ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 80, No. 4 (July, 2012), 1387–1432

CONTRACT PRICING IN CONSUMER CREDIT MARKETS

LIRAN EINAV
Stanford University, Stanford, CA 94305, U.S.A. and NBER

MARK JENKINS
Wharton School, University of Pennsylvania, Philadelphia, PA 19104, U.S.A.

JONATHAN LEVIN
Stanford University, Stanford, CA 94305, U.S.A. and NBER

Econometrica, Vol. 80, No. 4 (July, 2012), 1387–1432

CONTRACT PRICING IN CONSUMER CREDIT MARKETS

BY LIRAN EINAV, MARK JENKINS, AND JONATHAN LEVIN[1]

We analyze subprime consumer lending and the role played by down payment requirements in screening high-risk borrowers and limiting defaults. To do this, we develop an empirical model of the demand for financed purchases that incorporates both adverse selection and repayment incentives. We estimate the model using detailed transaction-level data on subprime auto loans. We show how different elements of loan contracts affect the quality of the borrower pool and subsequent loan performance. We also evaluate the returns to credit scoring that allows sellers to customize financing terms to individual applicants. Our approach shows how standard econometric tools for analyzing demand and supply under imperfect competition extend to settings in which firms care about the identity of their customers and their postpurchase behavior.

KEYWORDS: Contract pricing, subprime lending, credit markets, asymmetric information.

1. INTRODUCTION

THE DRAMATIC CREDIT CYCLE of the last decade has brought renewed attention to consumer credit markets. In this paper, we develop an econometric model of consumer lending, and use it to investigate the pricing and demand for subprime credit. Our main goal is to show how different elements of credit offers affect the quality of the borrower pool and the subsequent prospects for repayment, and draw out some implications of these findings.

We analyze data from a large subprime lender that was active during the credit boom. The company specialized in financing auto sales to consumers with low incomes or poor credit histories. The subprime auto loan market is a useful setting to understand high-risk lending. Consumer liquidity is low, default rates are high, and there is substantial heterogeneity across borrowers in the likelihood of repayment. This heterogeneity is captured only partially by credit histories. Our unusually rich data include not only contract terms and repayment outcomes for a large pool of borrowers, but also detailed information on loan applicants who subsequently declined borrowing. Sharp variation in the company's pricing schedule over time allows us to identify the effects of changing down payment requirements and markups on cars. Using this variation, we find that down payment requirements play an important role in screening out risky borrowers and limiting loan sizes, while higher markups on cars primarily induce larger loans.

We develop our analysis in several steps. We first set out the problem of loan pricing in a way that links it to the traditional problem of pricing against

[1]We have benefited from the suggestions of the co-editor and three referees, and from the comments of many seminar participants. We acknowledge the support of the Stanford Institute for Economic Policy Research, the National Science Foundation (Einav and Levin), and the Alfred P. Sloan Foundation (Levin).

 DOI: 10.3982/ECTA7677

a downward sloping demand curve. As in the traditional case, pricing incentives depend on a small number of elasticities. Estimating these elasticities requires a model of consumer demand that incorporates both borrowing and repayment decisions. After describing the empirical setting, we develop such a model. Our model starts with a standard consumer theory framework in which heterogeneous consumers choose whether to purchase and how much to borrow, and then make decisions over time about whether to continue making loan payments or to default. We use the model to derive a set of linear estimating equations that capture purchasing, borrowing, and repayment decisions. The linear specification makes the identification of the key pricing quantities particularly transparent.

To estimate the demand model, we use a series of discrete changes in the loan terms offered by the lender. Our estimates reveal several striking features of consumer demand. The first is that purchasing decisions are highly sensitive to down payment requirements. A natural interpretation, and one that is consistent with our theoretical model, is that consumers are constrained by cash on hand. This makes them sensitive to current expenditures, especially compared to price increases that can be financed. Indeed, we find that price increases have relatively little effect on purchasing or down payment decisions; instead, they translate mostly into larger loans. We also find that larger loans decrease the likelihood of repayment substantially.

The central part of our analysis uses the demand estimates to analyze contract design: how different aspects of the financial contract affect the quantity and quality of loan originations. Down payment requirements generate a trade-off between loan volume and loan quality. Tighter down payment requirements do not matter for some buyers who intend to make a substantial down payment. For others, however, a more stringent borrowing limit means coming up with extra cash or foregoing the purchase. Our estimates suggest that these "marginal" buyers are relatively likely to default, so screening them out improves the composition of borrowers. If the firm has predictive information about individual borrowers, its optimal pricing schedule has the intuitive property that safer borrowers are allowed to borrow more, a feature that is shared by the firm's observed contracts.

Changes in the markup on cars, which are similar to interest rate changes in that, for a fixed down payment, they affect the resulting repayment obligation, are different. We estimate a relatively small effect of car prices on sales volume and on down payments. Instead, increased prices primarily translate into larger loans and higher monthly payments. This means more revenue while borrowers are making payments, but also a higher rate of default. We show how optimal prices resolve this trade-off and have the somewhat surprising feature that it may be desirable to grant discounts to high-risk borrowers whose ability to repay is more sensitive to the size of the payment obligation.

We also use our model, combined with detailed cost data, to assess whether the observed pricing policies are optimal. We find that optimal pricing would have relied more heavily on credit scoring information: lowering down payment requirements for better risks and raising them for higher risks. We also find that there is a sense in which the lender maintained a relatively conservative origination policy, in that there appear to be marginal loans that would have been profitable had they been originated. When we back out an implied cost of lending that rationalizes the firm's "average" prices, we find it to be roughly $1500 per loan. One interpretation is that this cost reflects a concern about taking on aggregate portfolio risk, a justifiable concern given the events that occurred subsequent to our data (which ended in spring 2006).

The last part of the paper applies the model to quantify the value of using credit scores to set down payment requirements. The question is of interest because the advent of sophisticated credit scoring has been a major change in consumer lending over the last quarter century. We find that relative to setting a uniform down payment requirement, risk-based financing can increase profits by 22 percent. Having perfect information about consumer characteristics would increase expected profits by 96 percent. We also illustrate how proprietary information can create a strategic benefit by leaving rival lenders with an adversely selected pool of borrowers.

Our analysis in this paper builds on Adams, Einav, and Levin (2009, henceforth AEL), in which we used the same data to document some interesting features of subprime borrowing behavior.[2] There we focused on two points: first, that subprime purchasing is highly sensitive to "cash on hand"; second, that lenders face considerable adverse selection in borrowing decisions. Sharp evidence for cash sensitivity comes from the seasonal nature of demand. In AEL, we showed that demand spikes in early February as tax refund checks become available from tax preparers and that the spike is attributable to customers who are eligible for large refunds. We also described the sensitivity of purchasing to down payment requirements, which reappears here as part of our demand model. As evidence of adverse selection, we showed that borrowers who make lower down payments are more likely to default, even conditional on having the same future payment obligation, indicating the information content in down payment decisions.

Here we link these findings in a unified model of consumer purchasing, borrowing, and repayment behavior derived from consumer theory. The combined model allows us to analyze how high down payment requirements improve the composition of the borrower pool, and then to examine pricing decisions by the lender, the trade-offs involved in setting down payment requirements and car prices, the relationship between observed and optimal prices, and the extent to which credit scoring facilitates profitable lending. The analysis of pricing decisions and trade-offs is the main contribution of this paper. In estimating

[2]We have used data from the same lender in two other papers: Einav, Jenkins, and Levin (2011), which studies the implementation of credit scoring and risk-based pricing, and Jenkins (2009), which looks at the difficulty of recovering value after subprime defaults.

demand, which is a key input to the pricing analysis, we reference some of the institutional details provided in AEL, and particularly the extensive set of checks provided there to support the identification of various demand and repayment elasticities.

Our finding that borrowing limits can play a central role in curbing consumer demand and maintaining the quality of the borrower pool can be related to some recent work that studies the turn-of-the-century credit boom. For example, the idea that households may respond sharply to relaxed lending policies underlies many explanations of the run-up in mortgage borrowing and housing prices (e.g., Mian and Sufi (2009)). Of course, while the consumer response to easier credit might be similar across asset purchases, there are differences between auto and housing loans. In particular, cars tend to depreciate relatively predictably and there are no inherent limits in their supply, both of which mute the link between credit conditions and asset prices that is emphasized in work on housing.

Our analysis also relates to studies that document the presence of asymmetric information in consumer credit markets (Ausubel (1991, 1999), Edelberg (2004, 2006), Karlan and Zinman (2008, 2009), Adams, Einav, and Levin (2009)) and insurance markets (e.g., Einav, Finkelstein, and Levin (2010)). Similar to recent work in this literature, our demand estimates quantify the extent of asymmetric information between borrowers and lenders, and also borrower responsiveness to contract terms. We also go on, however, to analyze the implications for profitability and pricing decisions in a way that typically has not been pursued.[3]

2. SIMPLE ECONOMICS OF CONTRACT PRICING

We begin by describing the simple economics of contract pricing in a setting where a seller has market power, and cares about the type of buyers he attracts and how they behave subsequent to signing the contract.

There is a unit mass of potential buyers. Each is described by a vector of characteristics ζ, and we let F denote the population distribution of characteristics. The seller can offer a contract described by a vector of terms $\phi \in \Phi$. Let $U(\phi, \zeta)$ denote the value a type-ζ consumer assigns to a contract ϕ and let $\overline{U}(\zeta)$ denote the value the consumer assigns to not accepting the contract. A type-ζ consumer will select the contract if $U(\phi, \zeta) \geq \overline{U}(\zeta)$. If this happens, subsequent events will lead to an outcome $y(\phi, \zeta)$, at which point the seller realizes a profit $\pi(\phi, y)$.[4]

[3]The contract design aspect of our analysis bears some relation to work on monopoly regulation under asymmetric information (e.g., Wolak (1994) or Perrigne and Vuong (2011)).

[4]Note that the model allows for the consumer to be initially uncertain about some of her own characteristics, such as her future income, that subsequently will affect outcomes. The key point is that U and $\overline{U}$ must be measurable with respect to the consumer's initial information, so if ζ

In our empirical application, a contract offer consists of a car, a price, a required down payment, an interest rate, and a loan repayment period. The contract outcome will be the duration of time over which the consumer makes loan payments and the recovery in the event of default. The relevant consumer characteristics include individual demographics, location, and calendar dates. A slight complication is that the buyer is able to choose how much to borrow, subject to the down payment requirement, as opposed to having to accept or reject a given offer. In Section 4, we show how the contract value U can be derived from a stochastic model of car use and repayment, and in Section 5, we develop the components of the profit function π in detail.

We now illustrate the seller's pricing problem, assuming that the seller offers a single contract to the entire buyer population. To the extent the seller can condition his offer on consumer characteristics such as a credit score, we can think of an analogous problem applying to each subpopulation.

If the firm offers a contract ϕ, its total sales will be

$$Q(\phi) = \int \mathbf{1}\{U(\phi, \zeta) \geq \overline{U}(\zeta)\}\, dF(\zeta) \tag{1}$$

and its expected profits are

$$\Pi(\phi) = \int \pi(\phi, y(\phi, \zeta))\mathbf{1}\{U(\phi, \zeta) \geq \overline{U}(\zeta)\}\, dF(\zeta). \tag{2}$$

Alternatively, we can express expected profits as total sales times per-sale expected profits and write the seller's problem as

$$\max_{\phi \in \Phi} \Pi(\phi) = Q(\phi)R(\phi), \tag{3}$$

where $R(\phi) = \mathbb{E}_\zeta[\pi(\phi, y(\phi, \zeta))|U(\phi, \zeta) \geq \overline{U}(\zeta)]$ represents the seller's expected profit (or net revenue) conditional on sale.

To focus ideas, suppose the seller optimizes over a single dimension of the contract, for instance, the required down payment, the car price, or the interest rate. Assume the functions defined above are differentiable, that consumer utility U is decreasing in ϕ, and that the seller's problem is concave. An optimal contract offer will satisfy the first-order condition

$$\frac{R(\phi)/\phi}{dR(\phi)/d\phi} = \frac{-Q(\phi)/\phi}{dQ(\phi)/d\phi}, \tag{4}$$

which is a generalization of the standard monopoly formula that equates the markup to the inverse demand elasticity. In the standard problem, $\pi(\phi, \zeta) =$

includes characteristics that are initially unknown to the buyer, then U and $\overline{U}$ will not vary with these characteristics, although y and π may.

$\phi - c$, where ϕ is the product price and c is the unit cost, and the left-hand side of equation (4) is the markup $(\phi - c)/\phi$. Here the inverse revenue elasticity on the left-hand side of equation (4) captures two further effects. First, the realized profits may depend on the type of consumer: there can be selection effects. Second, the terms of the contract may affect consumer repayment: there can be behavioral effects. In our application, we will see both. For instance, allowing borrowers to take more leveraged loans will worsen the characteristics of the loan pool and reduce the repayment rate for a given borrower.

To develop this point further, suppose that ζ is single-dimensional and $U - \overline{U}$ is increasing in ζ. We then can write the marginal profit to increasing ϕ as

$$\frac{d\Pi(\phi)}{d\phi} = \frac{dQ(\phi)}{d\phi} \cdot \mathbb{E}_\zeta\big[\pi(\phi, y(\phi, \zeta))|U(\phi, \zeta) = \overline{U}(\zeta)\big] + Q(\phi) \cdot \mathbb{E}_\zeta\left[\frac{d\pi(\phi, y(\phi, \zeta))}{d\phi}\big|U(\phi, \zeta) \geq \overline{U}(\zeta)\right]. \tag{5}$$

The first term in equation (5) reflects the loss of "marginal" buyers.[5] The foregone profit depends on their characteristics: below we find that marginal borrowers are generally riskier than average borrowers. The second term captures how a change in the contract will affect inframarginal buyers. For instance, we find below that increasing the markup on autos leads to a larger loan liability, which increases the size of the monthly payments and raises the default rate.

Equation (5) highlights the empirical objects of interest from the perspective of pricing. They are the sensitivity of demand and conditional profit to the relevant contract terms, and the relationship between marginal and inframarginal buyers. Moreover, because changes in pricing have both a selection effect (altering the composition of buyers) and a direct effect (altering the behavior of borrowers), it is natural to consider an empirical strategy based around traditional selection methods, making use of plausibly exogenous variation in pricing to identify the relevant parameters. This is the approach we pursue.

