

The Capitalization of Consumer Financing into Durable Goods Prices

BRONSON ARGYLE, TAYLOR NADAULD, CHRISTOPHER PALMER,
and RYAN PRATT*

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% to 90% of the monthly payment increase. Most of this is driven by shifting toward lower-quality cars, but affected borrowers offset 20% to 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 car's price by 2.8%.

THE RELATIONSHIP BETWEEN CREDIT AND asset prices has been an important area of study in postcrisis finance. Recent work, which focuses 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 that is driven by differences in borrowers' access to

*Bronson Argyle is with the Marriott School of Business, Brigham Young University. Taylor Nadauld is with the Marriott School of Business, Brigham Young University. Christopher Palmer is with the Sloan School of Management, MIT, and NBER. Ryan Pratt is with the Marriott School of Business, Brigham Young University. Bronson Argyle, Taylor Nadauld, Christopher Palmer, and Ryan Pratt have read *The Journal of Finance*'s disclosure policy and have no conflicts of interest to disclose. 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; John Mondragon; 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, and the SFS Cavalcade; and Natalie Bachas, Effi 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.

Correspondence: Bronson Argyle, Brigham Young University, 689 TNRB, Provo, UT 84602;
e-mail: bsa@byu.edu.

DOI: 10.1111/jofi.12977

© 2020 the American Finance Association

credit, which isolates 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 unit price and the quantity purchased of the good. Conceptually, this corresponds to the standard interpretation in 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 supply and demand framework suggests that affected borrowers would respond to a decrease in available credit by substituting toward lower-quality cars, an effect that 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 selected that car because of a credit shock or not. In contrast, we find that 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 only 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 in effect cars are treated as if they fully 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, maturity policies that feature such 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. The pre- and postevent dates together with the 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.

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

² Englmaier, Schmöller, and Stowasser (2018) present evidence that in the European used-car market, cars are treated as aging relative to the year in which they were first registered.

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 using a battery of manufacture year-make-model-trim (“YMMT”) \times purchase-month fixed effects (e.g., 2014 Honda Accord LXs purchased in February 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. In particular, 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 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 documents 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) in our data set has as large an impact on the monthly payment of an average vehicle as would a seven-percentage-point increase in interest rate. Still, we are interested in the effect of credit terms more broadly, and hence 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 variation in interest rates. Using a two-stage least squares (2SLS) procedure in which we simultaneously instrument for both maturity and interest rates, we find that maturity accounts for 70% to 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, which corresponds to an elasticity of price with respect to monthly payment size of -0.19. This finding implies a lower bound on the consumer discount rate needed to rationalize these results of around 11.6%.⁴

³ To assess the magnitude of this estimate, note that estimates of gross margins on used cars are 5% to 20% (Gavazza, Lizzeri, and Roketskiy (2014), Huang, Luo, and Xia (2019), and Larsen (2020)).

⁴ By contrast, Busse, Knittel, and Zettelmeyer (2012) estimate that in the context of fuel economy, the average buyer acts as if she has a discount rate roughly equal to her loan interest rate.

We further estimate that for a one-percentage-point increase in interest rate, 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 irrespective of whether that change came from interest rate or maturity variation is consistent with monthly payments playing a central role in borrowers' decision making. Finally, we find that for the same *YMMT* in the same month, borrowers with a 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 choose 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 II. In one test, we examine the subset of our data that consist 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. To further address concerns about unobserved heterogeneity in car quality, we examine 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 (2019) 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? Commuting zone fixed effects rule out the possibility of 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.

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

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 off load some of the cost of poor financing terms onto the seller, effectively leveraging their lower reservation price to negotiate a better transaction price.

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* a 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, which suggests that the lower transaction prices are likely the result of more successful negotiations by treated borrowers.

While our data set focuses 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 our conceptual motivation. Section II discusses the data set. Section III outlines the empirical strategy that we employ, while Section IV presents the results and mechanism. Section V concludes.

I Related Literature and Conceptual Motivation

A Related Literature

The empirical literature studying the causal link between credit and prices has concentrated on the housing market. See, for example, Mian and Sufi (2009), Glaeser, Gottlieb, and Gyourko (2012), Adelino, Schoar, and Severino (2012), Favara and Imbs (2015), Landvoigt, Piazzesi, and Schneider (2015), Zevelev (2020), Bhutta and Ringo (2017), Di Maggio and Kermani (2017), Verner and Gyöngyösi (2020), and Davis et al. (2020), none of whom study loan maturity.⁷ Favara and Imbs (2015) exploit state-level exposure to bank

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

⁷ A recent empirical macro literature also studies the causes and consequences of credit shocks for house prices (Jordà, Schularick, and Taylor (2015)), price-discrimination markups (Cornia, Gerardi, and Shapiro (2012)), 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 (2017) for a survey of recent work on credit-driven business cycles.

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 antipredatory 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 (2012) 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 also 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 the 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, this is 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) show 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 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 that 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

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

⁹ Hertzberg, Lieberman, and Paravisini (2018) argue the shared sentiment that the role of maturity has been understudied relative to interest rates in this literature.

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

(1991) and Müller 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 firms' financial flexibility, which strengthens a firm's bargaining position. We find 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 look 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 the capitalization of cost shocks into vehicle prices. For example, Busse, Silva-Risso, and Zettelmeyer (2006) examine the effects of manufacturer cash rebates for new cars and show 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% to 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 and contribute the lack of pass-through to Toyota's concerns about future demand given the dynamics of buyer price beliefs. Busse, Knittel, and Zettelmeyer (2012) find that resale prices capitalize exposure to gasoline taxes.¹² We complement this literature by studying the transmission of credit supply shocks with cross-sectional identification, and we further emphasize that disaggregate credit shocks can have disaggregate price effects.

Finally, we highlight the contribution of this paper relative to Argyle, Nadauld, and Palmer (2020a) and Argyle, Nadauld, and Palmer (2020b). Although each of these papers uses the same broad set of auto data, the three papers differ in economic focus. Argyle, Nadauld, and Palmer (2020b) document the relevance 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 (2020a) document that even relatively unconstrained borrowers make loan decisions based on monthly payment amounts. In particular, many consumers anchor

¹¹ See 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.

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

around hundred-dollar multiples (e.g., \$200, \$300, \$400), which supports the conjecture that cognitive frictions like mental accounting influence consumers' high-stakes debt decisions.

The present paper builds on the other papers in this literature 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 to this finding, they do not imply the cross-sectional pricing effects of individual credit shocks that are the contribution of this paper. Our empirical strategy also differs from these papers in that we exploit vehicle-level discontinuities in offered maturity across the car-age space, while the companion papers exploit borrower-level discontinuities across the credit score spectrum.

B Conceptual Motivation

The central question we examine in this paper is how do 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: 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 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 meet, 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 hence bargaining ensues, with the negotiated price dividing the bilateral surplus. Costly search therefore 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 in which 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 off-load some of the cost onto the seller by negotiating a lower price. We aim to quantify the extent to which borrowers substitute versus pay lower prices for equivalent goods as well as 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 (2020b) 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 observation. 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 data set that consists of over four million auto loans originated by 372 unique lenders covering all 50 states. The data include only those loans originated directly through the lending institution. The data exclude 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 case, 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 case, 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 the financing terms.

