

Credit Constraints in the Market for Consumer Durables: Evidence from Micro Data on Car Loans*

Orazio P. Attanasio
UCL, IFS and NBER

Pinelopi K. Goldberg
Yale University and NBER

Ekaterini Kyriazidou
UCLA

March 2007

Abstract

We investigate the empirical significance of borrowing constraints in the market for consumer loans. Using micro data from the Consumer Expenditure Survey (1984-1995) on auto loan contracts we estimate the elasticities of loan demand with respect to loan interest rate and maturity. The econometric specifications we employ account for important features of the data, such as selection and simultaneity. We find that – with the exception of high income households – consumers are very responsive to maturity changes and less responsive to interest rate changes. Both maturity and interest rate elasticities vary with the level of household income, with the maturity elasticity *decreasing* and the interest rate elasticity *increasing* with income. We argue that these results are consistent with the presence of binding credit constraints in the auto loan market, and that such constraints significantly affect the borrowing behavior of some groups in the population, low income households in particular.

*We would like to thank Costas Azariadis, Richard Blundell, Bo Honoré, Tom MaCurdy, Costas Meghir, Derek Neal, Josef Perktold, Frank Vella, Frank Wolak, and several seminar participants in the U.S. and Europe for useful conversations and suggestions. The current version has benefited substantially from the detailed comments of the editor and three anonymous referees. Kyriazidou and Goldberg thank the Sloan Foundation Faculty Research Fellowship Program for financial support. In addition, Goldberg thanks the NSF (Grant SBR-9731979 through the NBER). Eduardo Fajnzylber provided excellent research assistance. Jeremy Nalewaik provided invaluable help in using the NBER tax simulation program to compute marginal tax rates. The usual disclaimer applies.

1. Introduction

The existence of borrowing constraints in the market for consumer loans has important implications at both the micro and macro levels. At the micro level, credit constraints can affect both the intra- and intertemporal allocations of resources and have important consequences for the effects of policy measures. At the macro level, liquidity constraints, as borrowing restrictions are often characterized, have been invoked to explain the observed correlation between expected consumption and income growth, and the rejection of the permanent income hypothesis. Moreover, the possibility that individual agents have limited means of smoothing consumption over time has been for a long time considered as a justification for a Keynesian consumption function (see for instance Flemming, 1973). But despite the importance of the topic, and the substantial amount of theoretical and empirical research that has been devoted to it, there is still no conclusive evidence on the economic significance of borrowing constraints.

A potential explanation for this lack of consensus is the fact that most empirical work on the subject has utilized only consumption data, and not data on loans. The majority of this work has been framed in terms of a test of the life cycle - permanent income hypothesis, focusing on the excess sensitivity of consumption to expected labor income (see, for example, Hall and Mishkin (1982), Altonji and Siow (1987), Zeldes (1989), Runkle (1991)). The problem with this approach is that the interpretation of the results critically depends on explicit or implicit assumptions about the utility function. In particular, the inference of the existence of credit constraints often rests on the assumption of separability between consumption and leisure, which has been empirically rejected (Browning and Meghir (1991)).

Departing from this tradition, Jappelli (1990) relied on survey questions to identify individuals who have been denied credit, or feel that they would have been denied, had they applied for it. Given that liquidity constraints are primarily restrictions placed on borrowing, it is rather surprising that none of the above papers have utilized data on borrowing behavior to examine the empirical relevance of credit rationing.¹

This paper attempts to fill in this gap by proposing and implementing a novel approach for testing for borrowing constraints that exploits micro data on car loans. Our basic idea is that borrowing restrictions have specific implications for certain features of the demand for loans, and in particular for its interest rate and maturity elasticities. By empirically exploring these implications, one can shed light on the empirical significance of credit restrictions. The strength of this approach is that it does not rely on functional form assumptions concerning the utility function. It is particularly promising

¹More recently, another set of papers has tried to exploit the idea that in the presence of (at least partly) *collateralizable* loans (that are often used to finance durables), liquidity constraints introduce distortions in the intratemporal allocation of resources between durables and non-durables (Brugiavini and Weber (1992), Chah, Ramey and Starr (1995), Alessie, Devereux and Weber (1997)). But this idea was again implemented using only data on aggregate or household consumption.

if information on loan contracts is combined with data on socioeconomic characteristics to identify households that are a-priori more likely to face liquidity constraints.

Our approach rests on the idea that credit constraints introduce kinks and convexities in the intertemporal budget set. Liquidity constrained individuals are the ones who are either at a kink, or in the steeper portion of the budget set. This leads to the following implications which will be discussed in more detail in the next section. The demand for loans of unconstrained individuals, consuming at the flatter portion of the budget set, should be a function of the price of the loan (the primary interest rate), but independent of the loan maturity; liquidity constrained consumers, on the other hand, should respond less to changes in the primary interest rates, and more to changes in the borrowing limit. In consumer loan markets, changes in the borrowing limit are primarily achieved through changes in loan maturities; a longer maturity decreases the size of the monthly payment, allowing the consumer to assume a larger amount of debt. The implicit assumption here is that debt repayment, rather than finance charges, dominates the size of the monthly payments. This is probably a realistic assumption for the credit markets for durables which are characterized by short term contracts.² Hence, one can assess the empirical relevance of credit rationing by estimating the elasticities of loan demand with respect to interest rate and maturity, and examining how consumers respond to changes in these loan terms. A particularly interesting exercise is to estimate these responses for different consumer groups, which, based on their characteristics, have a different likelihood of being liquidity constrained, and examine whether consumers who are more likely to be constrained exhibit a larger maturity and a lower interest rate elasticity than the other groups.