3. SUBPRIME AUTO LOANS

Our study uses data from a company that operates used car dealerships across the United States. The company sells to individuals with low incomes or poor credit histories. Customers who arrive at a dealership fill out a loan application and are matched to a car that fits their needs. The dealer quotes a price and offers financing options that reflect the buyer's creditworthiness. Virtually all buyers finance a large fraction of their purchase. The loans are risky.

[5]A similar expression can be derived if ζ is multi-dimensional, and is identical but for the fact that marginal types must be weighted to account for their flow into and out of purchasing in response to an increase in ϕ. See Veiga and Weyl (2012) for a general treatment.

Defaults are common and recoveries typically amount to only a small fraction of the company's cost of the car. For this reason, both customer selection and the structure of financing are of central importance, making this an attractive setting to study the pricing and design of consumer credit contracts.

3.1. *Data and Environment*

For the present study, we use data on all loan applications and sales from June 2001 through December 2004. We observe well over 50,000 loan applications (the exact number is proprietary).[6] For each loan application, we observe dozens of demographic and credit history measures. We utilize a subset of these in our empirical analysis. Table I reports summary statistics. On average, applicants have a household income of just under \$29,000 a year. Many applicants appear to have little access to savings or credit: almost a third have no bank account, just 17 percent have a FICO score above 600—a typical cutoff for obtaining a bank loan—and 18 percent have no FICO score at all.

The company obtains its inventory at used car auctions. For each car, it sets a list price based on a formula that relates markup to the cost of the car, and then there is some potential for negotiation at the dealership. The average sale price is just under \$11,000. Buyers are required to make a minimum down payment, which depends on their credit category and can vary from a few hundred up to \$2000. Credit categories are assigned using a proprietary credit scoring algorithm. Buyers are offered terms at which they can finance the remainder of the car price, although they do not need to take the maximum loan. The offered interest rates are typically 25–30 percent on an annual basis, with the rate often equal to the state limit. Loans are expected to be repaid over 3–4 years. Our data include all of this information: the financing terms offered at each point in time, the company's cost for each car, the markup schedule determining list prices, the negotiated price for successful sales, and the loan sizes actually taken.

Just over one-third of the customers who arrive at a dealership and fill out a loan application purchase a car. Of these, 43 percent make exactly the minimum down payment. The remaining buyers make somewhat larger down payments, but frequently not by much. Fewer than 10 percent make a down payment that exceeds the minimum by \$1000. The average down payment is around \$1000, so that after taxes and fees, the average loan size is a bit under \$11,000. This translates into monthly payments on the order of \$400, a fairly significant fraction of household income.

A large portion of borrowers subsequently default. Our data end before the last payments are due on some loans, but of the loans with uncensored payment periods, 61 percent end in default. Defaults tend to come early in the

[6]We report summary statistics based on the full sample of applicants and loans, but to reduce computational time, we use a random subsample of 45,000 applicants to estimate the model.

TABLE I

SUMMARY STATISTICS[a]

	Obs.	Mean	Std. Dev.	5%	95%
Applicant Characteristics					
Age	N	32.8	10.7	19	53
Monthly income ($U.S.)	N	2414	1074	1299	4500
Home owner	N	0.15	—	—	—
Live with parents	N	0.18	—	—	—
Bank account	N	0.72	—	—	—
Risk category			—	—	—
Low	N	0.27	—	—	—
Medium	N	0.45	—	—	—
High	N	0.29	—	—	—
Car purchased	N	0.34	—	—	—
Buyer Characteristics					
Age	$0.34N$	34.7	10.8	20	55
Monthly income	$0.34N$	2557	1089	1385	4677
Home owner	$0.34N$	0.17	—	—	—
Live with parents	$0.34N$	0.16	—	—	—
Bank account	$0.34N$	0.76	—	—	—
Risk category					
Low	$0.34N$	0.35	—	—	—
Medium	$0.34N$	0.47	—	—	—
High	$0.34N$	0.17	—	—	—
Car Characteristics					
Acquisition cost ($U.S.)	$0.34N$	5090	1329	3140	7075
Total cost ($U.S.)	$0.34N$	6096	1372	4096	8212
Car age (years)	$0.34N$	4.3	1.9	2.0	8.0
Odometer (miles)	$0.34N$	68,775	22,091	31,179	102,299
Time on lot (days)	$0.34N$	33	44	1	122
Car price ($U.S.)	$0.34N$	10,777	1797	8095	13,595
Transaction Characteristics					
Min. down payment (applicants)	N	648	276	400	1200
Min. down payment (buyers)	$0.34N$	750	335	400	1400
Interest rate (APR)	$0.34N$	26.2	4.4	17.7	29.9
Loan term (months)	$0.34N$	40.5	3.7	35.0	45.0
Down payment	$0.34N$	942	599	400	2000
Loan amount	$0.34N$	10,740	1801	7982	13,559
Monthly payment	$0.34N$	395	49	314	471
Loan Outcomes (uncensored sales only)					
Default	$0.13N$	0.61	—	—	—
Fraction of payments made	$0.13N$	0.58	0.38	0.04	1.00
Loan payments	$0.13N$	7972	5635	491	16,587
Nonzero recovery (all defaults)	$0.08N$	0.78	—	—	—
PV of recovery (all recoveries)	$0.06N$	1579	1328	231	4075
Gross operating revenue	$0.13N$	9614	5192	2169	17,501
Net operating revenue	$0.13N$	3333	5020	−3906	10,284

[a]Loan payments, gross operating revenues, and net operating revenues are in present value (PV) terms. Gross operating revenue is equal to the sum of the down payment and the present value of the loan payments and recovery amount, assuming an internal firm discount rate of 10 percent. Net operating revenue is equal to gross operating revenue minus total car cost (including sales tax). To preserve the confidentiality of the company that provided the data, we do not report the exact number of applications.

loan period. Nearly 50 percent of the defaults occur before a quarter of the payments have been made, and nearly 80 percent occur in the first half of the loan length. A notable feature of the data is that the value recovered following a default is rather low. For 22 percent of defaults we observe, no recovery is made at all, sometimes because the car has been in an accident or stolen. Even when the recovery value is positive, the average present value of the recovery is less than \$1600, compared to an average car cost of around \$6000. Jenkins (2009) explored this aspect of subprime lending in more detail.

From the lender's perspective, there is a large difference between loans that pay off and loans that end in default. Figure 1 plots the distribution of per-sale profits, computed by adding up the firm's revenue from the down payment, loan payments, and recoveries, and subtracting off the cost of the car and collections. The distribution is highly bimodal. Paid loans are quite profitable; defaulted loans are not. Note that these profit calculations, and those we make subsequently in the paper, are made as if loans are held to maturity. In fact, loans during the sample period were securitized and sold into the secondary market. We return to this point when it becomes relevant in our discussion of optimal pricing.

Figure 1 raises the question of whether it is possible to predict likely defaulters. To provide a rough assessment, we group buyers into "high," "medium," and "low" risk using the company's proprietary credit categories. In the sample of uncensored loans, the default rate for high-risk buyers is 71 percent,

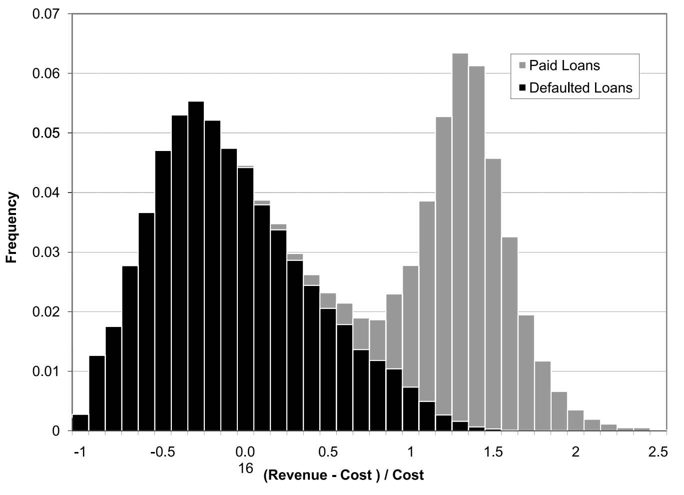

FIGURE 1.—Rate of return histogram. A unit of observation is a loan, and the figure is based on data from uncensored loans only. Revenue is calculated as the sum of the down payment, the present value of the loan payments, and the present value of the recovery amount, assuming an internal firm discount rate of 10 percent. Cost includes the purchase cost at the auction and reconditioning cost.

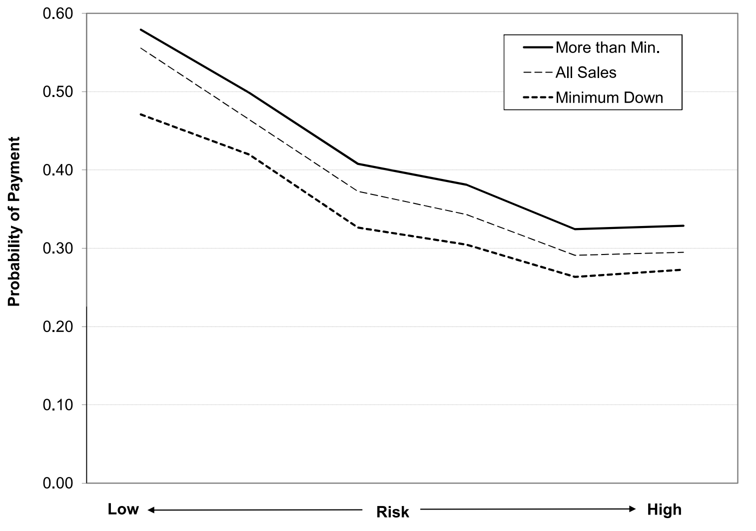

FIGURE 2.—Probability of payment by risk type and down payment. The figure is based on data from uncensored loans only. The horizontal axis represents a (discrete) measure of buyer riskiness based on the internal company credit scoring. The vertical axis represents the probability that buyers in that risk category repay the loan in full. The "All Sales" line (long-dashed) shows the relationship between risk category and the probability of payment for all buyers. The "More than Min." line (solid) shows the relationship for buyers who put down more than the required minimum down payment. The "Minimum Down" line (short-dashed) shows this relationship for buyers who put down exactly the required minimum down payment.

compared to 44 percent for the low-risk buyers. Our later estimates, which include more detailed controls, bear out the finding that credit score is highly predictive. Figure 2 illustrates this graphically, plotting the likelihood of default against the underlying credit score, again using the sample of uncensored loans.

Figure 2 also addresses a more nuanced question of whether, conditional on observed credit risk, a buyer's actions at the dealership reveal information about the subsequent probability of default. The two other lines in Figure 2 plot default rates for two groups of buyers: those who make a minimum down payment and those who voluntarily made a larger down payment. The default rate for the former group is 67 percent compared to 56 percent for the latter group. There are two explanations for the correlation between financing decisions and default rates (which, as we will see, survive the addition of detailed controls for buyer and car characteristics, as well as fixed effects for dealership and time periods). The first explanation is selection: high-risk buyers choose to make low down payments and leverage aggressively. The second explanation is causal: buyers who leverage more aggressively end up with larger debt burdens, and this makes default more likely. Our empirical strategy below disentangles these channels and verifies the importance of both.

3.2. *Changes in Pricing*

Our empirical strategy takes advantage of discrete changes over time in the firm's pricing. Because these changes play an important role in identifying demand responses, it is useful to provide some detail on how the company sets down payment requirements and car prices. More details about the pricing changes and why they provide useful variation for demand estimation can be found in Adams, Einav, and Levin (2009).

As described above, the company sets a separate down payment requirement for each applicant credit category. The company adjusted this schedule on more than 20 different occasions during the sample period. Figure 3 shows the time series for three representative credit categories. Some of the changes are seasonal; the company typically increases down payment requirements in the February "tax rebate" season. The company reports that the other changes arose from shifts in emphasis toward higher or lower risk borrowers, arising for strategic or back-end financing reasons, and because over time, headquarters became more confident in the credit scoring algorithm, which led to the increasing spread observed in Figure 3. Below, we use year and month controls to capture seasonal shifts and slow-moving trends in demand, and rely on

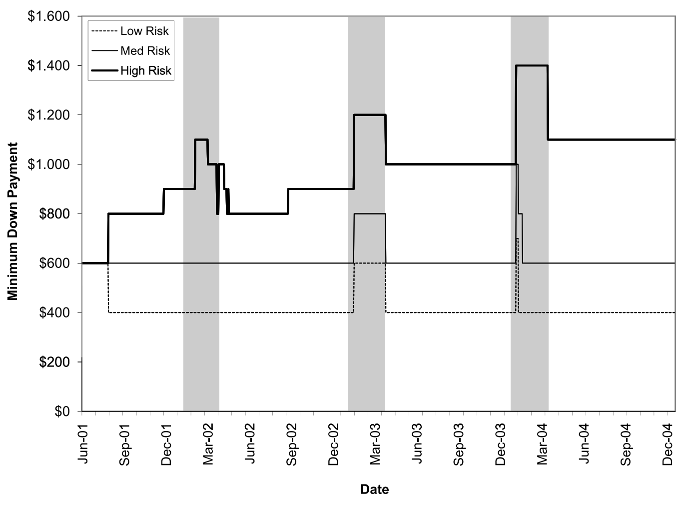

FIGURE 3.—Minimum down payments over time. Shaded areas represent "tax season" (mid-January through end of March) during which volume is higher and down payment requirements increase. Each line represents the time series of the required down payment for a *representative* high-risk, medium-risk, and low-risk applicant, based on the firm's credit scoring. Other credit categories (not shown in the figure) follow a similar time series.

the abrupt variation at the price changes to identify the demand response to different down payment requirements.[7]

The policy for setting car prices is more complicated because, as noted above, the company sets a "list price" for each car but there can be subsequent negotiation at the dealership. There were two major changes in the margin formula used to generate list prices.[8] To account for the possibility that negotiated discounts might be correlated with unobservable features of demand such as the buyer's liquidity or car value, we rely only on the list price variation for identification (that is, we use a car's list price as an "instrumental variable" for its negotiated price). The discrete changes in the list price schedule create variation over time in the price of cars with equivalent cost, and also variation over time across cars with different costs. Because we control for cost in a continuous way, we also gain some additional identifying variation from the fact that margins increase in discrete jumps, so cars with near-identical costs can have discretely different list prices.

