

NBER WORKING PAPER SERIES

AUTO DEALER LOAN INTERMEDIATION:
CONSUMER BEHAVIOR AND COMPETITIVE EFFECTS

Andreas Grunewald
Jonathan A. Lanning
David C. Low
Tobias Salz

Working Paper 28136
<http://www.nber.org/papers/w28136>

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

The views expressed are those of the authors and do not necessarily reflect those of the Consumer Financial Protection Bureau, the Federal Reserve Bank of Chicago, the Federal Reserve System, the United States, or the National Bureau of Economic Research. Andreas Grunewald and Tobias Salz thank the Institute for Social and Economic Research and Policy at Columbia University and the Russel Sage Foundation for financial support. We thank Brian Bucks, Glenn Ellison, Matt Gentzkow, Kate Ho, Éva Nagypál, Michael Riordan, Mike Whinston, Maximilian Voigt, and audiences at Berkeley, BRIQ Bonn Structural Behavioral Conference, Boston College, IO workshop at CREST, CFPB, University of Cologne, University of Copenhagen, DICE Duesseldorf, EIEF Junior Conference, Frankfurt School, German Economist Abroad Conference, Goethe University Frankfurt, Harvard, University of Mainz, University of Marburg, University of Minnesota, University of Konstanz, NYU, University of Oregon, University of Pennsylvania, Penn State, Solvay Brussels School / Ecare, Stanford, Toulouse School of Economics, and Yale for helpful comments and suggestions. Guy Aridor, Jasper Clarkberg, Thi Mai Anh Nguyen, and Yining Zhu provided outstanding research assistance.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Andreas Grunewald, Jonathan A. Lanning, David C. Low, and Tobias Salz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Auto Dealer Loan Intermediation: Consumer Behavior and Competitive Effects
Andreas Grunewald, Jonathan A. Lanning, David C. Low, and Tobias Salz
NBER Working Paper No. 28136
November 2020
JEL No. G41,G51,L0,L13,L5,L62

ABSTRACT

This paper studies the intermediation of auto loans through auto dealers using new and comprehensive administrative data. The arrangements between auto dealers and lenders incentivize dealers to increase loan prices. We leverage details of the corresponding contracts to demonstrate that many consumers are less responsive to finance charges than to vehicle charges. Taking this behavior into account, we estimate an equilibrium model of dealer price setting and lender competition. We explore counterfactuals where dealers have no discretion to price loans and final rates are set by lenders instead. We find large gains in consumer surplus from such a policy.

Andreas Grunewald
Frankfurt School of Finance and Management
Adickesallee 32-34
Frankfurt, Hess 60322
Germany
gruni06@googlemail.com

Jonathan A. Lanning
Federal Reserve Bank of Chicago
2018 W Churchill St
Chicago, IL 60647
jonlanning@gmail.com

David C. Low
CFPB
1700 G St. NW
Washington, DC 20002
david.low@cfpb.gov

Tobias Salz
MIT Department of Economics
77 Massachusetts Avenue, E52-404
Cambridge, MA 02139
and NBER
tsalz@mit.edu

1 Introduction

In retail financial markets consumers make infrequent, high-stakes decisions about complex products. Often, such decisions are made in negotiations with professional intermediaries, brokers, or sales agents who are incentivized to extract revenue from consumers. A key question is therefore to what extent the pricing discretion of these agents can harm consumers. In particular, the dispersion in markups enabled by such discretion likely reflects differences in consumers' knowledge about the terms of financial products and their ability to assess prices correctly. While regulations, such as the Equal Credit Opportunity Act, are meant to curb certain sources of price dispersion, in practice different consumers may achieve very different outcomes in such markets.

In this context, the market for financed vehicles has received considerable attention from regulators and consumer advocacy groups. Auto dealers act as intermediaries between lenders and consumers for the financing of the car. Specifically, lenders suggest loan prices but allow dealers to increase them; dealers often do so and are compensated with a share of the resulting profit through incentive contracts. Dealer loan intermediation is of outsized importance for American consumers as they hold more than \$1 trillion in auto debt, and more than eighty percent of auto loans are obtained “indirectly” through dealers.¹ Several regulators, including the Federal Trade Commission (FTC), the Consumer Financial Protection Bureau (CFPB), and the U.K.'s Financial Conduct Authority (FCA) have stated that the arrangements between lenders and dealers can harm consumers.² In particular, if some consumers respond less to loan price markups than to financially-equivalent vehicle price markups, dealers can tailor “price bundles” to extract more revenue from consumers who are less responsive to loan prices. Understanding the market for financed vehicles therefore requires understanding how lenders incentivize dealers to price loans and how these incentives affect dealer and consumer behavior.

This paper uses new and comprehensive administrative data to study the agreements between dealers and lenders and how they affect the market for financed vehicles. First we show that consumers are substantially more sensitive to changes in car prices than loan prices. To this end, we leverage the incentive structure under which dealers make those decisions, which we observe in our data. The intuition for our approach is simple: dealers keep nearly all profit from car price markups, but share profits from loan price markups with lenders. Thus, if a consumer is equally responsive to car and loan markups, a dealer maximizes profit by charging the highest car price that consumer will accept and setting the loan markup to zero. Con-

¹For an estimate of total auto loan debt outstanding, see https://www.newyorkfed.org/medialibrary/interactives/householdcredit/data/pdf/HHDC_2019Q4.pdf. We estimate the percent of auto loans obtained indirectly in Appendix D.3; see also Davis (2012). The affordability of auto loans is increasingly a concern. See, for example, <https://www.wsj.com/articles/the-seven-year-auto-loan-americas-middle-class-cant-afford-their-cars-11569941215> and <https://www.npr.org/2019/10/31/773409100/the-7-year-car-loan-watch-your-wallet>.

²See CFPB (2016), FCA (2019), Reynolds and Cox (2020) and Sullivan *et al.* (2020).

versely, if consumers are sufficiently unresponsive to loan markups, then dealers should mark up loans even though they keep only a fraction of the profits that loan markups generate. Loan markups therefore contain information about consumers' sensitivity to loan prices relative to vehicle prices. This argument requires only mild assumptions about consumer preferences and market structure. Furthermore, since we observe the markup on every loan in our data, we can derive consumer-specific estimates of the wedge between the responsiveness to loan and vehicle prices.

We find that, on average, consumers behave as if they would pay a dollar more in finance charges to reduce the vehicle price by \$0.86. Hence, consumers respond substantially more to vehicle prices than to loan prices. We also find that on average consumers act as if they perceive finance charges to be at least \$380 less than actual finance charges. Since we obtain these measures at the consumer level, we can study the heterogeneity in this wedge. We find that there are large differences across consumers. Consumers in the 10th percentile are willing to pay a dollar in finance charges to reduce the vehicle price by just 77 cents, but at the 90th percentile consumers are essentially indifferent between finance and vehicle charges.

We consider several potential explanations for our findings, including time discounting and credit constraints, vehicle taxes, loan prepayment, and repeated dealer-lender interactions. These factors cannot account for our results either qualitatively or quantitatively. However, our findings are consistent with evidence that several regulators such as the FTC, the FCA, and the CFPB have collected indicating that most consumers do not fully understand that loan terms are negotiable or do not correctly assess loan prices (Center for Responsible Lending *et al.*, 2012, CFPB, 2016, FCA, 2019, Reynolds and Cox, 2020, Sullivan *et al.*, 2020).³ Our results are also consistent with the literature on add-on pricing, which argues that firms can charge higher prices on less transparent dimensions of a product (Lal and Matutes, 1994, Verboven, 1999, Ellison, 2005).

In the second part of this project, we explore the equilibrium implications of dealer discretion, taking into account our estimates for consumer behavior from the previous part. We assume that dealers play a differentiated-product pricing game (Berry *et al.*, 1995) by setting both prices and interest rates strategically to compete for consumers. Lenders compete in interest rate auctions for the business of dealers. We use the model to explore the quantitative effects of two counterfactuals.

In our first counterfactual, we recompute equilibria under the assumption that consumers respond equally to car and loan prices. This counterfactual delineates how the consumer behavior that we document is reflected in the market equilibrium. We find that total prices

³Consumers may fail to minimize total expenses for a number of reasons. They might lack financial literacy in general (Lusardi and Mitchell, 2014), the car price might be more salient to them (Chetty *et al.*, 2009, Kőszegi and Szeidl, 2012, Bordalo *et al.*, 2013), dealers might exploit the complexity of financial contracts to shroud or obfuscate crucial information (Gabaix and Laibson, 2006, Ellison and Ellison, 2009), or consumers might neglect some dimensions of the price because their capacity to process information is limited (Sims, 2006, Mackowiak *et al.*, 2018).

would be slightly more than half a percent lower and consumer surplus 4.26 percent higher. One important insight from this counterfactual is that eliminating the wedge in consumer responsiveness makes the entire market more competitive. If consumers are equally responsive to car and loan prices, dealers have no incentive to mark up loans and so only compete over vehicle prices. Since consumer demand is elastic with respect to vehicle prices, total prices fall and consumer surplus rises.

Second, in our main counterfactual, we compute the market equilibrium when dealers are prohibited from marking up loans. This counterfactual closely approximates policies that regulators have considered or enacted. For example, the FCA has announced that it will ban dealer markups beginning in 2021.⁴ In this counterfactual, lenders fully determine loan prices, but they cannot price discriminate as effectively as dealers can because dealers interact directly with consumers. As in the first counterfactual, eliminating dealer markups increases the competitiveness of the market. If dealers have discretion to price both the loan and the car, they can tailor the bundle of prices to specific consumer types. This leads to high interest rates for consumers that are less responsive to loan charges and low interest rates for consumers that are more responsive to loan charges. Under such targeted interest rates, consumer demand is on average less elastic to increases in prices. In contrast, without discretion over interest rates, dealers cannot evade competition with tailored price bundles. As a consequence, consumer surplus increases by between 3.24% and 5.63%. Interestingly, we find that consumer surplus increases in this counterfactual irrespective of whether the wedge reflects suboptimal consumer behavior or actual preferences.

Our work contributes to the literature on retail financial markets and to the literature on auto markets. A common theme of closely related papers is that intermediaries market products to consumers who are either financially unsophisticated, face substantial search costs, or are unaware of sellers' conflict of interest (Woodward and Hall, 2010, Allen *et al.*, 2014, Guiso *et al.*, 2018, Egan, 2018, Egan *et al.*, 2019, Robles-Garcia, 2019, Benetton *et al.*, 2019, Bhattacharya *et al.*, 2019). Our findings imply that policies which simplify the decision problem of consumers in markets for complex goods may benefit consumers. More specifically, we show that eliminating intermediaries' price discretion can increase consumer welfare. Our results suggest that the detrimental effects of intermediation on consumer surplus might be particularly severe in markets where the intermediary is also the seller of the base good.

There are a number of other papers that connect the literatures on cars and loans. In the subprime market, many consumers are liquidity-constrained (Adams *et al.*, 2009) and sensitive to monthly payments (Attanasio *et al.*, 2008, Argyle *et al.*, 2018). In part because

⁴See <https://www.bbc.com/news/business-53567495>. In the U.S. in 2013, the Consumer Financial Protection Bureau issued guidance urging auto lenders to consider alternatives to the current system, including eliminating dealer markup, in order to avoid violations of the Equal Credit Opportunity Act. In 2018, this guidance was rescinded. See https://files.consumerfinance.gov/f/201303_cfpb_march_-Auto-Finance-Bulletin.pdf. In one survey in the U.S., 93% of respondents favored requiring dealers to disclose the lowest interest rate borrowers qualified for (Center for Responsible Lending *et al.* (2012)).

of binding credit and liquidity constraints for subprime consumers, loan performance can be improved by down payment requirements (Einav *et al.*, 2012), credit scoring (Einav *et al.*, 2013), and post-default wage garnishment (Brown and Jansen, 2019). Our study focuses on prime borrowers with negligible default risk. The necessity to account for default risk of risky consumers opens the door to contract-specific pricing even for consumers that actually pose no risk. Our work suggests that this leads to loan price heterogeneity even in subpopulations with minimal default risk, harming these consumers on average.

This paper also contributes more broadly to the literatures on auto markets. While dealer revenues from auto-loan intermediation are substantial (Davis, 2012), the literature on competition in auto markets has largely abstracted from dealer loan intermediation and the implied joint pricing of cars and loans (Berry *et al.*, 1995, Morton *et al.*, 2001, 2003, Gavazza *et al.*, 2014, Nurski and Verboven, 2016, Murry, 2017, Biglaiser *et al.*, 2019).⁵ We complement these studies by accounting for the strategic considerations that arise from the joint pricing of cars and car loans. Accounting for joint pricing not only contributes to our understanding of dealer competition in the auto market but is also essential to evaluating how dealer markups affect market outcomes.

There is a growing literature estimating empirical models with some form of suboptimal consumer decision-making.⁶ When relying only on observational data, this type of work faces a major challenge when identifying parameters that control the deviations from more standard models. In particular, it is often not clear how such parameters are separately identified from those that capture true preferences. Several studies in the literature have overcome this challenge in the context of health insurance plans (Handel, 2013, Abaluck and Gruber, 2011), cell phone plans (Grubb and Osborne, 2015), electricity providers (Hortaçsu *et al.*, 2017), and extended warranties (Abito and Salant, 2017). We complement this research by showing that, in the prime car market, many consumers do not minimize the cost of financed vehicles. Moreover, previous studies either leverage some form of exogenous variation or auxiliary data in addition to prices and quantities. In our case, we instead leverage the incentive structure that the *supply side* faces to derive insights about the *demand side*. An important advantage of our approach is that it does not require specifying the exact functional form of consumer preferences, nor does it require assumptions about the exact structure of the market. Our quantitative results on consumer responsiveness to financial charges are therefore less prone to model misspecification.

Finally, there are several papers at the intersection of financial decision-making and behavioral economics that establish mechanisms for our results by documenting that consumers make flawed decisions in retail financial markets. For a recent survey of this literature see

⁵Davis (2012) estimates that franchise dealerships selling to private consumers generate more than half of their profit through their Finance and Insurance departments.

⁶Overviews on structural work in behavioral economics in general and within the field of industrial organization can be found in DellaVigna (2018) and Grubb (2015), respectively.

Madrian *et al.* (forthcoming).⁷ We contribute to this literature by showing that consumers are less sensitive to loan prices than to car prices. To the extent that this wedge in sensitivities is driven by suboptimal financial decision-making, we also demonstrate how market outcomes depend on such anomalies. For example, auto loans are a major factor behind the overall increase in household debt (Mian and Sufi, 2016, Schlagenhauf and Ricketts, 2016). Our results suggest that some of the expansion in auto loan credit may be driven by a combination of the failure of car buyers to minimize total costs and dealers’ corresponding incentive to charge markups. Hence, our results may also be of broader interest to the household finance literature.

2 Institutional Details

Car prices and loan terms are negotiated between dealers and customers, so actual prices and interest rates often differ substantially from advertised ones. Typically, a sales agent and the consumer negotiate vehicle specifics, e.g., trim level, options, and an initial price quote.⁸ Then, the consumer arranges financing with the dealer’s “Finance and Insurance” (F&I) department at which stage they may negotiate loan terms and the final car price.⁹ To get rate quotes, the F&I agent typically submits information on the customer and vehicle into at least one of three major systems: DealerTrack, RouteOne, and/or Credit Union Direct Lending. Dealers may select specific lenders from which to solicit bids, or send the application to all lenders; DealerTrack advertises access to more than 1,500 lenders. Dealers appear to work with significantly fewer lenders than they potentially have access to. On average, dealers appear to maintain active relationships with about 4.35 lenders and are likely to request quotes from several of them for any given deal.

Each solicited lender submits a *buy rate*, which is the minimal interest rate at which it is willing to make the loan.¹⁰ For prime borrowers, this process happens very quickly.¹¹ The dealer may then add a markup to the buy rate. “Markdowns”, in which the dealer pays a fixed

⁷For example, present bias has been used to explain why consumers are willing to hold high-interest debt (Ausubel, 1991) and why some consumers with substantial illiquid assets borrow at high interest rates (Laibson *et al.*, 2015, Stango *et al.*, 2017) use RAND’s American Life Panel to relate various behavioral anomalies to each other and to financial decision making.

⁸Sales agents are often paid based on commission and sales targets. The commission is often a function of the sale (e.g. a “flat” commission) and / or the simple profit to the dealer. The sales targets are often discontinuous, and can result in direct bonuses and / or increases in the commission percent earned.

⁹F&I agents are compensated by the dealer via a commission (or commission-like) mechanism based on the profit generated by F&I products, including loan markups.

¹⁰Technically, the dealer originates the loan and then sells it to the lender. The “buy rate” is then the lowest interest rate at which the lender will buy the loan from the dealer. Dealers know the buy rate in advance and sell the loan almost instantly, so for all practical purposes the lender originates the loan.

¹¹Super prime quotes (e.g. credit score 740+) are often fully automated and virtually instantaneous. “High prime” deals (e.g. credit score 700+) are sometimes automated and are usually handled quickly. One bank that “manually” underwrites prime loans targets a decision time of two minutes. “Near/low prime” deals (e.g. credit score 620+) often require manual pricing and / or underwriting and therefore take longer. Subprime (e.g. credit score < 620) quotes are often available from only a few lenders and often require phone calls, and so take even longer.

fee to decrease the contract rate below the lender’s buy rate, are allowed by most lenders but occur rarely. Most lenders allow markups of at most 200-250 basis points, but markups are otherwise discretionary. Lender-imposed caps on markups arose after a series of class-action lawsuits against auto lenders that settled between 2003 and 2006. Before, many markups were even higher (see Cohen, 2012). The additional revenue generated by markups is split between dealers and lenders according to pre-specified contracts. The dealer’s share of the markup revenue is included in a one-time, upfront payment that lenders give to dealers for making the loan called the “dealer reserve”.

Since dealers only act as intermediaries for financing, loans are generally transferred directly to the balance sheet of the lender.¹² Loans present both default risk and prepayment risk. For prime consumers, default risk is minimal (less than one percent). Default risk is much higher for subprime consumers, and so in the subprime market contracts often discipline markups and split default risk between dealers and lenders. In both the subprime and prime markets, prepayment risk is substantial. Dealers typically assume all prepayment risk for the first three to six months. If the loan is prepaid during this time period, the dealer often returns the entire dealer reserve to the lender. Lenders assume all prepayment risk after this time period.

We estimate that at least 80 percent of auto loans are “indirect”, i.e. obtained through auto dealers as described above (see Section D.3 for details). However, consumers can also get loan quotes directly from lenders. They can obtain these quotes either before negotiating with the dealer or afterward. If the consumer finances the vehicle directly, the dealer receives no revenue from the financing.

3 Data and Descriptive Evidence

This section first describes the datasets and then provides descriptive evidence on markups. We pay special attention to the agreements between lenders and dealers and to the dealer reserve that results from these agreements.

3.1 Data

This project uses three different datasets. First, and most importantly, we use a new administrative dataset of auto loans from various financial institutions. The data records car, loan, and buyer characteristics and includes several million transactions from 2010 to 2014. We

¹²Banks do sometimes sell loans on the secondary market, but this is rare. Captive finance companies (finance companies owned by car manufacturers) and Buy-Here-Pay-Here dealerships (which finance sales themselves) do so much more frequently, but neither kind of lender is in our data.

observe the make and models of cars, whether they were new or used, their mileage and model year, and in some cases the price of add-ons. In terms of buyer characteristics, we observe a buyer's zip code, income, and credit score. For each transaction, we also observe the encrypted numeric identifier of the lender. There are several lenders in our data and an average of over 7,000 dealers per lender. The loan characteristics we observe include the interest rate, the term length, the down payment, and the trade-in value for the old car.¹³ Crucially, we observe the buy rate, the markup, and the dealer reserve. Recall that the buy rate is the rate at which the lender is willing to finance the transaction. The markup is the discretionary interest that the dealer adds to the buy rate, and the dealer reserve is the payment that the dealer obtains from the lender for originating and marking up the loan.

The administrative data provides detailed information on observed transactions, but it does not cover the entire market. We therefore use complementary commercial data from AutoCount[®], from which we can observe market shares of lenders and dealers for the majority of states in the U.S.¹⁴ These data also record whether dealers are franchised (i.e., associated with car manufacturers) or independent. However, it does not include information on buy rates or markups.

The third dataset we use is the CFPB's Consumer Credit Panel (CCP). The CCP is a longitudinal, nationally-representative sample of approximately five million de-identified records from one of the three nationwide consumer reporting agencies. The CCP also does not contain data on buy rates or markups, nor does it include data on the vehicle securing the loan. We mainly use it to study loan performance after origination.

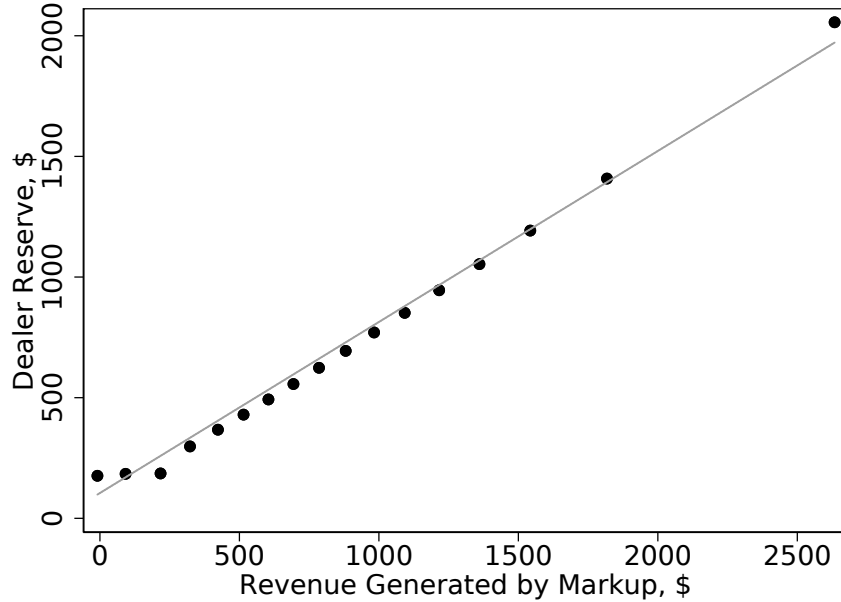
The subprime auto loan market is more complex than the prime market in a number of ways, the most important of which is default risk. To abstract from these concerns, we restrict attention everywhere in the paper to "prime" consumers, i.e., those with credit scores above 720. Because we frequently use model fixed effects and need them to be reliable, we drop observations from models that appear less than 50 times in the data. Price, loan amount, income, and down payment are winsorized by model at the 99.9 percent level.

Because our administrative dataset does not come from the universe of auto lenders, it is not nationally representative. To examine how similar our dataset is to the national market, Table A12 presents summary statistics of several variables from our data and from the 2011 commercial data. Our administrative data appear broadly comparable to nationally representative data.

