

SUPPLEMENT TO “CONTRACT PRICING IN CONSUMER
CREDIT MARKETS”

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

BY LIRAN EINAV, MARK JENKINS, AND JONATHAN LEVIN

This appendix provides a more detailed analysis of a parameterized version of the consumer model presented in Section 4.1. We first describe our parameterization of the model. Second, we illustrate some of the theoretical properties of the model and show that the model can incorporate several important features observed in our data. In particular, the model allows for adverse selection (borrowers who put less down are more likely to default) and behavioral responses to contract terms (borrowers are more likely to default on larger loans), and generates a possibly large mass of borrowers who make exactly the minimum down payment. Third, we calibrate the parameters of the model to match several key moments in the data. We then use the calibrated model to further explore the features of the model, and, in particular, how purchase, financing, and repayment outcomes change with the two primary contract terms we focus on: car price and minimum down payment. Finally, we report how well this model can be approximated by the linearized version in Section 4.3.

APPENDIX A: THE CONSUMER OPTIMIZATION MODEL
AND ITS CALIBRATION

A.1. *A Model of Borrower Behavior*

WE BEGIN WITH THE CONSUMER’S value function shown in equation (6) in Section 4.1:

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

To calibrate the model, we make functional form assumptions on the utility function and the laws of motion for v_t and y_t . We assume that flow utility in the event of payment is given by $u(v_0, y_0 - D) = v_0 + \ln(y_0 - D)$ in the period of purchase and by $u(v_t, y_t - m) = v_t + \ln(y_t - m)$ in all subsequent periods. This functional form assumes that the borrower’s flow utility from payment is increasing in the value of the car, and increasing and concave in net income. Moreover, the log functional form dictates that consumption utility approaches $-\infty$ as income approaches the payment amount. After the car is paid off in period T , $u(v_t, y_t) = v_t + \ln(y_t)$.

The borrower’s flow utility in the period of default is given by $u_d(y_t) = -\varphi + \ln(y_t)$. This utility function assumes that when the borrower defaults, she no longer derives utility from the car, but can consume her entire income y_t . The utility function also includes a one-time, nonmonetary cost φ , which is assumed to be constant over time. This parameter captures the utility lost due to default costs such as a decline in credit score or the hassle of dealing with repossession.

After the period of default, the borrower receives utility from consuming her income, or $\bar{u}(y_t) = \ln(y_t)$.

We assume that both liquidity y and car value v follow a first-order Markov process, and that liquidity evolves stochastically according to $y_t \sim F(y_t|y_{t-1})$. Specifically, we assume that liquidity follows a random walk $y_t = y_{t-1} + \varepsilon_t$, where ε_t is an independent and identically distributed (i.i.d.) normally distributed liquidity shock with variance σ_ε^2 . In contrast, car value depreciates deterministically according to $v_t = (1 - \delta)v_{t-1}$.

A.2. Properties of the Model

In this section, we make several observations about the solution to the borrower's problem. The repayment problem is a standard one, in which the individual trades off having more money today (if she defaults) versus obtaining the utility flow from the car today and the option value to keep the car later (if she pays). These trade-offs and the resulting outcomes depend on the borrower's characteristics y_0 and v_0 , contract terms p , d , and m , and model primitives φ , β , δ , and $F(y_t|y_{t-1})$. To simplify the discussion, we use the following shorthand notation. Define the borrower's expected future value from default (or nonpurchase) in period t to be $\text{EV}D_{t+1} = \mathbb{E}[\bar{U}(y_{t+1})|y_t]$, and the expected future value from payment (or purchase) in period t to be $\text{EVP}_{t+1} = \mathbb{E}[U_{t+1}(v_0, y_{t+1}; L)|y_t]$ if $t < T$ and $\text{EVP}_{T+1} = \mathbb{E}[U_{T+1}(v_0, y_{T+1}; L)|y_T]$ if $t = T$.

Some Properties of the Value Function

The borrower's value function is increasing in liquidity y , weakly increasing in car value v , and weakly decreasing in the payment size m (which is increasing in the size of the loan L) and cost of default φ . This can be seen through backward induction. Consider the borrower's value function in the last period of the loan. Since EVP_{T+1} and $\text{EV}D_{T+1}$ do not depend on m or φ , $\text{EV}D_{T+1}$ does not depend on v and EVP_{T+1} is increasing in v , and the borrower's value function in the last period is (weakly) increasing in v and decreasing in m and φ . Furthermore, since both $\text{EV}D_{T+1}$ and EVP_{T+1} are increasing in y , the value function is also increasing in y .

Now consider the borrower's value function in period $T - 1$. Since EVP_T is increasing in v and decreasing in m , and since EVP_{T+1} is decreasing in φ and $\text{EV}D_{T+1}$ does not depend on φ , it follows that EVP_T is decreasing in φ , and since EVP_{T+1} and $\text{EV}D_{T+1}$ are increasing in y , EVP_T is increasing in y . A similar argument holds for all previous repayment periods, meaning the value function in each of these periods is increasing in y , weakly increasing in v , and weakly decreasing in m and φ . Finally, consider the borrower's value function at the time of purchase. This equation shows that the borrower's value function at the time of purchase is increasing in her initial liquidity y_0 and car value v_0 . Second, it is clear that the borrower's value function is decreasing in

the price of the car p , since the future value from purchase $EV P_1$ is decreasing in m , and is weakly decreasing in the minimum down payment d , since a larger d reduces the borrower's choice set with no offsetting benefit to utility.