Juster and Shay (1964) were the first to stress the implications of borrowing restrictions for the interest rate and maturity elasticities of the demand for loans. They used experimental data to assess the responsiveness of loan demand to interest rate and maturity in 1960. In contrast to them, we do not have experimental data, but micro data on auto loan contracts from the Consumer Expenditure Survey (1984-1995). Such contracts are an important, and fast growing component of consumer installment credit - Sullivan (1987), for example, reports that 39% of consumer credit is auto credit. We see the main strengths of our data set as being threefold: First, our information refers to actual household behavior rather than responses to hypothetical questions. Second, there is substantial time variation in interest rates and maturities that we exploit to identify the parameters of the loan demand equation. This variation can be exploited to identify credit constraints. Third, the information on demographics allows us to split the sample into various subgroups, some of which are more likely to be credit rationed than others (for example young or low income households), and test for the presence of credit rationing

²One could argue that downpayment requirements have a similar function, as they effectively limit the amount that can be borrowed. In the U.S., however, downpayment requirements are unlikely to be binding in the automobile loan markets, as most consumers use the receipts from trade-in allowances, to satisfy them. In addition, such requirements have, in many markets, dropped to zero in recent years.

separately in each of them. We are particularly interested in comparing the relative sizes of interest rate and maturity elasticities across groups.

With all its advantages, however, our data also pose several challenges: First, there is potential selection bias - observations on financing are available only for consumers who purchased a car and decided to finance such a purchase. Second, an important feature of auto loan contracts is that financing is bounded between 0 and the value of the car. Third, simultaneity issues are potentially important; the observed interest rate and maturity of a realized loan are likely to be endogenous, both in the economic and econometric sense (the loan rate and maturity lenders offer typically depend on the amount borrowed; and loan rate and maturity are likely to be correlated with unobserved consumer heterogeneity). Finally, normality assumptions often used in the estimation of empirical models seem particularly inappropriate in our framework. If one considers the loan terms facing an individual consumer to be the result of a search process (this would, for example, be the case if the consumer chooses the lowest interest rate and the maximum maturity among various offered alternatives), then the corresponding loan variables observed in our data would not be distributed normally, even if the original distribution of interest rates and maturities were.

We employ an estimation approach that deals with each of these issues. We first specify an empirical model which - while not directly derived from a full structural model - is informed by the discussion of the next section. We next estimate this model using two different semiparametric approaches, each of which exhibits different strengths and weaknesses. The results across the two methods are similar. Both approaches rely on the same identification strategy which involves two sets of important assumptions. First, regarding selection, we assume that vehicle stock variables (e.g., number of cars, age of cars in existing car stock, etc.) and population size of the town of residence affect selection into our sample (that is the decisions whether or not to buy a car, and whether or not to finance), but not the size of the loan. Second, regarding the endogeneity of interest rate and maturity, we exploit the tax reform of 1986 that gradually phased out the tax deductibility of consumer credit interest, and the increased durability of cars during our sample period to construct instruments for the interest rate and maturity.

In terms of empirical results, we find that the aggregate demand for loans is highly sensitive to maturity: increasing maturity by one year, increases loan demand by approximately 88.5% according to our estimates. In contrast, we cannot reject the hypothesis that the elasticity of loan demand with respect to interest rate is zero. These estimates look however quite different when we perform the estimation for different subgroups in the population. While, contrary to our expectations, we do not find any evidence that younger consumers are more constrained than older consumers, our results provide strong support for the hypothesis that low-income consumers are substantially more constrained than high-income consumers. In particular, we find that low-income consumers are less

sensitive to interest rates and more responsive to maturity changes. Interestingly, the high-income group is the only one for which we cannot reject the hypothesis of a zero maturity elasticity. However this group constitutes only a small fraction of our sample, about 15% of the observations who finance a car. At the same time this group exhibits high interest rate sensitivity: an interest rate increase of 1% reduces the loan demand of this group by 14% according to our estimates. These results suggest that the high-income group is not liquidity constrained in the sense used in this paper. Given, however, that this group is small, credit constraints appear to have a large effect on borrowing behavior in the aggregate.

The remainder of the paper is organized as follows: In the next section, we discuss more extensively the implications of credit rationing for the interest rate and maturity elasticities of loan demand. We use this discussion to motivate our main identification assumptions. Section 3 presents the empirical model and estimation approach; section 4 describes our data and offers some preliminary descriptive results, and section 5 discusses the results from the estimation of the model. Section 6 concludes.

2. A Theoretical Framework

2.1. The Demand for Loans With and Without Liquidity Constraints

This sub-section discusses the relationship between the presence of liquidity constraints, and the interest rate and maturity elasticities of loan demand. We argue that these elasticities are interesting because their magnitude is informative about the relevance of binding credit constraints, especially if one examines how they vary across different groups in the population, which have a-priori a different likelihood of being liquidity constrained.