Our sample includes loans originated between 2005 and 2017, though over 80% of the sample loans were originated between 2011 and 2017. The growth in originations over time is driven mostly by growth in our data provider's client base, although it also partly reflects increased reporting of loan originations within lender over time as our data provider's products became more integral to the lenders' businesses. Moreover, aggregate auto loan originations have

Table I
Summary Statistics

The table shows means and standard deviations in parentheses for the overall estimation sample (1), as well as 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. Credit Score 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.

	Overall Estimation Sample (1)	Control Sample (2)	Treatment Sample (3)	Difference (2) - (3)
Credit Score	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	

increased substantially over our sample period, with outstanding auto debt in the United States increasing 56% between 2010 and 2017. Similar data are used in Argyle, Nadauld, and Palmer (2020a, 2020b).

The data set, anonymized of any personally identifiable information, includes loan contract features such as 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 vehicle's manufacture year, make, model, and trim (*YMMT*). Borrower information includes credit 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 × car age × month cell, as described 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 × car age × month cells with the fewest observations). We lose an additional quarter of the remaining observations that are missing vehicle trim information, which we use in constructing *YMMT* fixed effects that hold observable vehicle quality fixed.

These two restrictions leave us with a final data set of 972,621 loans originated by 112 unique lenders. Table I reports summary statistics broken down 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 do 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, the average car in our sample is 3.9 years old and sold for \$20,341. The average loan-to-value (LTV) ratio is 90.7%, with the average loan having a maturity of 61.3 months and an interest rate of 4.1%. We compare treatment and control groups after defining them in Section III.A.

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? Moreover, 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 addressing 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 that are 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 such a setting by exploiting maturity rules used by lenders that are based on the age of cars. Conversations with lenders suggest that the maximum maturity borrowers are offered on an auto loan is frequently a function of car age, a practice motivated by lenders' risk management concerns. A longer maturity increases the likelihood that the loan balance will exceed the collateral value during the life of a loan, exposing lenders to a loss in the case of default. For older cars with shorter remaining expected life, lenders offer a 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 Figure 2 illustrates. 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 become 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 car 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

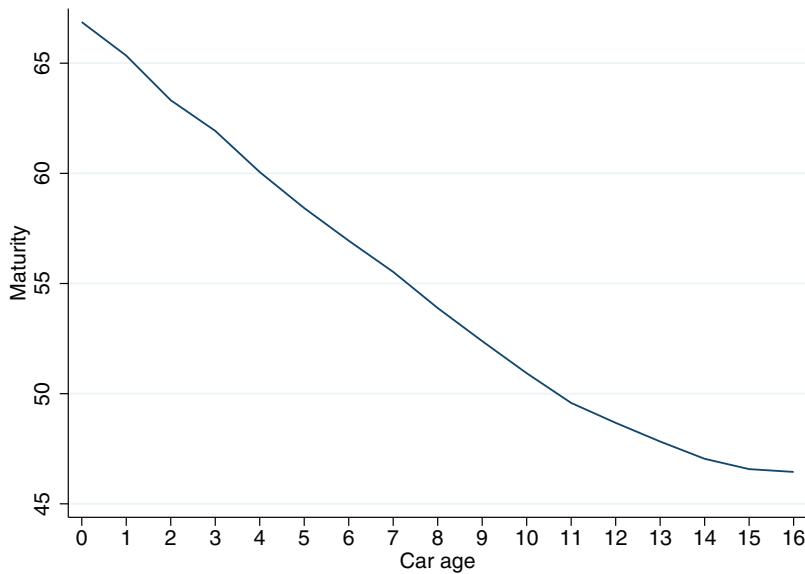


Figure 1. Average maturity by car age. The figure plots the average maturity of an auto loan in our data as a function of the age of the underlying collateral. (Color figure can be viewed at wileyonlinelibrary.com)

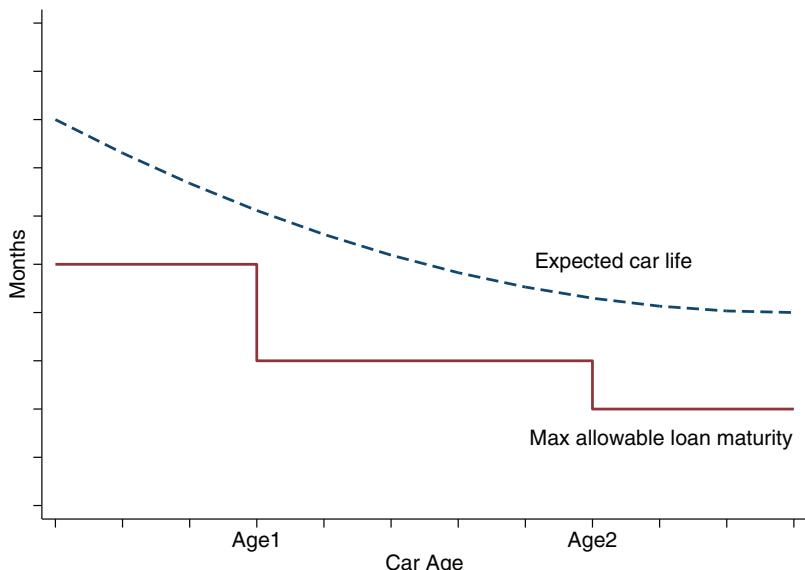


Figure 2. Hypothetical lender maximum maturity rule. The figure depicts a hypothetical discontinuous lender rule. While the conditional expected remaining life of a car decreases smoothly with car age, many lenders have discrete maximum allowable loan maturity rules that discontinuously decrease maximum allowable maturity once a car age reaches certain cutoffs. (Color figure can be viewed at wileyonlinelibrary.com)

institution without such a discrete drop for the car age being purchased) 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.

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 data set is that we can hold the quality of the goods fixed 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 the majority of the variation in car quality, as we discuss the interpretation of our tests, we take care to address the possibility that our results are affected by 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 that we face is the fact 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 with respect 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

Table II
Maturity Distribution within Lender × Car Age × Month

The table shows the maturity distribution relative to the maximum maturity within lender × car age × 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.

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

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 × car age varies a lot from month to month—the standard 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 at 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 Section II of the Internet Appendix (see Internet Appendix Tables IA.V and IA.VI) that we can reliably detect age-based maturity policies at lower percentiles and our results are not sensitive to the set chosen.¹³

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. 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

¹³ The Internet Appendix may be found in the online version of this article.