Optimal Repayment Behavior

Our choice of logarithmic utility of money implies a decreasing marginal value of money, which (coupled with the monotonicity we assume about the income process) leads to an optimal default strategy that can be described using a cutoff rule. That is, the optimal default strategy can be characterized by a vector $(y_1^*(v_1, m), y_2^*(v_2, m), \dots, y_T^*(v_T, m))$, such that the individual defaults in period t if and only if $y_t < y_t^*(v_t, m)$. Moreover, these cutoffs are decreasing in v_t and increasing in m . In each payment period, the borrower defaults if the marginal utility she derives from consuming her loan payment exceeds the benefit she receives from continuing in the loan or if

$$(A.2) \quad \ln(y_t) - \ln(y_t - m) > v_t + \varphi + \beta[EV P_{t+1} - EV D_{t+1}].$$

The cutoff value $y_t^*(v_t, m)$ is the value of y_t that makes the borrower indifferent between payment and default, or that makes the two sides of the above inequality equal. To derive comparative statics results for the cutoff levels, we can first observe that the last period cutoff y_T^* will be increasing in m and decreasing in both v_T and φ . Proceeding iteratively, the continuation value to paying rather than defaulting in a given period t will be higher if v_t or y_t is higher, or if m or ϕ is lower, and the cutoffs $y_t^*(v_t, m)$ will be decreasing in v_t and increasing in m , and will also be higher if ϕ is higher. A lower car depreciation rate δ also will translate into lower cutoffs, as it increases the continuation value of repayment $EV P_{t+1}$. Finally, the numerical results in the next section will show that the cutoff values decrease with t , provided the depreciation rate of car utility is not too large. This is due to the increased continuation value from payment that occurs as borrowers get closer to paying off their loans in full and permanently owning the car.

The above results imply that assessed at time of purchase, the probability of repayment and the expected fraction of payments made for an individual loan will be decreasing in payment size m and increasing in the initial value of the car v_0 and the cost of default φ . The probability of default also depends on the initial income y_0 and the persistence in the income process $F(y_t|y_{t-1})$. Here the analysis is analogous to traditional first passage time models of corporate default (e.g., Merton (1974)). The probability of default and the speed of time into default are both decreasing in y_0 and increasing in the volatility of the income process. The model also makes predictions about how the probability of repayment changes with the term of the loan T , but we defer these to the numerical section below.

Optimal Down Payment

A borrower's optimal down payment in the absence of minimum down payment requirements is given by

$$(A.3) \quad D^*(y_0, v_0) = \arg \max_D \{ \ln(y_0 - D) + v_0 + u_0 + \beta \text{EV} P_1(m(p - D)) \}.$$

The first-order condition equates the marginal utility of consuming an additional dollar today with the marginal effect on expected continuation value arising from reducing the loan size by \$1.

In the calibrated version of our model, the borrower's optimal down payment is increasing in v_0 and decreasing in φ , p , and the volatility of the income process. In other words, the borrower puts more down when her likelihood of repayment is higher. The optimal down payment is also increasing in the borrower's initial liquidity y_0 . This occurs for two reasons. First, higher initial liquidity lowers the marginal cost of putting an additional dollar down. Second, since higher initial income leads to lower likelihood of default, it increases the marginal benefit of lowering loan size. With a minimum down payment requirement, these relationships hold weakly instead of strictly, since borrowers who are constrained by the minimum down payment may be unaffected by changes in other parameters. The optimal constrained down payment is also weakly increasing in d .

Optimal Purchase Behavior

An applicant purchases a car if and only if y_0 and v_0 are sufficiently high that the utility from entering into the loan exceeds the utility from the outside option, or

$$(A.4) \quad \ln(y_0 - D^*(y_0, v_0)) + v_0 + \beta \text{EV} P_1(D^*(y_0, v_0)) \geq \ln(y_0) + \beta \text{EV} D_1.$$

The comparative statics of the purchase decision are similar to those of the optimal default policy. Let y_0^* be the initial liquidity that makes the above inequality bind, and define the probability of purchase as the fraction of applicants in the population for which $y_0 > y_0^*$. Since $\text{EV} P_1$ is increasing in v_0 and decreasing in φ , p , and the volatility of the income process, while $\beta \text{EV} D_1$ is independent of these parameters, the probability of purchase is increasing in v_0 and decreasing in φ , p , and the volatility of the income process. We can also condition on y_0 and interpret the probability of purchase as the fraction of applicants in the population for which $v_0 > v_0^*$. In this case, the probability of purchase is increasing in y_0 and decreasing in φ , p , and the volatility of the income process. Taken together, the purchase probability depends on the joint distribution of individuals over the two-dimensional space of y_0 and v_0 , as illustrated below. The probability of purchase is also decreasing in the minimum down payment d .

Additional Model Features

In addition to the properties detailed above, the model has several properties that match important features of our data.

First, the model gives rise to selection effects. As shown above, the borrower's optimal down payment and expected repayment duration are both increasing in her initial liquidity y_0 . This induces a positive correlation between observed down payments and the number of subsequent loan payments made, thus giving rise to adverse selection on loan size. This correlation between down payment and loan size—namely that borrowers who put more down at the time of purchase repay more and default less—is a notable feature of our data. We note that while the model above implies this must hold, the linearized version of the model that we take to the data (described in Section 4.3) does not impose it directly.

The model also allows for selection in the purchase/no purchase decision. An increase in the required down payment screens out low-liquidity buyers who are relatively likely to default in the future. So a higher required down payment improves the composition of buyers. There is also a selection effect with respect to car price. An increase in the car price tends to screen out buyers with lower liquidity and/or lower car utility than the average buyer, meaning that they are more likely to default. In the data, however, we find that borrowers in general are insensitive to price, which suggests that these selection effects may be small.

The second important feature of the model is moral hazard, by which we mean that for a given borrower, a larger debt obligation makes her/him less likely to repay. That loan size has a positive causal effect on default is another central feature of the data. In the model above, this effect may arise due to a “mechanical” effect (i.e., borrowers faced with larger payment obligations may simply not have the cash to make payment) or a “strategic” effect (i.e., borrowers with larger payment obligations are less inclined to make a payment even if it is feasible). The former can be thought of as occurring whenever $y_t < m$, while the latter occurs if $y_t > m$ but also $y_t < y_t^*$.

A third feature of the model is that many borrowers optimally choose to put exactly the required minimum down. In the data, over 40 percent of the borrowers chose to put down exactly the required minimum, and this can be generated by the model given that the constraint in the borrower's optimal down payment problem ($D \geq d$) may be binding for many levels of initial liquidity y_0 . Moreover, since the optimal down payment is increasing in y_0 , borrowers who put down the minimum are those who have the least initial liquidity and are therefore those who are most likely to default (see Figure 2).

Finally, the model can match key properties of the default timing distribution. In particular, the frequency of early defaults and relatively infrequent late defaults. Early defaults are prevalent in the model for two reasons. First, as long as car utility does not depreciate too quickly, the minimum level of

income required for payment increases with time. This occurs because the future value from payment increases as the borrower gets closer to the end of the loan due to the possibility of owning the car after paying off the loan in the last period. Second, borrowers who make payments are positively selected over the loan term in the sense that surviving borrowers will have, on average, higher liquidities and car values than borrowers who have defaulted, and are therefore less likely to be close to the default margin.