¹³ Additionally, we observe whether loans are "subvented", i.e., subsidized by car manufacturers to increase vehicle demand. Subvented loans are typically from captives, but non-captive lenders do sometimes have agreements with car manufacturers to extend subvented loans.

¹⁴ In the 2011 data, banks, captives, credit unions, finance companies, and Buy-Here-Pay-Here companies respectively had 46.1 percent, 26 percent, 14.1 percent, 9 percent, and 4.8 percent of the market.

Figure 1: Markup Revenues and Dealer Reserve



NOTE: The figure is a binscatter plot. The x-axis is the revenue generated by dealer markup over the lifetime of the loan, and the y-axis is dealer reserve. The graph shows that the relationship is nearly linear.

3.2 Contracts Between Lenders and Dealers

We now use the administrative data to study the relationships between dealers and lenders. Markup revenue is shared between the dealer and the lender according to a contractually-specified formula. Figure 1 plots dealer reserve against loan markup revenue in a binscatter plot. On average, dealers receive a fixed payment of \$70 per loan plus about 75 percent of markup revenue. The dealer share of markup revenue varies slightly from lender to lender with a coefficient of variation of 0.079.¹⁵ Adjustments of those terms across dealers and geographic regions appear to be rare. Two points are worth emphasizing. First, contracts are almost linear.¹⁶ Hence, dealers receive a pre-specified amount of every dollar of extra revenue they generate through loan markups. Second, the slope is well below one. A dealer therefore receives only a fraction of markup revenue. We build on these observations in our model in Section 4.

¹⁵For one lender we directly observe the dealers' share. It varies somewhat across observations, so we take the median dealer share for this lender for the calculation of the coefficient of variation.

¹⁶If we do not pool lenders but instead plot the same graph for lenders separately, contracts are still nearly linear. To protect the confidentiality of lenders in our data, we cannot show these graphs.

Table 1: Summary Statistics of Markup and Dealer Reserve

	CREDIT SCORE		INCOME		CAR PRICE	
	MEAN	STD DEV	MEAN	STD DEV	MEAN	STD DEV
MARKUP (%-POINTS)						
1st Quartile	1.20	0.83	1.20	0.82	1.24	0.87
2nd Quartile	1.14	0.82	1.13	0.81	1.13	0.82
3rd Quartile	1.09	0.82	1.09	0.81	1.08	0.80
4th Quartile	1.08	0.81	1.08	0.80	1.09	0.77
MARGIN OVER BUY RATE						
1st Quartile	0.39	0.32	0.42	0.32	0.39	0.33
2nd Quartile	0.42	0.33	0.40	0.32	0.43	0.34
3rd Quartile	0.43	0.36	0.40	0.32	0.43	0.35
4th Quartile	0.45	0.37	0.43	0.35	0.46	0.36
DEALER RESERVE						
1st Quartile	694	556	581	453	386	293
2nd Quartile	652	529	611	475	554	370
3rd Quartile	601	489	637	508	661	445
4th Quartile	585	470	754	661	950	708

NOTE: Summary statistics of markups (upper panel), margins (middle panel), and dealer reserve (lower panel), by credit score, income, and vehicle price quartile. Margin refers to markup as a fraction of the buy rate.

3.2.1 Markups and Dealer Compensation

In our data 77.8 percent of loans are marked up and 0.7 percent of loans are marked down. The average markup in the dataset is 1.08 percentage points, which is substantial at approximately 43 percent of the buy rate. While markups are common for all types of buyers, their size varies systematically with buyer observables. Table 1 shows how markup and dealer reserve vary with the buyer’s credit score, income, and the price of the vehicle. Markups are higher for buyers with low credit scores. As a consequence, dealer reserve is also higher on average for buyers with low credit scores. Markups are also higher for low-income buyers. However, dealer reserve is *lower* for low income buyers, because they typically buy cheaper cars with smaller loans that generate less revenue for a given markup.

The results from Table 1 also inform the model below. In particular, the descriptive evidence suggests that dealers often mark up loans even though they must share the resulting profit with lenders. This is true regardless of buyer income, risk characteristics, and vehicle segment.

4 Measuring The Wedge in Price Responsiveness

Building on the institutional details described above, we now outline our empirical strategy to quantify consumer sensitivity to changes in finance charges relative to the vehicle price. Our estimates of this section are also an input to the full equilibrium model in Section 6. Our basic identification argument is the following. Dealers share a substantial fraction of loan markup revenue with lenders, but keep all vehicle markup revenue. If customers are indifferent between loan charges and vehicle charges, dealers should extract all possible revenue from the consumer by maximizing the price of the car, which would let them keep all the profit. If instead dealers set positive markups, it must be true that consumers are less sensitive to loan charges than to vehicle charges.¹⁷

This basic insight allows us to estimate consumers' relative price sensitivity under very mild assumptions on the dealers' pricing problem. We specify neither the exact market structure under which dealers compete nor a particular model of consumer behavior that gives rise to this wedge in price sensitivities. Our approach is therefore reminiscent of recent structural work that does not require a full specification of the model (Pakes *et al.*, 2015). It also shares certain advantages with the sufficient statistics approach (Chetty, 2009) in that it is less prone to errors due to model misspecification.

This section develops the identification in a one-period model. Section 5 argues that alternative explanations for our findings that assume standard optimal consumer decision-making are implausible. It shows that time preferences and credit constraints (Section 5.1)¹⁸, prepayment risk (Sections 5.2 and long-term relationships between dealers and lenders (Section 5.3) do not account for our results.

4.1 One-period model

Denote the buy rate for consumer i by b_i and one plus the sales tax rate by τ_i . The car dealer has a linear markup revenue sharing agreement with the lender (see Section 3.2). This agreement is characterized by slope α and intercept β , both of which are lender-specific. The dealer chooses the vehicle price p_i and the loan interest rate $r_i \geq b_i$, taking as given the cost to the dealer of supplying the vehicle c_i , the consumer's down payment $\kappa_i \geq 0$ and the total utility that must be delivered to the consumer \bar{u}_i . Note that we assume κ_i to be exogenous; we discuss this assumption in Section 5.1. We are agnostic regarding how \bar{u}_i is determined.¹⁹

¹⁷This argument implicitly assumes that dealers perceive the pricing of the car and the loan as a joint maximization problem. Argyle *et al.* (2018), Argyle *et al.* (2019), and Brown and Jansen (2019) provide strong evidence that loan and vehicle prices are jointly determined. In Appendix D.2, we provide further evidence that they are jointly determined using subvented loans in our data.

¹⁸The basic intuition for why time preferences do not matter is that auto loans fully amortize, so the allocation of total charges between finance charges and car charges does not affect payments in any way.

¹⁹If the dealer makes a take-it-or-leave-it offer to the consumer, and the consumer can make no credible commitments, then \bar{u}_i is the value of the consumer's outside option. If \bar{u}_i is determined by Nash bargaining, it may be as high as the maximum utility the dealer can

We only impose that $\bar{u}_i < 0$ and $|\bar{u}_i| > \kappa_i$.

The consumer can accept the dealer's offer but still get a loan *directly* from a lender with interest rate $r^L > 0$. To get a loan directly, the consumer must search for one. r^L has a continuously differentiable density function $g_i(\cdot)$ with compact support. The search cost is $s_i \geq 0$.²⁰

Our approach allows us to impose only mild assumptions on how loan and car charges affect consumers' utility. We assume that the corresponding disutility is additive in the two components and given by $p_i + M_i(x)$, where x is the amount of loan charges and $M_i(x) \in C^2 \forall i$. If $M_i(x) = x$ the consumer treats financial and vehicle costs identically. If $M_i(x) \neq x$ for some x , the consumer treats the two dimensions differently.

Assumption 1. *Consumer utility is additive in the disutility of vehicle charges p and the disutility of finance charges $M_i(x)$, where $M_i(x) \in C^2$.*

Importantly, we impose no further structure on $M_i(x)$. Therefore, our model nests many others that create explicit microfoundations for $M_i(\cdot)$.²¹ The dealer's maximization problem is given by:

$$\begin{aligned} \max_{r_i, p_i} \quad & p_i + (\tau_i \cdot p_i - \kappa_i) \cdot (r_i - b_i) \cdot \alpha - c_i + \beta & (1) \\ \text{s.t.} \quad & -\tau_i \cdot p_i - M_i((\tau_i \cdot p_i - \kappa_i) \cdot r_i) \geq \bar{u}_i \\ & -M_i((\tau_i \cdot p_i - \kappa_i) \cdot r_i) \geq -\int M_i((\tau_i \cdot p_i - \kappa_i) \cdot r^L) \cdot g_i(r^L) dr^L - s_i, \\ & r_i \geq b_i, p_i \geq 0 \end{aligned}$$

To understand the basic intuition of our argument, first consider the case of $M_i(x) = x$. In this case, the optimal policy for the dealer is a corner solution. If the incentives provided by the lender are not too high (formally, $\alpha < \frac{1}{(1+b_i) \cdot \tau_i}$), the dealer sets the interest rate as low as possible ($r_i = b_i$). The reason is that the consumer dislikes loan and vehicle price markups equally, but the dealer prefers vehicle price markups because he keeps a greater share of the revenue they generate. Note that the α above which dealers may mark up loans is strictly below one. This is because a higher car price increases the size of the loan and therefore finance charges, which dealers share with lenders. Thus, dealers do not keep all the revenue they generate from vehicle price markups.

Proposition 1. *Suppose $M_i(x) = x$ and $\alpha < \frac{1}{(1+b_i) \cdot \tau_i}$. Then any solution to the maximization problem in (1) features $r_i^* = b_i$.*

deliver to the consumer while earning nonnegative profits.

²⁰This formulation is also compatible with a model where people search multiple times.

²¹For our argument we implicitly assume that the dealer is aware of the consumer's utility function $M_i(\cdot)$. However, our estimation procedure does not rely on this assumption. If the dealer does not know $M_i(\cdot)$ but has a point belief about it, our estimates reveal this point belief.

We relegate all proofs to Appendix A. Proposition 1 has an important implication: if $M_i(x) = x$, the dealer does not mark up loans with $\alpha < \frac{1}{(1+b_i)\tau_i}$. This inequality holds for 97.5 percent of observations in our data, and so the substantial markups documented in Table 1 are incompatible with $M_i(x) = x$, i.e. equal consumer sensitivity to loan price and car price.

Using the size of markups in the data we are then able to quantify the extent to which finance and vehicle charges are treated differently.

Proposition 2. *For any optimal offer (r_i^*, p_i^*) with $r_i^* > b_i$, $p_i^* > 0$ it holds that*

$$M'_i(r_i^*(\tau_i p_i^* - \kappa_i)) \leq \frac{\alpha \cdot \tau_i}{1 - \alpha \cdot \tau_i \cdot b_i}, \quad (2)$$

where (2) holds with equality if only the first constraint in (1) is binding at the optimum.

Proposition 2 allows us to use the prices set by the dealer to infer consumers' sensitivity to additional charges at the equilibrium contract. We observe all variables on the right hand side in (2), as well as the argument $r_i^* \cdot (\tau_i \cdot p_i^* - \kappa_i)$ of function $M_i(\cdot)$. Therefore, the data provide individual-specific estimates of the marginal disutility $M'_i(\cdot)$ at observed finance charges. Furthermore, $M'_i(r_i^* \cdot (\tau_i \cdot p_i^* - \kappa_i))$ has a transparent economic interpretation: it is the marginal disutility in dollars that a consumer associates with an increase in the loan price. If $M'_i(\cdot)$ is smaller than one, for example, the consumer is less sensitive to an increase in the loan price than to an equal increase in the car price. Whether Equation 2 holds with equality for any given consumer depends on unobservables. Hence, the right hand side of Equation (2) can be interpreted as a sufficient statistic establishing an upper bound for $M'_i(\cdot)$ at the observed finance charges.

Equation 2 may not hold with equality because of the second constraint in the dealer's optimization problem. When this constraint binds, it is because the dealer cannot charge a higher interest rate without causing the consumer to obtain the loan directly from a lender, even though the consumer would otherwise agree to pay a higher interest rate in exchange for a lower vehicle price. In these cases, our estimates for $M'_i(\cdot)$ are greater than the actual $M'_i(\cdot)$. This situation may arise because our model does not allow dealers to condition the price of the vehicle on whether or not the consumer also obtains financing from the dealer. If it did, the second constraint in the dealer's problem would never bind, and our estimates for $M'_i(\cdot)$ would be tight.²²

Our $M'_i(\cdot)$ estimates measure consumers' marginal valuation of finance charges. However, each marginal valuation is estimated at the specific financial charges in the data. As a result, two individuals with the same finance charges and the same estimated $M'_i(\cdot)$ may still differ in $M_i(\cdot)$. To obtain estimates comparable across individuals, we next calculate bounds on the difference between actual finance charges and those that rationalize the consumers' choices.

²²Anecdotally, dealers often do condition the price of the vehicle on the source of financing.

Proposition 3. Suppose (r_i^*, p_i^*) and $p_i^* > 0$ is an optimal offer.

(i) If $M'_i(x) \leq 1 \quad \forall x$ and $r_i^* \geq b_i$ it holds that

$$\begin{aligned} (\tau_i \cdot p_i^* - \kappa_i) \cdot (r_i^* - b_i) - [M_i(r_i^* \cdot (\tau_i \cdot p_i^* - \kappa_i)) - M_i(b_i \cdot (\tau_i \cdot p_i^* - \kappa_i))] \\ \geq (\tau_i \cdot p_i^* - \kappa_i) \cdot (r_i^* - b_i) \cdot [1 - \tau_i \cdot \alpha \cdot (1 + b_i)] = B_i^M \end{aligned} \quad (3)$$

(ii) If $M'_i(x) > 0 \quad \forall x$ and $r_i^* > b_i$, it holds that

$$(\tau_i \cdot p_i^* - \kappa_i) \cdot r_i^* - M_i(r_i^* \cdot (\tau_i \cdot p_i^* - \kappa_i)) \geq (\tau_i \cdot p_i^* - \kappa_i) \cdot r_i^* \cdot \left[1 - \frac{\alpha \cdot \tau_i}{1 - \alpha \cdot \tau_i b_i} \right] = B_i^O. \quad (4)$$

Proposition 3 establishes two different lower bounds. Both can be estimated directly from our data. Part (i) of Proposition 3 derives an individual-specific lower bound B_i^M for the difference between the costs of a markup and the disutility that a consumer associates with these costs. Intuitively, this result leverages the fact that dealers choose the markups seen in the data rather than setting them to zero. This behavior can only be optimal if consumers' sensitivity to markups is not too high. Part (ii) of Proposition 3 derives a lower bound B_i^O for the difference between overall finance charges and those that rationalize the observed markup.²³ For this result, we have to make the additional assumption that $M_i(\cdot)$ is weakly convex. If $M_i(\cdot)$ is convex, our estimates of $M'_i(\cdot)$ in combination with the observable amount of finance charges provide an upper bound on $M_i(\cdot)$ at the observed finance charges.

While the two estimated bounds are highly correlated and the underlying ideas are similar, we still consider them to be important complements. B_i^M is derived under milder assumptions, but only considers finance charges due to markups. Convexity of $M_i(\cdot)$ is a stronger assumption but allows us to consider all finance charges. B_i^O is also important from the dealer's perspective as it characterizes the extent to which the dealer can profit from differences in price sensitivities.

4.2 Results

We now estimate the measures of consumer price sensitivities given by Propositions 2 and 3. To simplify the analysis, we first use loan term lengths to convert annual buy rates and annual interest rates in the data to two-period interest rates. We use these two-period interest rates through the rest of this paper.²⁴

²³Note that both B_i^M and B_i^O are measured in dollars over the life of the loan. Interest rates were very low during the sample period and measuring these bounds in their present discounted value at the time of the car sale would yield very similar numbers. We consider consumers' discount rates in Section 5.1.

²⁴Finance charges are equal to the monthly payment times the loan term in months, minus the loan amount. The two-period interest rate is then finance charges over the loan amount.

Note that we can only estimate $M'_i(\cdot)$ and B_i^O for the approximately 77.8 percent of our sample with positive markup; there are several potential explanations for zero markups (e.g., the consumer has a good outside option or low loan search costs, or the consumer is indifferent between vehicle and finance charges) that we cannot disentangle. However, we can estimate B_i^M for the whole sample as it is zero for observations with zero markup.

Table 2: Summary Statistics of Estimates

MEASURE	MEAN	PIO	P25	P50	P75	P90
$M'_i(\cdot)$	0.86	0.77	0.80	0.86	0.91	0.95
B_i^O (\$)	380	106	187	324	511	722
B_i^M (\$)	96	0	17	72	145	228

NOTE: Selected summary statistics of consumers' sensitivity to finance charges. For $M'_i(\cdot)$ and B_i^O , statistics condition on a positive markup, while B_i^M are derived for the full sample.


Our procedure yields individual-specific estimates. For each contract in our data, we obtain a specific $M'_i(\cdot)$, B_i^M , and B_i^O estimate by applying equation (2), (3), and (4), respectively. Recall that our sample is restricted to consumers with a credit score above 720 to ensure that our results are not driven by default risk. Table 2 summarizes our estimates, revealing substantial variation in the wedge between car and loan price sensitivities. In particular, the mean $M'_i(\cdot)$ is only 0.86. This finding implies that, at the margin, the average consumer would be willing to pay a dollar more in finance charges to reduce the price of the car by 86 cents. For consumers at the 10th percentile, this value drops to 77 cents.²⁵

Figure 2 shows the distributions of the two lower bounds B_i^M and B_i^O . Three points are worth noting. First, the bounds are substantial. On average, marked-up consumers act as if the disutility of finance charges is at least \$380 less than actual finance charges. Similarly, on average consumers act as if the disutility of markup costs is at least \$96 less than actual markup costs. These estimates correspond to approximately 20.3 percent of total finance charges and about 19.2 percent of markup charges. Second, we estimate B_i^M to be equal to zero for consumers with zero markup and negative for the few consumers with marked-down loans. Third, there is significant heterogeneity in the bounds. The bound on the difference between finance charges and the disutility of finance charges at the 90th percentile is seven times larger than at the 10th percentile.

Our estimates also allow us to explore the average curvature of the function $M_i(\cdot)$. Specifically, we can investigate the effect of higher finance charges on our estimates of $M'_i(\cdot)$. Finance charges are obviously endogenous to $M_i(\cdot)$. However, the buy rate offered by the lender

²⁵Table A13 presents our average estimates for B_i^M and $M'_i(\cdot)$ for the 20 most common models in our data, showing that there is some variation across different models.

Figure 2: Histogram of Estimated Bounds



NOTE: The left panel plots the distribution of B_i^M , the bound on the difference between markup charges and consumers' sensitivity to markup charges. The right panel plots the distribution of B_i^O , the bound on the difference between finance charges and consumers' sensitivity to finance charges.

is strongly correlated with finance charges. Moreover, our data include virtually all the information the lender has. Conditional on this information, the buy rate is therefore exogenous to any unobservable consumer characteristics such as the consumer's utility function $M_i(\cdot)$ and so is a valid instrument. The results of the corresponding 2SLS regression are summarized in Table A10 in Appendix D.1. The estimated shape suggests that $M_i(\cdot)$ is on average a convex function. Hence, consumers respond more if more is at stake. This finding is reminiscent of recent work studying attention allocation (Bordalo *et al.*, 2012, 2013, Kőszegi and Szeidl, 2012). A common theme of this work is that consumers pay more attention to dimensions with more pronounced differences across products.

So far we have imposed only minimal assumptions on $M_i(x)$. For the rest of this paper, we have to impose more structure on $M_i(x)$ for two reasons. First, we study heterogeneity in our estimated bounds across consumers, which, as we discuss below, requires additional assumptions to be informative. Second, in our counterfactual policy experiments finance charges might change and we therefore need to take a stance on precisely how $M_i(\cdot)$ changes as x changes. We therefore assume the following functional form:

$$M_i(x) = x - \rho_i \cdot \log(x), \quad (5)$$

where ρ_i specifies the size of the wedge of consumer i . A larger ρ_i corresponds to a larger wedge (Figure A4 depicts an example of the functional form). This functional form has two important advantages. First, it is consistent with the evidence above on the curvature of $M_i(\cdot)$. In particular, it generates a marginal wedge in sensitivities that is large for small x but converges to 0 as x increases. Second, it enables us to study the heterogeneity in the bounds B_i^O . Theoretically, studying heterogeneity in bounds has one potential caveat: the difference

between these bounds and the objects they bound could be correlated with observables. As we explain in more detail in Section D.1, we can eliminate this concern if we impose the utility function in Equation (5). The results in this section show that under the above functional form an ordering of B_i^O always corresponds to an ordering of the underlying wedge.

5 Explanations for Differential Price Responsiveness

Above we show that consumers are less responsive to changes in loan prices than vehicle prices. In this section, we discuss potential explanations for this finding. We can rule out liquidity constraints, prepayment risk, and dealer-lender relationships as explanations for our findings. Instead, the most plausible explanation for our results is that the wedge does not reflect consumer preferences but rather suboptimal decision making: consumers seem to either systematically underestimate financial charges or are under the incorrect impression that interest rates are non-negotiable. However, in our quantitative policy experiments (Section 9) we show that consumers are harmed by dealer loan intermediation regardless of the specific mechanism that underlies the wedge and even if the wedge reflects actual consumer preferences.

5.1 Time Structure

It may appear intuitive that in a multi-period model consumers who are impatient or liquidity-constrained are willing to pay higher interest rates, and that this could partly explain our results. However, this intuition is largely incorrect, because auto loans have fixed monthly payments that fully amortize. If the down payment is fixed, an increase in the total price of a financed vehicle by \$1 for a loan lasting T months increases monthly payments by $\frac{1}{T}$. This is true whether the \$1 total increase comes from an increase in the price of the vehicle or the loan. Thus, there is a one-to-one mapping between monthly payments and the total price that does not depend on the specific allocation of charges. More formally, Section B.1.1 develops a model with consumer preferences defined over multiple periods. In this multi-period model, we show that our estimates for $M_i'(\cdot)$, B_i^O and B_i^M are not affected by intertemporal considerations if down payments do not depend on the allocation of charges between the price of the car and the loan.²⁶

If a higher car price would lead to a larger required down payment, however, intertemporal considerations could affect our results because impatient or constrained consumers might be more willing to pay a higher interest rate in order to reduce the car price and hence the down

²⁶Down payments do not depend on the allocation of charges if, for example, the down payment is determined before the price is determined, is zero, or is chosen by the consumer (which is common for prime auto loans).

payment. Our discussions with industry experts suggest that down payments for prime consumers are generally not determined in this way. This is not surprising since collateral value – not the consumer-specific price markup – influences both the consumer’s future default decisions and the creditor’s loss given default. Empirically, our estimates of $M'_i(\cdot)$ and B_i^O vary little with the size of the down payment. In particular, our estimates for consumers with zero and negative down payments are similar to those for other consumers (see Figure A2 in the appendix).²⁷ Thus consumer impatience or credit constraints cannot explain our results.