In addition to the list price schedule and down payment requirements, we also observe some variation in interest rates and loan lengths. For instance, the length of loans expanded somewhat over time, and state interest rate caps created some variation in the company's rates. We control for this variation, but feel less confident in our ability to identify convincingly how changes in these financing terms affect the quality and quantity of demand. This is a main reason that we focus on the down payment and car price, despite the fact that interest rates and payment terms are an important part of the loan contract.

3.3. *Making Use of the Data*

In making use of the data to estimate demand and to analyze pricing, we face several issues that require particular modeling decisions. Again, some additional institutional background is useful to understand these choices.

A first issue concerns the process by which customers arrive at a dealership and enter our data. The company attracts customers primarily through street presence and referrals. Referrals often come from other dealerships that are unwilling to offer financing to borrowers with very poor credit histories. This

[7]The results we present use the full time period for estimation, but we obtain similar estimates when we restrict the sample to narrow windows around price changes. See Section 4, and especially the Online Appendix of Adams, Einav, and Levin (2009) for more discussion. It is also possible to exploit additional variation by continuously controlling for the underlying credit score, but not credit category per se, and using regression discontinuity to compare buyers with credit scores just above and below the threshold for different credit categories. In Adams, Einav, and Levin (2009) we provided a more detailed discussion of this variation and presented estimates based on it.

[8]We focus on variation in the margin formula rather than in the list price of the car because the car cost (to the company) is used as a right-hand-side variable that we view as an excellent proxy for the car quality.

market segmentation, and the fact that the company's financing terms are not publicly posted and are tailored based on a proprietary credit score, means there is little reason for referring dealers to be aware of, or respond to, pricing changes at the firm whose data we study. As a result, our later analysis treats both the customer arrival process and the outside options of customers as independent of the company's pricing decisions, conditional on location, year, and month dummies.

When customers do arrive at a dealership, they face decisions about whether to purchase, what car to purchase, and how much of the transaction they should finance. In this paper, we focus on the purchase and financing decision, rather than on car selection. A main reason for this is that we are interested primarily in the financial aspects of the market rather than consumer preferences over used vehicles. A second reason is that many customers do not face a rich set of alternatives. The company has an algorithm to determine whether a buyer is eligible to receive financing for a given car, which often leads to tight restrictions. So while one could enrich the demand analysis to incorporate car choice, our view is that the additional complexity would not be matched by sharper insights into selection, consumer behavior, and optimal pricing. Adams, Einav, and Levin (2009, pp. 64–65) provided evidence and further discussion on this point.

In estimating consumer demand, we also face a standard problem related to nonpurchasers. Despite the extremely detailed nature of the data, we do not observe the exact car on the lot that a given nonpurchaser might have chosen, or the exact discount he or she might have negotiated. The obvious remedy, and the one we adopt, is to impute this missing data. For each nonpurchasing applicant, we randomly select an applicant in the same credit and income category who purchased a car in the same week at the same dealership, and we assign the nonpurchaser the same car and negotiated price.[9]

Finally, while we focus below on a particular specification and estimation strategy, we note that we have subjected at least some of our modeling decisions to a range of robustness checks. Many of these variations are detailed in Adams, Einav, and Levin (2009). We have not found our results to be particularly altered by changes in the exact data sample or specification.

4. LOAN DEMAND AND REPAYMENT

This section develops our empirical model of consumer demand. We start with a theoretical model of consumer behavior from which we derive a set of linearized estimating equations. We then describe how the variation in pricing identifies the key parameters and our resulting estimates. In the next section,

[9]By imputing prices for nonpurchasers, we potentially lose some of the true price variation. This may appear to weaken our ability to identify the demand response to car prices. We note, however, that (i) there is no obvious alternative without additional data and that (ii) our econometric model treats negotiated prices as endogenous, and the identifying variation (in list prices) is not compromised.

we use the estimates to analyze the trade-offs in high-risk lending, and the role that down payment requirements play in screening risky borrowers and limiting loan defaults.

4.1. *A Model of Consumer Borrowing and Repayment*

Our model begins at the point when a consumer enters a dealership, and is offered a car and possible financing. Let p denote the car price and let d denote the required down payment. If the consumer chooses to purchase, she can choose any down payment $D \in [d, p]$ and borrow the remaining $L = p - D$. Loans carry a monthly interest rate z and a loan length of T months. For a loan of size L, the monthly payment is $m(L) = (zL)/(1 - (1+z)^{-T})$. In each subsequent month, the individual chooses whether to make the monthly payment. Failure to pay means default, and the car is repossessed. In practice, repossession may not be immediate, but we view this as a reasonable modeling simplification.

Individuals derive utility from their purchased car and money, and maximize expected discounted utility, applying a monthly discount factor β. Each individual has an initial car use value, denoted v_0, that declines deterministically over time. Individuals also have an initial monthly income, denoted y_0, that evolves stochastically in subsequent periods so that $y_t \sim F(\cdot|y_{t-1})$. We assume that individuals cannot borrow additionally beyond the auto loan nor can they save. This is primarily to keep the model tractable, but arguably not that unrealistic in our setting.

In a given period, individual utility can be one of three objects. If the individual owns a car of current value v, has income y, and makes a payment m, utility is $u(v, y - m)$. If the individual defaults on the loan, utility is $u_d(y)$ in the default period and in any subsequent period is $\overline{u}(y)$.[10] We allow u_d to differ from $\overline{u}$ because of potential costs associated with default. In the initial period of purchase, we also allow for individuals to obtain some incremental utility u_0 from buying the car, to reflect the idea that there might be transitory needs that triggered the individual to be in the market at that point. Therefore, at the time of purchase, individuals can differ along three dimensions: their initial income or liquidity y_0, their initial monthly use value v_0, and their incremental benefit from purchasing u_0.

We can analyze the optimal repayment policy as a (finite horizon) dynamic programming problem. Suppose a consumer initially took a loan of size L, and is current on payments up to month t. Her value function is

$$U_t(v_0, y_t; L) = \max\big\{u(v_t, y_t - m) + \beta\mathbb{E}[U_{t+1}(v_0, y_{t+1}; L)|y_t], \quad u_d(y_t) + \beta\mathbb{E}[\overline{U}(y_{t+1})|y_t]\big\}, \tag{6}$$

[10] The assumption that the costs associated with default are all realized in the default period (rather than spread over multiple periods) is not important. We make it for convenience, as it simplifies the derivations below.

where $\overline{U}(y_{t+1}) = \mathbb{E}[\sum_{\tau=1}^{\infty} \beta^{\tau-1}\overline{u}_{t+\tau}(y_{t+\tau})|y_{t+1}]$ is the postdefault value function. Note that in writing equation (6), we have suppressed the dependence of v_t on the initial car value v_0 and the dependence of the monthly payment m on the initial loan size L.

After month T, the loan is paid off and the terminal value function is

$$U_{T+1}(v_0, y_{T+1}) = \sum_{\tau=1}^{\infty} \beta^{\tau-1}\mathbb{E}[u(v_{T+\tau}, y_{T+\tau})|y_{T+1}]. \tag{7}$$

The optimal repayment policy at t depends on the consumer's car value v_t (a deterministic function of v_0), her current income y_t, and her loan size L. We therefore let $S(v_0, y_1, \dots, y_T; L)$ represent the number of months for which the consumer optimally makes payments as a function of her realized income process $y_1, \dots, y_T$.[11]

Next, consider the initial borrowing decision. A purchaser with initial type v_0, y_0, u_0 would choose her down payment to solve

$$\max_{D\geq d} u(v_0, y_0 - D) + u_0 + \beta\mathbb{E}[U_1(v_0, y_1; p - D)|y_0], \tag{8}$$

where p is the car price and d is the minimum down payment. It is useful to let $D^*(y_0, v_0; p)$ denote the unconstrained solution to this problem, that is, the solution that ignores the constraint $D \geq d$. Then making the assumption that the objective is quasiconcave, the optimal down payment is

$$D(v_0, y_0, p, d) = \max\{D^*(y_0, v_0, p), d\}. \tag{9}$$

Finally we can write the individual's value from purchasing as

$$\begin{aligned} U(v_0, y_0, u_0; p, d) = {} & u(v_0, y_0 - D(v_0, y_0; p, d)) + u_0 \\ & + \beta\mathbb{E}\big[U_1(v_0, y_1; p - D(v_0, y_0; p, d))|y_0\big]. \end{aligned} \tag{10}$$

It is optimal to purchase if $U(v_0, y_0, u_0; p, d) \geq \overline{U}(y_0)$.

4.2. *An Illustrative Example*

We now describe some properties of the consumer model based on a parameterized version that we develop fully in Appendix A in the Supplemental Material (Einav, Jenkins, and Levin (2012)). In that appendix, we make functional form assumptions on the utility functions and on the stochastic income process, and fit the parameters by matching the model output to aggregate moments in the data. One feature of the fitted model is that the income process is

[11]Note that $S \leq T$ because the model ends at period T. To simplify notation, we omit the dependence of S on T and z.

persistent: a higher income in period t implies (stochastically) higher income in period $t+1$.

The parametrized model displays a number of intuitive properties. At month t, a consumer with an active loan will make her loan payment if y_t is above a threshold $y_t^*(v_0, L)$. A higher car value and a smaller loan both make payment more likely. As a result, consumers with higher initial car values or incomes and with smaller loans have longer expected repayment. Moreover, consumers with higher initial car values or incomes are more likely to purchase, and consumers with lower incomes that do purchase choose to make lower down payments, leading to larger loans and increasing the likelihood that they default. As a result, the model naturally generates a negative correlation between the choice of down payment and default, which as noted above is an important feature of the data.

Figure 4 provides a graphical illustration of these points, using the parameterization and calibration described in Appendix A. Panel (a) displays purchase and down payment decisions in the space of initial income y_0 and car value v_0 (fixing $u_0 = 0$). Each point represents a simulated consumer at the time of purchase. The solid black lines divide individuals into three groups: nonpurchasers, purchasers who make a minimum down payment, and purchasers who make a larger down payment. The gray lines show predicted default rates with each line corresponding to an isodefault curve.

As the figure shows, individuals with low car value (low v_0) or little available cash (low y_0) do not purchase. In the latter case, that can be because they simply have too little liquidity, $y_0 < d$, or because the marginal utility from using cash for other purposes is too high, or because they realize they will be unlikely to make later payments. Individuals in the middle region purchase and make the minimum down payment. Many of these individuals have a relatively high car value, but are sufficiently illiquid that they find it costly to come up with additional cash, and are suboptimal given their high likelihood of default. The final group of individuals have higher car value and income. These individuals make down payments above the minimum and have the lowest rate of subsequent default.

Panels (b) and (c) of Figure 4 illustrate how changes in the minimum down payment and the price of the car affect individual behavior. Again, the scatterplots in both figures represent draws of individuals based on the calibrated parameters. In panel (a), we consider a substantial \$500 increase in the required down payment. This shifts both solid curves up and to the right. Fewer individuals purchase, and the individuals who cease to purchase (the "marginal" purchasers) are almost entirely individuals with low initial liquidity who would make a \$1000 minimum down payment but are unwilling to make the higher down payment. These individuals represent high default risks relative to an average buyer. So the model predicts that an increase in the minimum down payment will reduce the number of loans being made, but improve loan performance.

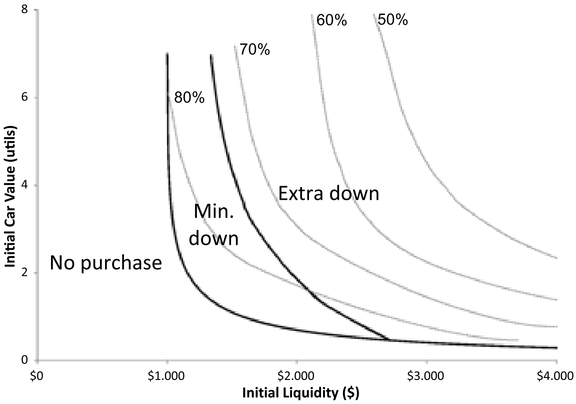

(a)

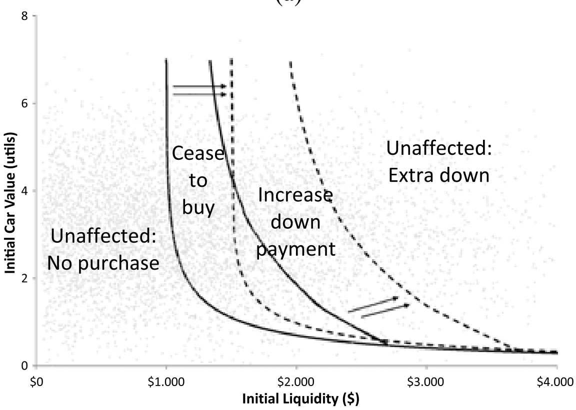

(b)

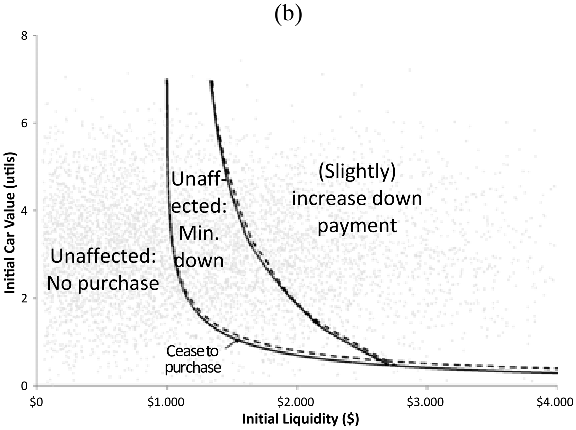

(c)

FIGURE 4.—Model illustration. Panel (a) shows consumer purchasing and down payment behavior. Panel (b) shows the distribution of consumers as a scatterplot, and the effect of changes in the minimum down payment. Panel (c) shows the effect of an increase in the car price.

Panel (c) shows the effect of a substantial \$2000 increase in car price. While the solid curves shift up in the same direction, they shift much less, indicating that the increase in car price has only a small effect on which consumers choose to purchase and on which consumers pay more than the minimum down. An important point here is that to rationalize the car price and minimum down payment demand elasticities (two of the moments used to calibrate the model), the model needs a relatively high distribution of car values and a relatively low distribution of consumer liquidity. An implication is that in response to a price increase, consumers primarily respond by taking extra financing because of their tight liquidity position. The larger loans that result increase the rate of subsequent default. So the model predicts that raising car price generates a trade-off: higher monthly payments from borrowers, but higher likelihood of default.