The term “liquidity constraints” is used here in two senses, one more stringent than the other. According to the stronger interpretation, a consumer is liquidity constrained if she cannot borrow as much as she would like in order to finance present consumption using resources that would accrue to her in the future. A weaker definition considers the consumer liquidity constrained if the interest rate at which she can borrow is greater than the rate at which she can lend, or, more generally, if the interest rate is increasing in the amount borrowed (Pissarides (1978)). The first definition is a subcase of the second if one considers the interest rate past a certain level of borrowing to be infinite.

The study of the demand for loans and the presence of liquidity constraints requires the formulation of a dynamic model in which consumers allocate resources over time. The standard model in the literature is the life-cycle/ permanent income one. Unfortunately, a version of the life-cycle model that incorporates uncertainty and different assets and liabilities with different interest rates yields a closed form solution for consumption (and therefore loan demand) only under very special circumstances. Nonetheless, we can use the first order condition for the standard life-cycle intertemporal optimization

problem, to obtain some intuition for the demand of loans in the presence of liquidity constraints. This condition can be written as:

$$\lambda_t \geq E_t \left[(1 + r_{t,t+k} + \frac{\partial r_{t,t+k}^i}{\partial A_t} A_t^i) \lambda_{t+k} \right] \quad (2.1)$$

where λ_t is the marginal utility of wealth at time t , $r_{t,t+k}$ is the rate of return between period t and $t+k$, A_t^i are the net holdings of asset i at time t , and E_t denotes the conditional expectation with respect to the information set available in period t . When condition 2.1 holds as an equality, it is the first order condition for an asset for which the consumer is not at a corner. The fact that the left hand side in 2.1 can be larger than the right hand side incorporates our first definition of liquidity constraints, while the term $\frac{\partial r_{t,t+k}^i}{\partial A_t}$ allows for our second definition.

Consider first the loan demand of a consumer who is liquidity constrained according to our first definition: the left hand side of equation 2.1 is strictly larger than the right hand side. This consumer would like to move resources from some future period to the present, so as to consume more in the periods between t and $t+k$; in other words, she is at a kink of her intertemporal budget constraint. In practice, such a consumer is either denied the desired loan amount by the lender, or feels constrained by the payments she will be required to pay in each period during the loan term. If these periodic payments were equal to the user cost of car services (the rental cost of a car), then the loan market would replicate a rental market, and the consumer's intertemporal condition would be undistorted. However, the economic life of a car is typically considerably longer than the duration of a car loan, so that the consumer is effectively required to pay now for a service she will receive later.

In this case, small changes in the interest rate will either not affect her demand for loans, or they will affect it minimally. Instead, an increase in the maturity of the loan, by reducing the monthly payments during the duration of the loan, will effectively allow the consumer to borrow more. These payments will depend to a certain extent on the interest rate (since lower interest rates decrease the payments for a given loan), but they will depend much more on the maximum available maturity. This is because given the typical structure of auto loan repayment, which takes the form of a series of monthly repayments of equal size, the size of each repayment is affected substantially more by the number of periods over which the total payment is spread, than by the interest rate.³

Similar considerations apply to the weaker definition of liquidity constraints, in which case the interest rate depends on the amount borrowed. In such a situation the consumer is on a steep portion of the intertemporal budget constraint and her demand for loans will therefore be less elastic compared to a situation in which the interest rate does not increase with the amount borrowed. If we express

³We should emphasize that the above argument applies to *consumer loans* only, which are typically short-term (3-5 years). It does not necessarily apply to housing loans which have substantially longer maturities, in which case the interest rate effect on the monthly payment is also significant.

the demand for loans as $A = f(r(A), z)$ where the vector z includes, among other things, current and future income, we have that:

$$dA/dr = \frac{f_1(r(A))}{1 - f_1(r(A))r'(A)}$$

where $f_1(r)$ is the derivative of f with respect to its first argument. If $r'(A) < 0$ (the interest rate a decreasing function of the net asset position)⁴, and $f_1(r) > 0$ (present and future consumption are normal goods so that lending (borrowing) is an increasing (decreasing) function of the interest rate), the denominator in the above expression is positive. Hence, the response of the net asset demand to interest rate changes will be smaller in absolute value than $f_1(r)$. But note that $f_1(r)$ would have been the response of demand for asset A to the interest rate, if the interest rate had not been a function of the asset position, in other words if liquidity constraints as defined based on the second, weaker, definition had not existed. In sum, no matter what definition one adopts, the general conclusion is that the loan demand of unconstrained consumers will be more responsive to interest rate changes than the demand of liquidity constrained consumers.