5.2 Prepayment Risk

Prepayment risk could appear to explain consumers’ relative insensitivity to finance charges documented in Section 4.2. Intuitively, consumers who are likely to prepay their loan may appear less sensitive to finance charges because they expect not to pay them. In practice, the potential effect of prepayment on our results is more subtle. Contracts between dealers and lenders include a “clawback” period, typically the first three to six months of the loan, during which the dealer bears prepayment risk. If the consumer prepays during the clawback period, the dealer refunds the entire dealer reserve to the lender. We show in Section B.2 that this “early” prepayment risk should bias the estimated bounds B_i^O and B_i^M downwards. Early prepayment makes loan markups less valuable to dealers and so even harder to explain. However, if the consumer prepays after the clawback period, the dealer keeps the entire dealer reserve, and so late prepayment should bias our estimates upwards.²⁸

To analyze the issue empirically, we create a proxy for prepayment risk by running a logit regression predicting prepayment in the Consumer Credit Panel (CCP).²⁹ Table A7 provides estimates from the supervisory data of the relationship between our bounds and estimated prepayment risk from the CCP, using a regression with a large number of controls. The conditional correlation between prepayment risk and our wedge estimates is *negative*. Hence, we estimate lower B_i^O ’s and B_i^M ’s for consumers with *higher* prepayment risk.³⁰ There are several potential explanations for this result. First, the theoretical effects of early prepayment risk may dominate those of late prepayment.³¹ Another potential explanation with substantial

²⁷To further address this point, in Section B.1.2 we recompute $M'_i(\cdot)$ under the (evidently counterfactual) assumption that dealers require a down payment that is a fixed percentage of the car price. The results are summarized in Table A1. Recall that we are restricting our sample to individuals with credit scores above 720. Hence, the sample consists of individuals that are unlikely to be either impatient or credit constrained and therefore should have reasonably high discount factors. Our results show quantitatively that even if the down payment were a fixed fraction of the price, our estimates would barely change for discount factors above 0.90.

²⁸In a sample of prime auto loans in the CCP with loan term length greater than two years, 5.7 percent were prepaid within the first 120 days and 27.0 percent were prepaid after 120 days but before two years.

²⁹The regression uses credit score, log loan amount, loan length, and state fixed effects to predict “late” prepayment (after 120 days but before two years). Coefficients are reported in Table A2. This is clearly a limited subset of the information available to both the borrower and the dealer, and yet it is remarkably predictive. 18.3 percent of those in the bottom decile of predicted risk prepay late, while 42.6 percent in the top decile do.

³⁰As shown in Table A7, when we separate prepayment risk into early prepayment risk and late prepayment risk, we still estimate lower B_i^O ’s and B_i^M ’s for consumers with higher late prepayment risk.

³¹Early prepayment risk and late prepayment risk are highly correlated, and it seems likely that dealers cannot predict them separately.

empirical support is that, instead of considering the cost of a financed vehicle through the life of a loan, many consumers care about monthly payments (Argyle *et al.*, 2019) and so do not care about prepayment. If both consumers and dealers do not care about prepayment, then prepayment risk does not affect our results, as in our main model. If consumers do not care about prepayment but dealers do, then dealers should be less willing to mark up loans for customers with a high chance of prepayment, and so our estimated bounds B_i^O and B_i^M are biased downward. This could explain the negative conditional correlation between prepayment risk and these estimated bounds.

5.3 Dynamic Relationships between Lenders and Dealers

Lenders and dealers interact repeatedly, so dealers may mark up loans to increase lenders' profits in exchange for favorable treatment on later deals, such as financing for a subprime consumer who might otherwise not qualify for a loan. Section B.4 explores whether markups are higher for loans originated by lenders that finance a higher portion of a dealer's sales. We find that markups for lenders that finance more than 20 percent of a dealer's sales are only three basis points higher than markups for lenders that finance less than one percent of a dealer's sales. Relative to the average markup in our data of 108 basis points, this statistical effect is negligible.

5.4 Suboptimal Decision Making

The most plausible explanation for our results seems to be that consumers are either financially unsophisticated or misperceive the nature of the negotiation process. Surveys have found that most respondents are not aware that dealer markups occur, that most respondents think that dealer markups are illegal, and that respondents who were told or believed that their dealer found them the best rate in fact obtained particularly high interest rates (Center for Responsible Lending *et al.*, 2012). In a recent study, the FTC found that few consumers were aware that interest rates are negotiable (Reynolds and Cox, 2020, Sullivan *et al.*, 2020). In a similar vein, the CFPB found that "many consumers did not fully explore their options for auto loans, and did not shop around and negotiate as much for the financing as they did for the vehicle itself" (CFPB, 2016). In response to similar concerns the U.K.'s Financial Conduct Authority has announced that it will ban dealer markups in 2021 in part because "consumers are not being provided with the right information about [auto loan] commissions at the right time" (FCA, 2019).³² The evidence collected by regulators³² is consistent with the view of

The correlation between our measures of early and late prepayment risk in the CCP is .63.

³²There is also strong evidence from the mortgage market that intermediaries exploit borrowers' confusion to increase prices (Woodward and Hall, 2012), that many borrowers do not believe there is price dispersion even though it is substantial (Alexandrov and Koulayev, 2018), and that mortgage knowledge varies substantially across prospective borrowers, is correlated with other socioeconomic characteris-

industry professionals. In an expert report, a former finance and insurance manager for an automotive dealership explained:

*“The standard industry practice is to prepare financing documents so that the customer is not alerted in any manner that the person with whom he is dealing has the ability to control the customer’s price of credit. [...] This type of pricing system is particularly successful when used in conjunction with the sale of an automobile, because the credit applicant’s attention is naturally focused on the price of the automobile [...].”*³³

Regional variation in our estimates support this explanation. Figure A3 in Appendix D shows how B_i^M varies across counties. To proxy for consumer financial sophistication, we use county-level data on education from the American Community Survey (ACS) and to proxy for the cost of acquiring information we use tract-level data on internet availability from the Federal Communications Commission (FCC). Table A9 shows how the bounds B_i^O and B_i^M correlate with education and internet availability at the regional level.³⁴ Columns (3) and (6) in Table A9 show that consumers that live in areas with less education and slower internet access display larger wedges in price sensitivities.

6 Equilibrium Model

This section describes a full equilibrium model to explore how the wedge between consumers’ sensitivities to finance charges and vehicle prices affects outcomes in a competitive environment. On the demand side we nest our estimates of the wedge in a random-coefficient logit demand model. On the supply side we use a differentiated-product Bertrand model to back out dealer costs, and a model of lender competition in the loan interest rate auction to back out lenders’ cost of funds. We first use the model to quantify the effect of the wedge on consumer behavior and on the pricing strategies of dealers. We then use these estimates to conduct two counterfactual experiments: one in which consumers treat all charges equally and one in which dealers cannot mark up loans and instead are required to pass through the buy rate.

tics, and has a strong relationship with ultimate interest rates obtained (Bhutta *et al.*, 2019).

³³Expert Report of Edward Ford Jr. in the matter of Addie T. Coleman et al. vs GMAC, U.S. District Court for the Middle District of Tennessee, August 21, 2003. McDonald, Kevin M., and Kenneth J. Rojc. “Automotive Finance: Shifting Into Regulatory Overdrive.” *The Business Lawyer* 69.2 (2014): 599-607.

³⁴We do not correlate $M'(\cdot)$ with observables because it is a marginal measure; two individuals with identical $M'(\cdot)$ could have been observed with different contract terms and so have very different $M(\cdot)$.

6.1 Consumer Utility and Dealer Market Shares

There is a set of markets \mathcal{M} and for every market $m \in \mathcal{M}$ a set of active dealers \mathcal{D}_m . Every dealer $d \in \mathcal{D}_m$ offers a specific set of models \mathcal{J}_d , which we take to be exogenous. Furthermore, each dealer works with an exogenous set of lenders \mathcal{L}_d . Consumers observe all dealer-model combinations available to them and choose the one that maximizes their utility. Dealers engage in differentiated-product Bertrand competition, setting both interest rates and car prices optimally, given consumer tastes, competing dealers' behavior, their marginal cost of lending, and the revenue sharing agreement with lenders.

To specify consumer i 's indirect utility function, we denote her travel costs to dealer d by $g(d, i)$. Furthermore, consumer i has marginal utility of money γ_i , which may depend on income y_m measured at the county level. In line with the findings from Section 4.2, consumers may treat different components of the overall price differently. To parameterize the utility function $M_i(\cdot)$, we use the utility function introduced in the last section, $M_i(x) = x - \rho_i \cdot \log(x)$. Our estimates of $M_i'(\cdot)$ from Section 4 allow us to back out ρ_i immediately. Summary statistics of the corresponding estimates can be found in Table A14.

Given the state tax rate τ_d , car price p_{jd} and interest rate r_{jd} for a car j at dealer d , consumers act as if the effective overall transaction price \tilde{p}_{jd} is:

$$\tilde{p}_{jd} = \tau_d \cdot p_{jd} + (\tau_d \cdot p_{jd} - \kappa_{jd}) \cdot r_{jd} - \rho_i \cdot \log((\tau_d \cdot p_{jd} - \kappa_{jd}) \cdot r_{jd}) \quad (6)$$

Besides the travel distance and car price, consumers' indirect utility depends on other observable loan and car attributes which we collect in the vector \mathbf{z}_j , which enters according to coefficient ζ . Furthermore, there are two types of unobservable taste shocks in the model. First, the relative attractiveness of purchasing model j at dealer d is captured by a dealer-model specific scalar variable ξ_{jd} . Second, every consumer's choice is affected by her i.i.d. taste shock ϵ_{ijd} for a dealer-model combination. Overall, consumer utility is thus given by:

$$u_{ijd} = g(d, i) - \gamma_i \cdot \tilde{p}_{jd} + \zeta \cdot \mathbf{z}_j + \xi_{jd} + \epsilon_{ijd}, \quad (7)$$

where

$$\gamma_i = \phi_0 + \phi_y \cdot y_m + v_i \quad \text{where } v_i \sim \mathcal{N}(0, \sigma_v). \quad (8)$$

In practice, we collapse $g(d, i)$ at the market level as a population-weighted average over the travel distances to dealer d from different zip-code centroids so that we have $g(d)$. As is customary in the literature we refer to mean utility as $\delta_{jdm} = g(d) - \gamma \cdot \tilde{p}_{jd} + \zeta \cdot \mathbf{z}_j + \xi_{jd}$ and denote the vector of those mean utilities for an entire market as $\boldsymbol{\delta}_m$.

Let \mathbf{p}_{dm} and \mathbf{r}_{dm} respectively denote the vectors of prices and interest rates that dealer d charges. Similarly, \mathbf{p}_{-dm} and \mathbf{r}_{-dm} are the vectors of prices and interest rates of all other

dealers in the market. With this notation we now describe the customers' choices, which determine the overall market share of a dealer-model combination. For this purpose, we integrate over the unobserved customer type v . Exploiting the logit structure of the idiosyncratic taste shocks, the share of model j at dealer d is given by:

$$s_{jd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm}) = \int \frac{\exp(g(d) - \gamma_i \cdot \tilde{p}_{jd} + \zeta \cdot \mathbf{z}_j + \xi_{jd})}{\sum_{d' \in \mathcal{D}_m} \sum_{k \in \mathcal{J}_{d'}} \exp(g(d') - \gamma_i \cdot \tilde{p}_{kd'} + \zeta \cdot \mathbf{z}_k + \xi_{kd'})} f(v) \cdot dv$$

6.2 Dealer Maximization Problem

Dealers decide on prices for all models that they sell. They also set final interest rates through discretionary markups. The buy rates that dealers face are determined in equilibrium by the set of bidders \mathcal{L}_d for loans at dealer d . We describe the lender's problem in more detail in subsection 6.3. Given a down payment κ_{jd} and contractual share of markup revenue α_{jd} , which are both taken to be exogenous, dealer d 's maximization problem is

$$\max_{\mathbf{p}_{dm}, \mathbf{r}_{dm}} \sum_{j \in \mathcal{J}_d} s_{jd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm}) \cdot (p_{jd} + \alpha_{jd} \cdot (p_{jd} - \kappa_{jd}) \cdot (r_{jd} - b_{jd}) - c_{jd}). \quad (9)$$

6.3 Lender Maximization Problem

In modeling the lenders' maximization problem, we make two simplifying assumptions. First, we assume that there is one second-price auction for each dealer-model combination. Hence, each of the lenders in \mathcal{L}_d associated with dealer d bids once for all the sales of model j . In Section D.4 we also show how one can estimate the model if we assume a first-price auction instead of the second-price auction.³⁵ While the results assuming a first-price and second-price auction are very similar, in the main text we use a second-price auction for simplicity. Second, we assume α is constant across lenders within a given dealer, which appears to be a good approximation in our data. Without this assumption we would have to model an asymmetric auction which would complicate the analysis considerably. Jointly, these two assumptions allow us to keep the setting tractable.

Assumption 2. *Auctions are held at the dealer-model level. All lenders $k \in \mathcal{L}_d$ are drawing from the same cost distribution $r_k^b \sim \mathcal{F}_l(\cdot)$ and have the same revenue-sharing contract α with a given dealer d .*

A lender wins the auction and issues the loans for a dealer-model combination if his buy

³⁵The first-price auction approach uses the key insight from Guerre *et al.* (2000) that there is a relationship between the observed distribution of buy rates and the unobserved distribution of costs. This insight allows us to substitute all unknown terms in the lender's first order condition with either observed or known objects based on our demand estimation (see Appendix D.4 for details).

rate is the lowest one submitted. He then issues the loan for the second-lowest buy rate in the same auction.

Let N_d be the number of lenders bidding for a contract and $r^l \sim \mathcal{F}_l(\cdot)$ the distribution of wholesale interest rates r , which characterize the lenders' cost of financing a loan. Lenders anticipate that the dealer markup $m_{jd}(b_{jd})$ will depend on the buy rate through the dealer's downstream decision. Since we assume a second-price auction, it is optimal for the lender to bid a buy rate such that expected profits are equal to zero.³⁶ Hence, lender k will bid a buy rate that satisfies:

$$(p_{jd}(b_{jd}^k) - \kappa_{jd}) \cdot \left[(1 - \alpha_{jd}) \cdot m_{jd}(b_{jd}^k) + b_{jd}^k \right] = (p_{jd}(b_{jd}^k) - \kappa_{jd}) \cdot r_{jd}^k. \quad (10)$$

With a second-price auction, the winning lender receives the buy rate and the associated markup of the second-lowest bidder who has the second-lowest cost. As the right-hand side of Equation (10) is observed in the data, we can directly recover the distribution of the second-lowest order statistic of $\mathcal{F}_l^{(2:N_d)}(\cdot)$. It is well known that each order statistic also identifies the underlying distribution via:

$$F_l^{(i:N_d)}(v) = \frac{N_d!}{(N_d - i)!(i - 1)!} \int_0^{F_l(v)} t^{i-1} (1 - t)^{N_d - i} dt.$$

We use this relationship to uncover lenders' distribution of costs from their second-order statistic. We perform this estimation and inversion conditional on the customers' state, and FICO-score bin in increments of 10 points.

6.4 Discussion of Model Specification

While the first part of the paper analyzes car sales at the contract level, this section introduces an approach that has traditionally been used to study posted-price markets. This modeling approach deserves some discussion, especially since some of our counterfactual results are driven by heterogeneity in the relative sensitivity of financial charges to vehicle-price charges. The modeling choice is made to account for important characteristics of the product we study, while keeping the setting tractable. In particular, we study a bundle of a retail financial product and an extremely differentiated durable. While retail financial products are typically modeled at the contract level (Allen *et al.*, 2013), where contract-specific pricing is taken into account, car markets have traditionally been treated as posted-price markets. For the equilibrium model we sacrifice some of the richness of the contract-level data and follow the latter approach, which allows us to build on established tools capturing the rich space

³⁶As our sample contains only prime consumers, we implicitly assume that there is no default risk associated with the loan. Therefore, the bids of competing banks also do not contain any additional information on expected profits for each bank.

of observed and unobserved attributes of cars and also allows us to rely on proven ways to account for price endogeneity. One alternative model specification that would preserve the simplicity of the “posted price model” and account for richer heterogeneity in consumer types, would be to estimate the demand model *conditional* on subsets of ρ_i . This would allow for within-county differences in prices. While this approach is certainly a feasible alternative path, we decided against it in the interest of model simplicity. At the same time, the more than one thousand counties in our data arguably capture much of the heterogeneity in both demographics and, crucially, in the relative sensitivity to car and financial charges.

Moreover, we assume that buyers who obtain loans indirectly are distinct from buyers who obtain loans directly, and the latter are outside our model.³⁷ In practice, the set of buyers who obtain loans directly includes both consumers who never obtained a rate quote from a dealer, and those who obtained but rejected a rate quote from a dealer. We expect this latter group to be quite small; dealers can obtain buy rate quotes from a very large number of lenders very quickly, so it is unlikely that they will be unable to beat an outside rate quote obtained by a consumer if they need to.

7 Estimation

To estimate the model, we use data from 30 states, covering 1134 counties.³⁸ We consider each county to be one market and estimate demand for the 70 most popular models, which account for the large majority of total sales. The remaining models are assigned to the outside good. We restrict our estimation to counties that have less than 45 dealers and in which at least one of the 70 most popular models is sold. After these data restrictions we are left with 917 markets which each include on average 7 dealers and a median of 5 dealers.

As described in Section 3, the commercial dataset covers the entire auto market in these thirty states while the supervisory data only cover a fraction of the market. We therefore define the total market as all new car purchases observed in the commercial data. Since these data are constructed from records from the states’ Departments of Motor Vehicles, they should be close to comprehensive. However, the commercial data lack information on dealer reserves and buy rates. To estimate the model we therefore impute buy rates in the commercial data from a saturated model, which is estimated using the supervisory data ($R^2 \geq 0.86$). Lastly, we use the Google distance API to build a proxy for the travel distance to each dealer. To generate this proxy, we subdivide each county into zip-code tabulation areas

³⁷One example of an empirical model of inter-modal competition is provided in Allen *et al.* (2014), where consumers can obtain loan quotes from their home bank or run an auction among non-home lenders. Another example is Salz (2017), in which buyers can search bilaterally or through brokers.

³⁸The data includes the following states: Alabama, Arkansas, California, Colorado, Connecticut, Florida, Georgia, Idaho, Illinois, Indiana, Kentucky, Maine, Maryland, Michigan, Mississippi, Missouri, Montana, Nebraska, New Jersey, New Mexico, North Carolina, North Dakota, Ohio, South Carolina, Texas, Utah, Vermont, Virginia, Washington, and Wisconsin.

and then query the travel distance from the centroid of each of those areas to the dealer using the Google distance matrix API. The overall travel distance to a dealer within a county is the population-weighted sum of the travel-distances from each of those centroids.

The price of the vehicle is potentially endogenous to the unobserved dealer-model specific demand shock ξ_{jd} . Failure to account for this potential endogeneity will result in biased price coefficients. We therefore interact ξ_{jd} with the following set of instruments: the average miles per hour of other models at the same dealer, the vehicle length of other models at the same dealer, the mileage of other models at the same dealer, the average travel distance to other dealers in the same market, the buy rate, and the average price of the same model in other markets. Except for the buy rate and the distance, these instruments are standard in the demand estimation literature (Nevo, 2000).

Following Berry *et al.* (1995) we use the observed aggregate market shares for a specific dealer-model combination, $s_m \forall m$ to compute a contraction mapping that recovers mean utilities.³⁹ We then regress those mean utilities on product attributes to uncover the linear parameters of the mapping (Nevo, 2000). Using δ_{jdm} we can compute $\xi_{jd} = \delta_{jdm} - \zeta_x \cdot \mathbf{x}$. Equation (11) provides the moment condition for estimation.

$$G(\theta) = \sum_{d \in \mathcal{M}_d} \sum_{j \in \mathcal{J}_d} \xi_{jd}(\theta) \cdot z_{jd} \quad (11)$$

We then solve the following minimization problem:

$$\arg \min_{\theta} G(\theta)' \cdot W \cdot G(\theta). \quad (12)$$

For W we use the optimal GMM weights from a first stage estimation of a standard logit model without random coefficients.

8 Estimation Results

8.1 Demand Results

Table 3 shows the coefficients from the demand model along with standard errors. The price interaction terms show that buyers with lower income are more price-elastic. The signs of most coefficients line up with intuition and the previous literature. The coefficients on horsepower and vehicle length are both positive but decreasing in income. However, the coefficient on miles per gallon is negative but increasing in income, suggesting that higher income

³⁹Unlike the nested logit model in Berry (1994), the random coefficient model does not allow for an analytic inversion. The mapping iterates on the following equation: $\delta_m^{h+1} = \delta_m^h + \log(s_m) - \log(\hat{s}(\delta^h, \sigma))$.

households place greater value on fuel efficiency. Consumers dislike travel distance to the dealer.

Table 3: Demand Model Coefficients

VARIABLE	COEFFICIENT	STANDARD ERROR
Price	-0.09	0.0002
Horsepower	0.21	0.0021
Vehicle Length	0.83	0.0114
MPG	-1.78	0.0728
Distance to Dealer	-0.89	0.0004
Income \times Price	0.89	0.0321
Income \times Horsepower	-2.56	0.3204
Income \times MPG	18.08	9.6959
Income \times Car Length	-8.07	1.6898
Standard Deviation Random Coefficient	0.10	< 0.0001

NOTE: The table shows the main coefficient estimates and the standard errors from the model. Standard errors are rounded to second digit after the comma. Travel time is in log-minutes and prices are in dollars. Income is measured in units of standard deviations.

The coefficients are difficult to interpret directly because of the interaction terms. To provide some intuition for the magnitudes of our estimates, we now discuss the average elasticities implied by our coefficient estimates.

Our model estimates consumers' price elasticities for the overall transaction price including the car price and all finance charges. Over all market-dealer-model combinations we obtain about 35,000 different price elasticities, with an average of -3.73. The own-price elasticities of the fifteen different makes that are included in the demand model vary from -5.4 for GMC to -2.5 for Hyundai (see Table A15 in Appendix D).⁴⁰ Our total price elasticities are close to previous estimates in the literature. Nurski and Verboven (2016), for example, find an average price elasticity of -3.14 for the Belgian market and Murry (2017) estimates an average own-price elasticity of -4.9 for the US market. Note that it is intuitive that we estimate demand to be slightly less elastic. Our estimates are relative to the overall price, including finance charges, while previous papers consider the car price alone. As increases in the car price are often accompanied by increases in the size of the loan and hence finance charges, ignoring finance charges should lead to lower estimates for demand elasticities.