A.3. Numerical Calibration and Illustration of the Model

The above results suggest that our model of borrower behavior makes several intuitively appealing predictions about the relationship between structural parameters and model outcomes. These predictions did not depend on the exact values of the contract terms or other parameters that affect the borrower's value function or on the distribution of applicant characteristics in the population (i.e., the joint distribution of y_0 and v_0). However, these qualitative features do not assure us that our model can match the important features of our data quantitatively. In this section, we calibrate our model by searching for parameters that allow the model to match several key moments in the data. We then examine the predictions of the calibrated model. This allows us to clarify some of the theoretically ambiguous predictions described above, such as the effect of loan term on purchase and repayment, and to consider the impact of pricing changes on borrower behavior.

Additional Parameterization of the Model

Before calibrating the model, we first make additional distributional assumptions. We assume that initial car utility v_0 and initial liquidity y_0 are possibly correlated and are drawn from a (truncated) bivariate normal distribution

$$(A.5) \quad \begin{pmatrix} v_0 \\ y_0 \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_v \\ \mu_y \end{pmatrix}, \begin{pmatrix} \sigma_v^2 & \rho_{vy}\sigma_v\sigma_y \\ \rho_{vy}\sigma_v\sigma_y & \sigma_y^2 \end{pmatrix} \right),$$

where—due to the log utility specification—we truncate this distribution at $y_0 = \$50$.¹ We normalize the one-time utility from purchase u_0 to zero.

Imposing Parameter Values on Some of the Parameters

While we choose some of the parameters to match moments in the data as described below, we impose assumptions on others. We set the offer terms to match the modal loan we observe in the data. Specifically, we set the purchase price $p = \$11,000$ (including taxes and fees), the required down payment $d = \$1000$, the annual interest rate on the loan $z = 29.9$ percent, and the

¹We also experimented with a log normal distribution, which would not require truncation, but our ability to fit the moments was much worse.

loan term $T = 42$ months. To solve the model more quickly, we assume that the loan is repaid in 42 equal monthly payments (periods) rather than the 90 payments that a typical (in our data) biweekly payment schedule implies. We also calibrate the car value to be discounted at an annual rate of 0.88, implying a per-period (monthly) depreciation rate $\delta = 0.989$, and that individuals discount utility at an annual rate of 0.75, implying a per-period (monthly) discount rate $\beta = 0.976$.

Calibration of the Remaining Parameters

We choose values for the remaining parameters of the model— $\mu_v, \mu_y, \sigma_v, \sigma_y, \rho_{vy}, \sigma_\varepsilon$, and φ —to match some key moments in the data. Specifically, we obtain values for these seven parameters so that the model predictions match the following seven (unconditional) moments we observe in the data: (i) the probability of sale, (ii) the probability of making exactly the minimum down payment conditional on sale, (iii) the average payment above the required minimum for those who pay more than the required minimum, (iv) the probability of default conditional on sale, (v) the average fraction of payments made conditional on default, (vi) the semielasticity of demand with respect to a \$100 increase in the minimum down payment, and (vii) the semielasticity of demand with respect to a \$1000 increase in car price.

Table A.I reports the values of these moments in the data and the corresponding values from the calibrated parameters. The observed moments in the table are qualitatively similar to analogous moments we obtain in the paper, although they are computed based on a more homogeneous subset of loans

TABLE A.I
MODEL FIT^a

		Actual	Model
(i)	Probability of purchase	0.345	0.341
(ii)	Probability of minimum down	0.427	0.416
(iii)	Average extra down payment if above minimum	\$577	\$543
(iv)	Probability of payment	0.392	0.412
(v)	Average fraction of payments made if default	0.344	0.292
(vi)	Change in close rate from \$100 increase in minimum down payment requirement	-0.022	-0.022
(vii)	Change in close rate from \$100 increase in car price	-0.002	-0.0003

^aThe “Actual” column lists the value of the moments that were used to calibrate the consumer behavior model described in this appendix. The corresponding “Model” column represents the model predictions for these moments at the best calibrated parameter values. Actual close rate is based on all applicants. Actual probability of minimum down and average extra down are conditional on sale. Probability of payment and fraction of payments conditional on default are based on uncensored sales only. Actual estimate for the change in close rate from \$100 increase in minimum down is based on OLS regression, using all applicants, of sale indicator on price, minimum down requirement, and other observables (the results of these regressions are qualitatively similar to those reported in Adams, Einav, and Levin (2009) and to those we obtain later in this paper).

in the data (42-month uncensored loans with biweekly repayment intervals) rather than all loans in our sample. This subset of loans was chosen to match the modal offer terms.

The observed moment corresponding to the semielasticity of demand with respect to a \$100 increase in minimum down payment was computed as the estimated coefficient on minimum down payment in a regression of a sale indicator on the minimum down payment and credit category fixed effects. The observed moment corresponding to the semielasticity of demand with respect to a \$100 increase in car price was computed analogously with car price as a regressor instead of the minimum down payment. Both estimates are similar in magnitude to those estimated in Adams, Einav, and Levin (2009), which were estimated using the full sample of applicants and a full set of controls.

Computational Details

The goal of the calibration exercise is to choose the parameters that minimize the distance between the seven empirical moments listed above and their simulated counterparts generated by the model. To compute these parameter values, we proceed in the following steps:

Step 1. Choose a set of parameter values $\mu_v, \mu_y, \sigma_v, \sigma_y, \rho_{vy}, \sigma_\varepsilon$, and φ .

Step 2. Compute optimal policy functions for all states in a discrete grid. The state variables of the model are stochastic liquidity y_t , deterministic car utility v_t , and constant payment size m . The borrower's policy functions can be summarized by an optimal down payment rule $D^*(y_0, v_0, p, d, \varphi, \sigma_\varepsilon)$ and a vector of default thresholds $(y_1^*(v_1, m, \varphi, \sigma_\varepsilon), \dots, y_T^*(v_T, m, \varphi, \sigma_\varepsilon))$. To compute these thresholds, we divide the state space into a discrete grid, with points on the grid separated by increments of \$100 for initial liquidity, \$10 for payment size, and 0.1 for initial car utility. For each discrete state, we compute the optimal policy by backward induction, starting from period T . In each period, the borrower's expectations about future liquidity depend on the Markov transition probabilities from her current discrete liquidity to each state in the next period. These probabilities are approximated by a cumulative normal distribution with mean equal to the current discrete liquidity state and standard deviation equal to the volatility of liquidity shocks σ_ε .

