturity and rates for used cars less than five years old (column 1) and five or more years old (column 2). Excluded instruments are the interactions of Post with Treatment dummies that identify treated lender × event age × event year combinations (as discussed in section III). Borrower controls include FICO scores, DTI, and Treatment dummies. Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** p<0.01, ** p<0.05, * p<0.1

Table AIV: Price Effects Adjusted for Potential Omitted Variable Bias

log(Price)	(1)	(2)	(3)	(4)	(5)	(6)
Estimated Coefficient	-0.026	-0.027	-0.009	-0.006	-0.007	-0.007
Omitted Variables Bias-Adjusted	-0.068	-0.071	-0.029	-0.010	-0.009	-0.007
Includes Zero?	No	No	No	No	No	No
Adjusted Coefficient within Original Confidence Interval?	No	No	Yes	Yes	Yes	Yes
Car Age FE		Yes				
Car Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
Commuting Zone FE					Yes	Yes
Lender FE						Yes

Notes: Table reports coefficients adjusted for omitted variable bias for the reduced-form price regressions performed in III. The first row reports the uncorrected coefficients found in Table III. The second row reports coefficients adjusted for potential omitted variable bias following Oster (2017). The third row reports whether the interval between the unadjusted and adjusted coefficients includes zero. Finally, the fourth row reports whether the adjusted coefficient is included in the 95% confidence interval of the observed regression. *MMT* signifies combinations of make-model-trim; *YMMT* signifies combinations of manufacture year-make-model-trim.

Table AV: January Discontinuities by Percentile

Quantile	Count of January Discontinuities	January Share of Negative Discontinuities (%)
Max	12	70.6
P90	40	93
P80	38	95
P70	47	87
P60	49	90.7
P50	37	94.9
P40	32	91.4
P30	22	81.5
P20	13	61.9
P10	3	25

Notes: Table shows the number of negative January discontinuities that we identify at each decile, as well as the percentage of the negative discontinuities detected using that quantile that occurred in January. All percentiles from the 90th to the 40th show more than 30 January discontinuities and have a percentage occurring in January near or above 90%. For reference, the table shows 293 total January discontinuities; no other month had more than six.

Table AVI: Robustness of Difference-in-Differences Maturity Results to Alternative Discontinuity Detection Criteria

Maturity	(1)	(2)	(3)	(4)	(5)	(6)
Treatment \times Post	-2.423*** (0.476)	-2.312*** (0.275)	-1.971*** (0.235)	-2.072*** (0.272)	-2.185*** (0.221)	-2.216*** (0.253)
Treatment	-0.389 (1.098)	-0.421 (0.423)	-0.392 (0.413)	-0.421 (0.371)	0.720** (0.326)	0.686** (0.336)
Post	-0.792*** (0.176)	0.971*** (0.184)	0.827*** (0.094)			
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes
Car Age FE		Yes				
Car Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
Commuting Zone FE					Yes	Yes
Lender FE						Yes
Observations	1,091,578	1,091,578	1,091,578	1,091,578	1,091,578	1,091,578
R-squared	0.003	0.127	0.202	0.340	0.398	0.437

Notes: Table reports difference-in-differences regressions of loan maturities measured in months over an event year from July to June. Post is a dummy equal to one for observations after January 1st. Treatment is a dummy equal to one for loans that originate from a lender whose maximum maturity policy changed discontinuously for the transacted car on January 1, as discussed in section III. Treated observations are those with discontinuities at any of the 50th, 60th, 70th, 80th, or 90th percentile of the maturity distribution, whereas Table III uses maturity discontinuities at only the 70th, 80th, or 90th percentiles. Borrower controls include FICO (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). *MMT* signifies combinations of make-model-trim; *YMMT* signifies combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table AVII: Robustness of Difference-in-Differences Price Results to Alternative Discontinuity Detection Criteria