4.3. *Econometric Specification*

We now link the consumer model to the observable data. Recall that in the data, we observe the characteristics of each loan applicant, the offer they received, their decision about whether to purchase, their choice of down payment and loan size, and their subsequent length of repayment. In the model, the characteristics of loan applicants and the offers they received are related to the observed outcomes by the purchase decision $Q = \mathbf{1}\{U(v_0, y_0, u_0; p, d) \geq \overline{U}(y_0)\}$, the down payment decision $D(v_0, y_0; p, d)$ and loan size $L = p - D$, and the repayment length $S(v_0, y_1, \dots, y_T, L)$.

One empirical approach would be to parameterize and estimate the primitive elements of the consumer model, as in Appendix A. However, this approach is computationally intensive, which makes it difficult to incorporate covariates, and it requires functional form assumptions about which the data provide little direct insight. Instead, we make functional form assumptions that are more closely related to the observed outcomes.

Linear Specification

Specifically, we assume the consumer's value from purchasing, her unconstrained optimal down payment, and her repayment length are well approximated by linear specifications:[12]

$$
\begin{aligned}
(11)\qquad & U(v_0, y_0, u_0; p, d) - \overline{U}(y_0) \\
& \quad \approx \alpha_v v_0 + \alpha_y y_0 + \alpha_p p + \alpha_d I(D^* \leq d) d + u_0,
\end{aligned}
$$

[12]There is one exception to this linear approximation. Because the down payment requirement d is a constraint that does not appear to be binding for a significant fraction of the buyers, it seems useful to interact its effect with an indicator variable that is turned on when the constraint is binding (i.e., when the unconstrained down payment D^* is below the required level).

$$D^*(y_0, \nu_0, p) \approx \beta_v v_0 + \beta_y y_0 + \beta_p p, \tag{12}$$

$$\ln S(v_0, y_1, \ldots, y_T, L) \approx \min\{\gamma_v v_0 + \gamma_L L + s, \ln T\}. \tag{13}$$

In the third equation, s is a scalar random variable, and it is natural to assume $S \leq T$ (i.e., no one makes payments after their loan is fully paid).[13]

We adopt this linear specification partly for convenience and because, as we show later, it provides a good fit to the observed data. It is also possible to assess how well it matches the policy functions derived from making assumptions on model primitives and solving the consumer problem to obtain U, D^*, and S. In Appendix A, we show that for the calibrated version of the model discussed above, these computed functions do not display significant curvature in either the required down payment or the car price, the two pricing variables on which we focus.

Incorporating Covariates

Next we incorporate individual and contract characteristics. As in Section 2, we describe an applicant by a vector of characteristics $\zeta = (x^a, x^d, \varepsilon_u, \varepsilon_v, \varepsilon_y, \varepsilon_s)$. Here x^a represents observed individual characteristics including age, income, credit category, and proxies for wealth, and x^d includes dealership and time dummies. The scalar characteristics ε_u, ε_v, ε_y, and ε_s are not observed. We summarize the offered contract by $\phi = (x^c, p, d)$. The vector x^c includes the characteristics of the applicant's preferred car on the lot, including its cost, the offered interest rate, and the loan length. The remaining terms are the car price p and the required down payment d. It is useful to let $x = (x^a, x^d, x^c)$ denote the complete vector of observed characteristics other than price and minimum down payment.

We parameterize individual types as (linear) combinations of observed and unobserved characteristics, as

$$\begin{aligned} v_0 &= x'\xi_v + \varepsilon_v, \\ y_0 &= x'\xi_y + \varepsilon_y, \\ u_0 &= x'\xi_u + \varepsilon_u, \\ s &= x'\xi_s + \varepsilon_s. \end{aligned} \tag{14}$$

Moreover, we assume that $(\varepsilon_u, \varepsilon_v, \varepsilon_y, \varepsilon_s)$ are drawn from a multivariate normal distribution $N(0, W)$, independently of (x, p, d). We discuss the independence assumption later.

[13]Note that we could have constants α_0, β_0, and γ_0 in the equations, but this is equivalent to allowing v_0 to have an arbitrary mean, which we do below. Also, note that in moving from $S(v_0, y_1, \ldots, y_T, T)$ to its approximation, there might appear to be a loss of dimensionality from the vector of income realizations $y_1, \ldots, y_T$ to the single dimensional s. However, the reduction of dimensionality comes because we are interested in S rather than in the underlying $y_1, \ldots, y_T$, and S is already one dimensional.

We can now combine our parametric assumptions to obtain

$$\begin{aligned}(15)\qquad & U(v_0, y_0, u_0; p, d) - \overline{U}(y_0) \\ & \quad \approx x'(\alpha_v \xi_v + \alpha_y \xi_y + \xi_u) + \alpha_p p \\ & \qquad + \alpha_d I(D^* \leq d)d + \alpha_v \varepsilon_v + \alpha_y \varepsilon_y + \varepsilon_u,\end{aligned}$$

$$(16)\qquad D^*(y_0, \nu_0, p) \approx x'(\beta_v \xi_v + \beta_y \xi_y) + \beta_p p + \beta_v \varepsilon_v + \beta_y \varepsilon_y,$$

$$(17)\qquad \ln S(v_0, y_1, \ldots, y_T, L) \approx \min\{x'(\gamma_v \xi_v + \xi_s) + \gamma_L L + \gamma_v \varepsilon_v + \varepsilon_s, \ln T\}.$$

We can define new parameters and random variables:

$$\begin{aligned}(18)\qquad & \alpha_x \equiv \alpha_v \xi_v + \alpha_y \xi_y + \xi_u, \quad \varepsilon_Q \equiv \alpha_v \varepsilon_v + \alpha_y \varepsilon_y + \varepsilon_u, \\ & \beta_x \equiv \beta_v \xi_v + \beta_y \xi_y, \quad \varepsilon_D \equiv \beta_v \varepsilon_v + \beta_y \varepsilon_y, \\ & \gamma_x \equiv \gamma_v \xi_v + \xi_s, \quad \varepsilon_S \equiv \gamma_v \varepsilon_v + \varepsilon_s.\end{aligned}$$

Because $(\varepsilon_u, \varepsilon_v, \varepsilon_y, \varepsilon_s)$ is drawn from $N(0, W)$ independently of (x, p, d), it follows that $(\varepsilon_Q, \varepsilon_D, \varepsilon_S)$ is also joint normal, mean zero, and independent of (x, p, d).

4.4. *Estimating Equations*

With these changes of variables, and imposing the approximation exactly, we have our three estimating equations. The first is the purchasing equation. The model implies that individuals decide to purchase ($Q = 1$) if and only if $U(v_0, y_0, u_0; p, d) \geq \overline{U}(y_0)$. From above,

$$(19)\qquad Q = \mathbf{1}\{x'\alpha_x + \alpha_p p + \alpha_d I(D^* \leq d)d + \varepsilon_Q \geq 0\}.$$

The parameters α_p and α_d play an important role in the pricing analysis. They define the sensitivity of purchasing decisions to changes in car price and in the required down payment (when binding). One should note that the indicator term $I(D^* \leq d)$ is a function of the desired down payment, which is given by $D^* = x'\beta_x + \beta_p p + \varepsilon_D$ as derived above.

The second equation specifies the choice of down payment, which is restricted by the required down payment d, so can be expressed by

$$(20)\qquad D = \max\{D^* = x'\beta_x + \beta_p p + \varepsilon_D, d\}.$$

A buyer also cannot make a down payment larger than the purchase price p, but we omit this in presenting the model. Here the parameter β_p plays a key role. It defines how consumer financing decisions react to price increases. A low value of β_p means that price increases translate mainly into larger loans or at least larger desired loans.

The third equation specifies the repayment outcome, which is given by

$$\frac{S}{T} = \min\{\exp(x'\gamma_x + \gamma_L L + \varepsilon_S), 1\}. \tag{21}$$

The parameter γ_L is an important one in this specification, as it defines the sensitivity of the repayment duration to changes in loan size. Increases in car price and changes in the down payment will both affect the loan size $L = p - D$. The larger is the value of γ_L, the greater is the causal link between financing decisions and subsequent repayment. Of course, correlation between ε_Q, ε_D, and ε_S will also connect choices at the time of purchase and loan performance.

Let us note several technical aspects that apply to the repayment decision. First, because loan lengths in the data vary somewhat across observations, we normalize the repayment duration by the loan length, so the dependent variable is $\frac{S}{T}$ (rather than S) and is capped at 1 (rather than T) for fully paid loans. Second, buyers in the data make payments on a regularly scheduled basis, most often biweekly, but sometimes more or less often. Because there is some variation in this schedule and because we sometimes observe deviations such as off-schedule or partial payments, it is convenient to work with a continuous model of repayment. Finally, for loans that occur later in our sample, there is some additional censoring because we do not observe the full repayment period. We account for it in estimating the model, but we defer a complete discussion of this detail to Appendix B, where we also provide additional details about estimation.

Discussion

The parameters of the demand system (19)–(21) are linear composites of the parameters of the value and policy functions (15)–(17). Note, however, that what matters for the pricing decisions we analyze below is how purchasing, down payments, and repayment vary with required down payments and car prices. In this sense, our estimating equations capture exactly the “right” information for the questions we want to pursue. Note also that while we derived (19)–(21) from a specific model of consumer optimization, one alternatively can view our estimating equations simply as a convenient linear model of purchasing and repayment that might approximate other choice models. To the extent that specific assumptions such as rational expectations or time consistency can be questioned, we view this flexibility as appealing.

Price Negotiation

As mentioned earlier, the company sets a list price for each car, but customers have some ability to negotiate at the dealership. Rather than formally model the bargaining process, we integrate price determination into the model by specifying a simple relationship between the negotiated price p and the list price l:

$$p = x'\lambda_x + \lambda_l l + \varepsilon_p. \tag{22}$$

A rough way to view the pricing equation is as a "first stage" regression, where list price l is used as an instrumental variable for the possibly endogenous covariate p. Here the possible endogeneity is captured by allowing ε_p, the unobservable aspect of negotiation, to be correlated with ε_Q and ε_D, the buyer's unobserved information at the time of purchase, or even with ε_S, which is unknown by the buyer but may be inferred by the dealer selling the car. In estimating the pricing equation, the parameter λ_l is of particular interest. When we consider optimal price-setting, we consider that the company influences prices through its choice of list price, and λ_l defines the rate at which list price changes pass through to actual transaction prices.

Stochastic Assumptions

To close the model, we specify a stochastic structure for the unobservables $(\varepsilon_Q, \varepsilon_D, \varepsilon_S, \varepsilon_p)$. From our earlier assumptions, $(\varepsilon_Q, \varepsilon_D, \varepsilon_S)$ are normally distributed, with mean zero. We further assume that ε_p is normally distributed, so that

$$\begin{pmatrix} \varepsilon_Q \\ \varepsilon_D \\ \varepsilon_S \\ \varepsilon_p \end{pmatrix} \sim N(0, \Sigma) \quad \text{with} \tag{23}$$

$$\Sigma = \begin{pmatrix} \sigma_Q^2 & \rho_{QD}\sigma_Q\sigma_D & \rho_{QS}\sigma_Q\sigma_S & \rho_{Qp}\sigma_Q\sigma_p \\ \rho_{QD}\sigma_Q\sigma_D & \sigma_D^2 & \rho_{DS}\sigma_D\sigma_S & \rho_{Dp}\sigma_D\sigma_p \\ \rho_{QS}\sigma_Q\sigma_S & \rho_{DS}\sigma_D\sigma_S & \sigma_S^2 & \rho_{Sp}\sigma_S\sigma_p \\ \rho_{Qp}\sigma_Q\sigma_p & \rho_{Dp}\sigma_D\sigma_p & \rho_{Sp}\sigma_S\sigma_p & \sigma_p^2 \end{pmatrix}.$$

The correlation parameters ρ_{QS} and ρ_{DS} have important economic meaning. They characterize the relationship between the applicant's unobserved motives for purchasing and down payment, and her subsequent repayment behavior. If both are zero, the purchase and financing decisions would reveal no new information about later default. Based on the descriptive evidence in Section 3, one expects that all else equal, purchasers and especially purchasers making larger down payments are better risks. This would imply that $\rho_{QS}, \rho_{DS} > 0$.

The correlation parameters ρ_{Qp}, ρ_{Dp}, and ρ_{Sp} play a role in identification. If all three were zero, the negotiated price could be treated as exogenous in modeling the purchase, financing, and repayment decisions. Finally, the variance parameters σ_Q, σ_D, σ_S, and σ_p capture the importance of unobserved characteristics relative to observed characteristics in negotiation and customer decisions.

4.5. *Estimation and Identification*

The demand model consists of the three equations for consumer purchasing, borrowing, and repayment ((19), (20), and (21)), and the pricing equation (22). These equations map the observed and unobserved characteristics of each applicant, along with the firm's list price and required down payment, into purchase, borrowing, and repayment decisions, Q, D, and S. These equations, combined with the stochastic assumption on unobservables in equation (23), lead to a likelihood function for the observed decisions. In Appendix B in the Supplemental Material (Einav, Jenkins, and Levin (2012)), we write out the full likelihood function and provide computational details on the parameter estimation, which we perform using maximum likelihood. Maximizing the likelihood is fairly standard, although somewhat cumbersome due to the need to integrate over multiple unobservables. The reported standard errors are computed using a bootstrap method.

Now consider how the variation in the data described above identifies the key parameters of the demand model. The important quantities for pricing include the sensitivity of purchasing and borrowing decisions to the down payment requirement and car price, and the sensitivity of repayment to loan size. As we discussed in Section 3.2, the company made multiple discrete changes in the required down payment schedule. In estimating demand, we include dummy variables for credit category, and for calendar year and month. The remaining time variation in the schedule identifies the purchasing and borrowing response to the required down payment. Moreover, the induced variation in down payments creates variation in observed loan sizes that identifies the sensitivity of repayment to loan size. As noted above, our estimates use the full time period, but are robust to focusing on narrower windows around the pricing changes.

The sensitivity of consumer demand to car prices is identified by the variation in the company's markup formula described in Section 3.2. There we also made the point that negotiated prices might be correlated with unobserved buyer characteristics. In the model, this is captured by allowing the unobservable ε_p in the pricing equation (22) to be correlated with the other unobservables. With this allowance, identification of the demand response to price changes comes from variation in prices induced by shifts in the list price schedule. To the extent that this same variation in car prices is passed through into loan sizes—and we will see that the majority is—it creates additional variation that identifies the sensitivity of repayment to loan size.