Step 3. Simulate initial applicant characteristics y_0 and v_0 based on $\mu_v, \mu_y, \sigma_v, \sigma_y$, and ρ_{vy} , and liquidity shocks ε_t based on σ_ε . The number of simulations we use for this exercise is 100,000.

Step 4. Compute optimal purchase and down payment decisions for each simulated applicant based on her initial liquidity and car value, and the optimal down payment function determined in Step 2. In addition, compute optimal repayment decisions for each simulated applicant based on her liquidity shocks and the default thresholds determined in Step 2. Each simulated applicant is then characterized by a purchase decision, a down payment decision, and a time to default.

Step 5. Compute the values for the seven moments we are trying to match using the simulated data calculated in Step 4.

Step 6. Compute the minimum distance function equal to the (weighted) sum of squared differences between the simulated moments and their empirical analogs. In particular, let g be a vector of empirical moments and let \hat{g} be the corresponding vector of simulated moments. The minimum distance objective function is given by $\Omega = [g - \hat{g}]'W[g - \hat{g}]$, where W is a weighting matrix.

Step 7. Repeat Steps 1–6 until the objective function Ω is minimized. In practice, this minimization is completed using a grid search over the parameter space followed by Nelder–Mead simplex minimization routines run in the regions of best fit.

Parameters and Model Fit

The calibrated model is able to fit these moments quite well, with parameter values that seem overall sensible. Table A.I reports the seven empirical moments used to calibrate the model. The model quite closely matches the probability of sale, probability of observing exactly the minimum down, the average down payment above the minimum, the probability of default, and the semielasticity of demand with respect to the minimum down payment. The model slightly underpredicts the fraction of payments made conditional on default and the semielasticity of demand with respect to a \$100 change in car price. However, we take the overall fit of the calibrated model to be an indication that our model of borrower purchase and repayment behavior is reasonable.

The calibrated model also fits other moments in the data (which were not used for the calibration) reasonably well, such as the distributional patterns of observed down payments and repayment duration. Using the calibrated models, Figure A.1(a) shows the distribution of down payment conditional on paying more than the minimum and Figure A.1(b) shows the distribution of default timing. Both these figures are plotted against the analogous distributions of the data and—despite the fact that this dimension of the data was not used to calibrate the model—they seem to be qualitatively quite similar.

The resulting calibrated parameters of the model are shown in Table A.II. The five parameters of the applicant characteristic distribution— μ_v , μ_y , σ_v , σ_y , and ρ_{vy} —are identified primarily from the five moments observed at the time of purchase. The calibrated mean flow utility (in dollars) from the car is \$17,744, with a standard deviation of \$2452.² While this parameter value is extremely high, this magnitude of car utility is needed to account for the fact that

²We do not calibrate this number directly. To arrive at this number we use our calibrated value of $\mu_v = 2.931$ and transform it to the dollar value x by solving for $\mu_v + \ln y = \ln(y + x)$ (recall that flow utility is $v + \ln y$), and assuming $y = \$1000$, which is the minimal liquidity required to purchase a car. The implied car value would be even higher if we used higher levels of y .

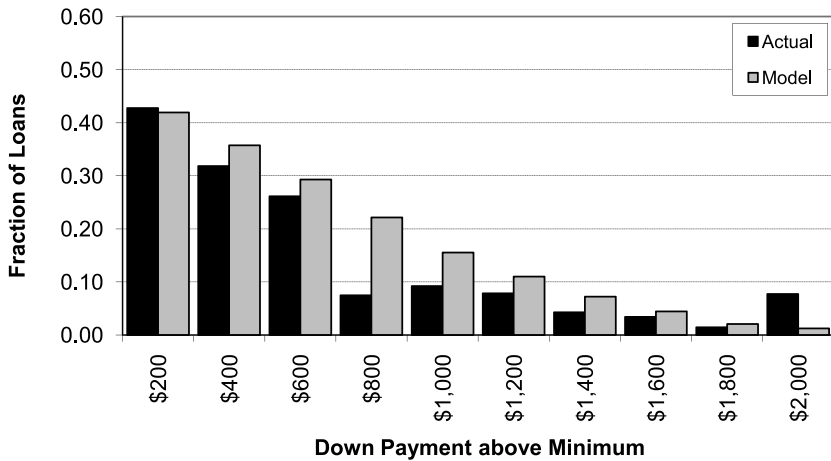


FIGURE A.1(a).—Distribution of down payments above minimum. Each bin shows the fraction of borrowers who made a down payment that exceeded the required minimum by between $(X - 199)$ and X . The bin labeled \$2000 includes all buyers whose down payments exceeded the required minimum by at least \$2000. Black bars represent observed down payments from all borrowers in our sample. Gray bars represent down payments generated by the model of consumer behavior described in this appendix.

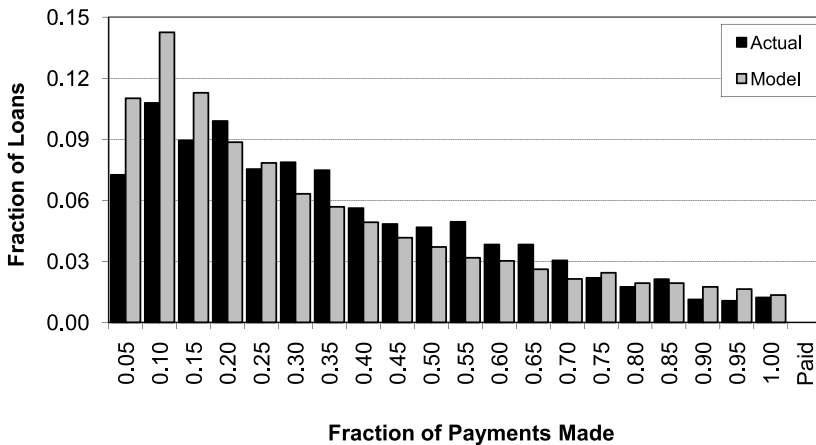


FIGURE A.1(b).—Distribution of default timing conditional on default. Each bin shows the fraction of defaulters who defaulted after making a fraction of their payments between X and $X - 0.05$. Fraction of payments made is defined as total number of payments made divided by the total number of payments due. Black bars represent observed fractions of payments made for all defaulters with uncensored loans. Gray bars represent fractions of payments made generated by the model of borrower behavior described in this appendix.