log(Price)	(1)	(2)	(3)	(4)	(5)	(6)
Treatment \times Post	-0.033 (0.025)	-0.024*** (0.005)	-0.008*** (0.001)	-0.007*** (0.001)	-0.009*** (0.003)	-0.009*** (0.003)
Treatment	-0.024 (0.055)	-0.025*** (0.007)	-0.016** (0.007)	-0.014* (0.008)	0.006 (0.005)	0.008 (0.006)
Post	-0.050*** (0.005)	0.060*** (0.006)	0.059*** (0.005)			
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes
Car Age FE		Yes				
Car Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
Commuting Zone FE					Yes	Yes
Lender FE						Yes
Observations	1,091,578	1,091,578	1,091,578	1,091,578	1,091,578	1,091,578
R-squared	0.057	0.365	0.868	0.905	0.908	0.911

Notes: Table reports difference-in-differences regressions results of log(prices) over an event year from July to June. Post is a dummy equal to one for observations after January 1st. Treatment is a dummy equal to one for loans that originate from a lender whose maximum maturity policy changed discontinuously for the transacted car on January 1, as discussed in section III. Treated observations are those with discontinuities at any of the 50th, 60th, 70th, 80th, or 90th percentile of the maturity distribution, whereas Table IV uses maturity discontinuities at only the 70th, 80th, or 90th percentiles. Borrower controls include FICO (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). *MMT* signifies combinations of make-model-trim; *YMMT* signifies combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table AVIII: Difference-in-Differences Results on LTVs

Loan-to-Value	(1)	(2)	(3)	(4)	(5)	(6)
Treatment \times Post	-0.022*** (0.004)	-0.023*** (0.003)	-0.020*** (0.003)	-0.024*** (0.004)	-0.023*** (0.004)	-0.021*** (0.004)
Treatment	0.004 (0.009)	0.006 (0.009)	0.006 (0.009)	0.010 (0.008)	0.012*** (0.003)	0.010*** (0.003)
Post	-0.002 (0.005)	-0.000 (0.005)	-0.005 (0.004)			
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes
Car Age FE		Yes				
Car Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
CZ FE					Yes	Yes
Lender FE						Yes
Observations	972,621	972,621	972,621	972,621	972,621	972,621
R-squared	0.007	0.014	0.053	0.224	0.255	0.270

Notes: Table reports difference-in-differences regressions of loan-to-value ratios over an event year from July to June. Post is a dummy equal to one for observations after January 1st. Treatment is a dummy equal to one for loans that originate from a lender whose maximum maturity policy changed discontinuously for the transacted car on January 1, as discussed in section III. Borrower controls include FICO (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). *MMT* signifies combinations of make-model-trim; *YMMT* signifies combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table AIX: Two-Stage Least Squares Effects of Maturity, Rate, and LTV on Log Price

log(Price)	(1)	(2)	(3)	(4)	(5)	(6)
Maturity	0.0102** (0.0040)	0.0090*** (0.0019)	0.0035*** (0.0010)	0.0020*** (0.0006)	0.0021*** (0.0006)	0.0020*** (0.0006)
Interest Rate	-1.4475** (0.6463)	-2.1823*** (0.6001)	-1.4860*** (0.4434)	-0.8581*** (0.3004)	-0.7828** (0.3016)	-0.8324*** (0.3065)
Loan to Value	-0.7908** (0.3183)	-0.2432* (0.1360)	-0.0139 (0.1253)	0.0346 (0.0637)	0.0429 (0.0654)	0.0386 (0.0642)
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes
Car Age FE		Yes				
Car Age \times MMT FE			Yes			
YMMT \times Month FE				Yes	Yes	Yes
Commuting Zone FE					Yes	Yes
Lender FE						Yes
Observations	972,621	972,621	972,621	972,621	972,621	972,621
R-squared	0.050	0.409	0.871	0.913	0.917	0.919

Notes: Table reports two-stage least-squares regressions of log transaction prices on loan maturity, interest rate, and LTV. Excluded instruments are the interactions of Post with Treatment dummies that identify treated lender \times event age \times event year combinations (as discussed in section III). Borrower controls include FICO scores, DTI, and Treatment dummies. *MMT* signifies combinations of make-model-trim; *YMMT* signifies combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$