The other key parameters are the correlations between the unobservables that affect repayment, ε_S, and the unobservables that affect decisions at the time of purchase, ε_Q and ε_D. These correlations affect our inference about the extent of adverse selection, which translates to the implied difference between the average and the marginal customer, which are key for optimal pricing as described in Section 2. Here the variation described above is also useful: as the minimum down payment requirement changes, we observe the change in average per-loan profits and (implicitly) back out the implied profitability of the marginal applicant. Of course, our normality assumption on the unobservables is very useful here. It helps us to back out the profitability of the marginal borrowers from changes in the average profitability of inframarginal ones. Because normality is not implied by economic theory, we have experimented with other distributional assumptions and are reassured that the key parameters are not overly sensitive to our choice of distribution.

We should emphasize that our strategy for estimating demand does not impose any assumptions about prices being optimal or profit-maximizing. Indeed, by assuming that price changes are orthogonal to individual level unobservables conditional on year and month dummies, we implicitly take the view that prices were not exactly optimal on a daily basis. This view is consistent with discussions with company personnel that indicate that pricing was revisited only periodically. One advantage of our approach is that it allows us to analyze differences between the observed prices and those that maximize profit against the estimated demand system. Later, we also investigate what can be learned from imposing weak assumptions that pricing "on average" was optimal. The empirical validity of our identifying assumptions and a variety of robustness checks are discussed in more detail in Adams, Einav, and Levin (2009).

4.6. *Demand Estimates*

Table II reports summary statistics for the data next to averages predicted by the model, showing that we are able to fit the key moments in the data fairly well. Figure 5 shows the distributions of down payments and repayment lengths observed in the data and predicted by the model. The model does well in matching the distribution of down payments and repayment length. This suggests that the distributional assumptions imposed by the model—truncated normal in the case of down payments and truncated log normal in the case of repayment length—are not particularly restrictive. In fact, we tried to estimate versions of the model with an additional parameter that tilted the repayment distribution and were unable to reject the baseline specification.

The second and fourth columns of Table III report our estimates of purchasing and borrowing behavior. The second column reports the marginal effects of the variables on the probability of sale. The probability that an applicant purchases—the "close rate"—is very sensitive to the required down payment and much less sensitive to changes in the car price. A \$100 increase in the required down payment lowers the probability of sale by 2.3 percentage points, which is equivalent to a 6.7 percent reduction in volume. In contrast, a \$100 increase in car prices has essentially no economically meaningful effect on the probability of sale. An increase in car prices also has a relatively

TABLE II

MODEL FIT[a]

	Raw Data	Demand Model
Close Rate		
All applicants	0.343	0.339
Low risk	0.451	0.440
Medium risk	0.398	0.391
High risk	0.249	0.250
Probability of Minimum Down Payment		
All buyers	0.431	0.461
Low risk	0.234	0.335
Medium risk	0.428	0.463
High risk	0.570	0.566
Average Loan Size		
All buyers	$10,709	$10,674
Low risk	$11,047	$10,837
Medium risk	$10,660	$10,636
High risk	$9992	$10,093
Probability of Payment		
All buyers	0.390	0.405
Low risk	0.559	0.551
Medium risk	0.363	0.378
High risk	0.289	0.330
Fraction of Payments Made		
All buyers	0.594	0.620
Low risk	0.715	0.740
Medium risk	0.576	0.599
High risk	0.521	0.554

[a] Raw data moments are computed directly from the estimation sample (see notes to Table III for details). Demand model moments are computed based on the econometric model described in Section 4 and the parameter estimates presented in Table III. The close rate, or probability of sale, is computed using data on all applicants. The probability of making minimum down payment and average loan size are conditional on sale. The probability of payment and the fraction of payments made are computed using data on uncensored loans only.

small effect on a buyer's desired down payment. As reported in the fourth column of Table III, we estimate that a $100 increase in car prices raises a buyer's desired down payment by about $11. Therefore, it appears that the primary effect of higher car prices is to increase the size of loans that buyers take.

These findings—the substantial effect of down payment requirements and the tendency of buyers to finance price increases—are consistent with the idea that customer liquidity is an important factor in explaining purchasing and borrowing decisions. Applicant characteristics are consistent with this hypothesis. Applicants with higher income and applicants with a bank account are more likely to purchase, and those with bank accounts tend to make larger down

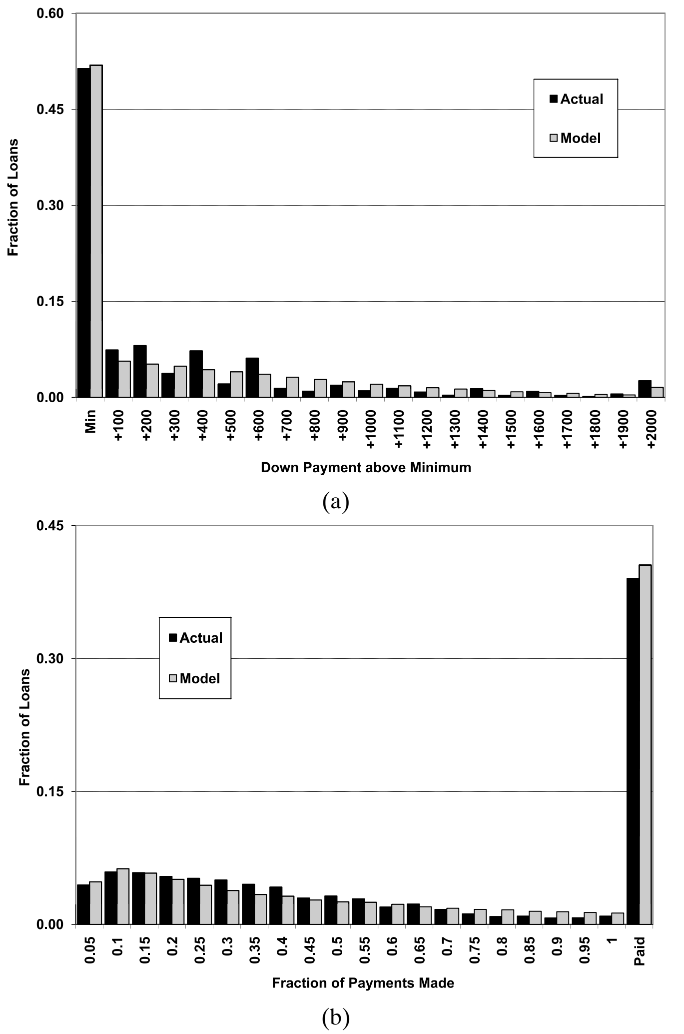

FIGURE 5.—Model fit. Panel (a) shows the model fit for down payment decisions. The last bin includes all down payments of at least $2000 above the minimum down. Panel (b) shows the model fit for the fraction of loan payments made. Note that the model slightly overstates the fraction of consumers who default just prior to paying off their loan.

payments.[14] Adams, Einav, and Levin (2009) provided additional evidence for consumer liquidity effects by analyzing an annual spike in applications that occurs each February when consumers become eligible for tax rebates. We ac-

[14]Applicants who own a house are less likely to purchase, a finding which may reflect their potential access to better credit terms at other lenders.

TABLE III

DEMAND ESTIMATES[a]

	Probability of Sale		Down Payment		Payments Made	
Dependent Variable	Marg. Eff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
Offer Variables						
Minimum down ($1000s)	−0.231	(0.011)	—	—	—	—
Negotiated price ($1000s)	−0.025	(0.006)	0.106	(0.086)	—	—
Maximum interest rate (%)	0.002	(0.002)	0.024	(0.004)	−0.030	(0.006)
Term (years)	−0.060	(0.011)	−0.375	(0.100)	0.070	(0.101)
Loan amount ($1000s)	—	—	—	—	−0.386	(0.041)
Car Characteristics						
Car cost ($1000s)	0.011	(0.007)	0.162	(0.087)	0.389	(0.043)
Premium (Cost > $7500)	0.013	(0.009)	0.458	(0.036)	0.059	(0.070)
Car age (years)	0.001	(0.001)	0.000	(0.005)	−0.049	(0.009)
Odometer (10,000s)	−0.003	(0.001)	0.000	(0.004)	0.005	(0.008)
Lot age (months)	0.001	(0.002)	−0.041	(0.017)	−0.079	(0.011)
Individual Characteristics						
Income ($1000s/month)	0.021	(0.002)	−0.002	(0.009)	0.078	(0.014)
Age	0.007	(0.002)	−0.015	(0.004)	0.014	(0.007)
Age squared	−6E−05	(2E−05)	2E−04	(5E−05)	−1E−04	(8E−05)
Bank account	0.034	(0.005)	0.036	(0.017)	0.214	(0.031)
House owner	−0.032	(0.006)	0.006	(0.019)	0.008	(0.042)
Lives with parents	0.003	(0.007)	0.043	(0.019)	−0.111	(0.036)
Credit Grade Fixed Effects						
Representative low-risk grade	−0.012	(0.017)	0.349	(0.060)	0.869	(0.101)
Representative medium-risk grade	0.004	(0.014)	0.260	(0.056)	0.316	(0.081)
Representative high-risk grade	Omitted		Omitted		Omitted	
Seasonal Effects						
Tax season	0.137	(0.010)	0.545	(0.027)	0.023	(0.065)
Other Fixed Effects	Year, Month, City, Credit Grade					

Covariance Matrix	ε_Q		ε_D		ε_S		ε_p	
ε_Q (purchase equation)	1.000	—	0.062	(0.007)	0.003	(0.014)	−0.018	(0.006)
ε_D (down payment equation)	0.062	(0.007)	0.557	(0.027)	0.022	(0.014)	−0.024	(0.058)
ε_S (repayment equation)	0.003	(0.014)	0.022	(0.014)	2.334	(0.047)	0.078	(0.015)
ε_p (car price equation)	−0.018	(0.006)	−0.024	(0.058)	0.078	(0.015)	0.283	(0.020)

[a] All estimates are based on the demand model described in Section 4. The sample for the purchase and down payment equations is a random sample of all applicants; the sample size is $0.1N$, where $N \gg 50{,}000$ (see Table I). The sample for the fraction of payments made equation is all sales; sample size is $0.034N$. Reported estimates in the second column show the marginal effects of a 1 unit change in each of the explanatory variables on the probability of sale. For dummy variables, this is computed by taking the difference between the probability of sale when the variable is equal to 1 and the probability when the variable is equal to 0 (holding other variables fixed). For continuous variables, this is computed by taking a numerical derivative of the probability of sale with respect to the continuous variable. Estimates in the fourth column show the effects of a 1 unit change in each explanatory variable on the desired down payment (in $1000s). For instance, a $1000 increase in price raises the desired down payment of the average applicant by $106. Estimates in the sixth column show the effects of a 1 unit change in each explanatory variable on the log of fraction of payments made. For example, a $1000 increase in loan amount decreases the fraction of payments made by $1 - \exp(-0.386)$, or 32 percent. Standard errors are based on 60 bootstrap samples.

count for this in our demand specification by including month dummies, reporting the coefficient on the February tax season dummy in Table III. All else equal, the close rate in February is 14 percentage points higher, or 40 percent, and desired down payments are $545 higher.

Our estimates of repayment behavior are reported in the sixth column of Table III. Loan size is a primary determinant of payment duration and hence the likelihood of default. All else equal, a buyer who takes a $1000 larger loan (which translates into monthly payments being about $35 higher) makes about 32 percent fewer payments. Payment duration varies with individual and car characteristics in ways that are largely predictable. All else equal, buyers are less likely to default on higher quality cars, and buyers with greater income or with a bank account are more likely to make payments.

As discussed earlier, the company's credit score is a strong predictor of expected repayment. Table III reports the coefficients for two representative categories—a low-risk category and a medium-risk category—relative to the omitted category, which is a representative high-risk category. A representative low-risk buyer is expected to make 87 percent more payments than a representative high-risk buyer and is 24 percent less likely to default. Individual credit scores are also predictive of decisions at the time of purchasing. High credit risks appear to have the highest desire to borrow, while medium-risk applicants are the most likely to purchase. One interpretation for the nonmonotonicity in purchase probability is that low-risk buyers may have better outside opportunities.

The bottom of Table III reports the estimated variances and covariances of the unobserved individual characteristics. Consistent with our earlier discussion of Figure 2, unobserved drivers of purchasing and down payment are positively correlated with the fraction of loan payments made. All else equal, a buyer who is inclined to make a $500 larger down payment is expected to make 21.6 percent more of her payments.

To put this in context, a thought experiment is useful. Suppose two buyers who are identical on observables arrive at the lot and are offered the same car and the same financing terms. Both choose to purchase, but the first applicant chooses to make a down payment that is $500 more than the second applicant. From this decision, the company can infer that the first applicant is in fact a better credit risk. If the applicants had identical loans, the first would be expected to make 2.3 percent more payments. But the additional down payment also has a second, larger, effect, which is to reduce the first applicant's loan principal. The $500 reduction in loan size increases the expected fraction of payments by 19.3 percent. So the fact that the first applicant volunteers a larger down payment has both a signalling aspect and a direct repayment effect, and both are important.

An alternative way to think about buyers selecting into larger and smaller loans, and one that is useful for pricing calculation, is to compare the predicted repayment of the average buyer with the predicted repayment of a marginal buyer who is just indifferent between purchasing and not purchasing, and of an average nonbuyer. The estimates imply that the average buyer, given an average loan size, will make 62 percent of her payments on average and has a 59 percent chance of default. A marginal buyer, who puts down exactly the required minimum given the same loan size, would be expected to make 55 percent of her payments and would have a 66 percent chance of default. An average nonbuyer, given the same loan size, would be in between and would be expected to make 58 percent of her payments and would have a 62 percent chance of default. So we can view both selection into purchasing and into larger loans as advantageous: the average buyer represents a substantially better risk than either the average nonbuyer or the buyers who demand larger loans. However, relaxing the down payment requirement selects relatively bad risks even when compared to the general pool of nonbuyers.