The fact that car loans are loans against a specific durable commodity raises a couple of additional issues. First, it is possible that the presence of liquidity constraints induces a consumer not to buy a car at all. That is, consumption today is reduced by sacrificing the consumption of car services. While the empirical approach we develop in the next section does control for selection into the sample of consumers who bought and financed a car, our inference of liquidity constraints is based on analyzing the behavior of those consumers who did take car loans. Accordingly, we cannot identify liquidity constraints among those who did not buy or finance a car. Second, as has been noted in the literature, in the presence of liquidity constraints, the fact that it is typically easier to borrow against a car than against non-durable consumption distorts the intratemporal allocation between non-durables and cars, making cars effectively cheaper (see Brugiavini and Weber (1992), Chan, Ramey, and Starr (1995), and Alessie, Devereux and Weber (1997)). We do not consider this aspect of the problem, since it does not affect the implications of liquidity constraints for the interest rate and maturity elasticities of loan demand we discussed earlier.

In sum, liquidity constrained consumers will fall into two categories. Some of them will be rationed out of the car market, in which case they will show no sensitivity to either loan interest rate or maturity. Others will buy and finance a car, in which case their loan demand will be highly sensitive to the loan maturity, but less sensitive to the interest rate.

Now consider the case of a consumer who is not liquidity constrained in either of the above senses.

⁴We argue that the case $r'(A) < 0$ is realistic from an empirical point of view, especially when one focuses on *negative* asset positions (that is loans), as we do in this paper. But the assumption $r'(A) < 0$ is also more interesting from a theoretical point of view; if $r' > 0$, individuals could make money by borrowing at low levels of r , and saving at a higher r .

When deciding whether or not to finance a car purchase, such a consumer compares the interest rate she is offered on a car loan with her own after-tax interest rate. If the interest rate on the car loan exceeds the latter, then the consumer will not finance any amount.⁵ If it is lower, she will finance the full amount. Loan demand is in this case highly sensitive to the interest rate: if the auto loan rate increases, some (liquidity unconstrained) consumers will switch from financing to not financing. The loan maturity is here only relevant to the extent that the yield curve is typically upward sloping, so that depending on the maturity of the loan offered, the alternative rate that the auto loan competes with could be lower or higher. However, the yield curve during our sample period was almost flat for most years, while auto loan rates have always been an increasing function of loan maturity. Given this, we would expect little sensitivity to maturity among unconstrained consumers, while there should be high sensitivity to interest rates, at least at the margin. The story is slightly more complicated if the interest rate offered depends on the amount borrowed, as in this case liquidity unconstrained consumers may find it optimal to finance at intermediate levels, but the essence of the argument why loan demand would be highly sensitive to interest rates would remain the same.

In sum, consumers who are not liquidity constrained will also fall into one of two categories: some of them may find it optimal to never finance, so that they will show no sensitivity to either maturity or interest rate; others will find it optimal to finance, in which case their loan demand will exhibit high sensitivity to interest rate changes, but limited (if any at all) sensitivity to maturity changes.

While the above discussion demonstrates why the elasticities of loan demand with respect to interest rate and maturity can give insight into the significance of liquidity constraints, it also illustrates the limitations of this approach. Our inference is based on analyzing the behavior of consumers who *did* obtain a car loan. But as was shown above, both liquidity constrained and liquidity unconstrained consumers may find it optimal not to finance, though for entirely different reasons. Our approach that rests on estimating the elasticities of loan demand based on *observed* loan contracts obviously cannot identify liquidity constraints in this case. Accordingly, our results need to be interpreted with caution. Because the focus on those who financed may understate the overall importance of liquidity constraints, lack of evidence that such constraints exist among consumers who financed does not constitute conclusive evidence that liquidity constraints are immaterial, since such constraints could exist among households who did not finance. On the other hand, evidence that liquidity constraints exist among consumers who financed is more compelling.

The discussion so far has treated interest rate and maturity as exogenous variables. The premise

⁵Even in this case, it is conceivable that some (liquidity unconstrained) consumers decide to finance for other reasons, for example the desire to build up a credit history. Though the overall importance of such motives for car financing is probably limited, it is worth noting that they imply a high interest rate sensitivity of loan demand; an interest rate increase would be equivalent in this case to an increase in the “price” of obtaining a credit history, and some consumers would find it desirable to switch to no-financing.

of exogenous loan terms is unreasonable, for a variety of reasons. We take up the issue of their endogeneity next.

2.2. Supply Side and Identification

Assuming that all consumers face a common interest rate and maturity that are exogenously predetermined is invalid for two reasons. First, a distinguishing feature of the market for car loans is that loan contract terms vary significantly across finance sources (banks, dealers, credit unions, finance companies, etc.), and borrowers have - to a certain extent - the option to choose among them. If interest rate and maturity are the result of a search process, they should clearly be treated as endogenous, as they are likely correlated with unobserved consumer characteristics. Moreover, the existence of multiple loan sources may be itself indicative of credit rationing that arises as an equilibrium phenomenon in the presence of asymmetric information. This point is exemplified in Jaffee and Russell's (1976) model of consumer loan markets.

A second feature of car loans is that even *conditional* on the loan source, consumers face different loan "packets" in which loan size, interest rate and maturity are bundled together. For example, banks offered during our sample period loan packages that had four dimensions: loan size; new or used car; interest rate; and maturity. Larger loans, used cars, and longer maturities are typically associated with higher interest rates. As Jaffee and Stiglitz (1990) show, the availability of such multiple contracts can be thought of as a response of credit institutions to the presence of either adverse selection or moral hazard. Furthermore, car salesmen may tailor financing packages to individual households that appear to be credit constrained, so as to make the car purchase affordable; in such instances, the loan terms are clearly endogenous.