At an average of -0.16 consumers appear not to be too sensitive to changes in distance to different dealer locations.⁴¹

⁴⁰These elasticities are computed as averages over all models within a make.

⁴¹Murry (2017) also documents distance elasticities using the distance in miles instead of average travel times as a measure for the disutility associated with traveling to a dealership. He finds that buyers are more elastic to distance with estimated elasticities ranging from -1.1 to -1.8. This might be a result of a different subset of markets that Murry (2017) focuses on.

8.2 Dealer Cost

In order to conduct the counterfactual experiments, we use the first-order conditions from the dealer's pricing problem to recover dealer costs. We also estimate the parameter α , which determines the contractually-specified transfer (as a function of the markup) that dealer d receives from a lender when selling a car of model j . A dealer maximizes her profits with regard to both interest rate and car price. Let κ_{kd} be the down payment for model k at dealer d . We obtain a set of $2 \times |\mathcal{J}_d|$ first-order conditions, which are shown in the following two sets of equations.

$$s_{jd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm}) \cdot (1 + \alpha_{jd} \cdot (r_{jd} - b_{jd})) + \sum_{k \in \mathcal{J}_d} \frac{\partial s_{kd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm})}{\partial p_{jd}} \cdot (p_{kd} + \alpha_{kd} \cdot (p_{kd} - \kappa_{kd}) \cdot (r_{kd} - b_{kd}) - c_{kd}) = 0 \quad \forall j \in \mathcal{J}_d \quad (13)$$

$$s_{jd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm}) \cdot \alpha_{jd} \cdot (p_{jd} - \kappa_{jd}) + \sum_{k \in \mathcal{J}_d} \frac{\partial s_{kd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm})}{\partial p_{jd}} \cdot (p_{kd} + \alpha_{kd} \cdot (p_{kd} - \kappa_{kd}) \cdot (r_{kd} - b_{kd}) - c_{kd}) = 0 \quad \forall j \in \mathcal{J}_d \quad (14)$$

In Appendix C we show how one can combine this set of equations to recover the cost of dealers. The cost estimates that we obtain from this procedure are shown in Table 4. We estimate an average cost of \$17,600, ranging from \$7,600 in the 10th percentile to \$27,600 in the 90th percentile. The average markup is estimated at 25% and the implied average Lerner index is 0.294.⁴²

Table 4: Lerner Index, and Cost, Summary Statistics

VARIABLE	MEAN	P10	P25	P50	P75	P90
Lerner Index	0.294	0.102	0.147	0.228	0.383	0.606
Cost (\$ 1000)	17.6	7.6	11.9	17.5	23.6	27.6

NOTE: The table shows summary statistics for the Lerner index across all estimated markets. In each market we weight the index according to the market shares of the respective model.

8.3 Lender Cost Estimates

Lenders compete with other lenders that have an established relationship with the dealer. The mean and median number of lenders through which a dealer extends loans is 4.35 and 4,

⁴²The Lerner index is equal to the price of a good minus its marginal cost, normalized by the price of the good. It is often used as a measure of market power, with values near 0 indicating a competitive market and values near 1 indicating a concentrated market.

respectively. If we do not count lenders that originate less than five percent of a dealer’s loans, those numbers drop to 3.35 and 3, respectively. On average, we estimate lenders’ cost


Table 5: Lender Cost, Buyrate, Interest Rate

VARIABLE	MEAN	P10	P25	P50	P75	P90
Lender Cost	12.1	5.4	7.4	10.5	15.5	21.2
Buy Rate	8.7	5.5	6.5	7.8	9.7	12.7
Interest Rate	13	6.9	9.5	12.3	15.7	19.7

NOTE: The table shows the lender cost, buy-rate, interest rate. All quantities are two-period interest rates expressed in percent.

of funds to be 12.1%. A complete overview of our lender cost estimates, buy rates and interest rates is given in Table 5. Figure 3 depicts the estimated lender cost distribution together with the observed distributions of interest rates and buy rates. Note that the distribution of buy rates (except for very small values) stochastically dominates the distribution of lender costs. Lenders anticipate dealer markups and are therefore willing to offer buy rates below cost.

Figure 3: CDFs of interest rates, buy rates, and lender costs.



NOTE: This graph shows the CDFs of the interest rate, buy rate, and lender cost distributions.

9 Counterfactual Experiments

We present results from two different counterfactual experiments. We start by investigating the market equilibrium under the assumption that buyer utility is equally responsive to finance charges and vehicle charges, that is $\rho_i = 0 \forall i$. We refer to this as the NO WEDGE counterfactual. While this simulation does not directly reflect any feasible policy, it delineates how the wedge in consumer preferences affects overall market outcomes. In a second counterfactual,

we remove dealers' ability to mark up loans. The FCA has announced this policy will be enacted in the U.K. in 2021; it is perhaps the most obvious policy to counteract the potential adverse effects of dealer loan intermediation. Without dealer loan markups, lenders set the final interest rates for consumers and dealers only determine the price of the car. We refer to this as the NO DISCRETION counterfactual.

We compare outcomes under these counterfactual experiments to the outcomes from our estimated model, which we refer to as BASELINE. Our main outcome measures will be the effects on prices, consumer surplus, and producer surplus. To simplify the presentation of our results and to make interest rates comparable across contracts with different term lengths, we calculate total finance charges and show the implied two-period interest rate. We also compute two different measures for consumer surplus. First, we assume that the estimated $\hat{\rho}_i$'s accurately reflect consumer preferences and compute average consumer surplus for all scenarios under $\hat{\rho}_i$. Second, we compare consumer surplus across scenarios under the paternalistic welfare criterion that true preferences are given by $\rho_i = 0 \forall i$. The latter measure of consumer surplus is correct if consumers actually prefer to minimize total costs but fail to do so because, for example, they believe that interest rates are non-negotiable.

9.1 No Wedge Counterfactual

The first counterfactual investigates the market equilibrium if all consumers are indifferent between vehicle and finance charges. Comparing the estimated results from BASELINE to the counterfactual results under NO WEDGE allows us to analyze how the wedge in preferences changes the market equilibrium. As in the previous sections, lenders compete to provide loans in a second-price auction.

In NO WEDGE total prices are on average \$170 lower than in BASELINE, a reduction of 0.55%. This overall effect is due to opposing changes in the two prices. On the one hand, dealers increase average car prices from \$27,071 to \$27,524. On the other hand, final interest rates decrease. In NO WEDGE dealers do not mark up loans since consumers fully respond to finance charges. Lenders anticipate this behavior and so set higher buy rates than in BASELINE, but competition among lenders prevents buy rates (which are final interest rates without dealer markups) from rising to the level of final interest rates in BASELINE. As a result, two-period interest rates fall from 12.61% to 9.99%.

Table 6: Overview Counterfactual Results

	BASELINE	NO WEDGE	$\Delta\%$	NO DISCRETION	$\Delta\%$	NO DISCRETION COMPENSATED	$\Delta\%$
<i>Panel A: Market Variables</i>							
Average Total Price, $p \cdot (1 + r)$ (\$)	30,688	30,518	-0.55	30,406	-0.92	30,171	-1.68
Average Car Price, p (\$)	27,071	27,524	1.67	27,862	2.92	27,072	0
Average Interest Rate (%)	12.61	9.99	-20.76	9.15	-27.4	11.7	-7.22
Median Interest Rate (%)	11.16	8.84	-21.79	8.73	-20.87	11.24	0.63
<i>Panel B: Consumer Surplus</i>							
Total Cons. Surplus, $\hat{\rho}$ (Billion \$)	41.54	42.23	1.67	41.79	0.62	43.23	4.08
Total Cons. Surplus, $\rho = 0$ (Billion \$)	36.97	38.55	4.26	38.17	3.24	39.05	5.63
<i>Panel C: Profits</i>							
Total Dealer Profits (Billion \$)	3.61	3.19	-11.58	3.48	-3.67	3.53	-2.15
Total Lender Profits (Billion \$)	6.62	6.6	-0.3	5.57	-15.85	5.48	-17.21

NOTE: This table shows results for the three different counterfactual scenarios. In scenario NO WEDGE we set $\rho = 0$, which means that consumers treat all financial charges equally. In scenario NO DEALER DISCRETION lenders set interest rates directly and dealers compete downstream in prices taking them as given. In scenario NO DEALER DISCRETION COMPENSATED we compensate dealers for their lost revenue share through an increase in fixed payments from lenders to reduce a double marginalization effect. In all scenarios we compute consumer surplus under the assumption that the wedge reflects sup-optimal consumer decision making ($\rho = 0$) and under the assumption that the wedge reflects true preferences ($\rho = \hat{\rho}$). The first three rows (Total Price, Car Price, Interest Rate) are averages across all markets. The last four rows (Consumer Surplus and Dealer/Lender Profits) are totals over all markets.

Due to the drop in total prices, under the paternalistic welfare criterion, consumer surplus increases by 4.26%. Hence, eliminating the wedge in price sensitivities amplifies the competitiveness of the market. To study this force in more detail, we compare demand elasticities under BASELINE and NO WEDGE. For this purpose, we compute elasticities at the prices and interest rates in BASELINE. For comparability, it is necessary to take into account that dealers split increases in the total price differently in the two scenarios. While dealers in BASELINE distribute total price increases across finance and car charges, in NO WEDGE they only increase car prices. The corresponding demand elasticity is -3.7 under BASELINE and -4.08 under NO WEDGE, a 10% decrease in demand elasticity. Partly, this result arises because finance charges depend also on vehicle prices, since an increase in the price of a vehicle leads to increase in the loan amount. So making consumers more responsive to finance charges also makes them more responsive to vehicle prices.


Importantly, even if consumer utility is determined by $\hat{\rho}_i$ — i.e. the wedge reflects consumer preferences — consumer surplus increases by 1.67%. It may seem counterintuitive that consumers with preferences given by $\hat{\rho}$ benefit if dealers acted as if their preferences were given by $\rho = 0$ instead. One good way to understand this result is an analogy between ρ in (6) and the price coefficient γ_i in (7). Clearly, consumers would prefer firms to compete as if their γ_i were higher (in absolute value) than it actually is, because if they did the market would be more competitive and prices would be lower. Just as increasing γ_i would, decreasing ρ makes the market more competitive and decreases total prices.

The analogy between ρ and γ_i is, however, not perfect. If consumers act as if $\rho = 0$ when in fact $\rho = \hat{\rho}$, total prices are lower but the components of total prices are allocated suboptimally. The finding that lower total prices more than compensate consumers for suboptimally allocated prices is a quantitative result. Thus the NO WEDGE counterfactual increases consumer surplus, even if consumer utility is determined by $\hat{\rho}_i$. This result is important to understand why consumer surplus also increases in the NO DISCRETION counterfactual, regardless of whether we measure it assuming $\rho = \hat{\rho}$ or $\rho = 0$.

As we have pointed out, there is substantial heterogeneity in the wedge across consumers. This heterogeneity maps to heterogeneous price effects of the counterfactual. Figure 4 shows the change in total price by ρ quantile. The counterfactual total price change is largest for consumers who previously had the highest ρ . The counterfactual leads to a price increase of about \$130 for consumers with ρ in the bottom quartile, whereas for consumers in the highest quartile total prices decrease by more than \$400.

These findings illustrate how dealers compete for heterogeneous consumers in our posted-price setting. They offer low car prices to increase their market share, but only because they anticipate substantial loan markup profit from consumers with large ρ . NO WEDGE eliminates this cross subsidization, benefiting consumers with larger ρ and harming those with low ρ .

Figure 4: Price Effects of the NO WEDGE Scenario ρ -decile.



NOTE: The figure shows the estimated price effect of the NO WEDGE counterfactual by ρ quartile. In the counterfactual, prices fall more for consumers with larger wedges.

The increase in consumer price elasticity lowers dealers' profits. Since total prices decrease, aggregate dealer surplus falls from \$3.61 billion to \$3.19 billion, an 11.58% reduction. Hence, dealer profits would fall substantially if consumers were to treat finance charges and vehicle charges identically.

9.2 No Dealer Discretion Counterfactual

Next, in our key counterfactual, we explore what happens if dealers have no discretion to set interest rates. In this counterfactual, lenders set final interest rates directly and dealers take them as given. In line with the previous sections, we assume that lenders participate in a second-price auction to issue loans, i.e. interest rates are equal to the second-lowest bid of N_d bids, where N_d is the number of lenders that work with dealer d . To account for how contracts between dealers and lenders may change in the long term without dealer loan markups, we run two versions of this counterfactual. In the *first*, dealers receive no other payments from lenders to compensate for the loss of dealer loan markup. In the *second*, lenders pay dealers a fixed amount that fully compensates them for the loss of dealer loan markup. These two versions of the counterfactual provide bounds for intermediate cases in which dealers receive partial fixed compensation from lenders. We begin by describing the first version of the counterfactual in which dealers are not compensated.

Eliminating dealer discretion over loan prices has three important effects. First there is an effect on competition, in line with the NO WEDGE counterfactual. If dealers cannot adjust interest rates, they have to compete for customers only through car prices. Therefore they cannot attenuate competition by tailoring price bundles to specific consumer types. This

effectively makes consumers more price-elastic and therefore reduces total prices. Note however, that the mechanism is slightly different between the two counterfactuals. In NO WEDGE consumers are *actually* more price-elastic because their responsiveness to finance charges changes. In NO DISCRETION consumers do not change but are still *effectively* more price-elastic because dealers cannot tailor price bundles to consumer behavior. Conducting the same comparison of elasticities as in Section 9.1, we find that consumers are about 10% more price-elastic in NO DISCRETION than in BASELINE.

Second, there is an information effect. In the model, lenders have less information about consumers than dealers do and so they cannot price discriminate as effectively across consumer types. To test empirically how much information lenders have about consumer wedges, we estimate a fully-saturated regression model of markups, using virtually all information that lenders have. The adjusted R^2 is 0.11 (see Table 7). This finding suggests that lenders have very little information about markups and therefore consumer wedges.⁴³

Since we assume that lenders compete in a second-price auction, our model makes the extreme prediction that lenders base their bids only on cost-relevant consumer observables. Consumer wedges do not affect lender costs, so lenders would not price discriminate at all. This implicit assumption is in line with the evidence above that lenders have almost no information about consumer wedges, but it may be plausible that if dealer markups were eliminated lenders would attempt to collect more data to price discriminate themselves. However, lenders are not able to observe or interact with consumers directly, making it harder to gain access to the same soft information that dealers have. In practice, the Equal Credit Opportunity Act and other government regulations place strong restrictions on the information lenders can use to price loans. Overall, NO DISCRETION is best understood as a scenario in which lenders are permitted from using such additional information.

The information effect is illustrated by Figure 5, which shows a scatter plot together with a bin-scatter plot of interest rates against the wedge parameter ρ . Average interest rates increase with consumer wedges in BASELINE, but there is almost no relationship between the two variables in NO DISCRETION. Thus the effects of the counterfactual are heterogeneous. As in NO WEDGE, high- ρ consumers benefit more than low- ρ consumers if dealer discretion is banned.

There is also a third effect of eliminating dealer loan price discretion that results from the change in dealers' vertical incentives. In BASELINE, both dealers and lenders profit from dealer loan markups which effectively aligns their incentives. However, in NO DISCRETION, lenders prefer that consumer surplus be extracted through high loan prices while dealers prefer that

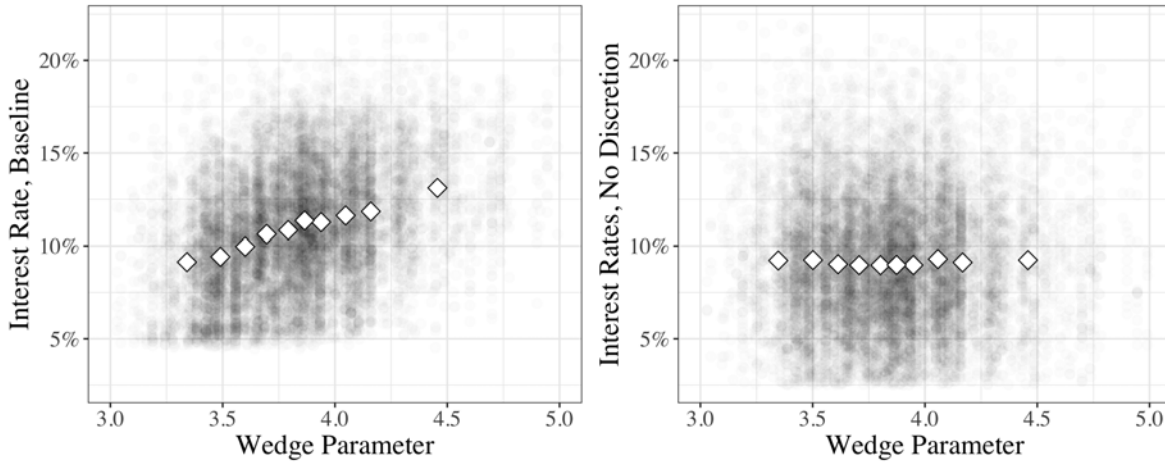
⁴³Our counterfactual might underestimate the importance of the information effect since we collapse consumer types at the county level whereas dealers can in practice tailor prices even more finely grained. Thus, consumers might on average benefit even more from banning dealer price discretion than suggested by our counterfactual simulations. As we pointed out in our model discussion, one way to capture richer dealer pricing would be to condition the demand model on subsets of ρ .

Table 7: Regressions of observables on markup (new cars only)

	LOAN MARKUP	LOAN MARKUP	LOAN MARKUP
Log Monthly Income	-0.0158*** (0.002)	-0.0242*** (0.002)	-0.0176*** (0.002)
Log Car Price	0.166*** (0.006)	0.397*** (0.007)	0.357*** (0.007)
Credit score, 100 points	-0.0989*** (0.003)	-0.0840*** (0.003)	-0.0860*** (0.003)
Mileage, Tens of Thousands	0.0228*** (0.006)	0.0286*** (0.005)	0.0217*** (0.006)
Log Loan Amount	-0.133*** (0.004)	-0.241*** (0.004)	-0.223*** (0.004)
Buy Rate	-0.0980*** (0.002)	0.000952 (0.002)	-0.000826 (0.002)
Model Fixed Effects	Yes	Yes	Yes
Term Fixed Effects	Yes	Yes	Yes
Month-Year Fixed Effects	Yes	Yes	Yes
Lender Fixed Effects	No	Yes	Yes
Zip Fixed Effects	No	No	Yes
Adjusted R^2	.049	.08	.11

NOTE: The table shows three different regressions of dealer markup on customer observables that are available to the lender at the time of the buy rate auction. Even though these regressions are very saturated, only a small percent of the variation in markups can be explained by the independent variables.

Figure 5: Scatter and Binscatter Plot of Interest Rate Change against ρ -estimate



NOTE: The binscatter plot shows the Interest Rate in BASELINE and NO DISCRETION against ρ . The gray dots are the scatter points and the white diamonds the binscatter points.

it be extracted through high car prices. Therefore, lenders increase buy rates from 8.7% in BASELINE to 9.15% in NO DISCRETION. Dealers then maximize their profits taking these buy rates (which, without dealer markups, are also the final interest rates) as given. This new mismatch between dealer and lender incentives thus leads to higher total transaction prices than if dealers set both prices. All else equal, this effect, which is reminiscent of double marginalization, decreases consumer welfare relative to the baseline model.

We find that the information and competition effect outweigh the double marginalization effect. Results are presented in Table 6. NO DISCRETION leads to a decrease in the average total price of \$283, a 0.92% decrease. As a result, consumer surplus increases. Even if we measure consumer welfare according to $\hat{\rho}$, consumer surplus increases by 0.62%. Hence, the increase in the competitiveness of the market is strong enough such that consumers benefit even if the wedge represents their true preferences. If we adopt the paternalistic welfare criterion, consumer welfare increases by 3.24%. While the effect on consumer surplus is on average positive, it varies substantially across different subgroups of the population. Consumers with $\hat{\rho}$ above the median experience a price decrease of \$384 whereas prices for individuals with a $\hat{\rho}$ below the median decrease only by \$182. For consumers that are equally responsive to finance and car charges, the policy may even lead to an *increase* in prices—total prices increase by \$310 for consumers with an estimated $\hat{\rho}$ in the lowest quartile. The size of the wedge is correlated with the level of education as well as consumers' income (see Section 5.4 and Table A9). Hence, eliminating dealer discretion to price loans, also leads to redistributive effects. While consumers with lower income and less education benefit on average, highly educated consumers may in fact suffer from the policy.

While NO DISCRETION assumed that contractual agreements between dealers and lenders

are equivalent to those in `BASELINE`, without dealer markups dealer-lender contracts might adjust to compensate dealers in other ways for intermediating loans. To illustrate how changes in the agreements between lenders and dealers might affect our results, we also run the second version of this counterfactual, which we call `NO DISCRETION COMPENSATED`. This counterfactual is identical to `NO DISCRETION` except for the fact that dealers receive an additional fixed payment by lenders for every loan they originate. To derive an upper bound for how readjustments of contracts may affect equilibrium prices, we assume that this payment on average fully compensates dealers under the old equilibrium prices, which amounts to a \$615 fixed payment per loan. Under this assumption, the effects on consumer surplus are considerably stronger (see also Table 6). Consumer surplus would then increase by 4.08 percent and 5.63 percent with respect to the non-paternalistic and paternalistic welfare criterion, respectively. In particular, this payment counteracts the quantity distortion as it induces dealers to pass through some of this compensation to consumers in order to increase market share.

10 Conclusion

In this paper we evaluate a key institution at the heart of a key retail financial market, which is the ability of dealers to directly price intermediated auto loans. Our analysis uses novel comprehensive administrative data with detailed information on the arrangements between lenders and dealers. Our analysis follows two steps. First we construct a measure of consumers' differential responsiveness to loan and vehicle charges. Our measure only requires that dealers set interest rate markups optimally but makes otherwise few assumptions on conduct or consumer behavior. Instead, the estimation strategy exploits the incentives that dealers face when they set interest rates. We find that consumers are significantly less responsive to finance charges than to car prices. Our analysis suggests that this wedge arises due to suboptimal consumer behavior.

We then study the equilibrium implication of dealers' loan price discretion by taking it away in a counterfactual. For this purpose we use our first-stage estimates of consumers' wedges in a differentiated-product model with Bertrand competition among dealers. In addition, we recover lender costs so we can model the interest rates set by lenders in this counterfactual. Removing dealers' discretion to price loans in the model benefits consumers. This result holds whether the wedge actually reflects true consumer preferences or suboptimal choices. The key mechanism is that, without loan markups, dealers compete only on car prices. Consumer demand is more elastic to car prices, so prices fall and consumers benefit.