A final point about the estimates pertains to the price negotiation. We find a relatively small correlation between the negotiated price and the (unobserved) credit risk. One interpretation of the low correlation is that dealership managers cannot infer much about credit risk beyond what is contained in the credit score, so that conditional on observables, the price negotiation outcome is driven mainly by factors that are unrelated to applicant creditworthiness. The small correlation that we do find is positive, which could be driven either by the higher-risk borrowers performing somewhat better in negotiation or because managers anticipate that they might benefit more from a slightly lower repayment burden (which is consistent with the results we report in the next section).

5. IMPLICATIONS FOR CONTRACT PRICING

We now turn from the demand side to consider contract pricing from the perspective of the firm. We start by defining profits and deriving conditions for optimal pricing. We then use these conditions to assess the optimality of pricing decisions and the implied costs of lending that rationalize the firm's observed policies. We explore the implications of the demand estimates for pricing decisions, and then estimate the value that can be derived from risk-based pricing and the extent to which better credit-scoring information can serve as a barrier to entry.

5.1. *Revenue Accounting*

The firm's profit from a given sale is the difference between the total revenue it receives and its cost. There are three sources of revenue: the initial down payment, the discounted value of the stream of loan payments, and the discounted recovery in the event of default. Let D denote the initial down payment, $p - D$ denote the amount that is borrowed, T denote the length of the loan, z denote the interest rate on the loan, and κ denote the firm's internal discount rate. As in the previous section, let S denote the time for which payments are made. Finally, let $k(S)$ be the expected recovery if default occurs at time S and let C denote the cost incurred in selling the car.

With this notation, the present value of a given sale is

$$\pi = D + \left(\frac{\frac{1}{\kappa}(1 - e^{-\kappa S})}{\frac{1}{z}(1 - e^{-zT})} \right) (p - D) + e^{-\kappa S} k(S) - C. \tag{24}$$

The down payment D is realized immediately. The second term is the present value of loan payments. The fraction in the expression represents the present value return on each dollar of loan principal. The third term is the value of the expected recovery, discounted back to the time of purchase. If the loan defaults (i.e., $S < T$), the expected value of the recovery is positive. If the loan is paid in full (i.e., $S = T$), there is no recovery, so $k = 0$. The final term in the profit expression is the company's cost of sale.

The demand model allows us to fill in the variables in the profit expression. To see how this works, consider an applicant with observed characteristics x and unobservables $\varepsilon = (\varepsilon_Q, \varepsilon_D, \varepsilon_S, \varepsilon_p)$. The interest rate z and length of the loan T are part of the characteristics vector x. Suppose the applicant faces a car price p and minimum down payment d. If she purchases a car, then according to the model, she will make a down payment $D(p, d, x, \varepsilon)$ and subsequently repay for length $S(p, d, x, \varepsilon)$.

The recovery value is not a component of the model in Section 4. We simplify computation by modeling and estimating this object separately from the demand model, and the recovery parameter estimates are taken as given when computing expected net revenues in supply-side estimation and counterfactuals. In estimating recoveries separately from the rest of the model, we ignore the possibility that unobserved heterogeneity in the recovery value might be related to other unobservables. Jenkins (2009) considered a unified analysis of default and recovery, and found little evidence that information or decisions at the time of purchase are systematically related to recovery value, conditional on observable characteristics and the time of default.

The recovery model consists of two separate equations. The first is a probit regression with a positive recovery indicator as the dependent variable, estimated using all loans that end in default. The second equation is a linear regression with recovery amount as the dependent variable, estimated using all observations with positive recoveries. Both equations use the same set of explanatory variables, which include car characteristics, applicant characteristics, time and city fixed effects, and the number of months that loan payments were made before default. The last variable is of particular importance, since it provides a link between recovery amount and the endogenous loan repayment

TABLE IV
RECOVERY PARAMETER ESTIMATES[a]

	Nonzero Recovery Ind. Probit		Recovery Amt. ($1000s) OLS	
Dependent Variable	dF/dx	Std. Err.	Coeff.	Std. Err.
Months paid	−0.006	(0.0002)	−0.045	(0.001)
Car Characteristics				
Car cost ($1000s)	0.027	(0.002)	0.526	(0.006)
Premium (cost > $7500)	−0.097	(0.008)	0.079	(0.024)
Car age (years)	−0.008	(0.001)	−0.068	(0.004)
Odometer (10,000s)	0.010	(0.001)	−0.029	(0.003)
Lot age (months)	−0.004	(0.001)	−0.131	(0.004)
Individual Characteristics				
Income ($1000s/month)	0.000	(0.002)	0.048	(0.006)
Bank account	0.001	(0.003)	0.078	(0.012)
House owner	0.001	(0.004)	0.118	(0.015)
Other fixed effects	Year, Month, City		Year, Month, City	

[a]The sample for the probit equation is all defaults; sample size is $0.18N$, where $N \gg 50{,}000$ (see Table I). Reported coefficients show the marginal effect of a 1 unit change in the explanatory variable on the probability of making a positive recovery. The sample for the ordinary least squares (OLS) equation is all recoveries; sample size is $0.14N$.

variable S. Credit category fixed effects are not included in either recovery equation since they are found to have very little explanatory power. Table IV shows the parameter estimates from both recovery equations. The results are fairly intuitive. An increase in the number of months before default decreases the probability of a nonzero recovery and also decreases the expected recovery amount, conditional on a nonzero recovery, by about $45 per month. Individual characteristics, such as higher monthly income, possession of a bank account, and home ownership, do not have a significant effect on the probability of recovery, but do significantly increase the expected recovery amount, conditional on recovery occurring.

The final components of profitability are the firm's internal discount rate κ and the cost of the sale. We assume an annual discount rate of 10 percent based on discussions with firm executives.[15] For the cost of sale, our data include detailed information on the cost of acquiring each car and transporting it to the lot. We treat these as the direct financial cost associated with a given sale. Naturally, we expect that sales may involve additional indirect costs, for

[15]Discussions with industry participants suggest that at the time of our study, lenders may have been using internal discount rates in the range of 8–12 percent. These relatively high rates, particularly compared to the low rates on Treasury bills at the time, reflect the risk involved in subprime lending. We also experimented by using values in the range of 5–15 percent, and found that our results were not much affected.

instance, from the additional collections effort. We address this in more detail in the next section.

5.2. *Optimal Pricing and Implied Costs of Lending*

We now examine the extent to which the firm's observed pricing can be viewed as profit-maximizing, and estimate the implied indirect costs of lending necessary to rationalize observed policies. At first glance, this problem appears straightforward. Given our estimates of the profit components, we can simply ask whether there are plausible cost assumptions that rationalize observed prices as expected profit-maximizing. The difficulty, however, is that the firm, in principle, could use an almost unlimited set of pricing policies, varying down payment requirements, prices, interest rates, and loan lengths with a host of underlying applicant and car characteristics. It seems unreasonable to expect the firm's pricing decisions to be designed optimally in every instance.

We therefore focus on a single dimension of pricing—the down payment requirement—and on alternative weak assumptions about optimality. The first is that over the sample period, the firm's down payment requirements were on average correct, in the sense that a uniform increase or decrease in the down payment required of every applicant in the data would not have improved expected profit. We explicitly do not assume that the structure of down payments was optimal for each risk category, allowing us to examine later whether the firm could have profitability shifted its lending by making down payment requirements more or less sensitive to assessed risk.

We also start with a very simple specification of the indirect costs of lending. We write the potential cost of a sale to a given applicant i as the sum of the financial cost of the offered car c_i and a constant per-loan indirect cost ψ. So total costs are

$$C_i = c_i + \psi. \tag{25}$$

We discuss the interpretation of the indirect cost below and also consider more flexible specifications in which the per-loan cost is not constant but varies over time or according to borrower risk.

To proceed with our strategy, we construct the expected profit from the observed pricing and from alternative policies. Consider an applicant in the data with observed characteristics (x, p, d). For any such applicant, we can use the demand and repayment estimates to compute the probability of purchase $\Pr_\varepsilon[Q(x, p, d, \varepsilon) = 1]$ and the expected profit conditional on purchase $\mathbb{E}_\varepsilon[\pi(x, p, d, \varepsilon; \psi)|Q(x, p, d, \varepsilon) = 1]$. The firm's expected profit is

$$\begin{aligned} &\Pi(x, p, d; \psi) \\ &\quad = \Pr_\varepsilon[Q(x, p, d, \varepsilon) = 1] \cdot \mathbb{E}_\varepsilon[\pi(x, p, d, \varepsilon; \psi)|Q(x, p, d, \varepsilon) = 1]. \end{aligned} \tag{26}$$

Summing over applicants, who we index by i, we obtain $\sum_i \Pi(x_i, p_i, d_i; \psi)$, which is the overall expected profit from the observed pricing according to the demand model.

We can do the same calculation for a counterfactual situation in which the firm required down payments that were higher or lower by some constant amount. We say the observed down payment requirements were "on average" optimal if

$$\sum_i \Pi(x_i, p_i, d_i; \psi) \geq \sum_i \Pi(x_i, p_i, d_i + a; \psi) \quad \text{for all } a \neq 0. \tag{27}$$

This revealed preference condition provides our first way to estimate implied lending costs. We simply calculate the (uniquely defined) indirect cost parameter ψ so that equation (27) holds.[16]

Our second approach to estimating implied lending costs also relies on a weak revealed preference assumption, but has a different motivation. As described above, at any point in time, applicants faced a required down payment that depended on their credit category. We observe 22 changes in the underlying schedule, and hence 23 pricing periods. Index these periods by $\tau = 1, \dots, 23$ and let I_τ denote the set of applicants in period τ. For any applicant i, $d_i = \delta(x_i, \tau_i)$, where τ_i is the pricing period and x_i includes i's credit category. Our alternative assumption about pricing is that while changes in the pricing schedule over time may not have been optimal, each change represents an improvement over the current status quo.

To capture this in our notation, let $d_i' = \delta(x_i, \tau_i - 1)$ denote the down payment requirement that applicant i would have faced under the prior pricing schedule. We assume that the observed prices should satisfy

$$\sum_{i \in I_\tau} \Pi(x_i, p_i, d_i; \psi) \geq \sum_{i \in I_\tau} \Pi(x_i, p_i, d_i'; \psi) \quad \text{for all } \tau = 2, \dots, 23. \tag{28}$$

In contrast to equation (27), there is not necessarily a value of the indirect cost parameter ψ that satisfies (28). Therefore, as our estimate of ψ, we use the value that minimizes the sum of squared violations of the 22 inequalities in equation (28).

We report our estimates of the implied lending costs in Table V. We find that for offered financing to be preferable to any uniform shift in the down

[16]There is always a unique value ψ that satisfies equation (27). To see why, observe that we can write $\Pi(x_i, p_i, d_i + a; \psi)$ as the product of the probability of sale $Q(x_i, p_i, d_i + a)$ and the expected profit conditional on the sale. Write the latter term as $R(x_i, p_i, d_i + a) - \psi$, where R is the expected revenue net of the direct financial cost of the car. Because Q is decreasing in a, $\Pi(x_i, p_i, d_i + a; \psi)$ has increasing differences in (a, ψ), and $a^*(\psi) = \arg\max_a \sum_i \Pi(x_i, p_i, d_i + a; \psi)$ is an increasing function. Condition (27) is equivalent to the requirement that $a^*(\psi) = 0$, which holds for a unique value of ψ.

TABLE V

ESTIMATES OF INDIRECT COSTS[a]

	Cost Estimate	Standard Error
Uniform	2481	(68)
Credit Category		
Low risk	3265	(111)
Medium risk	2373	(74)
High risk	1932	(88)
Seasonal		
Non-tax season	2498	(70)
Tax season	2285	(151)
Learning Assumptions		
Assumption (i)	1800	(332)
Assumption (ii)	2200	(240)
Assumption (iii)	2800	(222)
Assumption (iv)	2600	(207)

[a]All results are estimates of the firm's unobserved indirect cost of sale. The top six rows are based on the first-order condition for pricing described in Section 5. Estimates assume an internal firm discount rate of 10 percent. Standard errors are based on 60 bootstrap samples. The four estimates, which are based on "learning" inequalities, reflect the following conditions: (i) period $t+1$ pricing dominates period t pricing in period $t+1$; (ii) period $t+1$ pricing dominates period t pricing in period $t+1$ (weighted); (iii) period $t+1$ pricing dominates period t pricing in period t; (iv) period $t+1$ pricing dominates period t pricing in period t (weighted). In each case, the estimates are computed by minimizing the sum of squared violations of the inequality dictated by each condition. For indirect cost estimates (ii) and (iv), violations are weighted by the number of applicants in a given period. For example, to compute learning model (i), we compute expected profits per applicant in each pricing period using observed prices in that period and observed prices from the previous period. If the updated prices result in higher profits than the previous prices, then there is no violation of the learning inequality. If updated prices performed worse than previous prices, then we calculate the loss due to the pricing change. The sum of squared losses across all pricing changes is then minimized by searching over a grid of indirect costs, and the minimizing indirect cost is our estimate.

payment schedule, the unobserved component of costs must be fairly large, on the order of \$2500. Our alternative approach, based on the idea that changes in down payments are preferred to the status quo, produces similar estimates, ranging from \$1800 to \$2800. One way to understand these results is to observe that relative to the direct financial costs, the observed down payments are too high and the number of loan originations are too low to be optimal. Instead, the down payment requirements are set "as if" there is a fairly high indirect cost of extending credit.

It is natural to ask why exactly are the observed down payment requirements high relative to the direct financial costs or what might the implied indirect costs represent. Loan servicing is an obvious explanation. While we have no direct way to measure the cost of servicing each additional loan, the firm places this cost at around \$1000. This leaves another \$1500 or so to explain. One possibility is that expanding loan originations is costly due to various adjustment costs of increasing the scale of operations. An alternative explanation is that the company was concerned about the overall risk of its originated loans—in particular, the possibility of macroeconomic shocks. That is, while our default estimates are driven by realized loan outcomes up to April 2006, the company may have realized that there was a risk of an aggregate downturn, as occurred subsequently. Finally, a third possibility is that there are frictions between the lender and the secondary market that arise when loans are securitized.