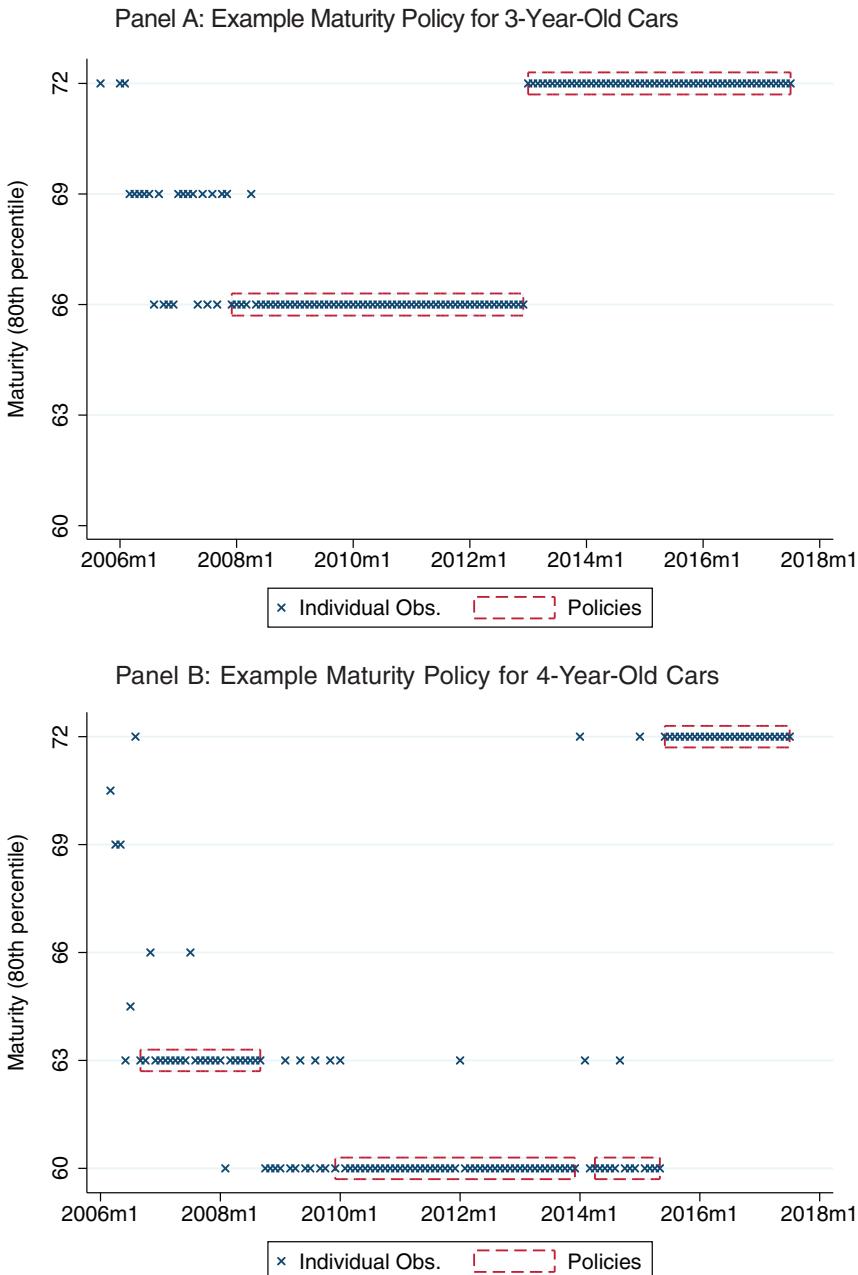


Figure 3. Example-lender maturity policies. The figure plots the maturity policies for the largest lender in our data. Individual observations x capture the 80th percentile of maturity for a given car age within a given month at this lender. Boxed areas correspond to the maturity policies identified by our algorithm (see Section III.A for details). Panels A and B show the policies for three- and four-year-old cars, respectively. (Color figure can be viewed at wileyonlinelibrary.com)

be part of a lender maximum maturity policy.¹⁴ 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 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.¹⁵

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 maturity at this particular lender.

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. Now consider the example of a 2012 Honda Civic LX 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.¹⁶

¹⁴ We require that the endpoints of the 13-month window 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 (2020a)).

¹⁵ The 80th percentile of loan maturities is more volatile in the early part of our sample because there are fewer loans during that period. The coverage of our data provider improves over the early part of our sample, even within lender.

¹⁶ Of 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 such cases as treated since they display a maturity shock.

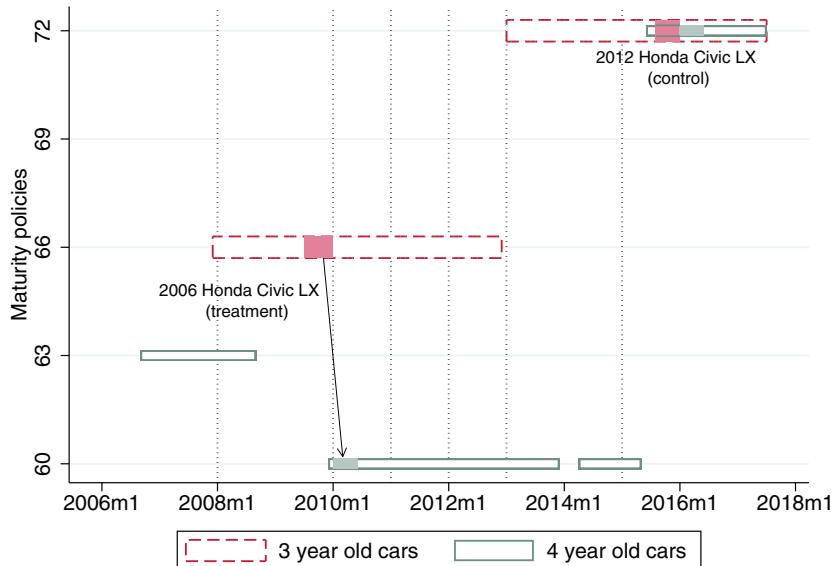


Figure 4. Example-lender maturity shocks. The figure illustrates our empirical design by plotting treatment- and control-group observations against identified maturity policies for the example lender used in Figure 3. Dashed rectangles correspond to the lender's maturity policy for three-year-old cars, 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 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. (Color figure can be viewed at wileyonlinelibrary.com)

Our requirement that we observe a stable maturity percentile for at least one year to infer the 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 the effectiveness of our approach at isolating exogenous variation in the supply of maturity arising from age-based maturity policies. To the extent that we correctly identify 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 mistakenly pick 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 by 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 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 that we identify represent actual policy updates versus false positives, at worst 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 reevaluate and update their maturity policies disproportionately 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 examine 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 highly persistent, we include all months from July through December in the preperiod and all months from January through June in the postperiod, although monthly event studies 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 as defined by our 106 January discontinuities. Table I reports summary statistics for the treatment and control samples. The groups are very similar on observables, including credit scores at origination, DTI ratios, and LTV ratios. Although the