Our results highlight a broader issue in retail financial markets. To allow for risk adjustment, many financial and insurance markets feature contract-specific pricing. However this allows sellers to price attributes besides risk, to the disadvantage of some consumers. Al-

though the Equal Credit Opportunity Act limits the attributes that can formally be used to price loans, in practice these limits may be less useful in situations in which buyers and sellers negotiate. Our results suggest that there is wide dispersion in consumer responsiveness to interest rates and that many consumers leave money on the table. Dealers are able to use this dispersion to their advantage.

While the considerable size of the auto market makes it a worthy subject of investigation in and of itself, there are many other markets in which products are sold along with financial contracts. For example, durables are often offered with extended warranties and installment plans and flights are sold with travel insurance. How the intermediation of financial contracts by sales agents affects consumer choices and welfare is therefore of broad interest. Our results indicate that this “dual role” of sales agents can hurt consumers by attenuating price competition in the respective markets.

References

- ABALUCK, J. and GRUBER, J. (2011). Choice inconsistencies among the elderly: evidence from plan choice in the medicare part d program. *American Economic Review*, **101** (4), 1180–1210.
- ABITO, J. M. and SALANT, Y. (2017). *The Effect of Product Misperception on Economic Outcomes: Evidence from the Extended Warranty Market*. Tech. rep., mimeo, University of Pennsylvania.
- ADAMS, W., EINAV, L. and LEVIN, J. (2009). Liquidity constraints and imperfect information in subprime lending. *The American Economic Review*, **99** (1), 49–84.
- ALEXANDROV, A. and KOULAYEV, S. (2018). No shopping in the u.s. mortgage market: Direct and strategic effects of providing information. *Consumer Financial Protection Bureau Office of Research Working Paper*, **2017** (1).
- ALLEN, J., CLARK, R. and HOUDE, J.-F. (2013). *The Effect of Mergers in Search Market: Evidence from the Canadian Mortgage Industry*. Tech. rep., National Bureau of Economic Research.
- , --- and --- (2014). Price dispersion in mortgage markets. *The Journal of Industrial Economics*, **62** (3), 377–416.
- ARGYLE, B., NADAULD, T. and PALMER, C. (2019). Monthly payment targeting and the demand for maturity. *Working Paper*.
- , ---, --- and PRATT, R. (2018). The capitalization of consumer financing into durable goods prices. *Working Paper*.
- ATTANASIO, O., GOLDBERG, P. and KYRIAZIDOU, E. (2008). Credit constraints in the market for consumer durables: Evidence from micro data on car loans. *International Economic Review*, **49** (2), 401–436.
- AUSUBEL, L. M. (1991). The failure of competition in the credit card market. *The American Economic Review*, pp. 50–81.
- BENETTON, M., GAVAZZA, A. and SURICO, P. (2019). *Mortgage Pricing and Monetary Policy*. Tech. rep., working paper.
- BERRY, S., LEVINSOHN, J. and PAKES, A. (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pp. 841–890.
- BERRY, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pp. 242–262.

- BHATTACHARYA, V., ILLANES, G. and PADI, M. (2019). *Fiduciary duty and the market for financial advice*. Tech. rep., National Bureau of Economic Research.
- BHUTTA, N., FUSTER, A. and HIZMO, A. (2019). Paying too much? price dispersion in the us mortgage market. *Working Paper*.
- BIGLAISER, G., LI, F., MURRY, C. and ZHOU, Y. (2019). Intermediaries and product quality in used car markets. *Available at SSRN 3003562*.
- BORDALO, P., GENNAIOLI, N. and SHLEIFER, A. (2012). Salience theory of choice under risk. *The Quarterly Journal of Economics*, **127** (3), 1243–1285.
- , --- and --- (2013). Salience and consumer choice. *Journal of Political Economy*, **121** (5), 803–843.
- BROWN, J. and JANSEN, M. (2019). Consumer protection laws in auto lending. *Working Paper*.
- CENTER FOR RESPONSIBLE LENDING, C. F. A. R., CONSUMER FEDERATION OF AMERICA, SAFETY, N. C. L. C., NATIONAL ASSOCIATION OF CONSUMER ADVOCATES and OF LA RAZA, N. C. (2012). Comments to the federal trade commission: Motor vehicle roundtables - project number p104811.
- CFPB (2016). Consumer voices on automobile financing.
- CHETTY, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annu. Rev. Econ.*, **1** (1), 451–488.
- , LOONEY, A. and KROFT, K. (2009). Salience and taxation: Theory and evidence. *American economic review*, **99** (4), 1145–77.
- COHEN, M. A. (2012). Imperfect competition in auto lending: Subjective markup, racial disparity, and class action litigation. *Review of Law and Economics*, **8** (1), 21–58.
- COLLARD-WEXLER, A., GOWRISANKARAN, G. and LEE, R. S. (2014). “nash-in-nash” bargaining: a microfoundation for applied work.
- DAVIS, D. (2012). Auto loans: The state of lending in America and its impact on U.S. households. *mimeo: Center for Responsible Lending*.
- DELLAVIGNA, S. (2018). Structural behavioral economics. *mimeo*.
- EGAN, M. (2018). Brokers vs. retail investors: Conflicting interests and dominated products. *Journal of Finance, Forthcoming*.

- , MATVOS, G. and SERU, A. (2019). The market for financial adviser misconduct. *Journal of Political Economy*, **127** (1), 233–295.
- EINAV, L., JENKINS, M. and LEVIN, J. (2012). Contract pricing in consumer credit markets. *Econometrica*, **80** (4), 1387–1432.
- , --- and --- (2013). The impact of credit scoring on consumer lending. *The RAND Journal of Economics*, **44** (2), 249–274.
- ELLISON, G. (2005). A model of add-on pricing. *The Quarterly Journal of Economics*, **120** (2), 585–637.
- and ELLISON, S. F. (2009). Search, obfuscation, and price elasticities on the internet. *Econometrica*, **77** (2), 427–452.
- FCA (2019). Motor finance discretionary commission models and consumer credit commission disclosure. Consultation Paper CP19/28.
- GABAIX, X. and LAIBSON, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics*, **121** (2), 505–540.
- GAVAZZA, A., LIZZERI, A. and ROKETSKIY, N. (2014). A quantitative analysis of the used-car market. *American Economic Review*, **104** (11), 3668–3700.
- GAYNOR, M. and TOWN, R. J. (2011). Competition in health care markets. In *Handbook of health economics*, vol. 2, Elsevier, pp. 499–637.
- GRUBB, M. D. (2015). Behavioral consumers in industrial organization: An overview. *Review of Industrial Organization*, **47** (3), 247–258.
- and OSBORNE, M. (2015). Cellular service demand: Biased beliefs, learning, and bill shock. *American Economic Review*, **105** (1), 234–71.
- GUERRE, E., PERRIGNE, I. and VUONG, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, **68** (3), 525–574.
- GUIO, L., POZZI, A., TSOY, A., GAMBACORTA, L. and MISTRULLI, P. E. (2018). The cost of steering in financial markets: Evidence from the mortgage market. *Available at SSRN 2951042*.
- HANDEL, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, **103** (7), 2643–82.

- HORN, H. and WOLINSKY, A. (1988). Bilateral monopolies and incentives for merger. *The RAND Journal of Economics*, pp. 408–419.
- HORTACSU, A., MADANIZADEH, S. A. and PULLER, S. L. (2017). Power to choose? an analysis of consumer inertia in the residential electricity market. *American Economic Journal: Economic Policy*, **9** (4), 192–226.
- KOSZEGI, B. and SZEIDL, A. (2012). A model of focusing in economic choice. *The Quarterly journal of economics*, **128** (1), 53–104.
- LAIBSON, D., MAXTED, P., REPETTO, A. and TOBACMAN, J. (2015). Estimating discount functions with consumption choices over the lifecycle.
- LAL, R. and MATUTES, C. (1994). Retail pricing and advertising strategies. *Journal of Business*, pp. 345–370.
- LUSARDI, A. and MITCHELL, O. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, **52** (1), 5–44.
- MACKOWIAK, B., MATEJKA, F. and WIEDERHOLT, M. (2018). Rational inattention: A disciplined behavioral model. *Working Paper*.
- MADRIAN, B. C., BESHEARS, J., J. CHOI, J. and LAIBSON, D. (forthcoming). *Behavioral Economics and Household Finance*, Elsevier.
- MIAN, A. and SUFI, A. (2016). Household debt and defaults from 2000 to 2010: The credit supply view. *evidence and innovation in housing law and policy*, p. 257.
- MORTON, F. S., ZETTELMEYER, F. and SILVA-RISSO, J. (2001). Internet car retailing. *The Journal of Industrial Economics*, **49** (4), 501–519.
- , --- and --- (2003). Consumer information and discrimination: Does the internet affect the pricing of new cars to women and minorities? *Quantitative marketing and Economics*, **1** (1), 65–92.
- MURRY, C. (2017). Advertising in vertical relationships: An equilibrium model of the automobile industry.
- NEVO, A. (2000). A practitioner’s guide to estimation of random-coefficients logit models of demand. *Journal of economics & management strategy*, **9** (4), 513–548.
- NURSKI, L. and VERBOVEN, F. (2016). Exclusive dealing as a barrier to entry? evidence from automobiles. *The Review of Economic Studies*, **83** (3), 1156–1188.

- PAKES, A., PORTER, J., HO, K. and ISHII, J. (2015). Moment inequalities and their application. *Econometrica*, **83** (1), 315–334.
- REYNOLDS, C. E. and COX, S. E. (2020). Buckle up: Navigating auto sales and financing.
- ROBLES-GARCIA, C. (2019). Competition and incentives in mortgage markets: The role of brokers. *Working Paper*.
- SALZ, T. (2017). Intermediation and competition in search markets: An empirical case study.
- SCHLAGENHAUF, D. and RICKETTS, L. (2016). The quarterly debt monitor: Trends in consumer debts in st. louis, little rock, louisville, memphis—and beyond. *In the Balance*.
- SIMS, C. (2006). Rational inattention: Beyond the linear-quadratic case. *The American Economic Review*, **96** (2), 158–163.
- STANGO, V., YOONG, J. and ZINMAN, J. (2017). *Quicksand or Bedrock for Behavioral Economics? Assessing Foundational Empirical Questions*. Tech. rep., National Bureau of Economic Research.
- SULLIVAN, M. W., JONES, M. T. and REYNOLDS, C. L. (2020). The auto buyer study: Lessons from in-depth consumer interviews and related research.
- VERBOVEN, F. (1999). Product line rivalry and market segmentation-with an application to automobile optional engine pricing. *The Journal of Industrial Economics*, **47** (4), 399–425.
- WOODWARD, S. E. and HALL, R. E. (2010). Consumer confusion in the mortgage market: Evidence of less than a perfectly transparent and competitive market. *American Economic Review*, **100** (2), 511–15.
- and --- (2012). Diagnosing consumer confusion and sub-optimal shopping effort: Theory and mortgage-market evidence. *American Economic Review*, **102** (7), 3249–3276.

APPENDIX

A Proofs of Propositions

Proof of Proposition 1

To solve the maximization problem in (1), ignore the second constraint and the constraint that $p_i \geq 0$ such that it becomes:

$$\begin{aligned} \max_{r_i, p_i} \quad & p_i + (\tau_i p_i - \kappa_i)(r_i - b_i)\alpha + \beta - c_i \\ \text{s.t.} \quad & -\tau_i p_i - (\tau_i p_i - \kappa_i)r_i \geq \bar{u}_i \\ & r_i \geq b_i. \end{aligned}$$

Any solution to this problem clearly fulfills its constraint with equality. Therefore, we can rewrite this problem to be:

$$\max_{r_i} -\frac{r_i \kappa_i - \bar{u}_i}{(1+r_i)\tau_i} + \left[\frac{r_i \kappa_i - \bar{u}_i}{1+r_i} - \kappa_i \right] (r_i - b_i)\alpha + \beta - c_i \quad r_i \geq b_i$$

The derivative with respect to r_i is:

$$\begin{aligned} & \frac{\kappa_i + \bar{u}_i}{(1+r_i)^2 \tau_i} + \frac{\kappa_i + \bar{u}_i}{(1+r_i)^2} \cdot (r_i - b_i) \cdot \alpha - \frac{\kappa_i + \bar{u}_i}{1+r_i} \alpha \leq 0 \\ \Leftrightarrow & \left[\frac{1}{\tau_i} + (r_i - b_i)\alpha - (1+r_i)\alpha \right] \leq 0 \\ \Leftrightarrow & 1 \geq \alpha(1+b_i) \cdot \tau_i \end{aligned}$$

Whenever the last equation is fulfilled, the derivative is negative such that the optimal solution features $r_i = b_i$, which corresponds to a zero markup. Also note that the derived price is positive because $|\bar{u}_i| \geq \kappa_i$ is assumed for all i . Hence, the constraint $p_i \geq 0$ is fulfilled. Finally, consider the second constraint. It will be fulfilled if:

$$-b_i(\tau_i p_i - \kappa_i) \geq -(\tau_i p_i - \kappa_i)\mathbb{E}[r^L] - s_i,$$

which is trivially fulfilled for $b_i \leq \mathbb{E}[r^L]$. If in contrast, $b_i > \mathbb{E}[r^L]$, we get that the second constraint is fulfilled if:

$$\begin{aligned} p & \leq \left(\frac{s_i}{r_i - \mathbb{E}[r^L]} + \kappa_i \right) \frac{1}{\tau_i} \\ \Leftrightarrow & \frac{r_i \kappa_i - \bar{u}_i}{(1+r_i)\tau_i} \leq \left(\frac{s_i}{r_i - \mathbb{E}[r^L]} + \kappa_i \right) \frac{1}{\tau_i} \\ \Leftrightarrow & s_i \geq (\mathbb{E}[r^L] - b_i) \left(\frac{\bar{u}_i + \kappa_i}{1+b_i} \right). \end{aligned} \tag{15}$$

As long as (15) holds, we will thus get that $r_i^* = b_i$ for any optimal offer by the dealer. Finally, suppose that there is an optimal offer featuring $r_i > b_i$ and (15) being not satisfied such that we have: $s_i < (\mathbb{E}[r^L] - b_i) \left(\frac{\bar{u}_i + \kappa_i}{1 + b_i} \right)$. From the proof of Proposition 2, we know that for any solution with $r_i^* > b_i$ the first constraint will be binding. Then from the arguments above it can not be the case that only the first constraint is binding, because $r_i = b_i$ would be optimal. Hence, both constraints have to be binding. Then, we get

$$p_i = \frac{r_i \kappa_i - \bar{u}_i}{(1 + r_i) \tau_i} = \left(\frac{s_i}{r_i - \mathbb{E}[r^L]} + \kappa_i \right) \frac{1}{\tau_i}$$

$$\Leftrightarrow s_i = -\frac{r_i - \mathbb{E}[r^L]}{1 + r_i} (\kappa_i + \bar{u}_i)$$

This however, is a contradiction to $s_i < (\mathbb{E}[r^L] - b_i) \left(\frac{\bar{u}_i + \kappa_i}{1 + b_i} \right)$, as $|\bar{u}_i| > \kappa_i$ and $\frac{r_i - \mathbb{E}[r^L]}{1 + r_i}$ is increasing in r_i . As a consequence, even if (15) does not hold, there cannot be an optimal offer featuring $r_i^* > b_i$, which concludes the proof.

Proof of Proposition 2

First note that the optimal solution to the maximization problem in (1) will, in this case, satisfy at least one of its constraints with equality. Otherwise the dealer could increase p_i and thereby increase profits. In the first step of the proof we argue that if the second constraint is binding at the optimum the first one is also binding. For this purpose, suppose in contradiction that there exists an optimal solution $(r_i^* > b_i, p_i^* > 0)$ such that only the second constraint binds. Then (r_i^*, p_i^*) is a maximizer of:

$$\max_{p_i, r_i} L = p_i + (\tau_i p_i - \kappa_i)(r_i - b_i)\alpha + \beta - c_i +$$

$$\mu \left(M_i(r_i(\tau_i p_i - \kappa_i)) - \int M_i((\tau_i p_i - \kappa_i)\tau_i r^L) g_i(r^L) dr^L - s_i \right)$$

Leading to the following first order conditions:

$$\frac{dL}{dp_i} = 1 + \tau_i \alpha (r_i - b_i) + \mu \left(M_i'(r_i(\tau_i p_i - \kappa_i)) \tau_i r_i - \int M_i'((\tau_i p_i - \kappa_i)r^L) \tau_i r^L g_i(r^L) dr^L \right) = 0$$

$$\frac{dL}{dr_i} = \alpha (\tau_i p_i - \kappa_i) + \mu M_i'(r_i(\tau_i p_i - \kappa_i)) (\tau_i p_i - \kappa_i) = 0$$

Since both of these hold in equilibrium we get

$$\frac{1 + \alpha \tau_i (r_i - b_i)}{M_i'(r_i(\tau_i p_i - \kappa_i)) \tau_i r_i - \int M_i'((\tau_i p_i - \kappa_i)r^L) \tau_i r^L g_i(r^L) dr^L} = \frac{\alpha}{M_i'(r_i(\tau_i p_i - \kappa_i))}$$

$$M_i'(r_i(\tau_i p_i - \kappa_i)) (1 + \tau_i \alpha (r_i - b_i)) = \alpha (M_i'(r_i(\tau_i p_i - \kappa_i)) \tau_i r_i - \int M_i'((\tau_i p_i - \kappa_i)r^L) \tau_i r^L g_i(r^L) dr^L)$$

$$M_i'(r_i(\tau_i p_i - \kappa_i)) (1 - \tau_i \alpha b_i) = -\alpha \int M_i'((\tau_i p_i - \kappa_i)r^L) r^L \tau_i g_i(r^L) dr^L$$

The last equality cannot be fulfilled since the left hand side is positive but the right hand side is negative, which establishes the contradiction.

Using the insights from the first step, there are only two possibilities left, either both constraints are binding or only the first one is. Suppose only the first one is binding such that the optimal solution to the maximization problem in (1) maximizes:

$$\begin{aligned} \max_{p_i, r_i} \quad & p_i + (\tau_i p_i - \kappa_i)(r_i - b_i)\alpha + \beta - c_i \\ \text{s.t.} \quad & -\tau_i p_i - M_i((\tau_i p_i - \kappa_i)r_i) = \bar{u}_i \\ & r_i > b_i, p_i \geq 0 \end{aligned}$$

The corresponding Lagrange function is given by:

$$\max_{p_i, r_i} L = p_i + (\tau_i p_i - \kappa_i)(r_i - b_i)\alpha + \beta - c + \mu(\bar{u}_i + \tau_i p_i + M_i(r_i(\tau_i p_i - \kappa_i))),$$

leading to the following first order conditions:

$$\begin{aligned} \frac{dL}{dp_i} &= 1 + \tau_i \alpha (r_i - b_i) + \mu \left(\tau_i + M'_i(r_i(\tau_i p_i - \kappa_i)) \tau_i r_i \right) = 0 \\ \frac{dL}{dr_i} &= \alpha (\tau_i p_i - \kappa_i) + \mu M'_i(r_i(\tau_i p_i - \kappa_i)) (\tau_i p_i - \kappa_i) = 0 \end{aligned}$$

Since both of these hold in equilibrium we get:

$$\begin{aligned} \frac{1 + \tau_i \alpha (r_i - b_i)}{\tau_i + M'_i(r_i(\tau_i p_i - \kappa_i)) \tau_i r_i} &= \frac{\alpha}{M'_i(r_i(\tau_i p_i - \kappa_i))} \\ \Leftrightarrow M'_i(r_i(\tau_i p_i - \kappa_i)) &= \frac{\alpha \tau_i}{1 - \tau_i \alpha b_i} \end{aligned}$$

In particular, any solution to the problem in which only the first constraint is binding will satisfy this equation. Next suppose that both constraints are binding. In this case, the corresponding Lagrange function is:

$$\begin{aligned} L &= p_i + (\tau_i p_i - \kappa_i)(r_i - b_i)\alpha + \beta - c \\ &+ \mu \left(\int M_i((\tau_i p_i - \kappa_i)r^L) g_i(r^L) dr^L + s_i - M_i(r_i(\tau_i p_i - \kappa_i)) \right) \\ &+ \lambda (-\bar{u}_i - \tau_i p_i - M_i(r_i(\tau_i p_i - \kappa_i))) \end{aligned}$$

Then from the Kuhn-Tucker Theorem, we know that there exists $\lambda > 0$ and $\mu > 0$ such that the following two equations are fulfilled:

$$\begin{aligned} \frac{dL}{dp_i} &= 1 + \tau_i \alpha (r_i - b_i) + \mu \left(\int M'_i((\tau_i p_i - \kappa_i)r^L) \tau_i r^L g_i(r^L) dr^L - M'_i(r_i(\tau_i p_i - \kappa_i)) r_i \tau_i \right) \\ &+ \lambda \left(-\tau_i - M'_i((\tau_i p_i - \kappa_i)r_i) \tau_i r_i \right) = 0 \\ \frac{dL}{dr_i} &= \alpha (\tau_i p_i - \kappa_i) - \mu M'_i(r_i(\tau_i p_i - \kappa_i)) (\tau_i p_i - \kappa_i) - \lambda M'_i(r_i(\tau_i p_i - \kappa_i)) (\tau_i p_i - \kappa_i) = 0 \end{aligned}$$

From the second equation, we get:

$$\frac{\alpha - \mu M'_i(r_i(\tau_i p_i - \kappa_i))}{M'_i(r_i(\tau_i p_i - \kappa_i))} = \lambda$$

Substituting this into the first equation yields:

$$\begin{aligned} & 1 + \tau_i \alpha (r_i - b_i) + \mu \left(\int M'_i((\tau_i p_i - \kappa_i) r^L) \tau_i r^L g_i(r^L) dr^L - M'_i(r_i(\tau_i p_i - \kappa_i)) \tau_i r_i \right) \\ & + \frac{\alpha - \mu M'_i(r_i(\tau_i p_i - \kappa_i))}{M'_i(r_i(\tau_i p_i - \kappa_i))} \left(-\tau_i - M'_i((\tau_i p_i - \kappa_i) r_i) r_i \tau_i \right) = 0 \\ \Leftrightarrow & 1 - \tau_i \alpha b_i + \mu \int M'_i((\tau_i p_i - \kappa_i) r^L) \tau_i r^L g_i(r^L) dr^L - \tau_i \frac{\alpha - \mu M'_i(r_i(\tau_i p_i - \kappa_i))}{M'_i(r_i(\tau_i p_i - \kappa_i))} = 0 \\ \mu = & \frac{\tau_i \alpha b_i + \frac{\tau_i \alpha}{M'_i(r_i(\tau_i p_i - \kappa_i))} - 1}{\tau_i + \int M'_i((\tau_i p_i - \kappa_i) r^L) r^L \tau_i g_i(r^L) dr^L} \end{aligned}$$

since μ is positive $\tau_i \alpha b_i + \frac{\tau_i \alpha}{M'_i(r_i(\tau_i p_i - \kappa_i))} - 1 > 0$ holds, which is equivalent to the statement that $M'_i(r_i^*(\tau_i p_i^* - \kappa_i)) < \frac{\tau_i \alpha}{1 - \tau_i \alpha b_i}$.