To explore the adjustment cost story, we reestimated indirect costs, allowing them to differ in tax season, and assuming that equation (27) holds separately for tax season and non-tax season applicants. To the extent that scale or capacity constraints are important, we should expect higher costs in tax season when demand is at a peak. As Table V shows, however, we do not observe this. The implied costs are similar and perhaps even slightly lower in tax season. We used a similar approach to estimate costs separately for each credit category and to explore whether the indirect costs estimate is driven primarily by a hesitancy to lend to the highest-risk borrowers. As reported in Table V, however, we find little evidence for the hypothesis that estimated costs are higher for better risks. These estimates have an interesting alternative explanation, which is that if the indirect costs are constant across risk categories, then the observed down payments should have been somewhat lower for low risks and higher for high risks. The general trend in the company's pricing, as shown in Figure 3, is consistent with them (the company) figuring this out slowly over time.

This leaves us with possible costs from securitization and aggregate risk exposure, which are connected if the secondary market pricing was distorted or sheltered the firm from risk exposure. We have no data on secondary market prices, but given the time period, we suspect that prices were, if anything, quite optimistic. Indeed a concern with mortgage lending during this period is that securitization gave lenders an excessive incentive to make marginal loans. We see little evidence of this in our estimates, and one explanation is that the securitization contracts in our setting called for substantial penalties in the event of excessive defaults. Because of this, we view aggregate risk exposure as a fairly plausible cost of lending and as a rationale for our estimates, although fees or other transaction costs of securitization might also explain some fraction of the residual implied cost.

5.3. *Down Payments and Markups*

In this section, we use our demand estimates to analyze the trade-offs involved in pricing. We do this for both the down payment requirement and the markup of car prices over cost so as to emphasize the very different nature of these contract terms.

Although substantively quite different, the mechanics are similar to the previous section. For each applicant in the data, we compute the probability of purchase $\Pr_\varepsilon[U(x, p, d, \varepsilon) \geq 0]$, the probability of default conditional on purchase $\Pr_\varepsilon[S(x, p, d, \varepsilon) < T | Q(x, p, d, \varepsilon) = 1]$, and the expected profit $\Pi(x, p, d)$ (specified in equation (26) above) for varying levels of the down payment requirement d. We include in the profit calculation our indirect cost estimate $\psi = 2481$ (see Table V). We then average over applicants to find the expected outcomes that would result from specifying a down payment requirement d for a given credit category. In this exercise, as above, we hold fixed the car assignment, offered price, and other contract characteristics for each applicant. We repeat this same exercise for car pricing. In this case, we vary the markup of price over cost for the applicants, holding fixed the other contract elements including the down payment requirement.[17]

Figure 6 illustrates the effect of changes in minimum down payment requirement: panel (a) does so for low-risk applicants and panel (b) for high-risk applicants. In both panels, the probability of sale line reflects the large estimated elasticity of purchase with respect to the down payment requirement. Our estimate, extrapolated a bit out of sample, suggests that absent a down payment requirement, the probability of purchase would reach 50%, but it sharply declines and comes close to zero as the requirement reaches \$2000 or \$2500. The second line in Figure 6 presents the probability of default and shows how the default probability decreases with required down payments. As discussed earlier, this happens for two reasons. First, an increased down payment requirement (holding car prices fixed) leads to smaller loans, and hence greater ability and incentive to repay. Second, an increased down payment requirement screens out the relatively high-risk marginal borrowers, so the remaining pool of purchasers is of better quality.

The qualitative pattern is similar for low- and high-risk borrowers. In both cases increasing the minimum down payment requirement trades off a reduction in loan originations with an improvement in loan repayments. However, the default probability is significantly higher for high-risk borrowers, implying that the resolution to this trade-off is quite different. Because high-risk applicants are less profitable, the company has a greater incentive to screen them out, leading to a relatively high optimal down payment requirement. In contrast, even marginal low-risk applicants are quite profitable, implying that expected profits are maximized at a much lower level of down payment requirement, which screens out only a small fraction.

Figure 7 repeats the exercise, focusing on changes in the markup of prices over cost and holding the down payment requirements fixed. Compared to Figure 6, the patterns are dramatically different. Here the probability of sale line is almost flat, reflecting our finding that purchase probability is hardly sensitive to the markup. What then is the downside to higher markups? The answer

[17]In analyzing price–cost margins, we continue to abstract from car choice. This involves a more substantive restriction than when we estimated the demand model alone, because a large out-of-sample change in the pricing policy might cause an applicant to substitute to a different preferred car. More precisely, the way to think about the current exercise is that we are considering a change in the markup for all cars on the lot and assuming that such a change does not affect the identity of a buyer's preferred car. This seems like a natural assumption for small price changes, but perhaps more questionable for very large ones.

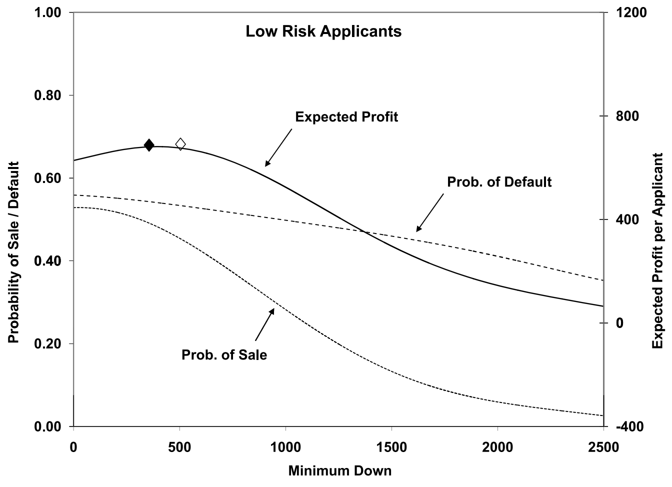

(a)

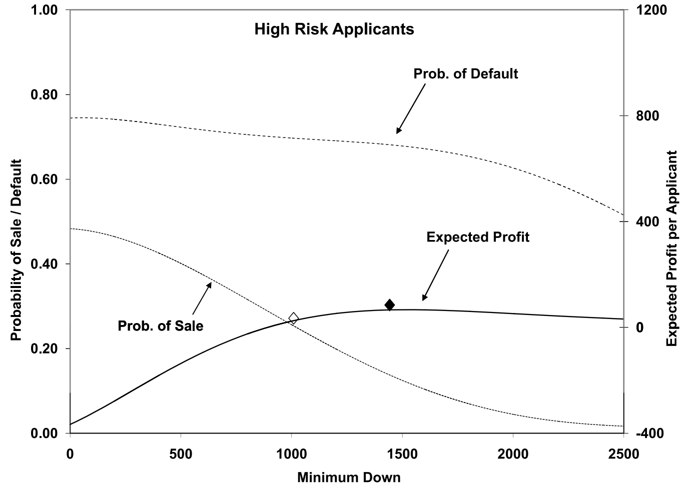

(b)

FIGURE 6.—Effect of minimum down payment changes. The figure is based on model estimates for all low-risk (a) and high-risk (b) applicants. The horizontal axis represents the required minimum down payment applied to applicants in each risk category. The left-hand y axis represents the probability of sale (for applicants) and the probability of default (for buyers). The right-hand y axis represents expected profit per applicant, calculated as the probability of sale times net operating revenue per sale, where net operating revenue is equal to the sum of the down payment and the present value of loan payments and recoveries, minus total cost (observed cost and unobserved cost). The unobserved cost is estimated as described in Section 5. Open diamonds show observed average minimum down payments for each credit category. Solid diamonds show optimal minimum down payments based on the model estimates.

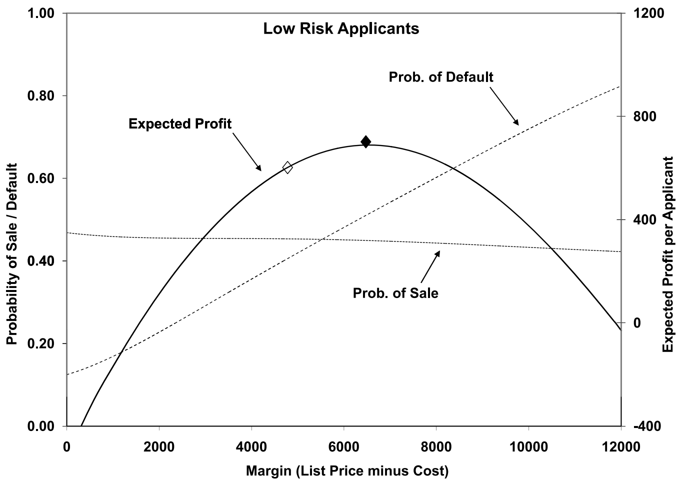

(a)

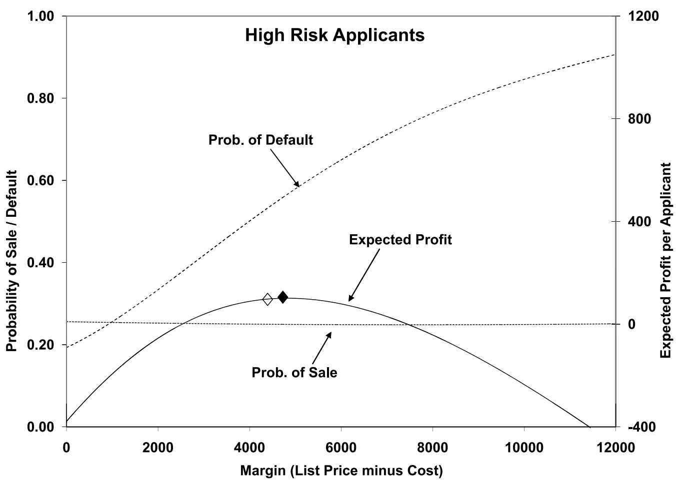

(b)

FIGURE 7.—Effect of price–cost margin changes. The figure is based on model estimates for all low-risk (a) and high-risk (b) applicants. The horizontal axis represents the target margin (list price minus cost) applied to applicants in each risk category. The left-hand *y* axis represents the probability of sale (for applicants) and the probability of default (for buyers). The right-hand *y* axis represents expected profit per applicant, calculated as the probability of sale times net operating revenue per sale, where net operating revenue is equal to the sum of the down payment and the present value of loan payments and recoveries, minus total cost (observed cost and unobserved cost). The unobserved cost is estimated as described in Section 5. Open diamonds show the observed average target margin for each credit category. Solid diamonds show the optimal target margin based on the model estimates.

becomes clear by looking at the probability of default. Increased car prices substantially raise the probability of default. So in looking at how expected profits vary with price, the curve in Figure 7 reflects the trade-off between the size of the payments and the probability they will be made, with customer selection playing essentially no important role. Comparing the two panels suggests that the optimal price is actually higher for low-risk applicants, who are less likely to default when margins are increased.

To summarize, the empirical findings in Figures 6 and 7 highlight the discussion in the earlier sections. Different elements of the loan contract introduce very different pricing trade-offs. Setting down payment requirements leads in large part to a trade-off between origination volume and borrower selection. Setting price–cost margins leads to a trade-off between the specified payments and the probability that payments will actually be received. Note that we might suspect interest rates to play a similar role, with higher rates leading to larger specified payments but more chance of default. The idea, therefore, that one could track a credit market by focusing on a single pricing dimension is clearly not the right approach to thinking about how markets may be functioning or evolving.

5.4. *Risk-Based Financing and the Value of Information*

The importance of credit scores in consumer lending is widely appreciated. Most consumers these days understand that their credit history may disqualify them from certain loan offers or qualify them for better terms. Our data, and our estimated model, provide an opportunity to assess quantitatively the value that this sort of risk-based pricing generates for a lender. The setting is interesting for this calculation, both because subprime lending is so risky and because we have already seen that the company's credit scores have substantial predictive power and are used heavily in setting down payment requirements.

To assess the value generated by risk-based pricing, we use our model to calculate the expected profitability of the company in different informational scenarios. The estimates are reported in Table VI. The table shows the company's per-loan profitability under the observed prices, the expected profitability if the company had set down payment requirements optimally using its existing credit categories, and its expected profit if it had not been able to use credit scores and instead set an optimal uniform down payment in each pricing period. Overall, our estimates imply that under the observed pricing, an applicant arriving at the lot was worth on average \$341. Setting down payments optimally given the credit categories, which as noted above would have meant somewhat lower down payments for low-risk applicants and somewhat higher ones for high-risk applicants, increases the expected profit per applicant by \$25 or 7 percent. In contrast, without the ability to categorize applicants by credit risk, optimal uniform pricing implies a per-applicant profits of \$299, which is 18 percent less than optimal risk-contingent profits.

TABLE VI

VALUE OF CREDIT SCORING[a]

	Low Risk	Med. Risk	High Risk	All Applicants
Minimum Down Payment				
Observed pricing	\$400	\$600	\$1000	—
Optimal credit-based pricing	\$0	\$700	\$1550	—
Optimal uniform pricing	\$800	\$800	\$800	\$800
Pricing with perfect information	—	—	—	—
Close Rate				
Observed pricing	0.451	0.398	0.249	0.343
Optimal credit-based pricing	0.505	0.355	0.119	0.305
Optimal uniform pricing	0.345	0.355	0.312	0.339
Pricing with perfect information	0.504	0.443	0.266	0.394
Profit Conditional on Sale				
Observed pricing	\$1967	\$797	\$106	\$996
Optimal credit-based pricing	\$1840	\$944	\$679	\$1201
Optimal uniform pricing	\$2137	\$944	−\$34	\$883
Pricing with perfect information	\$2292	\$1333	\$918	\$1490
Expected Profit per Applicant				
Observed pricing	\$886	\$317	\$26	\$341
Optimal credit-based pricing	\$930	\$335	\$81	\$366
Optimal uniform pricing	\$736	\$335	−\$11	\$299
Pricing with perfect information	\$1156	\$591	\$244	\$587

[a]All results are based on model estimates. Close rate is the probability that an applicant purchases a car. Profit conditional on sale is defined as net operating revenue (the sum of down payment and the present value of loan payments and recoveries, minus vehicle cost) minus our estimate of indirect cost in the top row of Table V. Expected profit per applicant is equal to the close rate times profit conditional on sale. Each counterfactual represents a different minimum down payment policy. List prices are held fixed at observed values in all counterfactuals. Observed pricing describes outcomes based on the company's observed minimum down payments, which vary both over time and across credit categories. Optimal credit-based pricing describes a counterfactual in which minimum down payments vary by credit category so as to maximize expected profit per applicant in each observed pricing period. Optimal uniform pricing describes a counterfactual in which a single minimum down payment, which is constant across credit categories, is chosen to maximize expected profit per applicant in each pricing period. Pricing with perfect information describes a counterfactual in which the firm can observe the borrower's unobservables at the time of purchase and sets a minimum down payment for each applicant equal to the maximum amount that the applicant is able to put down, and sells only to applicants with positive expected profits at this minimum down payment.