To address the endogeneity of the loan terms we exploit two developments of the last two decades to construct instruments for the interest rate and maturity.⁶

The first event is the tax reform of 1986 that phased out the deductibility of interest on consumer loans over a five year period,⁷ thus changing the after-tax interest rate consumers faced on their loans. To understand the role of the tax reform, note that a borrower's after tax nominal rate in year t is given by the formula:

$$r_t = r_p * (1 - t_f * \alpha_t * I - t_s * \alpha_t * I + t_f * t_s * \alpha_t * I)$$

⁶To the extent that shifts in interest rates and maturities over time were brought about by exogenous events, the time variation can be legitimately used as a source of identification. But if loan terms responded to changes in unobserved factors that also had an independent direct effect on loan demand, our estimates will suffer from the usual simultaneity bias. Though we make an attempt to minimize such bias by including variables in the estimation that proxy for current macroeconomic events as well as expectations about future developments (e.g., regional unemployment rates, regional per capita income), such proxies are imperfect.

⁷A detailed description of the reform can be found in Maki (2001).

The term r_p represents the pre-tax consumer loan rate. The second term in the parenthesis is the so-called tax price of the debt. The variables t_f and t_s represent the marginal federal and state tax rates respectively. I denotes a dummy that takes the value of 1 if the consumer itemizes deductions on his/her tax return. Finally, α_t denotes the proportion of consumer interest that is deductible from income. The subscript t captures the fact that the phaseout of interest deductibility was gradual. In particular, α was 1 prior to 1987, equal to 0.65 in 1987, 0.40 in 1988, 0.20 in 1989, 0.10 in 1990, and 0 after 1990. The very last term in the expression for the tax price of the debt $t_f * t_s * \alpha_t * I$, is needed because state taxes are deductible from federal taxes.

The tax reform of 1986 has two implications for our estimation. First, the nominal interest rate on car loans is modified according to the formula above to take the varying tax deductibility of consumer interest into account. Second, the formula for the after-tax rate suggests valid instruments for the purpose of identifying the parameters of the loan demand equation. Using the tax reform of 1986 to construct instruments presents two advantages: First, because the phaseout of deductibility was gradual, there is time variation in the effects of the reform. Second, because the effects depended on whether or not the consumer was an itemizer, the instruments we construct also exhibit a limited degree of cross-sectional variation. The details of the instrument construction are presented in the Appendix.

The discussion above suggests that one could use the tax price of the debt facing each individual (which we can construct after matching the data available in the Consumer Expenditure Survey with tax rate data), or various components of that variable (e.g., the fraction α interacted with the itemizer dummy I , the marginal federal tax rate t_f interacted with α and the itemizer dummy, the marginal state tax rate t_s interacted with α and the itemizer dummy, etc.) as instruments for the loan terms. The important identification assumption underlying this procedure would then be that the tax reform affected loan demand *only through* its effect on the after-tax interest rate facing car buyers, and did not have any direct effects on loan demand. However, the use of individual marginal tax rates and the itemizer dummy might well be correlated with the unobserved components of the loan demand equation. This would, for example, be the case if a consumer who took a large loan in order to finance a car purchase, decided to start itemizing. Similarly, marginal tax rates depends on capital income (see Maki (2001) for details). To the extent that capital income is endogenous to the borrowing decision, marginal tax rates should be considered endogenous too. As we discuss in the Appendix, these issues are not likely to be relevant in our data because of the way we construct the instruments (we use tax information from the year prior to the car purchase). Nevertheless, to be sure that our results are not driven by spurious correlation between the components of the tax price of the debt and loan demand, we compute averages of the above components, as well as the average of the tax price of the debt, by year and region, and use those averages as instruments.

While the tax reform provides us with powerful instruments for capturing exogenous variation in after-tax interest rates, we also need to instrument for the steady increase in maximum institutional maturities. To this end, we exploit an important development in the car market: the increased durability of cars. This development has been documented several times in the past (see various studies by Polk Corp. and Consumer Reports and Hendel and Lizzeri (2002)). Increased durability (i.e., slower depreciation) implies that lenders can offer longer maturities as cars retain their collateral value for a longer period of time. This argument suggests that one can use depreciation rates for cars over our sample period as instruments for maturity extensions. To construct these instruments we use the information on the vehicle stock provided by each household in the Consumer Expenditure Survey. For each vehicle owned by the household, the household provides information on the model year and the estimated value of the car. Based on this information, we constructed depreciation rates going *forward*. That is, for cars with model year 1990, we constructed a 1-year depreciation rate based on the values reported by households in 1990 and the values reported in 1991; a 2-year depreciation rate based on the values reported in 1990 and 1992, and so on. The advantage of this approach is that, since our sample ends in 1995, and we have access to Consumer Expenditure Survey data up to 1999, we could go as far as constructing 4-year depreciation rates. The information is averaged across all car models, to construct aggregate measures of durability. Even though we could have exploited additional information on car models to construct depreciation rates by car type (e.g., luxury cars, small cars, etc.) we chose not to do so, since the choice of the car type itself is plausibly correlated with loan demand. We should note that the depreciation rates constructed by this method reflect not only physical, but also economic depreciation of the durable, which is the relevant measure if one wants to capture the effect of durability on the collateralizable value of the car.