Proof of Proposition 3

For the first part assume that $M'_i(x) \leq 1 \quad \forall x$. Let $r_i^* > b_i, p_i^* > 0$ be the optimal offer. Denote by \tilde{p}_i the price that is optimal for the dealer given that he offers $r_i = b_i$. Note first that $\tilde{p}_i > p_i^*$ holds. In particular, both constraints in (1) have to be slack if $r_i = b_i$ and $p_i = p_i^*$. Thus, the dealer can profit from increasing the price and we get $\tilde{p}_i > p_i^*$. As a consequence, the first constraint in (1) will also be binding at b_i, \tilde{p}_i . To see this, recall that the first constraint is binding at r_i^*, p_i^* due to Proposition 2. We therefore get:

$$\bar{u}_i + \tau_i p_i^* = -M_i((\tau_i p_i^* - \kappa_i) r_i) \geq - \int M_i((\tau_i p_i^* - \kappa_i) r^L) g_i(r^L) dr^L - s_i.$$

As $\tilde{p}_i > p_i^*$,

$$\bar{u}_i + \tau_i \tilde{p}_i > \bar{u}_i + \tau_i p_i^* \geq - \int M_i((\tau_i p_i^* - \kappa_i) r^L) g_i(r^L) dr^L - s_i \geq - \int M_i((\tau_i \tilde{p}_i - \kappa_i) r^L) g_i(r^L) dr^L - s_i.$$

Thus the second constraint is slack and the first one has to be binding. Hence, \tilde{p}_i and p_i^* are given by:

$$\begin{aligned} -\tau_i \tilde{p}_i - M_i((\tau_i \tilde{p}_i - \kappa_i) b_i) &= \bar{u}_i \\ -\tau_i p_i^* - M_i((\tau_i p_i^* - \kappa_i) r_i) &= \bar{u}_i. \end{aligned}$$

The optimal offer has to imply a weakly higher profit than offering $r_i = b_i$. Hence,

$$\begin{aligned} p_i^* + (\tau_i p_i^* - \kappa_i) \alpha (r_i^* - b_i) + \beta - c_i &\geq \tilde{p}_i + \beta - c_i \\ \Leftrightarrow \frac{-\bar{u}_i - M_i((\tau_i p_i^* - \kappa_i) r_i)}{\tau_i} + (\tau_i p_i^* - \kappa_i) \alpha (r_i^* - b_i) &\geq \frac{-\bar{u}_i - M_i((\tau_i \tilde{p}_i - \kappa_i) b_i)}{\tau_i} \\ \Leftrightarrow M_i((p_i^* - \kappa_i) r_i) - M_i((\tilde{p}_i - \kappa_i) b_i) &\leq \tau_i (\tau_i p_i^* - \kappa_i) \alpha (r_i^* - b_i) \end{aligned} \tag{16}$$

Using that $\tilde{p}_i \geq p_i^*$, we therefore get that:

$$\begin{aligned}
& (\tau_i p_i^* - \kappa_i)(r_i^* - b_i) - \left[M_i((\tau_i p_i^* - \kappa_i)r_i) - M((\tau_i p_i^* - \kappa_i)b_i) \right] \\
&= (\tau_i p_i^* - \kappa_i)(r_i - b_i) - \left[M_i((\tau_i p_i^* - \kappa_i)r_i) - M((\tau_i \tilde{p}_i - \kappa_i)b_i) \right. \\
&\quad \left. + M_i((\tau_i \tilde{p}_i - \kappa_i)b_i) - M((\tau_i p_i^* - \kappa_i)b_i) \right] \\
&\geq (\tau_i p_i^* - \kappa_i)(r_i^* - b_i) - \tau_i(\tau_i p_i^* - \kappa_i)\alpha(r_i^* - b_i) - \left[M_i((\tau_i \tilde{p}_i - \kappa_i)b_i) - M((\tau_i p_i^* - \kappa_i)b_i) \right] \\
&\geq (\tau_i p_i^* - \kappa_i)(r_i - b_i) - \tau_i(\tau_i p_i^* - \kappa_i)\alpha(r_i^* - b_i) - ((\tau_i \tilde{p}_i - \kappa_i)b_i - (\tau_i p_i^* - \kappa_i)b_i) \\
&= (\tau_i p_i^* - \kappa_i)r_i - (\tau_i \tilde{p}_i - \kappa_i)b_i - \tau_i(\tau_i p_i^* - \kappa_i)\alpha(r_i^* - b_i), \tag{17}
\end{aligned}$$

where the last inequality follows from the fact that $M'(x) \leq 1 \forall x$. Moreover, (16) implies that:

$$\begin{aligned}
& p_i^* + (\tau_i p_i^* - \kappa_i)\alpha(r_i^* - b_i) + \beta - c_i \geq \tilde{p}_i + \beta - c_i \\
&\Leftrightarrow b_i[\tau_i p_i^* - \kappa_i + \tau_i(\tau_i p_i^* - \kappa_i)\alpha(r_i^* - b_i)] \geq [\tau_i \tilde{p}_i - \kappa_i]b_i
\end{aligned}$$

Substituting this into (17) yields:

$$\begin{aligned}
& (\tau_i p_i^* - \kappa_i)r_i - (\tau_i p_i^* - \kappa_i)b_i - (M_i((\tau_i p_i^* - \kappa_i)r_i) - M_i((\tau_i p_i^* - \kappa_i)b_i)) \\
&\geq (\tau_i p_i^* - \kappa_i)r_i - b_i[\tau_i p_i^* - \kappa_i + \tau_i(\tau_i p_i^* - \kappa_i)\alpha(r_i^* - b_i)] - \tau_i(\tau_i p_i^* - \kappa_i)\alpha(r_i^* - b_i) \\
&= (\tau_i p_i^* - \kappa_i)(r_i - b_i)[1 - \tau_i b_i \alpha - \tau_i \alpha],
\end{aligned}$$

which concludes the proof of the first part of the proposition. Part (ii) of the proposition follows from Proposition 2 and the assumed convexity of M . In particular, convexity implies that:

$$\begin{aligned}
& M((\tau_i p_i^* - \kappa_i)r_i) - M_i(0) \leq M_i'((\tau_i p_i^* - \kappa_i)r_i)(\tau_i p_i^* - \kappa_i)r_i \\
&\leq \frac{\tau_i \alpha}{1 - \tau_i \alpha b_i}(\tau_i p_i^* - \kappa_i)r_i
\end{aligned}$$

Which implies that a lower bound for the finance charges that agents do not respond to is given by:

$$\begin{aligned}
& (\tau_i p_i^* - \kappa_i)r_i - M_i((\tau_i p_i^* - \kappa_i)r_i) \geq (\tau_i p_i^* - \kappa_i)r_i - \frac{\tau_i \alpha}{1 - \tau_i \alpha b_i}(\tau_i p_i^* - \kappa_i)r_i = \\
& \qquad \qquad \qquad (\tau_i p_i^* - \kappa_i)r_i \left[1 - \frac{\tau_i \alpha}{1 - \tau_i \alpha b_i} \right]
\end{aligned}$$

B Explanations for Differential Price Responsiveness

B.1 Time Structure

In this section, we argue that the time structure of payments is unlikely to explain our results. For this purpose, we first show that a model with multiple periods would yield the same estimates for

consumers' differential price responsiveness. Second, we show how our estimates would change if car dealers were to require a fixed fraction of the car price as the down payment.

B.1.1 A Multi-Period Model with Exogenous Down Payment

In this subsection, we assume that consumer preferences are defined over T periods, where T is the loan term in months. We still maintain Assumption 1 from Section 4.1 that consumer utility is additive in the disutility of the car price p and the disutility of finance charges $M_i(x)$. We further assume that the within-period utility function is stable over time and that utility is additive across periods, which are both standard assumptions in multi-period models.

Suppose consumer i has down payment κ_i and discounts utility from month t by δ_t^i . We allow δ_t^i to vary with i and t in an arbitrary way. Since auto loans fully amortize, consumer i 's total loan payments will be

$$T \cdot \frac{r_i(1+r)^T}{(1+r_i)^T - 1} (\tau_i p_i - \kappa_i) \equiv (1 + \hat{r}_i)(\tau_i p_i - \kappa_i),$$

which equals T times the monthly loan payment. \hat{r} denotes the fraction of the loan amount that is due in interest rate charges over the term of the loan, which is equivalent to the two-period interest rate that we use in the main part of the paper.

First, consider Proposition 1. If $M_i(x) = x$, then consumer utility is given by:

$$u = -\kappa_i - \sum_{t=1}^T \frac{\delta_t^i}{T} [(1 + \hat{r}_i)(\tau_i p_i - \kappa_i)] = -\kappa_i - \left(\sum_{t=1}^T \frac{\delta_t^i}{T} \right) \cdot [(1 + \hat{r}_i)(\tau_i p_i - \kappa_i)]$$

Hence, consumer utility is the same as in the one-period model except for multiplication by a constant and the addition of an exogenous variable. Therefore, we can redefine the utility of the outside option such that the maximization problem is identical to the one in Section 4.1. In particular, the dealer's marginal incentives to allocate charges to the vehicle or the loan do not change and none of the relevant calculations are affected. Therefore Proposition 1 holds in this multi-period model.

The same argument also holds for the proof of Proposition 2. In our multi-period model, consumer utility is given by:

$$u = -\kappa_i - \sum_{t=1}^T \frac{\delta_t^i}{T} [(\tau_i p_i - \kappa_i) + M_i((\tau_i p_i - \kappa_i)\hat{r}_i)] = -\kappa_i - \left(\sum_{t=1}^T \frac{\delta_t^i}{T} \right) \cdot [(\tau_i p_i - \kappa_i) + M_i((\tau_i p_i - \kappa_i)\hat{r}_i)]$$

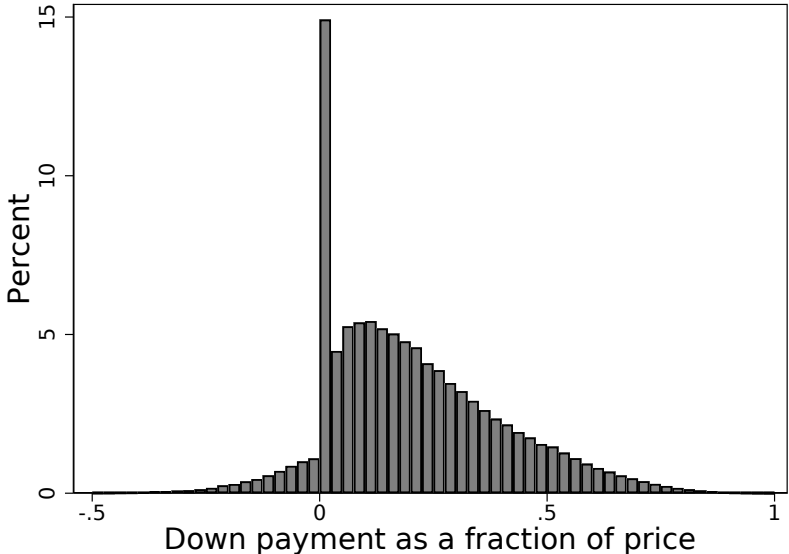
Hence, consumer utility again is the same as in the one-period model except for multiplication by a constant and the addition of an exogenous variable. We can therefore redefine the outside option and end up with the same maximization problem for the dealer. Therefore, since our estimates for $M_i(\cdot)$,

B_i^O , and B_i^M do not depend on \bar{u}_i , they are not affected by intertemporal considerations if the down payment is independent of the allocation of charges. Section B.1.2 examines our estimates if a larger car price causes a larger down payment.

B.1.2 Multi-Period Model with Down Payment as Fixed Fraction of Price

If a larger car price causes a larger down payment, intertemporal preferences could affect our results. Suppose, for example, the down payment is a fixed percentage of the car price. In this case, impatient or credit-constrained consumers may prefer to pay more for the loan to reduce the car price and hence the down payment. Figure A1 shows the distribution of down payments as a fraction of the car price. This figure, together with our discussions with industry experts, reveal that down payments for prime consumers are generally not determined in this way.⁴⁴ Variation of the relative size of the down payment is large and a marginal increase in the car price does not appear to increase the down payment.

Figure A1: Distribution of Down Payment as a Fraction of Car Price



NOTE: The figure shows the distribution of down payment as a fraction of the car price. The down payment is the price of the car minus the loan amount. Down payments can be negative if upfront fees, vehicle add-ons, or other costs are included in the loan. The figure shows that consumers in our data do not appear to be required to provide a specific fraction of the car price as a down payment.

⁴⁴Down payment requirements are often nonbinding for prime consumers. When they are nonzero, they are typically related to the collateral *value*, not its price, because collateral value is what affects both the consumer’s default decision and what the lender can recover if the consumer does default.

Still, we break down our estimates by the size of the down payment in Table A2. We find a considerable wedge between car and loan price sensitivities across the entire distribution of down payments. In particular, the wedge is substantial for zero and negative down payments, for which time preferences cannot explain any of our estimates. The median B_i^O for consumers in the 10th to 25th percentile of the down payment distribution is \$630 while it is \$563 for consumers between the 75th and 90th percentile. Overall, these results show that intertemporal preferences and credit constraints cannot explain the substantial dealer markups in our data.

To further address this point, we recompute $M'_i(\cdot)$ under the (evidently counterfactual) assumption that dealers require a down payment that is a fixed percentage of the car price, conditional on various discount factors for the consumer.⁴⁵ The results are summarized in Table A1. Recall that we are restricting our sample to individuals with credit scores above 720. Hence, the sample consists of individuals that are unlikely to be either impatient or credit constrained and therefore should have reasonably high discount factors. Our results show quantitatively that even if the down payment were a fixed fraction of the price, our estimates would barely change for discount factors above 0.90.

Table A1: Estimates if Down Payments Are a Fraction of the Car Price


	Mean	p10	p25	p50	p75	p90
M' with $\delta = 1$	0.80	0.64	0.69	0.84	0.89	0.93
M' with $\delta = 0.98$	0.80	0.64	0.69	0.84	0.89	0.92
M' with $\delta = 0.95$	0.81	0.66	0.70	0.85	0.91	0.94
M' with $\delta = 0.90$	0.84	0.67	0.73	0.86	0.94	0.98
M' with $\delta = 0.6$	1.08	0.73	0.87	1.03	1.27	1.51

NOTE: Table shows percentiles of the distribution of estimated M' , assuming that down payments are a fixed fraction of the car price and loan term length varies as in the data, but all consumers have an annual discount factor equal to the given value of δ .

The estimates in Table A1 are derived as follows. We make all the assumptions on preferences across time made in Section B.1.1. To derive a specific formula for M' we can take to the data, we further assume consumers discount the future at exponential rate δ . We assume that the down payment κ_t is given by $\kappa\tau_i p_t$, where $\kappa \in (0, 1)$ is an exogenous fraction. Moreover, we assume that the down payment is due immediately while all other payments are financed by a fully amortizing loan.

⁴⁵In doing so we first compute the fraction of the price that is equal to the observed down payment. We then assume that the dealer would require this fraction for any car price.

Figure A2: 10th, 50th, and 90th Percentiles of M' , B_i^O , and B_i^M by Down Payment Tercile



NOTE: For different downpayment categories the figure shows the 10th, 50th, and 90th percentiles of M' , B_i^O , and B_i^M in the upper, middle, and lower panel, respectively. The first down payment category consists of observations with zero or negative down payments. The other three down payment categories are terciles of down payment relative to the price of the vehicle, for observations with positive down payments. M' and B_i^O are only defined for the 77.8 percent of observations with positive markups. The figure shows that the distribution of our estimates vary little with down payment.

Denote by T the term of the loan in months, and by $\hat{r}_i = \frac{r_i(1+\frac{r_i}{12})^T}{(1+\frac{r_i}{12})^T - 1} T - 1$ the fraction of the payment that is due to finance charges if the yearly interest rate is r_i . Consumer utility is then given by:

$$u = -\kappa\tau_i p_i - \sum_{j=1}^T \frac{\delta^j}{T} [(1-\kappa)\tau_i p_i + M_i((1-\kappa)\tau_i p_i \hat{r}_i)] = -\kappa\tau_i p_i - \frac{\delta - \delta^{T+1}}{(1-\delta)T} [(1-\kappa)\tau_i p_i + M_i((1-\kappa)\tau_i p_i \hat{r}_i)]$$

The car dealer's optimization problem then becomes:

$$\begin{aligned} \max_{\hat{r}_i, p_i} \quad & p_i + (1-\kappa)\tau_i p_i (\hat{r}_i - \hat{b}_i) \alpha + \beta - c_i \\ \text{s.t.} \quad & -\kappa\tau_i p_i - \frac{\delta - \delta^{T+1}}{(1-\delta)T} [(1-\kappa)\tau_i p_i + M_i((1-\kappa)\tau_i p_i \hat{r}_i)] \geq \bar{u} \\ & -M_i((1-\kappa)\tau_i p_i \hat{r}_i) \geq -\int M_i((1-\kappa)\tau_i p_i \hat{r}_i^L) g(r^L) - s_i, \\ & r_i, p_i \geq 0 \end{aligned}$$

With the same arguments as above, we can concentrate on the case where the first constraint is binding, but the second not. The Lagrangian is then given by:

$$L = p_i + (1-\kappa)\tau_i p_i (\hat{r}_i - \hat{b}_i) \alpha + \beta - c_i + \mu \left[\bar{u} + \kappa\tau_i p_i + \frac{\delta - \delta^{T+1}}{(1-\delta)T} [(1-\kappa)\tau_i p_i + M_i((1-\kappa)\tau_i p_i \hat{r}_i)] \right].$$

The derivatives are given by:

$$\begin{aligned} \frac{dL}{d\hat{r}_i} &= (1-\kappa)\tau_i p_i \alpha + \mu \frac{\delta - \delta^{T+1}}{(1-\delta)T} M_i'((1-\kappa)\tau_i p_i \hat{r}_i) \tau_i p_i (1-\kappa) \\ \frac{dL}{dp_i} &= 1 + (1-\kappa)(\hat{r}_i - \hat{b}_i) \tau_i \alpha + \mu \left[\kappa\tau_i + \frac{\delta - \delta^{T+1}}{(1-\delta)T} \left[\tau_i(1-\kappa) + M_i'((1-\kappa)\tau_i p_i \hat{r}_i) \tau_i \hat{r}_i (1-\kappa) \right] \right]. \end{aligned}$$

Since both have to equal to zero simultaneously, we get:

$$\begin{aligned} \frac{\alpha}{\frac{\delta - \delta^{T+1}}{(1-\delta)T} M_i'((1-\kappa)\tau_i p_i \hat{r}_i)} &= \frac{1 + (1-\kappa)(\hat{r}_i - \hat{b}_i) \tau_i \alpha}{\kappa\tau_i + \frac{\delta - \delta^{T+1}}{(1-\delta)T} \left[\tau_i(1-\kappa) + M_i'((1-\kappa)\tau_i p_i \hat{r}_i) \hat{r}_i (1-\kappa) \tau_i \right]} \\ \Leftrightarrow \frac{\delta - \delta^{T+1}}{(1-\delta)T} M_i'((1-\kappa)\tau_i p_i \hat{r}_i) [1 + (1-\kappa)\tau_i \alpha (\hat{r}_i - \hat{b}_i) - (1-\kappa)\tau_i \hat{r}_i \alpha] &= \alpha \left[\tau_i \kappa + (1-\kappa)\tau_i \frac{\delta - \delta^{T+1}}{(1-\delta)T} \right] \\ \Leftrightarrow M_i'((1-\kappa)\tau_i p_i \hat{r}_i) &= \frac{\tau_i \alpha \left[\kappa + (1-\kappa) \frac{\delta - \delta^{T+1}}{(1-\delta)T} \right]}{\frac{\delta - \delta^{T+1}}{(1-\delta)T} [1 - \tau_i (1-\kappa) \alpha \hat{b}_i]} \end{aligned}$$

B.2 Early Prepayment Risk

Consumers often prepay auto loans, either by trading in their vehicle for a new one, by prepaying in cash, or (less commonly) through refinancing. The potential effects of prepayment on our results

depends on when it occurs. Contracts between dealers and lenders include a “clawback” period, typically the first three to six months of the loan, during which the dealer bears prepayment risk. If the consumer prepays during the clawback period, the dealer refunds the entire dealer reserve to the lender. Otherwise, the dealer keeps the dealer reserve. We define “early” prepayment as prepayment that occurs during the clawback period and “late” prepayment as prepayment that occurs after it.

First, we explore the role of early prepayment. Early prepayment is not uncommon; in a sample of prime auto loans in the CCP, 5.7 percent are prepaid within the first 120 days. To account for this, we add to our baseline model an early prepayment probability $(1 - \gamma)$. Hence, the dealer’s maximization problem becomes:

$$\begin{aligned} \max_{r_i, p_i} \quad & p_i + (\tau_i p_i - \kappa_i) \cdot (r_i - b_i) \cdot \gamma \cdot \alpha - c & (18) \\ \text{s.t.} \quad & -\tau_i \cdot p_i - \gamma \cdot M_i((\tau_i p_i - \kappa_i) \cdot r_i) \geq -\bar{u}_i \\ & -\gamma \cdot M_i((\tau_i p_i - \kappa_i) \cdot r_i) \geq -\gamma \cdot M_i((\tau_i p_i - \kappa_i) \cdot r_i^E) - s_i, \\ & r_i, p_i \geq 0 \end{aligned}$$

This problem can be solved along the same lines as the one in Proposition 2. On the one hand, the consumer is now less sensitive to interest rates so the dealer has a stronger incentive to mark up the loan. On the other hand, the dealer values markups less because he may not keep the dealer reserve. The following corollary shows that the latter effect dominates. As a consequence, our estimates for $M'_i(\cdot)$ are upper bounds because they do not account for early prepayment risk.