It is instructive to compare these estimates to the analysis in Einav, Jenkins, and Levin (2011). In that paper, we use different data from the same company to compare profits from before and after the company adopted credit-scoring technology and began to use risk-based pricing. Because we observe the full history of each loan in that paper, the statistical analysis is straightforward. We simply compute observed profits (revenues minus direct financial costs) and compare before and after the advent of credit scoring. The numbers we obtain there are quite comparable to what we get from the current modeling exercise if we compare profits per borrower with the observed prices to profits per borrower with a uniform down payment set at the level the company used prior to credit scoring. This provides an instructive out-of-sample check on our estimates.

An interesting question from a conceptual perspective is whether existing credit-scoring technology extracts most of the value that could be extracted with available information. A rough way to get at this is to ask how well the firm could have done with additional information, in particular, if it was able to set down payment requirement's based not just on the credit score, but also on each applicant's private information at the time of purchase (i.e., ε_Q and ε_D). Our results suggest that using this additional information would increase profits substantially: by 96 percent compared to the case of optimal uniform pricing, by 60 percent compared to the case of optimal pricing based on current risk categorization, and by 72 percent compared to observed pricing. These numbers suggest that despite dramatic improvements in risk classification and credit scoring, consumers still have significant private information that remains unpriced.

To understand exactly how credit risk information benefits the lender, rather than just the size of the benefit, we can use a graphical treatment to decompose the effect of risk-based down payment requirements. Panel (a) of Figure 8 illustrates a situation where the firm has no minimum down payment requirement.[18] Each bubble in the plot represents a credit category. The size of the bubble represents the number of loan originations that would have resulted according to the estimated model. The location on the horizontal axis represents the default rate and the location on the vertical axis represents the average down payment. Two features are notable. The first is the large number of high-risk loan originations. The second is the correlation across credit categories between down payment and default. Observably high-risk borrowers demand larger loans and subsequently default at much higher rates.

Panel (b) of Figure 8 illustrates how imposing the optimal risk-based down payment requirements affects the number of loan originations and the resulting loan sizes. The required down payment is highest for the riskiest borrowers, and loan originations for these applicants fall dramatically. Moreover, the high-risk applicants who do borrow are forced to make substantial down payments. So conditional on purchase there is a better matching between credit risk and loan size. Rather than making the smallest down payments, the highest risk borrowers are forced to make the largest. The final panel (c) in Figure 8 shows the effect on default rates. Default rates fall for all credit categories, but most dramatically for the highest risk category. As explained above, there are two forces at work: the down payment requirement screens out the risky borrowers within a credit category, and it forces others to take smaller loans, decreasing the chance of default.

[18]The analysis would be qualitatively similar if we started with a nonzero, but still uniform, down payment requirement. This version is useful also in highlighting the general importance of the down payment.

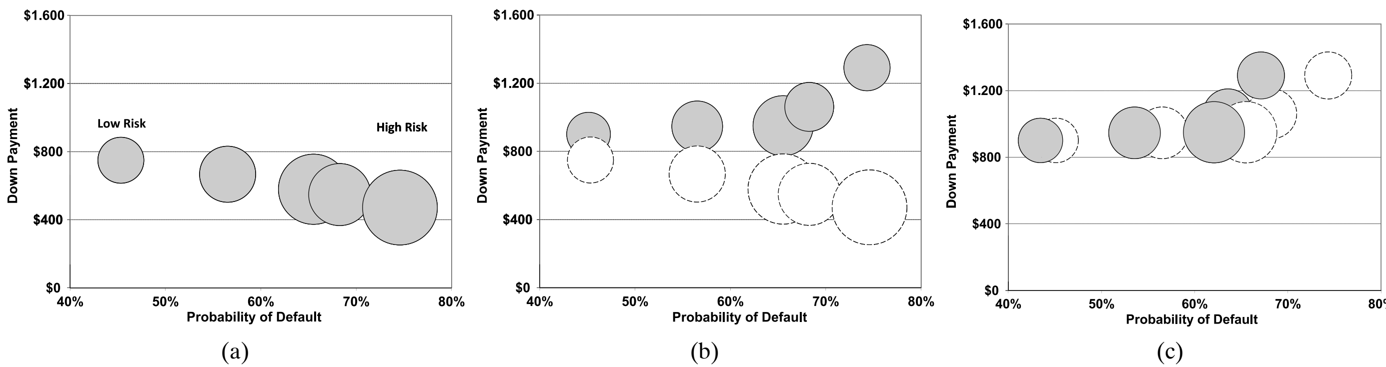

FIGURE 8.—The relationship between down payments and default rates under different conditions. The figure shows model-estimated average down payments and default rates for borrowers in different risk categories under two different pricing regimes: one with no minimum down payment and one with observed minimum down payments that vary by risk category. Each bubble represents an observable risk category defined by the company's proprietary credit score. The size of each bubble represents the estimated loan volume for a given risk category. Estimated average down payments, default rates, and loan volumes are calculated for all applicants based on the demand model estimated in the paper. Panel (a) shows estimated average down payments, default rates, and loan volumes by assuming that all applicants face observed car prices and no minimum down payment. Panel (b) shows estimated average down payments assuming applicants in each risk category face their observed average minimum down payment, keeping the default rates as in panel (a). Panel (c) adjusts panel (b) to include estimated default rates based on observed average minimum down payments (the average down payments and bubble sizes are the same as in panel (b)).

5.5. *Proprietary Information as an Entry Barrier*

A further potential effect of credit scoring is strategic. In consumer credit markets, having a better credit-scoring technology or an informational advantage can impose a cost on competitors who end up facing an adversely selected pool of borrowers. Our model provides a simple way to explore this idea, albeit one that takes us much further out of sample and relies on many stylized assumptions. Despite these limitations, we view the exercise as useful to illustrate the extent to which information may operate as an entry barrier.

For this exercise, we use the estimated model to compute per-applicant profits for a set of monopoly and duopoly scenarios. We examine cases where neither the incumbent firm nor the entrant has access to credit scores, where both have access to credit scores, and where only the incumbent has access to credit scoring. The results are reported in Table VII. Each cell presents the profits for the incumbent first and for the potential entrant second. In each scenario, we take our estimated demand system to be the market demand (as opposed to our earlier interpretation as residual demand), so the monopoly profits mirror those reported in Table VI. When an entrant is present, we find Nash equilibrium of the duopoly game, where the firms simultaneously set uniform or category-based required down payments, depending on the information they have available. Note that in our model, a fraction of the applicants are unaffected by the required down payment, and these applicants represent better-than-average risks. We assume that such applicants are randomly split between the two firms. We also make the strong assumption that in all other respects, including the car offered to each applicant and the car price, the competing firms are identical.[19]

The top row of Table VII implies that the value of information for a monopolist is \$67 per applicant. Information also has an entry-deterring benefit: better selection of applicants for an incumbent implies also that a competitor who does not have that information (and therefore has to price uniformly) would face worse selection of applicants. We compute that information reduces the potential competitor's profits by \$39 per applicant. That is, the break-even sunk entry cost that would justify entry needs to be 30% lower when an incumbent can offer contingent financing terms. Interestingly, Table VII also shows that the unilateral incentive to offer risk-based pricing is quite similar in the presence of competition, with per-applicant profits increasing by \$64 compared to \$67 when the company is a monopolist.

[19]We calculate a pure strategy Nash equilibrium (PSNE) for all cases assuming firms must choose from a discretized set of prices. As a general theoretical matter, the existence of a PSNE is not assured when one firm chooses a grade-based minimum and the other firm chooses a uniform minimum. Fortunately, we do not encounter a nonexistence problem for our particular calibrated model.

TABLE VII

CREDIT SCORING AS A BARRIER TO ENTRY[a]

	(Incumbent Profit per Applicant, Entrant Profit per Applicant)	
	Incumbent Prices Uniformly	Incumbent Prices by Risk Category
No entrant (monopoly)	($299, $0)	($366, $0)
Entrant prices uniformly	($132, $132)	($196, $93)
Entrant prices by grade	($93, $196)	($166, $166)
	Equilibrium Minimum Down Payments	
	Uniform	By Risk Category (Low, Med., High)
Monopoly	$750	$50, $750, $1450
Versus uniform	$450	$50, $400, $1000
Versus risk-based	$750	$0, $400, $1050

[a]All results are based on model estimates. Each cell in the first panel of the table presents the expected profits per applicant for an incumbent lender (first) and the profits per applicant for a potential entrant (second), calculated at the equilibrium minimum down payments shown in the corresponding cell of the second panel. The top row presents the case of no entrant, which is also presented in Table VI. The second row presents the case where an entrant prices uniformly (i.e., sets one minimum down payment for all risk categories), and the third row represents the case where an entrant prices by risk category. Each column represents the pricing strategy of the incumbent firm. In each scenario, we find the Nash equilibrium of the duopoly game in which each firm simultaneously either sets uniform or risk-based minimum down payments. We assume that applicants who choose to put down more than either firm's minimum down payment are randomly split between the two firms, and other applicants choose the lender with the lowest minimum down payment. We also assume that car prices remain the same in all scenarios. Expected profits conditional on sale are then calculated using the estimated repayment equation (Table III, sixth column) and estimated indirect costs (Table V).

6. CONCLUSIONS

This paper had two objectives: to analyze the demand and pricing for subprime loans, and to illustrate how empirical methods for studying supply and demand under imperfect competition might extend to markets where asymmetric information plays an important role. From a practical standpoint, one of our main findings is the central role that down payment requirements play in limiting loan originations and constraining borrower leverage. Our estimates show that even modestly relaxing these requirements can greatly expand and increase the riskiness of the borrower pool. This point seems particularly relevant in light of the events in the subprime mortgage market, in which lax down payment requirements allowed borrowers to become highly leveraged and, therefore, vulnerable in the face of declining house prices and underlying income or liquidity risk. Our estimates also reveal a high value, both direct and strategic, to innovations in credit scoring that allow offers to be based on the observed riskiness of loan applicants.

Our analysis of contract pricing decisions builds on an econometric model of consumer credit demand. In modeling demand, we started with a standard model of consumer choice, but then moved to a set of linearized estimating equations for estimation. We argued that one benefit of our approach is that the econometric model might be a valid approximation of a range of underlying models of consumer choice, and not just the standard model of intertemporal optimization. By focusing on the actions consumers take and not relying explicitly on a specific parameterization of an intertemporal choice model, we somewhat limit the range of counterfactuals we can perform and our ability to do welfare analysis. Nevertheless, we view this as an acceptable trade-off in modeling subprime borrowing, where standard assumptions about revealed preference, rational expectations, and consumer sophistication can be questioned.

We conclude by noting that economists increasingly have access to the type of data used in this paper, that is, detailed microdata from insurance, credit, and other contract markets. These data offer the promise of advancing our understanding of markets with asymmetric information, and providing a laboratory to test and apply the large theoretical literature on pricing and contract design. In this paper, we have tried to take a small step toward realizing this agenda. We hope the approach taken here will encourage future empirical work on pricing and contract design in settings of asymmetric information.

REFERENCES

ADAMS, W., L. EINAV, AND J. LEVIN (2009): "Liquidity Constraints and Imperfect Information in Subprime Lending," *American Economic Review*, 99, 49–84. [1389,1390,1397-1399,1410,1412]

AUSUBEL, L. (1991): "The Failure of Competition in the Credit Card Market," *American Economic Review*, 91, 50–81. [1390]

——— (1999): "Adverse Selection in the Credit Card Market," Report, University of Maryland. [1390]

EDELBERG, W. (2004): "Testing for Adverse Selection and Moral Hazard in Consumer Loan Markets," Finance and Economics Discussion Series 2004-9, Board of Governors of the Federal Reserve System, Washington. [1390]

——— (2006): "Risk-Based Pricing of Interest Rates for Consumer Loans," *Journal of Monetary Economics*, 53, 2283–2296. [1390]

EINAV, L., A. FINKELSTEIN, AND J. LEVIN (2010): "Beyond Testing: Empirical Analysis of Insurance Markets," *Annual Review of Economics*, 2, 311–336. [1390]

EINAV, L., M. JENKINS, AND J. LEVIN (2011): "The Impact of Information Technology on Consumer Lending," Report, Stanford University. [1389,1426]

——— (2012): "Supplement to 'Contract Pricing in Consumer Credit Markets'," *Econometrica Supplemental Material*, 80, http://www.econometricsociety.org/ecta/Supmat/7677_extensions-1.pdf; http://www.econometricsociety.org/ecta/Supmat/7677_extensions-2.pdf; http://www.econometricsociety.org/ecta/Supmat/7677_data and programs.zip. [1401,1409]

JENKINS, M. (2009): "Measuring the Costs of ex post Moral Hazard in Subprime Lending," Report, University of Pennsylvania. [1389,1395,1416]

KARLAN, D., AND J. ZINMAN (2008): "Credit Elasticities in Less Developed Countries: Implications for Microfinance," *American Economic Review*, 98, 1040–1068. [1390]

——— (2009): "Observing Unobservables: Identifying Information Asymmetries With a Consumer Credit Field Experiment," *Econometrica*, 77, 1993–2008. [1390]

MIAN, A., AND A. SUFI (2009): "The Consequences of Mortgage Credit Expansion: Evidence From the U.S. Mortgage Default Crisis," *Quarterly Journal of Economics*, 124, 1449–1496. [1390]

PERRIGNE, I. M., AND Q. H. VUONG (2011): "Nonparametric Identification of a Contract Model With Adverse Selection and Moral Hazard," *Econometrica*, 79, 1499–1539. [1390]

VEIGA, A., AND E. G. WEYL (2012): "Multidimensional Product Design," Report, University of Chicago. [1392]

WOLAK, F. A. (1994): "An Econometric Analysis of the Asymmetric Information, Regulator-Utility Interaction," *Annales d'Economie et de Statistique*, 34, 13–69. [1390]

Dept. of Economics, Stanford University, Stanford, CA 94305, U.S.A. and NBER; leinav@stanford.edu,

Wharton School, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19104, U.S.A.; mjenk@wharton.upenn.edu,

and

Dept. of Economics, Stanford University, Stanford, CA 94305, U.S.A. and NBER; jdlevin@stanford.edu.

Manuscript received January, 2008; final revision received August, 2011.