The reason depreciation rates are promising instruments for the purpose of identifying the loan demand parameters is illustrated in Figure 1 that plots the average car depreciation rate between the 2nd and 4th year (that is the depreciation rate between the time a car is 2 and 4 years old) against time. It is striking that this depreciation rate declines over time, just as the institutional maximum maturity norm increases. The patterns for other depreciation rates are similar. We focus on the 2nd-4th year depreciation rate because during our sample period maturities were extended first from two to four, and later to five years⁸; accordingly, we expect a particularly close link between the 2nd-4th year depreciation rate and the maturity extension. To investigate how “good” depreciation rates are as instruments, we also ran various partial regressions, in which the average maturity was successively regressed against the 0-1, 0-2, 1-4, and 2-4 year depreciation rates. In each of these cases, the coefficient on the depreciation rate was, consistent with the above argument on the effect of

⁸Unfortunately, we cannot construct the 2nd to 5th year depreciation rate without giving up the observations for 1995, since the last year for which we have the Consumer Expenditure Survey data is 1999.

increased durability on maturity, negative and highly significant, while the R-squares of the regressions were consistently above 0.30. The best fit was achieved by regressing the average maturity on the 2nd-4th year depreciation rate (R-square: 0.62). A regression of average maturity against all of the aforementioned depreciation rates, results in an R-square of 0.83. This suggests that depreciation rates can explain a substantial proportion of the time variation of maturity. Because interest rates and maturities are set jointly by lenders, we use the depreciation rates discussed above, as well as the variables related to the tax reform of 1986, as instruments for both interest rate and maturity.

Before moving to the discussion of the empirical specification, we need to mention an important caveat regarding the use of increased durability as an instrument for maturity. While from both a theoretical and an empirical point of view, durability seems to be an important determinant of maturity, it is possible that in the presence of liquidity constraints increased durability also has an independent role in the loan equation. This would be for example the case if consumers, anticipating lower repair and maintenance expenditures, would be able to afford higher monthly payments and borrowed more. Our identification strategy is based on the premise that this is not a quantitatively important channel in practice.

3. Empirical specification and econometric issues

To recap, the main goal of our paper is to estimate the elasticities of automobile loan demand with respect to interest rate and maturity, and examine how these elasticities differ across different groups in the population. The equation we want to estimate can be written as follows:

$$l^* = \ln(L^*) = x\theta_l + f(r, m) + \varepsilon \quad (3.1)$$

where L^* is defined as the desired loan amount. The dependent variable is expressed in logarithmic form to take into account the fact that loans cannot be negative. x is a vector of variables that capture demographic and life cycle effects, as well as macroeconomic effects on the loan amount. The variables r and m are the interest rate and the maturity of the loan, respectively, and are both considered endogenous. Finally, ε is an unobserved error term.

The exact functional form of the structural loan demand equation will be a complex expression, depending on both the functional form of the utility function and the structure of the monthly payments. Given that we do not want to make specific functional form assumptions about preferences or payment structure, we view (3.1) as an approximation to the true functional form.⁹

⁹In principle, the demand equation will be a nonlinear function of (x, r, m, ε) and it will contain interaction terms between x and (r, m) . Estimation of a specification with such interaction terms is challenging given the selection and endogeneity issues we face. In particular, none of the semiparametric methods we employ allows for such interaction terms. We do address this issue indirectly by estimating separate demand equations by different age and income groups.

When trying to estimate the loan demand equation (3.1), we are faced with a number of sample selection issues. First, loans are observed only for those households who decide to buy *and* finance a car. Because the decision to buy is most likely affected by the availability and cost of credit, one has to correct for the sample selection bias induced by the nature of our data. The two decisions, “buy vs. not buy”, and “finance vs. not finance”, can be either treated separately, or collapsed into one estimating equation in which the dependent variable is 1 if the individual buys and finances, and zero otherwise. As we could not think of any exclusion restrictions that would allow us to identify the coefficients of two separate equations (that is we could not find variables that would affect the decision to buy but not to finance) we chose the second approach. Second, the desired loan amount may equal or even exceed the value of the car. Lending institutions however will only finance up to 100% of the car value.

We model these selection issues as an ordered discrete choice model, where for each household in the sample we observe a discrete variable d which takes three possible values, say 0, 1 and 2 corresponding to “not buy/finance”, “finance less than 100% of the car value” and “finance 100% of the car value”, respectively. Let $(-\infty, a_1)$, (a_1, a_2) and (a_2, ∞) be a partition of the real line where a_1 and a_2 are unknown thresholds. The observed value of d (0, 1, 2) depends on whether the unobserved desired level of the loan is below a_1 , between a_1 and a_2 , or above a_2 . In principle, these thresholds may be endogenous and varying across households. In particular, the lower threshold may be affected by variables that enter the decision of whether to buy a car, but not the amount of the loan, for example whether the household owns a car or not. The upper threshold is the value of the car which is itself a choice variable and may depend on unobservable tastes. A reduced form of the model may be written as:

$$d = \begin{cases} 0 & \text{if } Z\beta_d + u_d < \tilde{a}_1 \\ 1 & \text{if } \tilde{a}_1 < Z\beta_d + u_d < \tilde{a}_2 \\ 2 & \text{if } Z\beta_d + u_d > \tilde{a}_2 \end{cases} \quad (3.2)$$

where \tilde{a}_1 and \tilde{a}_2 are now fixed constants, identical for all households. We include in the vector Z the variables in x as well as *all* observable variables that enter the equations for the interest rate and maturity. In addition, the vector Z includes controls that may affect the decision to purchase a car, but not necessarily the loan amount. These include the population size of the residence town, and several vehicle stock variables (number of cars in the stock, dummy for not currently owning a car, average age of the stock, age of the oldest car, and age of the newest car).¹⁰

The modelling approach above lumps together households who did not buy a car, households who

¹⁰A potential problem with the use of vehicle stock variables is that these are essentially functions of lagged dependent variables. To the extent that there is serial correlation in the error term, these variables are then not exogenous. However, we have no way of addressing the potential endogeneity of the stock variables.

did buy a car but paid cash, and households who bought a car but used leasing instead of financing.¹¹ We think of (3.2) as a reduced form model for the joint decision to buy *and* finance *and* of the extent of financing. This decision clearly depends on interest rate and maturity, as loan terms influence not only the amount a consumer borrows, but potentially also the decision whether or not buy a car, and finance the purchase. However, since loan terms are endogenous in our framework, we do not include them directly in the reduced form equation (3.2). We do include however, as pointed out above, *all* exogenous variables that affect interest rate and maturity.

Finally, as discussed above, interest rates and maturities are endogenous variables described by:

$$r^* = x\theta_r + W\delta_r + u_3 = X\beta_r + u_r \quad (3.3)$$

$$m^* = x\theta_m + W\delta_m + u_4 = X\beta_m + u_m \quad (3.4)$$

where $X \equiv \begin{bmatrix} x & W \end{bmatrix}$, $\beta_r \equiv \begin{bmatrix} \theta_r & \delta_r \end{bmatrix}'$ and $\beta_m \equiv \begin{bmatrix} \theta_m & \delta_m \end{bmatrix}'$. The vector x is the same as in the l^* equation since we could not think of any reason why a subset should be excluded from the r^* or m^* equations. The important identification assumption in our approach is that there exists a set of variables W , that affect interest rate and maturity, but which can be excluded from the loan demand equation. These variables were discussed in the previous section, and their construction is described in detail in the Appendix.

The variables r^* and m^* are observed only if the consumer actually takes a loan (i.e. for those that have $d = 1$ or 2 , that is, if $d > 0$). Let $1\{A\}$ denote the indicator function that is equal to 1 if A occurs and is zero otherwise. Then, the observed interest rate, r , and observed maturity m are given by:

$$r = 1\{d > 0\} \times r^* \quad (3.5)$$

$$m = 1\{d > 0\} \times m^* \quad (3.6)$$

¹¹We do not attempt a separate analysis of leasing in this paper, mainly because we do not have the data to do so. Leasing is almost non-existent in the early years of our sample; in the later years (mid-1990's), leasing becomes quantitatively more important in the U.S. market, however only ca. 2% of the car buyers in the Consumer Expenditure Survey report to have leased their cars. In addition, leasing is a complex phenomenon that requires a full empirical investigation in its own right, especially since the connection between leasing and liquidity constraints is not readily apparent. The popular perception is that leasing relaxes liquidity constraints, since it involves lower cash flows; the empirical implication of this statement is that leasing would be preferred by lower income households. However, Hendel and Lizzeri (2002) develop a theoretical model of leasing under adverse selection, in which leasing is shown to be preferred by higher income households. The simple intuition for this result is that in a leasing contract the price of the option of keeping the car is higher compared to the price of the same option in a selling contract. Car buyers with high valuations (income) do not value this option very much though, because they are not likely to keep the car anyway (they buy a new car every few years). Such buyers prefer therefore leasing. In contrast, buyers with lower evaluations (income) value the option of keeping the car more, and prefer a selling contract. This theoretical prediction is in fact supported by the empirical findings in Aizcorbe and Starr-McCluer (1997). This suggests that leasing may or may not be indicative of liquidity constraints, just as not-buying-a-car may or may not be a sign of such constraints.

Our empirical model thus consists of equations (3.1)-(3.6). It is possible to estimate the model by maximum likelihood. To illustrate the approach, suppose we assume a first-degree polynomial in r and m for $f(r, m)$, so that

$$l^* = \ln(L^*) = x\theta_l + \gamma_1 r + \gamma_2 m + \varepsilon \quad (3.7)$$