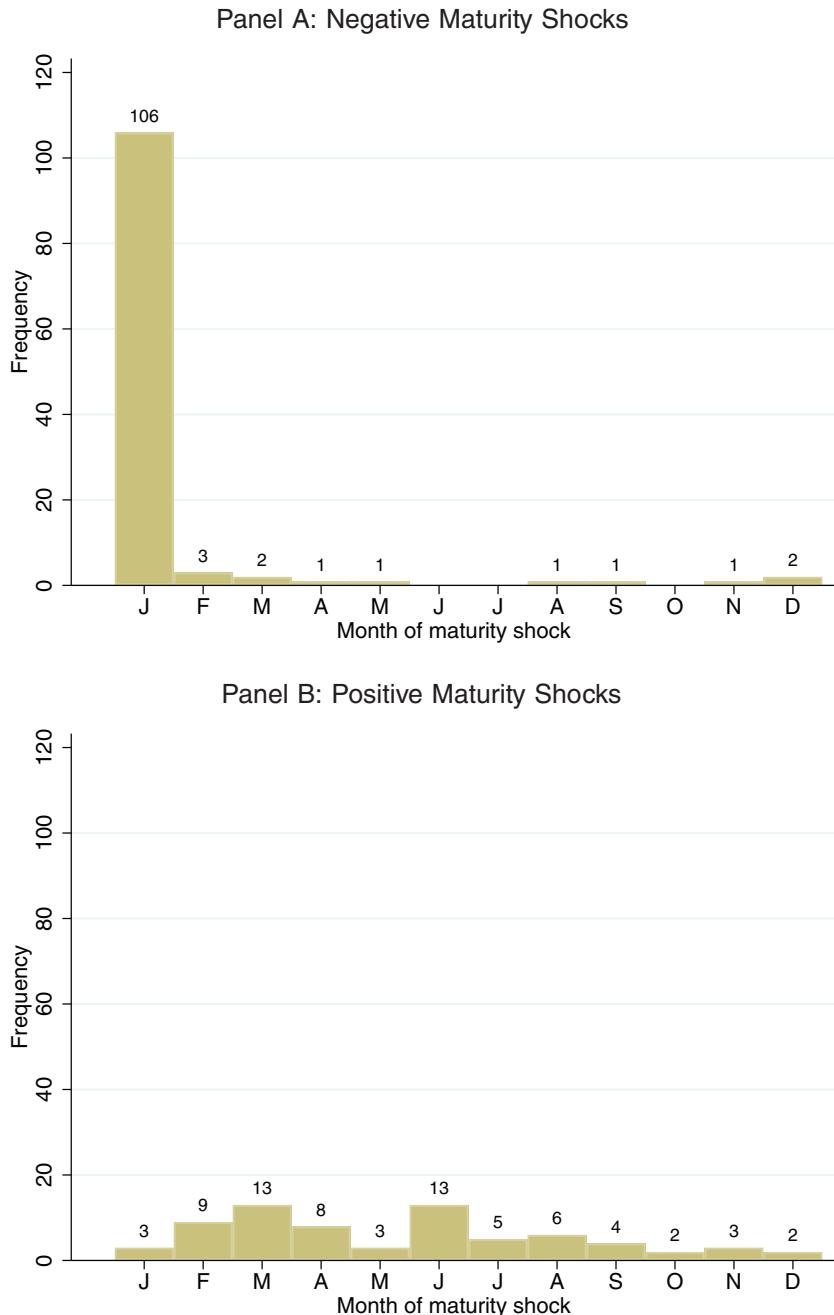


Figure 5. Distribution of maturity shock timing. The figure plots the number of discrete changes in maturity identified in our data that occur in each month for both negative (Panel A) and positive (Panel B) shocks to maturity. (Color figure can be viewed at wileyonlinelibrary.com)

differences in means are mostly statistically significant owing to the precision afforded by our large sample size, the typical difference is a 10th of a standard deviation, suggesting that treatment- and control-group borrowers are balanced for practical purposes. Consistent with being slightly older (4.29 versus 3.86 years), treated cars have slightly lower prices (\$18,821 versus \$20,432), shorter maturities (59.3 months versus 61.4 months), and higher interest rates (4.31% versus 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 postperiod.

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 rather from borrowers buying a car that has recently crossed a discontinuity in a lender's maturity policy. We estimate¹⁷

$$\text{Maturity}_{iglt} = \beta_1 \text{Treatment}_i + \beta_2 \text{Post}_t + \beta_3 \text{Treatment}_i \times \text{Post}_t + X'_{it} \gamma + \varphi_g + \psi_l + \varepsilon_{iglt}, \quad (1)$$

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. The treatment variable is a dummy equal to one for observations within a $\text{Lender} \times \text{EventAge} \times \text{EventYear}$ 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 we define EventYear as the calendar year of that January. Controls X_{it} consist of borrower characteristics (DTI and credit score) and various fixed effects that control for the quality of collateral, such as $\text{YMMT} \times \text{month}$ fixed effects. In some specifications, we also control for commuting zone fixed effects φ_g and lender fixed effects ψ_l . We double-cluster 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 $\text{Treatment} \times \text{Post}$ coefficient unchanged, suggesting that

¹⁷ Using the high-dimension fixed effect estimator of Correia (2017).

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

The 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 1. Borrower controls include Credit Score (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). *MMT* indicates combinations of make-model-trim; *YMMT* indicates combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double-clustered by month and commuting zone.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

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 ²	0.004	0.132	0.207	0.350	0.407	0.447

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 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 holds across all specifications, showing a stable effect on originated maturities of slightly more than two months.

As indicated above, auto loan maturities cluster significantly on multiples of 3, 6, or 12 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 a shorter maturity than the maximum allowable. Borrowers that demand loan maturities

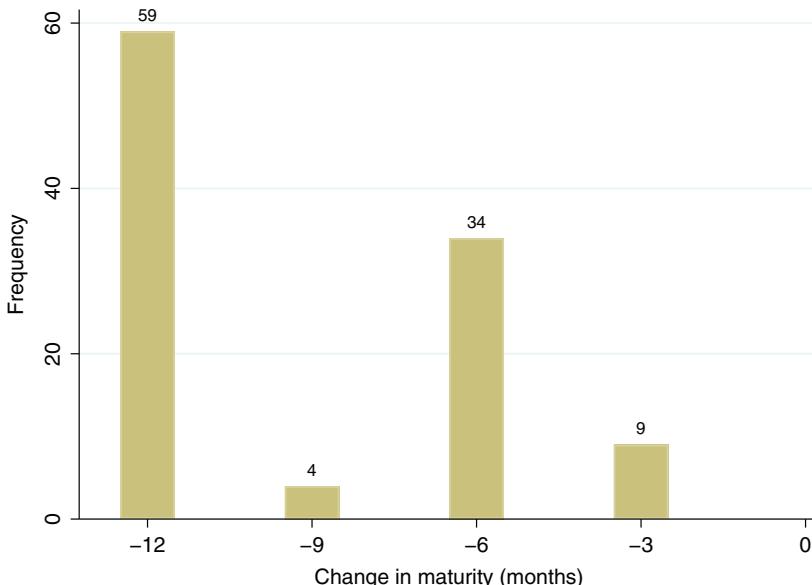


Figure 6. Distribution of maximum maturity shocks by size in months. The figure plots the number of occurrences of detected discontinuous maturity drops in January based on the size of the drop in maximum offered maturity. (Color figure can be viewed at [wileyonlinelibrary.com](#))

lower than the maximum allowable could be unaffected by changes in maturity policy. Our instrumental variables strategy below is designed precisely to address 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.