Corollary 1. *Suppose that consumer i repays the loan early with probability $1 - \gamma$. Then, the statements from Proposition 2 hold true with the only change that:*

$$M'_i(r_i^* \cdot (\tau_i \cdot p_i^* - \kappa_i)) \leq \frac{\alpha \cdot \tau_i}{1 - \gamma \cdot \tau_i \cdot \alpha \cdot b_i}.$$

Proof. With the same arguments developed in the proof of Proposition 2, we know that either only the first one or both constraints are binding at the optimum. If only the first constraint is binding, we get the following Lagrange function:

$$\max_{p_i, r_i} L = p_i + (\tau_i p_i - \kappa_i)(r_i - b_i)\alpha\gamma + \beta - c + \mu(\bar{u}_i + \tau_i p_i + \gamma M_i(r_i(\tau_i p_i - \kappa_i)))$$

Leading to the following first order conditions:

$$\begin{aligned} \frac{dL}{dp_i} &= 1 + \tau_i \gamma \alpha (r_i - b_i) + \mu \left(\tau_i + \gamma M'_i(r_i(\tau_i p_i - \kappa_i)) \tau_i r_i \right) = 0 \\ \frac{dL}{dr_i} &= \gamma \alpha (\tau_i p_i - \kappa_i) + \mu \gamma M'_i(r_i(\tau_i p_i - \kappa_i)) (\tau_i p_i - \kappa_i) = 0 \end{aligned}$$

Since both of these hold in equilibrium we get:

$$\frac{1 + \tau_i \gamma \alpha (r_i^* - b_i)}{\tau_i + \gamma M'_i(r_i^*(\tau_i p_i^* - \kappa_i)) \tau_i r_i} = \frac{\alpha}{M'_i(r_i^*(\tau_i p_i^* - \kappa_i))}$$

$$\Leftrightarrow M'_i(r_i^*(\tau_i p_i^* - \kappa_i)) = \frac{\tau_i \alpha}{1 - \tau_i \gamma \alpha b_i}$$

If both constraints are binding, we can again use the same arguments as in the proof of Proposition 2 to show that:

$$M'_i(r_i^*(\tau_i p_i^* - \kappa_i)) < \frac{\tau_i \alpha}{1 - \gamma \tau_i \alpha b_i}$$

□

B.3 Late Prepayment Risk

Lenders bear all prepayment risk after the end of the clawback period. If borrowers consider late prepayment when buying a car and if dealers can screen “late” prepayers from “early” prepayers, then the logic outlined in B.2 does not hold, and borrowers may appear less sensitive to finance charges because they expect not to pay them. Late prepayment risk is substantial and so this is an important concern; 27.0 percent of prime auto loans (with term greater than two years) in the CCP are prepaid after 120 days but before two years.

To create a proxy for late prepayment risk, we run a logit regression predicting late prepayment in the CCP, using credit score, log loan amount, loan length, and state fixed effects. Coefficients are reported in Table A2. This is clearly a limited subset of the information available to both the borrower and the dealer, and yet it is remarkably predictive. 18.3 percent of those in the bottom decile of predicted risk prepay late, while 42.6 percent in the top decile do.

While these variables are highly predictive of late prepayment and readily observable to lenders, recall that lenders do not condition α on them. This is the first indication that borrowers with predictably higher prepayment risk do not pay higher markups.

Still, because lenders do not condition α on this information set, borrowers with predictably higher prepayment risk may agree to higher markups because the effective cost is lower. Tables A3, A4, A5 and A6 provide percentiles of markup, $M'_i(\cdot)$, B_i^O and B_i^M , respectively, conditional on percentiles of predicted late prepayment risk. Table A7 provides estimates from a regression with a large number of controls. The conditional correlation between observable late prepayment risk and our wedge estimates is negative. Hence, we estimate lower B_i^O 's and B_i^M 's for consumers with higher prepayment risk. This is strong evidence that late prepayment risk does not drive our results.

Table A2: Logit Regression of Prepayment Risk

	Prepayment	Early Prepayment	Late Prepayment
730 ≤ Credit score ≤ 749	-0.000224 (0.013)	0.0155 (0.028)	-0.00338 (0.013)
750 ≤ Credit score ≤ 769	0.00974 (0.013)	0.134*** (0.027)	-0.0209 (0.013)
770 ≤ Credit score ≤ 789	0.0262** (0.013)	0.247*** (0.027)	-0.0333** (0.013)
790 ≤ Credit score ≤ 809	0.0256** (0.013)	0.370*** (0.026)	-0.0709*** (0.013)
810 ≤ Credit score ≤ 829	-0.0536*** (0.013)	0.312*** (0.027)	-0.142*** (0.014)
830 ≤ Credit score ≤ 849	-0.136*** (0.016)	0.258*** (0.032)	-0.219*** (0.017)
850 ≤ Credit score ≤ 869	-0.249*** (0.031)	0.131** (0.063)	-0.310*** (0.033)
Log loan size	-0.340*** (0.007)	-0.283*** (0.013)	-0.294*** (0.007)
30 ≤ Loan term ≤ 41	-0.911*** (0.065)	-0.0186 (0.101)	-0.837*** (0.060)
42 ≤ Loan term ≤ 53	-1.260*** (0.065)	-0.462*** (0.102)	-1.085*** (0.060)
54 ≤ Loan term ≤ 65	-1.443*** (0.064)	-0.212** (0.101)	-1.353*** (0.060)
66 ≤ Loan term ≤ 77	-1.376*** (0.065)	-0.214** (0.102)	-1.277*** (0.060)
78 ≤ Loan term ≤ 89	-1.278*** (0.067)	-0.349*** (0.108)	-1.140*** (0.063)
90 ≤ Loan term ≤ 101	-1.367*** (0.165)	-0.348 (0.356)	-1.241*** (0.171)
102 ≤ Loan term ≤ 113	-1.524*** (0.546)	0 (·)	-1.156** (0.545)
114 ≤ Loan term ≤ 125	-1.588*** (0.155)	-1.131*** (0.426)	-1.331*** (0.157)

NOTE: Table shows the results from a logit regression. The dependent variable for column (1) is one for consumers that pre-paid the loan in the first two years. The dependent variable for column (2) is one for consumers that prepaid the loan in the first 120 days, which is our proxy for prepayment during the clawback period. The dependent variable for column (3) is one for consumers that prepaid the loan after the first 120 days but within the first two years. All regressions condition on loans of length greater than two years.

Table A3: *Markups by prepayment probability percentile*

PREPAYMENT PERCENTILE	Markup Percentile				
	p10	p25	p50	p75	p90
Below 10th	0	0.22	1.01	2	2
10th-25th	0	0.21	1.25	2	2
25th-50th	0	0.25	1.23	2	2
50th-75th	0	0.36	1.2	2	2
75th-90th	0	0.48	1.19	1.98	2
Above 90th	0	0.5	1.16	1.91	2.1

NOTE: Table shows summary statistics for markups by late prepayment probability percentiles. Prepayment rates are estimated using the CCP.

Table A4: $M'(\cdot)_i$ by prepayment probability percentile

PREPAYMENT PERCENTILE	p10	p25	p50	p75	p90
Below 10th	0.78	0.81	0.87	0.91	0.96
10th-25th	0.78	0.80	0.85	0.91	0.96
25th-50th	0.77	0.80	0.85	0.90	0.96
50th-75th	0.77	0.80	0.86	0.91	0.95
75th-90th	0.77	0.79	0.86	0.91	0.94
Above 90th	0.77	0.79	0.87	0.90	0.94

NOTE: Table shows summary statistics for $M'(\cdot)_i$ by late prepayment probability percentiles. Prepayment rates are estimated using the CCP.

Table A5: B_i^O by Prepayment probability percentile

PREPAYMENT PERCENTILE	B_i^O Percentile				
	p10	p25	p50	p75	p90
Below 10th	107.34	211.51	388.05	607.28	835.87
10th-25th	103.38	221.04	376.07	559.68	747.19
25th-50th	108.87	208.68	348.28	523.88	718.63
50th-75th	108.65	186.61	309.21	481.85	683.92
75th-90th	104.7	165.72	274.74	437.95	641.47
Above 90th	96.23	155.62	262.49	430.58	650.86

NOTE: Table shows summary statistics for B_i^O by late prepayment probability percentiles. Prepayment rates are estimated using the CCP.

Table A6: B_i^M by prepayment probability percentile

PREPAYMENT PERCENTILE	B_i^M Percentile				
	p10	p25	p50	p75	p90
Below 10th	0	14.6	73.68	159.61	247.61
10th-25th	0	9.29	71.45	147.04	225.02
25th-50th	0	12.38	72.21	143.88	220.72
50th-75th	0	18.80	72.27	142.47	223.99
75th-90th	0	21.86	70.97	137.76	220.82
Above 90th	0	24.51	71.93	141.61	233.89

NOTE: Table shows summary statistics for B_i^M by late prepayment probability percentiles. Prepayment rates are estimated using the CCP.

Table A7: Regressions of Prepayment Risk on Estimated Bias

	B_{i^O}		B_{i^M}	
Log Monthly Income	-8.601***	-8.642***	-1.297***	-1.289***
	(0.263)	(0.264)	(0.136)	(0.136)
Credit Score, 100 points	-30.91***	-30.91***	-4.949***	-7.254***
	(0.406)	(0.518)	(0.202)	(0.235)
Mileage, Tens of Thousands	5.070***	5.068***	-0.580***	-0.547***
	(0.082)	(0.082)	(0.039)	(0.039)
New Car	-4.657***	-4.750***	-7.500***	-7.372***
	(0.415)	(0.415)	(0.235)	(0.236)
Log Loan Amount	390.5***	392.3***	81.44***	83.20***
	(1.208)	(1.214)	(0.407)	(0.415)
Average Years of Education	-3.742***	-3.733***	-1.044***	-1.047***
	(0.340)	(0.340)	(0.211)	(0.211)
Internet Access Quality	-8.743***	-8.716***	-8.469***	-8.452***
	(1.515)	(1.515)	(0.930)	(0.930)
Estimated Prepayment Probability	-33.72***		-60.34***	
	(5.991)		(1.739)	
Estimated Early Prepayment Probability		7.408		112.0***
		(21.399)		(8.704)
Estimated Late Prepayment Probability		-20.35***		-73.42***
		(6.602)		(1.972)

NOTE: Lender, model, and state fixed effects also included, but not shown. Standard errors clustered at zip code level. Early prepayment is defined as prepayment in the first 120 days. Late prepayment is defined as prepayment after 120 days but within the first 2 years. Estimated prepayment probabilities in the supervisory data are imputed using coefficient estimates from the CCP.

One potential explanation for this result is that dealers cannot distinguish late prepayment risk from early prepayment risk. Indeed, the correlation between our measures of early and late prepayment risk in the CCP is 0.63. In this case, our finding that higher prepayment risk predicts higher $M'_i(\cdot)$ estimates suggests that the early prepayment result in Section B.2 quantitatively dominates the opposing effect from late prepayment. In this case, ignoring prepayment biases against our results in Section 7.

Another possibility with substantial empirical support is that, instead of considering the cost of a financed vehicle through the life of a loan, consumers instead care about monthly payments (Argyle *et al.* (2019)). Because there is a one-to-one mapping between monthly payment and the total cost of a loan if it is paid on schedule, this model is equivalent to ours if we ignore prepayment risk.

B.4 Dynamic Relationships between Lenders and Dealers

The complementary commercial data includes important characteristics of dealer-lender relationships. This enables us to directly check the extent to which markups are higher in dealer-lender relationships that are more intense or that have lasted longer.

We proxy the intensity of the relationship between dealer i and lender j by the fraction f of loans intermediated by dealer i that are from lender j . Table A8 shows the results from an OLS regression where markups are the dependent variable and f is the independent variable of interest. Importantly, the regression includes dealer fixed effects, so that the estimates reveal if dealers choose higher markups when working with preferred lenders. We find that the intensity of a dealer-lender relationship has virtually no predictive effect for the markup. Markups for lenders that finance more than 20 percent of a dealer's sales are only three basis points higher than markups for lenders that finance less than one percent of a dealer's sales. Recall that the average markup in our data is 108 basis points. Even for lenders that finance less than five loans a year for a given dealer, we find that markups are on average 106 basis points. It is therefore implausible that dealers' markup decisions arise from their relationships with lenders.

B.5 Suboptimal Decision Making

This section shows how regional variation in our estimates provides suggestive evidence that the limited financial sophistication of some consumers drives our results.

Table A8: Dealer-Lender Relationships and Markups

	MARKUPS		$M'(\cdot)$	
Lender finances 5 – 20% of sales	0.0435*** (0.00331)	0.0405*** (0.00332)	-0.00123*** (0.000131)	-0.00126*** (0.000131)
Lender finances > 20% of sales	0.0126*** (0.00411)	0.00773* (0.00411)	0.000122 (0.000162)	0.0000675 (0.000163)
Credit Score	-0.000843*** (0.0000298)	-0.000840*** (0.0000297)	-0.0000805*** (0.000000883)	-0.0000803*** (0.000000883)
Log Monthly Income	-0.0260*** (0.00205)	-0.0261*** (0.00205)	-0.00186*** (0.0000694)	-0.00187*** (0.0000694)
Loan Term	0.00174*** (0.000133)	0.00177*** (0.000133)	0.000761*** (0.00000410)	0.000762*** (0.00000410)
New Car	-0.217*** (0.00278)	-0.218*** (0.00278)	-0.00727*** (0.0000790)	-0.00727*** (0.0000789)
Dealer market share 2.5 – 10%		-0.0155*** (0.00382)		-0.000895*** (0.000111)
Dealer market share > 10%		-0.0453*** (0.00499)		-0.000520*** (0.000161)
HH index 500 – 1500		-0.00813 (0.00495)		-0.000300* (0.000157)
HH index > 1500		0.00644 (0.00610)		-0.000897*** (0.000197)
Constant	1.608*** (0.0663)	1.625*** (0.0664)	0.777*** (0.00269)	0.778*** (0.00270)

NOTE: Lender, month, model, and zip code fixed effects also included, but not shown. Standard errors clustered at zip code level.

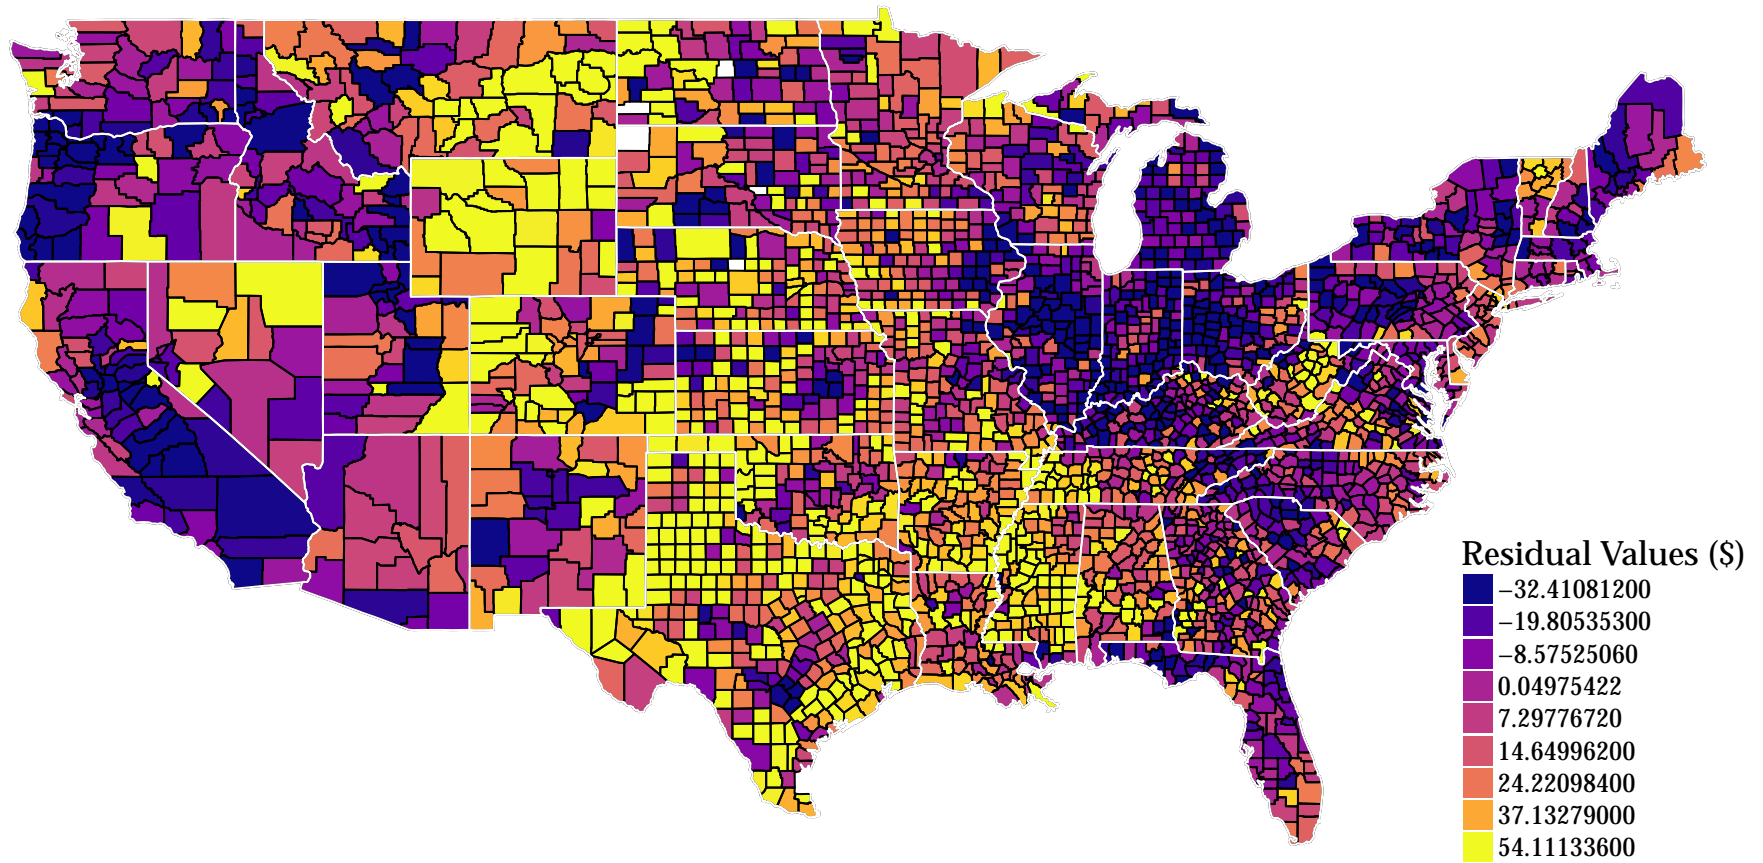


Figure A3: This map shows the residual of bound B^M (the bound due to markups) after removing lender fixed effects. To eliminate lender-specific variation, we plot residual variation in B_i^M after removing lender fixed effects. A county is lightly colored on the map if car buyers in that county exhibit larger wedges in sensitivities compared to buyers in other counties with loans from the same lender.

Table A9: Estimated Effect of Transaction and Regional Characteristics on Estimated Bounds

	B_i^O		B_i^M	
	(1)	(2)	(3)	(4)
Log Monthly Income	-9.277*** (0.262)	-8.688*** (0.261)	-1.847*** (0.135)	-1.528*** (0.135)
Credit Score, 100 points	-29.99*** (0.386)	-30.02*** (0.386)	-3.576*** (0.197)	-3.588*** (0.198)
Mileage, Tens of Thousands	5.054*** (0.081)	5.014*** (0.081)	-0.602*** (0.039)	-0.625*** (0.039)
New Car	-4.931*** (0.411)	-4.915*** (0.411)	-8.058*** (0.232)	-8.058*** (0.232)
Log Loan Amount	393.8*** (0.951)	393.4*** (0.950)	87.72*** (0.324)	87.51*** (0.324)
Average Years of Education		-3.734*** (0.336)		-1.028*** (0.208)
Internet Access Quality		-8.425*** (1.469)		-8.531*** (0.895)

NOTE: Table shows estimates from an OLS regression. Columns (1) and (3) control for transaction-specific observables. Columns (2) and (4) include county-level data from the ACS and tract-level data from the FCC. Lender, model, and state fixed effects also included, but not shown. Standard errors clustered at zip code level. B_i^M estimates are available for all loans in our sample. B_i^O is estimated only for loans with positive markups. We measure education by the percentage of county inhabitants with a college degree. The ACS also provides the percent of county residents with a high school degree. Our results also hold if we use this proxy for education or both proxies simultaneously.

C Cost Recovery from First Order Conditions of Dealers

This appendix chapter shows how dealer cost can be recovered from the dealer first order pricing conditions once the demand system is estimated. Recall that:

$$\tilde{p}_{jd} = p_{jd} + p_{jd} \cdot r_{jd} - \rho_i \cdot \log(p_{jd} \cdot r_{jd}).$$

Dealer d 's problem is:

$$\max_{\{p_{kd}, r_{kd}\}_{j \in \mathcal{J}_d}} \sum_{k \in \mathcal{J}_d} (p_{kd} + \alpha_{kd} \cdot (p_{kd} - \kappa_{kd}) \cdot (r_{kd} - b_{kd}) - c_{kd}) \cdot s_{kd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm}).$$

The FOCs for each $j \in \mathcal{J}_d$ are:

$$s_{jd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm})(1 + \alpha_{jd}(r_{jd} - b_{jd})) + \sum_{k \in \mathcal{J}_d} \frac{\partial s_{kd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm})}{\partial p_{jd}} \cdot (p_{kd} + \alpha_{kd} \cdot (p_{kd} - \kappa_{kd}) \cdot (r_{kd} - b_{kd}) - c_{kd}) = 0$$

$$s_{jd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm}) \cdot \alpha_{jd} \cdot (p_{jd} - \kappa_{jd}) + \sum_{k \in \mathcal{J}_d} \frac{\partial s_{kd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm})}{\partial r_{jd}} \cdot (p_{kd} + \alpha_{kd} \cdot (p_{kd} - \kappa_{kd}) \cdot (r_{kd} - b_{kd}) - c_{kd}) = 0.$$

Note that car prices affect shares only through perceived prices. Similarly, interest rates affect shares only through perceived prices. That is,

$$\frac{\partial s_{kd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm})}{\partial p_{jd}} = \frac{\partial s_{kd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm})}{\partial \tilde{p}_{jd}} \cdot \left(1 + r_{jd} - \frac{\rho_i}{p_{jd}}\right)$$

and

$$\frac{\partial s_{kd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm})}{\partial r_{jd}} = \frac{\partial s_{kd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm})}{\partial \tilde{p}_{jd}} \cdot \left(p_{jd} - \frac{\rho_i}{r_{jd}}\right)$$

Substituting these into the FOCs for j and rearranging give

$$\frac{1 + \alpha_{jd}(r_{jd} - b_{jd})}{1 + r_{jd} - \frac{\rho_i}{p_{jd}}} = \frac{\alpha_{jd} \cdot (p_{jd} - \kappa_{jd})}{p_{jd} - \frac{\rho_i}{r_{jd}}}.$$

This in turn gives

$$\alpha_{jd} = \frac{p_{jd} - \frac{\rho_i}{r_{jd}}}{(p_{jd} - \kappa_{jd}) \left(1 + r_{jd} - \frac{\rho_i}{p_{jd}}\right) - (r_{jd} - b_{jd}) \left(p_{jd} - \frac{\rho_i}{r_{jd}}\right)}.$$

After obtaining $\{\alpha_{jd}\}_{j \in \mathcal{J}_d}$, we can use the FOCs in interest rates to obtain costs (from a linear system):

$$\sum_{k \in \mathcal{J}_d} \frac{\partial s_{kd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm})}{\partial \tilde{p}_{jd}} \cdot c_{kd} \quad (19)$$

$$= \sum_{k \in \mathcal{J}_d} \frac{\partial s_{kd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm})}{\partial \tilde{p}_{jd}} \cdot (p_{kd} + \alpha_{kd} \cdot (p_{kd} - \kappa_{kd}) \cdot (r_{kd} - b_{kd})) \quad (20)$$

$$+ \frac{s_{jd}^m(\mathbf{p}_{dm}, \mathbf{r}_{dm}; \mathbf{p}_{-dm}, \mathbf{r}_{-dm}) \cdot \alpha_{jd} \cdot (p_{jd} - \kappa_{jd})}{p_{jd} - \frac{\rho_i}{r_{jd}}}. \quad (21)$$

Let Δ denote the matrix of the derivatives of shares in perceived prices and let \mathbf{c} denote the vector of costs of dealder d for different models. Let π denote the “adjusted revenue” vector. Finally, collect the last additive term, or “adjusted share”, from each j ’s equation into \mathbf{a} , we have

$$\mathbf{c} = \pi + \Delta^{-1} \cdot \mathbf{a}.$$

D Additional Facts, Graphs and Figures

D.1 Parametrization of Consumer Utility Function $M(\cdot)$

This section discusses our parametrization of the utility function $M(\cdot)$. First, we show evidence from a 2SLS regression suggesting that the curvature of $M(\cdot)$ is on average convex. Specifically, we can investigate the effect of higher finance charges on our estimates of $M'_i(\cdot)$. Finance charges are obviously endogenous to $M_i(\cdot)$. However, the buy rate offered by the lender is strongly correlated with finance charges. Moreover, our data include virtually all information the lender has. Conditional on this information, the buy rate is therefore exogenous to any unobservable consumer characteristics such as the consumer’s utility function $M'_i(\cdot)$ and so is a valid instrument. The results of the corresponding 2SLS regression are summarized in Table A10. The estimated shape suggests that $M_i(\cdot)$ is on average a convex function. Hence, consumers respond more if more is at stake.