TABLE A.II
MODEL PARAMETERS^a

Calibrated Parameters	Estimate
Initial car value (mean)	\$17,744
Initial car value (std. dev.)	\$2452
Initial liquidity (mean)	\$501
Initial liquidity (std. dev.)	\$1436
Corr. b/w initial car value and liquidity	-0.025
Volatility of liquidity shocks	\$544
Cost of default	\$2800
Other (Imposed) Parameter Values	Value
Exponential discount factor (annual)	75%
Car value depreciation rate (annual)	88%
Coefficient of relative risk aversion	1.00
Liquidity autocorrelation coefficient	1.00
Offer Terms	Value
Vehicle price including taxes and fees	\$11,000
Minimum down payment requirement	\$1000
Loan APR (annual)	29.9%
Loan term (months)	42

^aThe calibrated parameters in this table are those that minimize the distance between the observed and model-generated moments presented in Table A.I. The first five parameters are the parameters of the bivariate normal distribution of applicant characteristics (y_0, v_0) . The last two parameters—the volatility of liquidity shocks and the fixed cost of default—primarily affect the borrower's repayment problem. Imposed parameters were chosen to aid identification of the calibrated parameters, and offer terms were chosen to match the modal loan observed in the data.

numerous borrowers purchase even though they are likely to default on the loan (and thus will lose the car) and those who repay the loan will have to make large payments at a high interest rate. Moreover, while a lower mean car value together with a higher mean initial liquidity could also match the observed probability of sale, it would not allow us to match the extremely low semielasticity of demand with respect to price, since more applicants would be along the horizontal purchase threshold affected by price (see Figure A.2(c) below). In fact, such high (and possibly unrealistic) estimates are one reason that in the paper, we focus on estimating behavioral responses to pricing changes rather than on estimating utility parameters.

The calibrated mean initial liquidity for applicants in the model is \$501, with a standard deviation of \$1436. This parameter is driven by the facts that many borrowers choose not to purchase and that many borrowers change their purchase decision in response to changes in the minimum down payment. The latter suggests many borrowers are near the \$1000 threshold. The large standard

deviation on initial liquidity is due to the fact that many borrowers do make more than the minimum down payment, and high levels of initial liquidity are needed to explain these borrowers.

The parameters that affect borrower repayment behavior— σ_ε , and φ —are identified primarily from the two moments of the default timing distribution, namely the probability of default and the fraction of payments made conditional on default. The estimated standard deviation of liquidity shocks is \$544 per month. While this may seem high, it is important to note that this volatility captures not only changes in liquidity (which in turn may be driven by changes in income or prioritized expenses), but all changes in the marginal utility of owning the car not captured by the constant depreciation rate δ . The estimated one-time cost of default is \$2800, which may not be unreasonable for our sample population.

A.4. *Graphical Illustration of the Model*

Panel (a) of Figure 4—replicated in this appendix as Figure A.2(a)—provides intuition for the model implications by characterizing the purchase and down payment decisions in the space of initial liquidity y_0 and car utility v_0 . Since we only allow heterogeneity across individuals in these two parameters, a point in this space fully represents an individual at the time of purchase. The cloud of scatterpoints shows the distribution of the model’s simulated applicants, which are drawn from a bivariate normal distribution with the calibrated parameters μ_v , μ_y , σ_v , σ_y , and ρ_{vy} .

Figure A.2(a) allows us to classify individuals into three groups. The first group, in the bottom-left region of the figure, consists of individuals who decide not to purchase a car. These are individuals who either do not need the car as much (low v_0) or do not have enough liquidity to pay for it (low y_0). The latter may occur for several reasons. Individuals with very low liquidity ($y_0 < d$) simply cannot make the required down payment. Individuals with higher liquidity could make the down payment, but their liquidity is sufficiently low that the marginal utility of money after making the down payment is extremely high, making static considerations sufficient for a nonpurchase. Finally, a third effect is that low liquidity today is associated with low liquidity in the future, and, therefore, may lead to more frequent and faster default. This reduces the option value associated with purchase and would also make such individuals less likely to purchase.

The second group of individuals, in the middle region, is those who decide to purchase and pay exactly the minimum down payment. These are individuals who are constrained by the minimum down payment requirement, either because of low liquidity and high marginal utility of money, or because of high default probability (due to low liquidity or intermediate car value), which makes paying more up-front suboptimal. Yet, they sufficiently need the car to justify a purchase. As one can see, the boundary of this region is downward sloping; the