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.¹⁸ The figure shows stable maturities for the control group throughout the event year. The treatment group, in contrast, has stable maturities that are slightly higher than those of the control group from July through December, followed by a sharp drop in January that continues to February. Maturities in February through June are stable and significantly lower than those in the control group. It is difficult to say why exactly the drop in maturities spans January and February, but it may be 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

¹⁸ Specifically, 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.

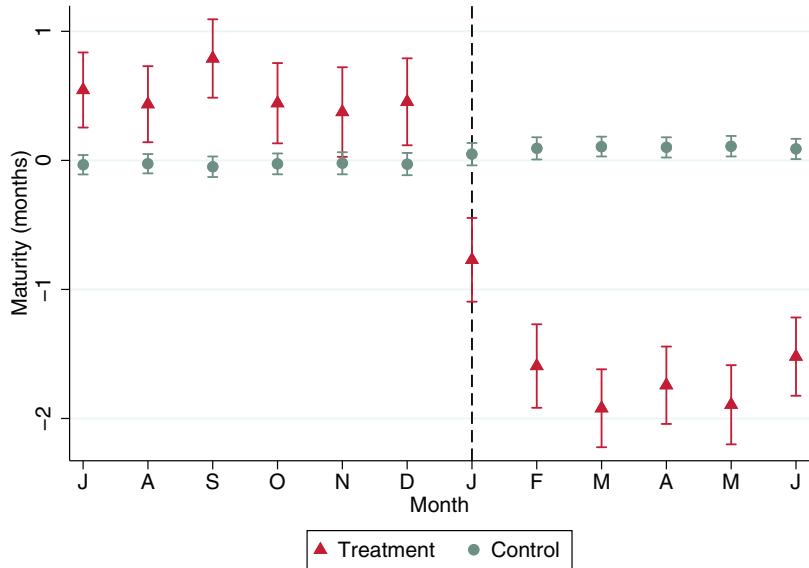


Figure 7. Average maturities around year-end by treatment. The figure plots the average conditional maturity by month around the year-end for both treatment and control groups. We first regress maturity on car age \times month-of-sale fixed effects and commuting zone fixed effects. We then plot average residuals within each month. (Color figure can be viewed at wileyonlinelibrary.com)

Treatment \times Post coefficients in Table III as causal effects of year-end discontinuous maturity policies.

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 car price. For consistency, we include the same borrower controls (DTI and credit score). Similarly, we include the same sets of fixed effects in each column as in Table III.

We report the reduced-form results in Table IV. In column (1), where we do not 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 magnitude of the coefficient. 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

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

The table reports difference-in-differences regression 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 1. Borrower controls include Credit Score (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). *MMT* indicates combinations of make-model-trim; *YMMT* indicates combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double-clustered by month and commuting zone.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

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 ²	0.060	0.369	0.872	0.909	0.911	0.914

price of a given car. To separate these two effects, we exploit the rich information on collateral quality available in our data set.

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 both bought, say, a four-year-old Honda Accord LX. The fixed effects soak up a large part of the variation, with the R^2 jumping to 0.87. The coefficient drops to 0.9%, indicating that a significant part of the effect on expenditure in column (2) is being driven by a shift in 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 than to shop for better financing.¹⁹ 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 data set. In column (4) we interact *YMMT* fixed effects with year-month fixed effects, in which case the coefficient gives us the difference in price paid by an affected borrower for the same *YMMT* purchased in the same month. These *YMMT* \times month fixed effects absorb any time-varying shocks to *YMMT* values,

¹⁹ In Section IV.B we test observable borrower characteristics for selection into treatment and find no economically meaningful results.

Table V
Two-Stage Least Squares (2SLS) Effects of Maturity on Log Price

The table reports difference-in-differences regressions of log prices for an event year running from July to June using 2SLS. The excluded instrument is Treatment \times Post, where Post is a dummy equal to one for observations after January 1. 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 Credit Score (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). *MMT* indicates combinations of make-model-trim; *YMMT* indicates combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double-clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

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 ²	0.210	0.411	0.874	0.911	0.913	0.916

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, in column (6) we add lender fixed effects to rule out the possibility that our results are being driven by 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 at around 0.6% to 0.7%. Recall that the magnitude of our estimate for the first-stage effect on average maturities is about 2.3 months. The estimates in Table IV thus 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 2SLS, Table V (with columns corresponding to those in Tables III and IV) shows a price effect of around 0.3% per month of maturity.²⁰

²⁰ See De Chaisemartin and D'Haultfoeuille (2017) and Hudson, Hull, and Liebersohn (2017) for a detailed 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.

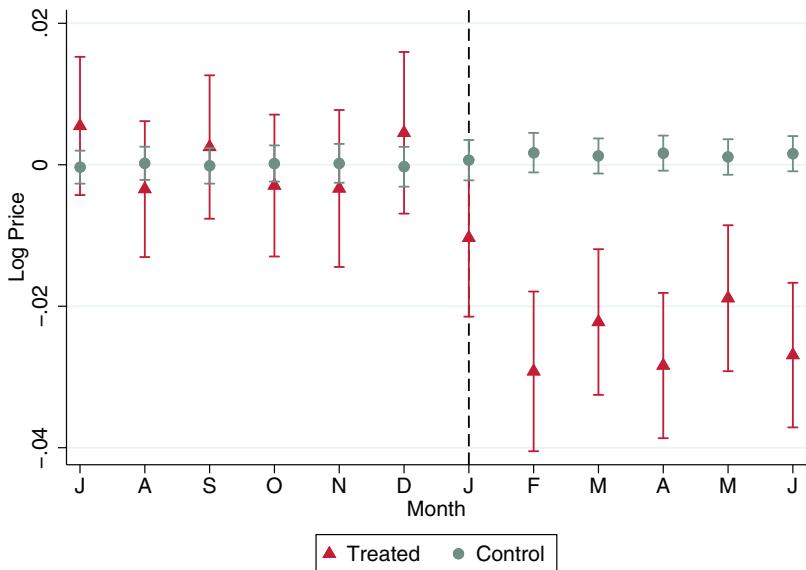


Figure 8. Average price around year-end by treatment. The figure shows the average conditional log price around year-end for both treatment and control groups. We first regress the log of car price on car age \times month-of-sale fixed effects and commuting zone fixed effects. We then plot the average residuals within each month. (Color figure can be viewed at wileyonlinelibrary.com)

The estimates in Table IV compare average prices during the postperiod (January through June) to average prices throughout the preperiod (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 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 \times month fixed effects and commuting zone fixed effects, as in Figure 7 for maturities. Control vehicles show a flat pattern over the event year, while treated vehicles' prices are largely flat except 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 many buyers would receive Earned Income Tax Credit (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—that is, the assignment of treatment and control—as we do to estimate potential pricing effects. Although we try 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 of the loans to a training sample and the remaining half to a hold-out sample. We use the training sample to identify lender maturity policies and in turn define treatment and control observations, following the procedure outlined in Section III.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 Internet Appendix Table IA.I, 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) to (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 versus \$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 off-load 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 focus on the maturity dimension of the financing contract, motivated by evidence in Argyle, Nadauld, and Palmer (2020a)

Table VI
Difference-in-Differences Results on Interest Rates

The 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 1. Borrower controls include Credit Score (the credit score of the borrower at loan origination) and DTI (the back-end debt-to-income ratio of the borrower at origination). *MMT* indicates combinations of make-model-trim; *YMMT* indicates combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double-clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

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 ²	0.426	0.443	0.499	0.604	0.640	0.664

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 in general, it would compromise the interpretation that estimated price effects are driven by changes in maturity. In this section, we turn our attention to disentangling effects due to changes in maturity from those due to changes in interest rate.