Second, we show that our parametrization indeed allows us to study how the bounds B_i^O vary across the population. In particular, given the parametrization we choose, the ordering of bounds corresponds to an ordering of the underlying wedge.

Table A10: IV Regression of M' on Finance Charges

Finance Charges, in thousands, instrumented by buy rate	0.00512*** (0.000)
Credit Score, 100 points	-0.00801*** (0.000)
Log Monthly Income	-0.00261*** (0.000)

NOTE: A regression of markup on log applicant income, log price, credit score bins (of width 20), mileage, log loan amount, and state, lender, model, and new car fixed effects yields an R^2 of .07. This goes up to .11 when using zip code fixed effects instead of state fixed effects, or .10 if restricting the regression to non-negative markups no higher than 2.5 (the 99th percentile in the data). We interpret this as evidence that lenders cannot predict markups.

Assumption 3. Suppose for every consumer i that: (i) the dealer can condition the price of the car on the source of financing, and (ii) $M_i(\cdot)$ can be written as $M_i(x) = x - h(\rho_i) \cdot \ln(x)$ for some positive, increasing, and continuously differentiable function $h(\cdot)$.

As discussed in Section 4, part (i) of Assumption 3 renders the bounds on $M'(\cdot)$ tight.⁴⁶ If $M(\cdot)$ is as described in part (ii) of Assumption 3, ρ_i specifies the size of the wedge between disutilities for consumer i , with larger ρ_i 's corresponding to larger wedges. Figure A4 depicts an example of the functional form. While the additional structure imposed on $M(\cdot)$ is a caveat, the functional form seems plausible and is consistent with evidence on the curvature of $M(\cdot)$ presented in Section 7. It generates a marginal wedge in sensitivities that is large for small amounts of financing but converges to zero as financing costs increase. This is consistent with models of bounded rationality in which consumers pay more attention to issues that are more important.⁴⁷

Proposition 4. Suppose that Assumption 3 holds. Then $B_i^O \geq B_j^O \Leftrightarrow \rho_i \geq \rho_j \quad \forall i, j$

Proof. We construct $M(x, \rho_i)$ such that the ordering of $B_i^O = [1 - M'(x, \rho_i)]x$ induces an ordering on ρ_i . Hence, $\rho_i > \rho_j \Leftrightarrow B_i^O > B_j^O$ for all i, j . As a direct consequence of continuity it then must hold that $\rho_i = \rho_j$ if and only if $B_i^O = B_j^O$. Consider the set of observations such that this is true. Hence:

$$[1 - M'(x, \rho)]x = C$$

⁴⁶Anecdotal evidence suggests that car prices frequently depend on whether or not the consumer finances the car through the dealership.

⁴⁷For papers that explore the behavioral implications of limited attention and the determinants of attention allocation see Kőszegi and Szeidl (2012), Bordalo *et al.* (2012), Bordalo *et al.* (2013), and Mackowiak *et al.* (2018)

Then with changing x , ρ needs to stay constant:

$$\begin{aligned}\frac{d\rho}{dx} &= \frac{1 - M'(x, \rho) - M''(x, \rho)x}{xM'_\rho(x, \rho)} = 0 \\ &\Leftrightarrow 1 - M'(x, \rho) - M''(x, \rho)x = 0\end{aligned}$$

This is a separable ODE that we can solve in the following way:

$$\begin{aligned}\frac{M''(x, \rho)}{1 - M'(x, \rho)} &= \frac{1}{x} \\ \Leftrightarrow \int \frac{1 - M'}{d} M' &= \int \frac{1}{x} dx \\ \Rightarrow -\ln(1 - M') &= \ln(x) + g(\rho) \\ \Leftrightarrow M' &= 1 - \frac{1}{xg(\rho)} \\ \Rightarrow M &= x - h(\rho)\ln(x),\end{aligned}$$

where $h(\rho)$ is some arbitrary function of ρ . Hence, every M function that satisfies the above induces the ordering on the bounds to be also an ordering on the rhos. Note that these functions are convex as long as $h(\rho) > 0$ and that higher ρ corresponds to a larger wedges in disutilities if $h'(\rho) > 0$.

□


Proposition 4 shows that, if Assumption 3 holds, an ordering of the bounds B_i^O induces an ordering of ρ_i 's. Hence, if $B_i^O \geq B_j^O$ for two consumers i and j then $\rho_i \geq \rho_j$. This implies we can study heterogeneity in B_i^O to understand heterogeneity in the difference between financing costs and the sensitivity to financing costs.⁴⁸

D.2 Subvented Loans

Our analysis requires that loan price and vehicle price be jointly determined, so this is worth verifying in our data. To do so we focus on subvented loans. Recall that vehicle manufacturers subsidize interest rates on subvented loans, so markups on these loans are typically restricted and often prohibited. As a consequence, in our data the mean markup for subvented loans is only 0.06 compared to 1.12 for non-subvented loans. Hence, dealers' profit for these types of sales depends much more on the vehicle price.

⁴⁸We can derive a qualitatively similar result for B_i^M . In particular, there exists a class of utility functions such that an ordering over B_i^M induces an ordering over ρ_i 's. Results in Section 5.4 are the same for B^M and B^O , suggesting that the correlations we document hold for both classes of utility functions.

Figure A4: Example of Assumed Functional Form For $M(\cdot)$



NOTE: This graph shows two examples of the functional form for $M(\cdot)$ that we assume from here on. Note that as finance charges increase, the ratio of finance charges and the disutility of finance charges converges to one.

Table A11 shows the estimates from a linear regression model that predicts the effects of subvention on the price of the car and the loan. We include fixed effects for the lender that issued the loan, the time the loan was issued, the purchased car model and make, and the buyer's zip code.

If car and loan prices are chosen independently in competitive markets, we would expect car prices not to be affected by loan subvention. Columns (1) and (2) suggest that this is not true. Subvention predicts two changes in prices. Subvented loan prices are on average about \$150 lower, but prices for cars with subvented loans are on average about \$850 higher.

These estimates are not causal and may reflect in part the effect of omitted variables. For example, cars financed by subvented loans may be more likely to come with valuable add-ons that increase their prices. Still, Table A11 provides suggestive evidence using our data that car and loan prices are jointly determined, consistent with results in Argyle *et al.* (2018), Argyle *et al.* (2019), and Brown and Jansen (2019).

Table A11: Predicted Effects on Price and Loan Outcomes

	CAR PRICE	FINANCE CHARGES	OVERALL CHARGES
Subvented	846.4** (2.13)	-148.1* (-1.91)	780.1* (1.74)
Credit Score	123.7*** (3.21)	-577.1*** (-57.59)	-443.1*** (-9.93)
Log Monthly Income	0.0759*** (15.97)	0.0108*** (14.76)	0.0924*** (16.90)
Loan Term	181.3*** (8.54)	970.2*** (203.27)	1166.5*** (48.05)

NOTE: Bank, month, model, and zip code fixed effects also included, but not shown. Standard errors clustered at zip code level.

D.3 Dealer Issued Loans

This section explains our procedure to estimate the fraction of auto loans originated “indirectly”, i.e. through auto dealers. The CCP includes data on hard credit inquiries, and so is one of very few datasets that covers search behavior outside of online markets. When an auto dealer intermediates a loan, the process should always begin with a hard credit inquiry so that the dealer can determine the kind of loan a borrower can qualify for. We take the percent of loans originated within a short window of at least one hard credit inquiry from a dealer as our proxy for the fraction of indirect loans. If the CCP included data on all hard credit inquiries, this would be straightforward to estimate. The main difficulty is that we only observe hard credit inquiries reported to the credit bureau our data is from; we do not see hard credit inquiries reported to the other two major credit bureaus.

We can deal with this difficulty if we assume a constant probability P_o that a hard credit inquiry is observed in the CCP.⁴⁹ Let $P_d(i)$ denote the probability that a loan is originated within a short time window of i hard credit inquiries from dealers, and let $P_{do}(i)$ denote the probability that a loan is originated within a short time window of i hard credit inquiries from dealers that we observe. Finally,

⁴⁹In particular, we assume that the probability of observing a given hard credit inquiry from a lender (which we can estimate) is the same as observing a given hard credit inquiry from a dealer (which we need). This is equivalent to assuming that the credit bureau’s market share for dealer inquiries is the same as its market share for lender inquiries. Unfortunately we cannot test this assumption, but our conversations with market experts lead us to believe it is reasonable. Dealers have an incentive to pull credit from the same credit bureau as the lenders they work with, so that they are operating with the same information.

assume that no more than N dealers perform a hard credit pull on a consumer within the time window. Then we need to estimate $P_d(0)$, and have the following equation:

$$P_{do}(i) = \sum_{n=i}^N P_d(n) \binom{n}{i} P_o^i * (1 - P_o)^{n-i} \quad (22)$$

First, we set $N = 3$. The next step is to estimate $P_{do}(i)$ and P_o . We do that by matching new auto loans to auto loan inquiries from the company the loan is from (for P_o) and to inquiries from auto dealers (for P_{do}). We match auto loans to auto loan inquiries if they are for the same consumer and if the inquiry date is no more than 14 days before or 7 days after the origination date of the auto loan. We restrict the sample during this step to consumers with credit scores above 680, to minimize the possibility of one dealer pulling credit records from multiple credit bureaus.⁵⁰

For a given guess of the vector $P_d(n)$, Equation 22 yields implied values of the vector $P_{do}(i)$. We take as our estimate of $P_d(n)$ the vector that minimizes the sum of squared deviations between implied and estimated values of $P_{do}(i)$. Using data for the U.S. as a whole, this yields an estimate of $P_d(0) = 0.17$. This implies that an estimated 83 percent of auto loans are opened a short time before or after a hard credit inquiry from an auto dealer, which we interpret to mean that roughly 83 percent of auto loans are indirect. A number of assumptions were required to obtain this estimate, and we do not view it as precise. However, we do interpret it as strong evidence that the vast majority of auto loans in the United States are indirect.

D.4 Lender Maximization Problem with first price auction

This section shows how we can use our data to estimate a model in which lenders compete in a first price auction instead of a second price auction. As in the main part, we make a number of simplifying assumptions. First, we assume that there is one auction for each model-dealer combination. Hence, each of the lenders in \mathcal{L}_d associated with dealer d bids once for all the sales of model j . Second, we assume α is constant across dealers and lenders. Empirically there is relatively little variation in α and allowing for heterogeneity would complicate the analysis considerably. Third, lenders take all buy rates arising in auctions for other models as given when making their bids for a particular contract.⁵¹

⁵⁰Lenders pay credit bureaus for every inquiry they make, so when deciding on the number of bureaus to pull from, they face a tradeoff between the cost of an additional pull and the benefit of obtaining more information. Because of the very large sums of money involved, mortgage lenders nearly always pull from all three major credit bureaus. Auto lenders typically only pull information from one credit bureau for borrowers who do not appear to be a credit risk, which is why we focus on consumers with good credit scores. Auto dealers have even less incentive to pull from multiple bureaus than auto lenders do, because auto dealers do not bear default risk.

⁵¹This simplification is reminiscent of the Nash-in-Nash approach that is frequently used in the analysis of bilateral oligopoly (see Horn and Wolinsky (1988), Gaynor and Town (2011), Collard-Wexler *et al.* (2014)). In this literature, the bargains between a particular

Assumption 4. (i) All lenders $k \in \mathcal{L}_d$ are drawing from the same cost distribution $r_k^b \sim \mathcal{F}$ and have the same revenue-sharing contract α with a given dealer d . (ii) In each auction (j,d) bidders take the buy rates arising in auctions for other makes as given.

The objective function is composed of the market share of the model for which the lenders are bidding, the loan amount $p_{jd} - \kappa_{jd}$, the probability of winning the auction, and the profit margin, which is composed of a share of the markup and a direct payment.

A lender wins the bid and issues the loans for a dealer-model combination if their buy rate is the lowest one submitted. Let N be the number of lenders bidding for a contract and $r^l \sim \mathcal{F}(\cdot)$ the distribution of wholesale interest rates r , which characterize the lenders' cost of financing a loan. Lenders anticipate that the likelihood of a sale depends on b_{jd} through the dealer's downstream decision on all prices and interest rates. Lender k 's objective function when bidding for a particular contract is therefore given by:

$$\max_{b_{jd}^k} s_{jd}^m(\mathbf{p}_{dm}(b_{jd}^k), \mathbf{r}_{dm}(b_{jd}^k); \mathbf{p}_{-dm}, \mathbf{r}_{-dm}) \cdot (p_{jd}(b_{jd}^k) - \kappa_{jd}) \cdot \left(1 - \mathcal{F}(\eta_{jd}^{-1}(b_{jd}^k))\right)^{N_d-1} \cdot \left[(1 - \alpha_{jd}) \cdot m_{jd}(b_{jd}^k) + b_{jd}^k - r_{jd}^k\right].$$

For notational convenience we omit some of the function's arguments and denote it by $s_{jd}(b)$. For the derivation, define $\psi_{jd}(b_{jd}) = s_{jd}^m(\mathbf{p}_{dm}(b_{jd}), \mathbf{r}_{dm}(b_{jd}); \mathbf{p}_{-dm}, \mathbf{r}_{-dm}) \cdot (p_{jd}(b_{jd}) - \kappa_{jd})$.

Lenders anticipate that dealers mark up a specific model by $m_{jd}(b_{jd})$. The markup function depends on observable attributes and unobservable attributes of customers that purchase model j as well as the other prices and interest rates in the market. Given our demand model, we can simulate \hat{m}_{jd} . We apply the key insight from Guerre *et al.* (2000) that there is a relationship between the observed distribution of buy rates and the unobserved distribution of costs. This insight allows us to substitute all unknown terms in the lender's first order condition with either observed or known objects based on our demand estimation.

Proposition 5. *Lenders' costs can be recovered as:*

$$r_{jd}^k = (1 - \alpha_{jd}) \cdot m_{jd}(b_{jd}^k) + b_{jd}^k + \frac{1 + (1 - \alpha) \cdot m'_{jd}(b_{jd}^k)}{\tilde{\psi}'(b_{jd}^k) - (N_d - 1) \cdot g(b_{jd}^k) \cdot (1 - \mathcal{G}(b_{jd}^k))^{-1}} \quad (23)$$

All expressions on the right hand side are observed or can be constructed from the demand side estimates. We can therefore back out the cost distribution of lenders in this market.

upstream supplier and a downstream firm take the outcome of all other bargains as given. Similarly, we assume that, when bidding for a particular contract, a lender takes all buy rates in other auctions as given.

Proof. The lender's maximization problem is given by:

$$\max_{b_k} s_{jd}(b) \cdot (p_{jd}(b_k) - \kappa_{jd}) \cdot \left(1 - \mathcal{F}(\eta_{jd}^{-1}(b_k))\right)^{N-1} \cdot \left[(1 - \alpha) \cdot m_{jd}(b) + b_k - r_k^l\right].$$

Taking the first order condition, we obtain:

$$\begin{aligned} & \left[(1 - \alpha) \cdot m'_{jd}(b_k) + 1\right] \cdot \psi(b_k) \cdot \left(1 - \mathcal{F}(\eta_{jd}^{-1}(b_k))\right)^{N-1} + \left[(1 - \alpha) \cdot m_{jd}(b_k) + b_k - r_k^l\right] \\ & \cdot \left[\psi'(b_k) \cdot \left(1 - \mathcal{F}(\eta_{jd}^{-1}(b_k))\right)^{N-1} - \psi(b_k) \cdot (N - 1) \cdot f(r) \frac{1}{\eta'(b_k)} \cdot \left(1 - \mathcal{F}(\eta_{jd}^{-1}(b_k))\right)^{N-2}\right] = 0 \end{aligned}$$

Which can be simplified to the following equation, using the insight that $G(b) = \mathcal{F}(\eta_{jd}^{-1}(b_k)) \quad \forall b$ and therefore $g(b) = \frac{1}{\eta'(r)} \cdot f(r)$.

$$\begin{aligned} & \left[(1 - \alpha) \cdot m'_{jd}(b_k) + 1\right] \cdot \psi(b_k) \cdot \left(1 - \mathcal{F}(\eta_{jd}^{-1}(b_k))\right) + \left[(1 - \alpha) \cdot m_{jd}(b_k) + b_k - r_k^l\right] \cdot \\ & \left[\psi'(b_k) \cdot \left(1 - \mathcal{F}(\eta_{jd}^{-1}(b_k))\right) - \psi(b_k) \cdot (N - 1) \cdot g(b_k)\right] = 0 \end{aligned}$$

Solving this equation for r and defining $\tilde{\psi}(b_k) = \frac{\psi'(b_k)}{\psi(b_k)}$ to simplify the exposition we get the following expression for the lender's financing costs

$$r_k = (1 - \alpha) \cdot m_{jd}(b_k) + b_k + \frac{1 + (1 - \alpha) \cdot m'_{jd}(b_k)}{\tilde{\psi}(b_k) - (N - 1) \cdot g(b_k) \cdot (1 - \mathcal{G}(b_k))^{-1}}.$$

□

D.5 Additional Estimations and Figures

Table A12: Summary Statistics in the complementary commercial and Administrative Data

	25TH PCTILE	50TH PCTILE	75TH PCTILE	MEAN
MONTHLY PAYMENT				
Administrative	284	366	475	402
Commercial	300	382	481	409
LOAN AMOUNT				
Administrative	15893	21345	28086	23145
Commercial	16193	21504	27361	22555
LOAN TERM				
Administrative	60.0	63.0	72.0	64.1
Commercial	58.0	60.0	69.0	60.7
INTEREST RATE				
Administrative	2.99	3.95	4.75	4.08
Commercial	2.85	3.85	4.95	3.94

NOTE: Summary statistics are shown for monthly payment, loan amount, loan term, and interest rate in the administrative data and in the commercial data from 2011. Statistics are shown for consumers with credit scores above 720 in both samples.

Table A13: $M'(\cdot)$, B^O , and B^M by Car Model

MODEL	$M'(\cdot)$	B^O	B^M	MODEL	$M'(\cdot)$	B^O	B^M
Ram	0.76	795.89	239.31	Chevrolet Equinox	0.81	499.41	169.74
Chevrolet Silverado	0.77	693.59	219.19	Ford Escape	0.79	493.89	170.68
Ford F-150	0.79	660.66	220.28	Ford Fusion	0.79	480.00	161.90
Kia Sorento	0.78	614.92	189.59	Toyota Camry	0.77	475.85	143.56
Ford Explorer	0.80	602.00	212.74	Hyundai Elantra	0.77	474.41	156.04
Jeep Gr. Cherokee	0.80	585.28	200.07	Chevrolet Malibu	0.80	450.21	154.90
Jeep Wrangler	0.81	576.16	201.04	Honda Civic	0.78	442.62	139.26
Hyundai Sonata	0.77	529.34	157.96	Chevrolet Cruze	0.81	437.70	144.71
Honda Odyssey	0.80	519.13	169.13	Honda Accord	0.78	437.70	150.26
Nissan Altima	0.79	507.10	159.52	Ford Focus	0.80	424.82	147.52

NOTE: The table shows for the 20 most common car models in the data the average estimates for the lower bound B_i^O , B_i^M and $M(\cdot)$.

Table A14: *Selected summary statistics of consumers' misperception of financial charges.*

	MEAN	P10	P25	P50	P75	P90
$\hat{\rho}$	0.77	0.67	0.74	0.78	0.82	0.83

Note: Table gives summary statistics on the distribution of estimated values of ρ .

Table A15: Demand Model Coefficients

MAKE	MEAN UTILITY SHIFT	STANDARD ERROR OF MEAN UTILITY SHIFT	MAKE-SPECIFIC OWN PRICE ELASTICITY
Chevrolet	-4.1	0.047	-4.2
Chrysler	-4.4	0.047	-4.2
Dodge	-4.8	0.044	-3.7
Ford	-4.0	0.047	-4.1
GMC	-2.7	0.073	-5.4
Honda	-3.5	0.047	-3.2
Hyundai	-4.0	0.047	-2.5
Jeep	-3.7	0.04	-4.1
Kia	-3.9	0.045	-2.7
Mazda	-4.1	0.048	-2.7
Nissan	-4.3	0.045	-2.9
RAM	-3.7	0.051	-4.9
Subaru	-3.8	0.044	-2.7
Toyota	-3.8	0.047	-3.4
Volkswagen	-4.2	0.052	-2.8

NOTE: The table shows how mean utility is shifted for different makes, the standard errors of those mean shifts, as well as own-price elasticities for different makes.