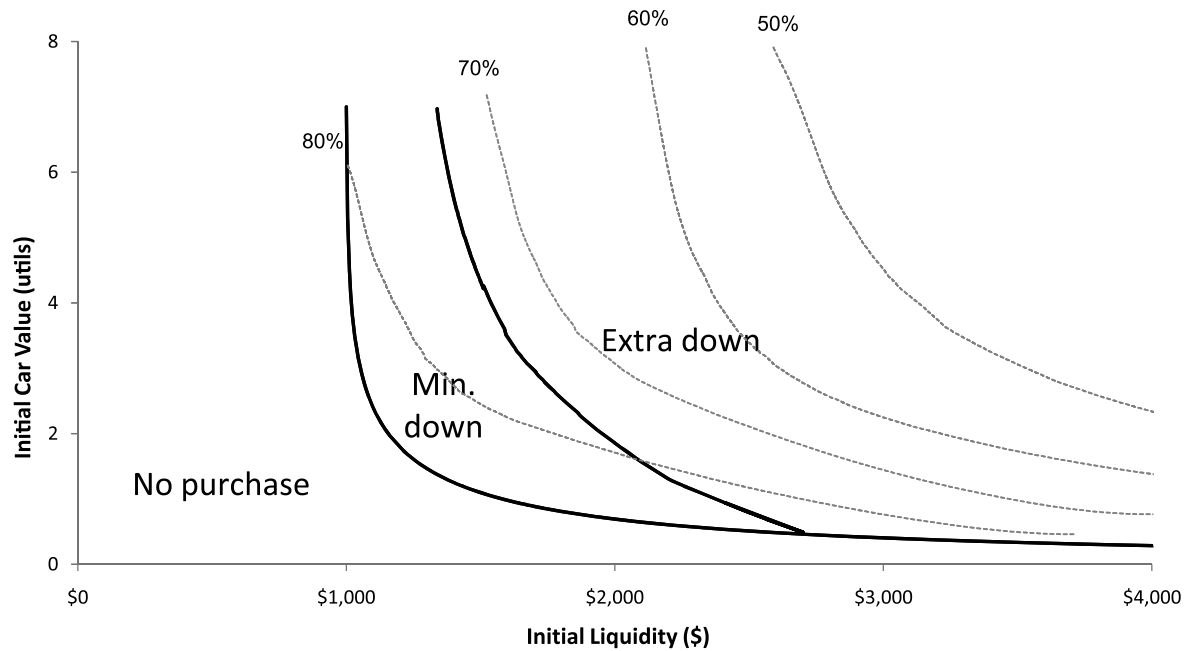


FIGURE A.2(a).—Model illustration. This figure duplicates panel (a) of Figure 4 and provides a graphical illustration of the consumer behavior model described in this appendix, assuming a required minimum down payment of \$1000 for all applicants. The horizontal axis shows the applicant’s initial liquidity (y_0), and the vertical axis shows the applicant’s initial car utility (v_0). Each applicant lies at a point in this (y_0, v_0) space. The region labeled “No purchase” contains all applicants who do not purchase. Applicants do not purchase for one of two reasons: either they do not have the liquidity to make the minimum down payment ($y_0 < 1000$) or they can make the minimum down payment but do not derive enough utility from the car to make it worthwhile. The region labeled “Min. down” contains all borrowers who make exactly the required minimum down payment. The region labeled “Extra down” contains all applicants who put down more than the required minimum. The dashed lines in the figure represent iso-default curves; for example, all applicants on the line labeled 50% have an expected default rate of 50% (based on model simulations).

liquidity level that makes an individual constrained is lower as the utility from the car is higher. This is because higher car utility is associated with less default, thereby increasing the incentive to pay more up front and reduce subsequent monthly payments.

The final group of individuals, located in the upper-right region of the figure, is those with high liquidity and high car utility. These individuals are not constrained by the minimum down payment and simply solve the unconstrained dynamic optimization problem. Indeed, these individuals also have the lowest default rate, which is consistent with the important negative correlation between down payment and default documented in Adams, Einav, and Levin (2009).

Panels (b) and (c) of Figure 4—replicated here as Figures A.2(b) and A.2(c)—illustrate how changes in the minimum down payment and the price of the car affect individuals. In Figure A.2(b), we consider a \$500 increase in the required down payment: Both the purchase threshold and the threshold for putting more than the minimum down shift up and to the right. The overall close rate is reduced, as individuals who were just to the right of the left curve no longer purchase. Given the calibrated values, these individuals have lower initial liquidity (and slightly lower initial car utility), and are, therefore, relatively high risk. This makes them less profitable than the average buyer. The second curve in Figure A.2(b) also shifts up given that the minimum down payment constraint is now binding for more individuals. In Figure A.2(c), we consider a \$2000 dollar increase in the price of the car. While again the curves shift up in the same direction, they shift much less, essentially leading to a minimal effect on the purchase decision and to only a small effect on the down payment decision.

A.5. *Linear Approximation*

As discussed in Section 4.3, we approximate the consumer model with a set of linearized estimating equations. Here we use the behavioral model developed above to illustrate how well these linear approximations approximate the policy functions derived from the calibrated version of the behavioral model.

As already described, Figures A.1(a) and A.1(b) present, respectively, the distribution of down payment conditional on paying more than the minimum and the distribution of default timing, using model simulations that are based on the calibrated parameters. Both these figures are plotted against the analogous distributions of the data and—despite the fact that this dimension of the data was not used to calibrate the model—they seem to fit qualitatively quite well. More importantly, it seems that both these figures could be well approximated by truncated normal and truncated log normal distributions, respectively, which is what we use for the estimation of the model.

In addition, in Table A.III and Figures A.3(a) and A.3(b), we use simulations from the calibrated model to report probit and tobit regressions of the key outcome variables (purchase, down payment, and repayment) on the key policy