We begin by examining whether the maturity shocks that we detect coincide with changes in interest rates by reestimating the difference-in-differences specifications from Table III, with the interest rate for each loan replacing maturity as the dependent variable. Table VI reports the results. With no fixed effects, the coefficient of interest on *Treatment* \times *Post* is six bps and statistically insignificant. As we control for increasingly fine collateral fixed effects in columns (2) to (4), the estimate remains insignificant, ranging in magnitude from four to nine bps. 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 bps, marginally significant at the 10% level. Finally, with the

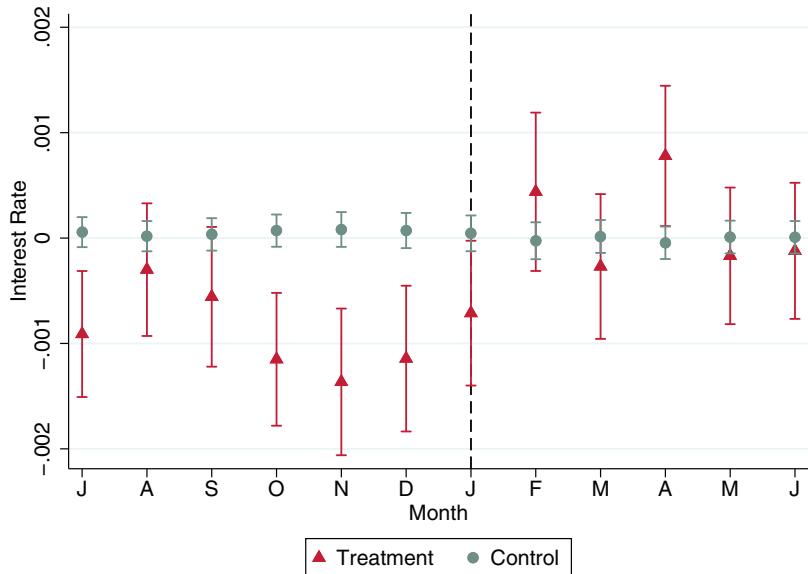


Figure 9. Average interest rate around year-end by treatment. The figure shows the average conditional interest rate around the new year for both treatment and control groups. We first regress the interest rate on car age \times month-of-sale fixed effects and commuting zone fixed effects. We then plot the average residuals within each month. (Color figure can be viewed at wileyonlinelibrary.com)

addition of lender fixed effects in column (6), the coefficient is 16 bps, 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-period months for treated cars, although 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, while 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 postperiod, which could potentially be driving some of the lower prices that we observe for treated borrowers in the postperiod.²¹

In an effort to pin down a causal estimate of maturity accounting for any interest rate impact, we estimate a 2SLS regression in which we instrument

²¹ Note that the finding 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.

for both maturity and interest rates. We estimate separate first stages for maturity and rates as follows:

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

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

As before, X_{iglt} 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 III.A. The key feature of this specification is that it allows unique magnitudes of the difference-in-differences coefficients π_k for each treated cell. Distinguishing the unique effect of maturity policy breaks from that of 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 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 bp increase in interest rate in this scenario. Meanwhile, a different lender's policy results in a 12-month decrease in maturity and a 10 bp 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 each endogenous variable, however, 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 there is sufficient independent variation to separately identify the price impact driven by each variable.

Table VII
Two-Stage Least Squares (2SLS) Effects of Maturity and Rate on Log Price

The table reports 2SLS 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 Credit Score, 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* indicates combinations of make-model-trim; *YMMT* indicates combinations of manufacture year-make-model-trim. Robust standard errors (in parentheses) are double-clustered by month and commuting zone. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

log(Price)	(1)	(2)	(3)	(4)	(5)	(6)
Maturity	0.0044*	0.0074***	0.0034***	0.0022***	0.0024***	0.0023***
	(0.0025)	(0.0015)	(0.0007)	(0.0004)	(0.0004)	(0.0004)
Interest Rate	-0.365	-1.951***	-1.467***	-0.920***	-0.863***	-0.905***
	(0.846)	(0.628)	(0.418)	(0.323)	(0.328)	(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 <i>p</i> -value	0.000	0.000	0.000	0.000	0.000	0.000
Maturity 1st Stage <i>p</i> -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 ²	0.164	0.421	0.875	0.911	0.914	0.916

Instruments for maturity and interest rates allow for a second-stage specification that builds on equations (2) and (3) and is given by

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

such that the η coefficients are semielasticities 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 the 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% to 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 offered maturity translates into a 23 bp higher price, significant at the 1% level. This compares to an estimate of 29 bps per month of maturity in the 2SLS specification without interest rates in Table V, which indicates that roughly 80% of the effect on prices estimated above is coming through the maturity channel. The coefficient on *Rate* indicates that for a one-percentage-point increase in interest rate, prices

fall by 90 bps. Given our first-stage estimate of a change in interest rate of 16 bps, these estimates imply that roughly 14 bps of price impact is driven by changes in interest rates, compared to the total price impact of 70 bps 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, which is enough to offset 60% of the increase in monthly payment. The estimates in columns (4) to (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 a borrower discount rate of 11.6%, sellers would bear the entire cost of shorter maturity. Buyers with discount rates below this internal rate of return of 11.6% would happily accept a lower maturity if doing so meant receiving 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) to (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 attribute 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 highly sensitive to available liquidity for a population of subprime borrowers. To assess this possibility, we might again rerun our main difference-in-difference specification with LTV as the dependent variable. However, Argyle, Nadauld, and Palmer (2020a) 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 Section III of the Internet Appendix, we present results indicating that affected borrowers have LTVs that are around two percentage points lower (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 rate (see Internet Appendix Tables IA.VIII and IA.IX).

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 heterogeneous 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 manifest in a lower relative price when the car is sold again. Relaxing our sample selection criteria for power considerations, our entire data set features 8,697 cars with at least two transactions. We require that the two transactions 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 each 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 , as follows:

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

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

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

This table reports difference-in-differences regression results of log price residuals for a second sale (1) and the initial sale (2) for those cars that we observe transacting twice in the data set, at least 18 months apart. These pricing residuals \hat{u} are calculated from equation (5) by controlling for manufacture year-make-model-trim \times month fixed effects, commuting zone fixed effects, and lender fixed effects. Post is a dummy equal to one for observations for which the first sale occurred after January 1. 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	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^2	0.001	0.001	

with $Treatment \times Post = 1$ by running a difference-in-differences regression in which $Treatment$ and $Post$ take their values as of the first sale at t_0 :

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

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, one concern with this exercise is that cars may only be observed selling twice in our data because 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 Internet Appendix Table IA.II, 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 bps) when sold a second time. Of course, this sample is different from our main sample, so in column (2) we 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.²² 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, although 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 for which 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 Internet Appendix Figure IA.1). To the extent that dealers may specialize in older or 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 reestimate 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 Internet Appendix Table IA.III. 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 are 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 Internet Appendix Figure IA.2 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 Oster (2019) and adjust our estimates for omitted variables bias. In Internet Appendix Table IA.IV, 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) and (2)), although 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) to (6)), where we see estimates very similar to the unadjusted coefficients reported earlier. The results of Internet Appendix Table IA.IV suggest that the scope for our

²² The only difference is that the fixed effects are accounted for in the creation of the pricing residuals rather than estimated directly in the regression.

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 (credit score 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 credit scores, regardless of the fixed effect structure. Figure 10 plots event studies of credit score and DTI by month of the year for treatment and control groups separately. The magnitudes of any differences are economically small with no consistent or statistically significant pattern that would suggest that the composition of treated borrowers changes in January relative to control-group borrowers.

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

C Mechanisms

What underlying mechanism would result in two buyers of observationally identical goods paying different prices? Our discussion of the economic framework, outlined in Section I.B, described the 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) to (2) of Table IV) to estimates with tight collateral controls (a 0.7% impact in columns (4) to (6)). Although collateral controls facilitate 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 that allow 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. Table IX reports the results. The limited loan application sample size does not allow for the YM_{MT} collateral controls featured in previous tables, so we control for collateral quality

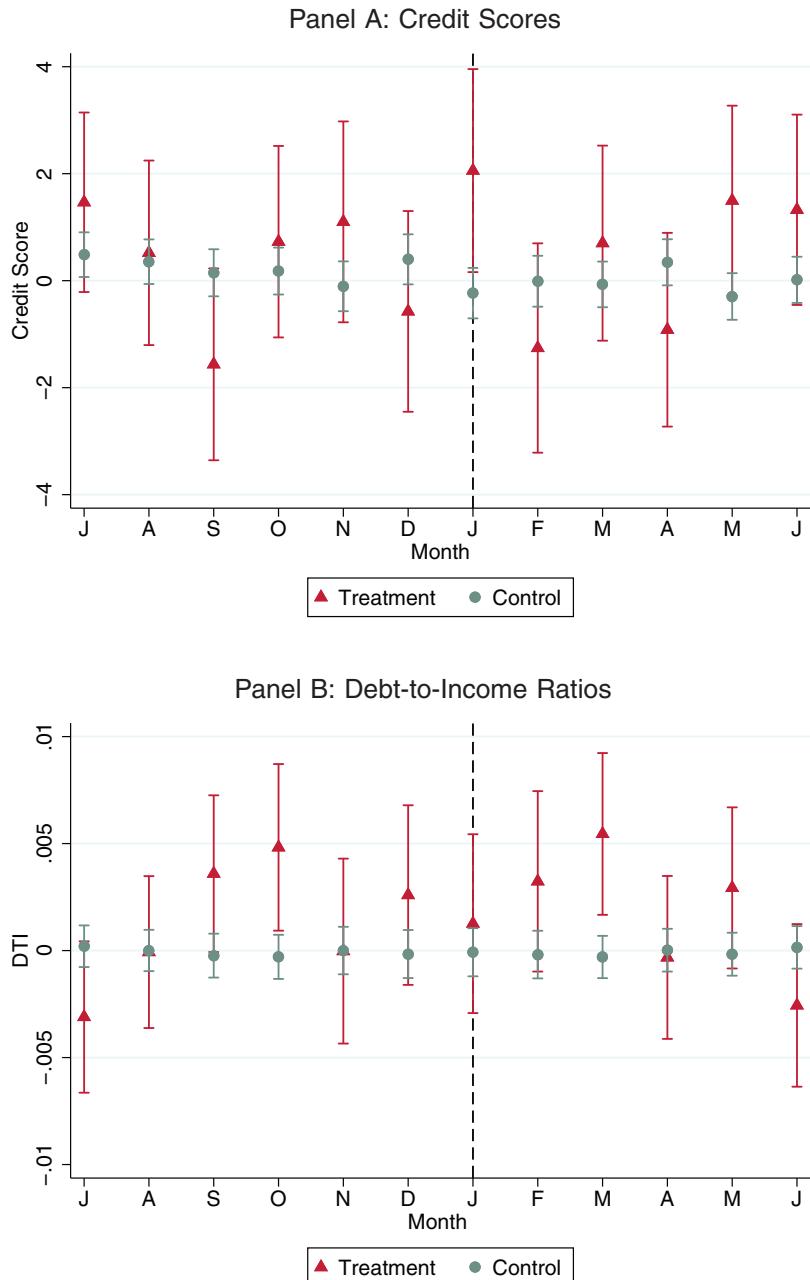


Figure 10. Average borrower characteristics around year-end by treatment. The figure plots an event study of average borrower characteristics by treatment and control group and month of year. Panel A plots average credit scores, and Panel B plots average debt-to-income ratios. We first regress each dependent variable on car age \times month-of-sale fixed effects and commuting zone fixed effects. We then plot the average residuals within each month. (Color figure can be viewed at wileyonlinelibrary.com)

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

The 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 Post is a dummy equal to one for observations after January 1. 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 Credit Score (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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(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 ²	0.01	0.01	0.05	0.05	0.05	0.06

with car age fixed effects. The negative and insignificant estimates indicate that *Treatment* \times *Post* observations are not associated with a meaningfully larger number of days between application and origination, which suggests 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, although our test is admittedly limited in addressing the possibility that treated borrowers search more intensely over the same number of days as nontreated borrowers.

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 (2020). 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. Taken together, the 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 into 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, if anything, the prices of cars bought by affected borrowers rebound when sold in later transactions, indicating that initial price differentials are 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 the question of whether our results would generalize to other markets is open, 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 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.²³

Our analysis also speaks to the transmission of policy actions through to final goods prices. For example, one goal of monetary policy is to influence consumer demand through the interest rate channel. Our results demonstrate that 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

²³ Note that our findings on the impact of consumer financing disruptions on sellers provides a positive economic rationale for vertical integration between lenders and dealers. See Murfin and Pratt (2019) for theoretical and empirical evidence on vendor financing in heavy construction equipment.

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 microeconometric 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 on 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 value most with potential implications for credit product design.

Initial submission: August 14, 2018; Accepted: November 14, 2019
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

REFERENCES

- Adams, William, Liran Einav, and Jonathan Levin, 2009, Liquidity constraints and imperfect information in subprime lending, *American Economic Review* 99, 49–84.
- Adelino, Manuel, Antoinette Schoar, and Felipe Severino, 2012, Credit supply and house prices: Evidence from mortgage market segmentation, NBER Working Paper 17832.
- Argyle, Bronson, Taylor Nadauld, and Christopher Palmer, 2020a, Monthly payment targeting and the demand for maturity, *Review of Financial Studies* 33, 5416–5462.
- Argyle, Bronson, Taylor Nadauld, and Christopher Palmer, 2020b, Real effects of search frictions in consumer credit markets, MIT Sloan NBER Working Paper 26645.
- 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, 317–365.
- Bhutta, Neil, and Daniel R. 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 E. V., and Philip Lowe, 2002, Asset prices, financial and monetary stability: Exploring the nexus, BIS Working Papers Series 114.
- Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer, 2012, Stranded vehicles: How gasoline taxes change the value of households' vehicle assets, Working paper, Northwestern University.
- Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer, 2006, \$1,000 cash back: The pass-through of auto manufacturer promotions, *American Economic Review* 96, 1253–1270.