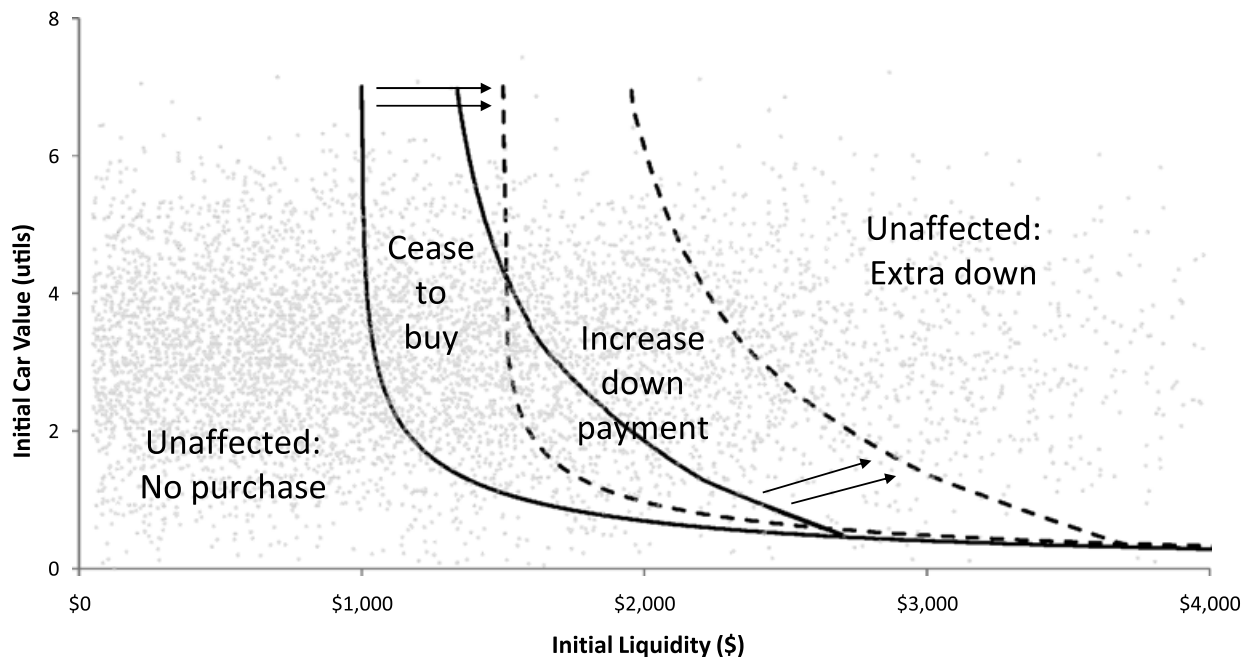


FIGURE A.2(b).—Model illustration of a \$500 increase in the minimum down payment requirement. This figure duplicates panel (b) of Figure 4 and illustrates the model prediction resulting from a \$500 increase in the required minimum down payment. Each scatterpoint on the chart represents one applicant with unobservable characteristics (y_0, v_0) , drawn from the calibrated distribution. Initial liquidity is drawn from a truncated normal distribution with truncation at \$50. The axes and the solid lines are the same as in Figure A.2(a). The dashed lines show how the curves shift as a result of the increase in the required down payment (from \$1000 to \$1500). The curves partition consumers to four groups: (i) those who did not purchase before the increase in the minimum down payment and obviously still decide not to purchase; (ii) high liquidity individuals who are unconstrained even by the higher required down payment; (iii) liquidity constrained individuals who were mostly putting the minimum down payment under \$1000 but decide to stop purchasing when faced with a \$1500 requirement; (iv) individuals with intermediate level of liquidity who still purchase, but increase their down payment to satisfy the higher requirement.




FIGURE A.2(c).—Model illustration of a \$2000 increase in car price. This figure duplicates panel (c) of Figure 4 and is similar to Figure A.2(b), but now illustrates the model prediction resulting from a \$2000 increase in the price of the car. The figure illustrates that when the car price increases by a large amount, the curves do not shift much. Individuals who did not purchase obviously do not purchase when prices go up. Most individuals who pay the minimum down payment keep doing the same, and simply take a larger loan in response to the price increase. There is only a small region in which individuals cease to purchase in response to the price change; moreover, the density of individuals in this region is relatively sparse (given our calibration exercise). These two effects allow us to match the low price elasticity in the data. Finally, all other individuals (including a small set of individuals who initially pay the minimum down payment requirement) respond to the price increase by a small increase in their down payments (which is much smaller than the price increase, thereby still leading to much larger loan amounts).

TABLE A.III
REGRESSION ESTIMATES WITH SIMULATED DATA

(A) Probit of Sale Indicator on Price ($N = 2,000,000$)				
Price	-0.0071 (0.000)	0.0017 (0.001)	-0.0026 (0.002)	-0.0025 (0.004)
Price ²	—	-0.0004 (0.000)	0.0001 (0.000)	0.0001 (0.001)
Price ³	—	—	0.0000 (0.000)	0.0000 (0.000)
Price ⁴	—	—	—	0.0000 (0.000)
log L	-1,338,220	-1,338,111	-1,338,107	-1,338,107
(B) Probit of Sale Indicator on Min. Down ($N = 2,000,000$)				
Mindp	-0.4660 (0.002)	-0.0546 (0.006)	0.2002 (0.013)	0.1320 (0.025)
Mindp ²	—	-0.2118 (0.003)	-0.5460 (0.016)	-0.3808 (0.054)
Mindp ³	—	—	0.1135 (0.005)	-0.0191 (0.041)
Mindp ⁴	—	—	—	0.0336 (0.010)
log L	-1,269,860	-1,267,101	-1,266,881	-1,266,875
(C) Tobit of Down Payment on Price ($N = 703,931$)				
Price	0.0070 (0.000)	-0.0002 (0.001)	0.0041 (0.001)	0.0000 (0.004)
Price ²	—	0.0003 (0.000)	-0.0001 (0.000)	0.0032
Price ³	—	—	0.0000 (0.000)	0.0006 (0.000)
Price ⁴	—	—	—	0.0000 (0.000)
log L	-620,486	-620,396	-620,391.31	-620,391
(D) Tobit of ln(fraction of payments) on Price ($N = 703,931$)				
Price	-0.0282 (0.000)	-0.0262 (0.002)	-0.0086 (0.004)	0.0000 (0.026)
Price ²	—	-0.0001 (0.000)	-0.0021 (0.009)	0.0091
Price ³	—	—	0.0001 (0.000)	0.0016 (0.001)
Price ⁴	—	—	—	0.0001 (0.000)
log L	-995,891	-995,890	-995,880	-995,870

instruments, minimum down payment, and price. Figures A.3(a) and A.3(b) show the relationships, and Table A.III reports regression results that use different degrees of polynomials of these key right-hand-side variables. As one can see, although the relationship has some curvature, a linear approximation

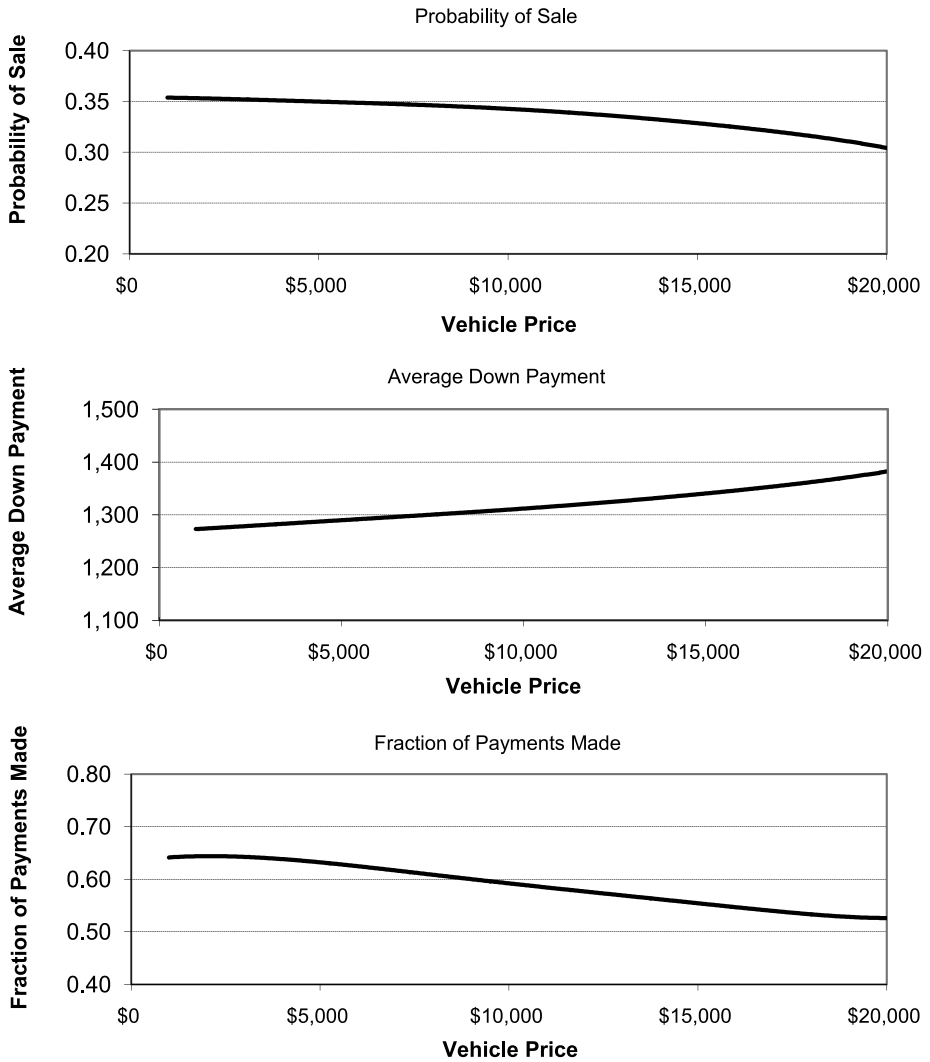


FIGURE A.3(a).—Simulated lending outcomes versus vehicle price. This figure illustrates the relationship between lending outcomes and car price for the model of borrower behavior described in this appendix. Lines are computed by fixing the pool of applicants and are simulated according to the calibrated parameter estimates presented in Table A.II, fixing the minimum down payment at \$1000 for all applicants, and varying car price. The probability of sale, average down payment, and fraction of payments made are then computed for each price based on the results of each applicant's optimization problem and simulated income draws for each borrower.

seems to fit quite well, with higher degree polynomials hardly changing the explanatory power.

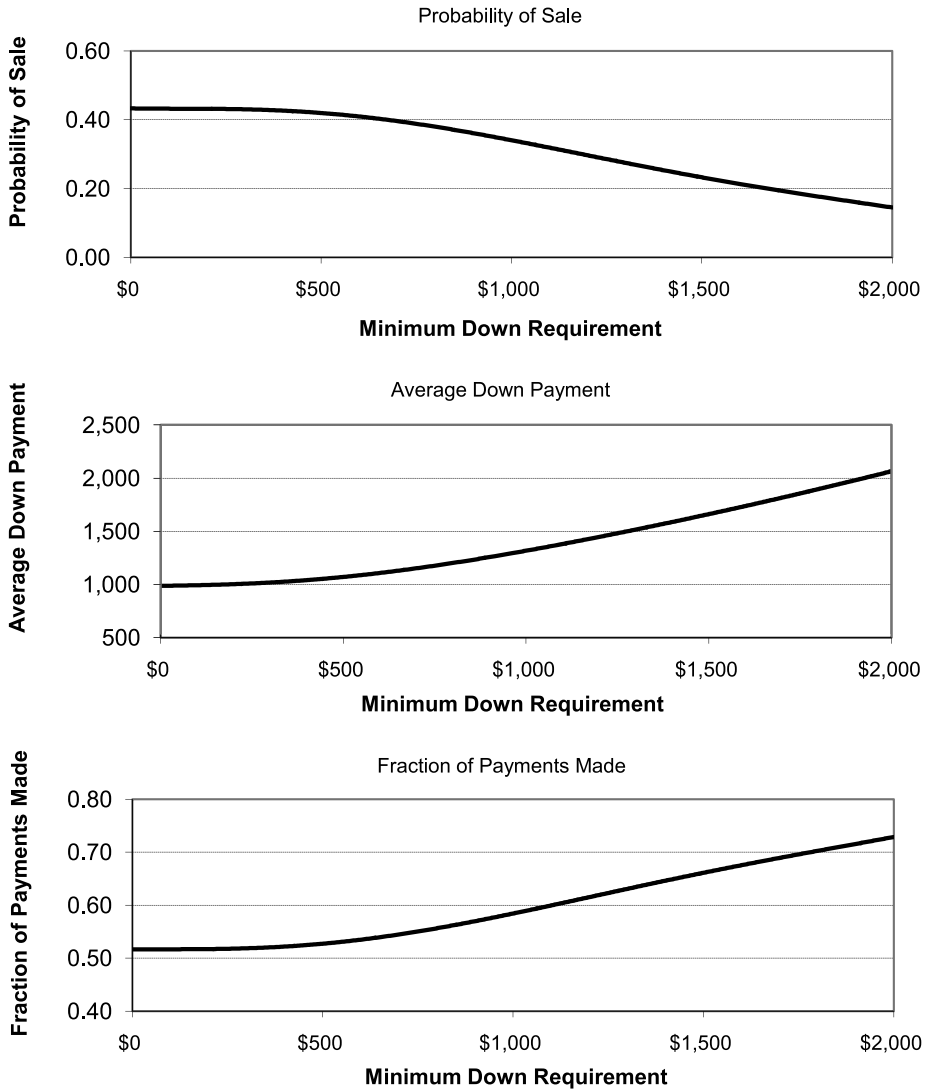


FIGURE A.3(b).—Simulated lending outcomes versus minimum down requirement. This figure illustrates the relationship between lending outcomes and minimum down payment for the model of borrower behavior described in this appendix. Lines are computed by fixing the pool of applicants and are simulated according to the calibrated parameter estimates presented in Table A.II, fixing the car price at \$11,000 for all applicants and varying the minimum down payment. The probability of sale, average down payment, and fraction of payments made are then computed for each price based on the results of each applicant’s optimization problem and simulated income draws for each borrower.

REFERENCE

MERTON, R. C. (1974): "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29, 449–470. [3]

*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.