- Cornia, Marco, Kristopher S. Gerardi, and Adam Hale Shapiro, 2012, Price dispersion over the business cycle: Evidence from the airline industry, *Journal of Industrial Economics* 60, 347–373.
- Correia, Sergio, 2017, reghdfe: Stata module for linear and instrumental-variable/gmm regression absorbing multiple levels of fixed effects, *Statistical Software Components* s457874.
- Davis, Morris A., Stephen D. Oliner, Tobias J. Peter, and Edward J. Pinto, 2020, The impact of federal housing policy on housing demand and homeownership: Evidence from a quasi-experiment, *Journal of Housing Economics* 48, 1–24.
- de Chaisemartin, Clément, and Xavier D'Haultfoeuille, 2017, Fuzzy differences-in-differences, *Review of Economic Studies* 85, 999–1028.
- Di Maggio, Marco, and Amir Kermani, 2017, Credit-induced boom and bust, *Review of Financial Studies* 30, 3711–3758.
- Diamond, Douglas W., 1991, Debt maturity structure and liquidity risk, *Quarterly Journal of Economics* 106, 709–737.
- Einav, Liran, Mark Jenkins, and Jonathan Levin, 2012, Contract pricing in consumer credit markets, *Econometrica* 80, 1387–1432.
- Englmaier, Florian, Arno Schmöller, and Till Stowasser, 2018, Price discontinuities in an online market for used cars, *Management Science* 64, 2754–2766.
- Favara, Giovanni, and Jean Imbs, 2015, Credit supply and the price of housing, *American Economic Review* 105, 958–992.
- Gavazza, Alessandro, Alessandro Lizzeri, and Nikita Roketskiy, 2014, A quantitative analysis of the used-car market, *American Economic Review* 104, 3668–3700.
- Glaeser, Edward L., Joshua D. Gottlieb, and Joseph Gyourko, 2012, Can cheap credit explain the housing boom? in Edward L. Glaeser, and Todd Sinai, eds.: *Housing and the Financial Crisis*. NBER Chapters (University of Chicago Press, Chicago, IL).
- Goolsbee, Austan, 1998, Investment tax incentives, prices, and the supply of capital goods, *Quarterly Journal of Economics* 113, 121–148.
- Hansman, Christopher, Harrison Hong, Wenxi Jiang, Yu-Jane Liu, and Juan-Juan Meng, 2018, Riding the credit boom, NBER Working Paper 24586.
- Hennessy, Christopher, and Dmitry Livdan, 2009, Debt, bargaining, and credibility in firm-supplier relationships, *Journal of Financial Economics* 93, 382–399.
- Hertzberg, Andrew, Andres Liberman, and Daniel Paravisini, 2018, Screening on loan terms: Evidence from maturity choice in consumer credit, *Review of Financial Studies* 31, 3532–3567.
- Hortacsu, Ali, Gregor Matvos, Chad Syverson, and Sriram Venkataraman, 2013, Indirect costs of financial distress in durable goods industries: The case of auto manufacturers, *Review of Financial Studies* 26, 1248–1290.
- Huang, Guofang, Hong Luo, and Jing Xia, 2019, Invest in information or wing it? a model of dynamic pricing with seller learning, *Management Science* 65, 5556–5583.
- Hudson, Sally, Peter Hull, and Jack Liebersohn, 2017, Interpreting instrumented difference-in-differences, MIT Working Paper.
- Israel, Ronen, 1991, Capital structure and the market for corporate control: The defensive role of debt financing, *Journal of Finance* 46, 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 23850.
- Landvoigt, Tim, Monika Piazzesi, and Martin Schneider, 2015, The housing market(s) of San Diego, *American Economic Review* 105, 1371–1407.
- Larsen, Bradley J., 2020, The efficiency of real-world bargaining: evidence from wholesale used-auto auctions, *Review of Economic Studies* forthcoming.
- Lee, Alice, and Daniel 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 Shen, 2018, Credit supply and the rise in college tuition: Evidence from the expansion in federal student aid programs, *Review of Financial Studies* 32, 423–466.
- MacCurdy, Thomas, 2015, How effective is the minimum wage at supporting the poor? *Journal of Political Economy* 123, 497–545.
- Matsa, David, 2010, Capital structure as a strategic variable: Evidence from collective bargaining, *Journal of Finance* 65, 1197–1232.
- Mian, Atif, and Amir Sufi, 2009, The consequences of mortgage credit expansion: Evidence from the U.S. mortgage default crisis, *Quarterly Journal of Economics* 124, 1449–1496.
- Mian, Atif, and Amir Sufi, 2011, House prices, home equity-based borrowing, and the U.S. household leverage crisis, *American Economic Review* 101, 2132–56.
- Mian, Atif, and Amir Sufi, 2018, Finance and business cycles: The credit-driven household demand channel, NBER Working Paper 24322.
- Mian, Atif, Amir Sufi, and Emil Verner, 2017, Household debt and business cycles worldwide, *Quarterly Journal of Economics* 132, 1755–1817.
- Müller, Holger, and Fausto Panunzi, 2004, Tender offers and leverage, *Quarterly Journal of Economics* 119, 1217–1248.
- Murfin, Justin, and Ryan Pratt, 2019, Who finances durable goods and why it matters: Captive finance and the coase conjecture, *Journal of Finance* 74, 755–793.
- Nadauld, Taylor, and Shane Sherlund, 2013, The impact of securitization on the expansion of subprime credit, *Journal of Financial Economics* 107, 454–476.
- Oster, Emily, 2019, Unobservable selection and coefficient stability: Theory and evidence, *Journal of Business & Economic Statistics* 37, 187–204.
- Rice, Tara, and Philipe Strahan, 2010, Does credit competition affect small-firm finance? *Journal of Finance* 65, 861–889.
- Sallee, James, 2011, The surprising incidence of tax credits for the Toyota Prius, *American Economic Journal: Economic Policy* 3, 189–219.
- Spiegel, Yossef, and Daniel Spulber, 1994, The capital structure of a regulated firm, *RAND Journal of Economics* 25, 424–440.
- Stigler, George J., 1961, The economics of information, *Journal of Political Economy* 69, 213–225.
- Verner, Emil, and Győző Gyöngyösi, 2020, Household debt revaluation and the real economy: Evidence from a foreign currency debt crisis, *American Economic Review* 110, 2667–2702.
- Zevelev, Albert, 2020, Does collateral value affect asset prices? Evidence from a natural experiment in Texas, SSRN Working Paper No. 2815609.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Internet Appendix.
Replication Code.