Then equations (3.5)-(3.6) can be substituted into equation (3.7) to obtain the following reduced form loan demand equation:

$$l^* = X\beta_l + u_l \quad (3.8)$$

where $\beta_l \equiv \left[\theta_l + \gamma_1 \theta_r + \gamma_2 \theta_m, \quad \gamma_1 \delta_r + \gamma_2 \delta_m \right]$ and $u_l \equiv \varepsilon + \gamma_1 u_r + \gamma_2 u_m$. Equations (3.2)-(3.8) constitute a reduced form system with unknown parameters of interest β_d , β_l , β_r , and β_m . Having obtained the ML estimates, we may recover the structural parameters of interest, γ_1 and γ_2 , by classical minimum distance. The implementation of ML requires to specify a distribution for the structural error vector $(\varepsilon, u_d, u_r, u_m)$, say a joint normal (which would also imply a joint normal distribution for the reduced form error vector, (u_l, u_d, u_r, u_m)). We believe however that in our context normality may be a particularly bad assumption. On one hand, it can be argued that, to the extent that the loan terms result from some search process by the agent, normality of the outcome variable of interest such as the terms of a loan (interest rate and maturity) is not justified. On the other hand, the estimates run the risk of being inconsistent if the functional form of the distribution is misspecified, which is a common problem in all parametric approaches. Furthermore, the joint normality imposes independence between error terms and covariates which is often undesirable. Finally, in our particular situation, assuming a more complicated functional form, say quadratic, for the function f via which the endogenous variables, r and m , enter the main equation of interest, would complicate estimation considerably, since it would imply conditional heteroskedasticity in the reduced form error of the loan demand equation, as well as a non-normal distribution, since it would involve a quadratic function of the two normal errors, u_r and u_m .

In the next section we will describe two estimation approaches that relax the assumption of a known joint distribution for the error vector. Because they are both multi-step procedures that rely on a first step estimation of the selection equation (3.2), and given that we are not interested in that equation, we will assume that its error term, u_d , is normally distributed. In other words we will assume that d follows an ordered probit model.¹² As we discuss next, it is possible to relax the normality assumption of a parametric approach for estimating the three main equations (3.1), (3.3) and (3.4) following either one of two estimation ideas from the semiparametric literature of sample selection models. The first one eliminates sample selection via a pairwise differencing scheme while the other

¹²See Powell (1994) for references on semiparametric estimation of discrete choice models.

one estimates the selection effect. The two methods also differ in how they deal with endogenous regressors although they are similar in the form of conditional heteroskedasticity that they permit.

3.1. Estimation: Two semiparametric approaches

The first approach follows Powell (2001)¹³. We start by noting that the equations describing the observed interest rate and maturity (equations (3.2)-(3.5) and (3.2)-(3.6)), respectively, constitute two standard sample selection models (Type 2 Tobit models, in the terminology of Amemiya (1985)). In other words, r^* and m^* are observed if $d > 0$, or equivalently if $u_d > \tilde{a}_1 - Z\beta_d$. Therefore,

$$E(r|X, Z, d > 0) = X\beta_r + E(u_r|X, Z, u_d > \tilde{a}_1 - Z\beta_d) \quad (3.9)$$

$$E(m|X, Z, d > 0) = X\beta_m + E(u_m|X, Z, u_d > \tilde{a}_1 - Z\beta_d) \quad (3.10)$$

The problem in estimating β_r and β_m lies in the fact that if the unobservables in the two equations, u_r and u_m , are correlated with the error term in the selection equation, u_d , then their conditional means (the ‘selection effects’ or ‘selection correction terms’ or ‘sample selection bias terms’) are not zero given the selection. In principle, these selection effects are functions of X , Z and the scalar ‘selection index’, $Z\beta_d$. Powell (2001) assumes that sampling across households is i.i.d. and that the conditional distribution of the errors in (3.3)-(3.4), u_r and u_m , given the covariates X and Z , only depends on the scalar index determining selection, $Z\beta_d$, albeit in an unknown manner. That is,

$$E(u_r|X, Z, u_d > \tilde{a}_1 - Z\beta_d) = \lambda_r(Z\beta_d)$$

$$E(u_m|X, Z, u_d > \tilde{a}_1 - Z\beta_d) = \lambda_m(Z\beta_d)$$

where $\lambda_r()$ and $\lambda_m()$ are unknown functions satisfying certain smoothness conditions. The idea then is that, for two households i and j that have approximately equal selection indices, i.e. $Z_i\beta_d \approx Z_j\beta_d$, the magnitudes of the selection bias terms are also approximately equal. Hence pairwise differencing eliminates the sample selection bias. As β_d is not known, it is estimated from the selection equation in a first step. In the second step the parameters of the continuous outcome equation are estimated by weighted least squares on the pairwise differenced selected sample, where the weight per pair varies inversely with the magnitude of the difference in the estimated selection indices for the pair. Formally, we estimate β_r and β_m by minimizing:

$$\sum_{i < j} 1\{d_i > 0\} 1\{d_j > 0\} K\left(\frac{(Z_i - Z_j)\hat{\beta}_d}{h_n}\right) [(r_i - r_j) - (X_i - X_j)\beta_r]^2$$

¹³The same approach is taken in Honore and Powell (2002) for the estimation of nonlinear models.

and

$$\sum_{i < j} 1 \{d_i > 0\} 1 \{d_j > 0\} K \left(\frac{(Z_i - Z_j) \hat{\beta}_m}{h_n} \right) [(m_i - m_j) - (X_i - X_j) \beta_m]^2$$

Here $K(\cdot)$ is a kernel “density” function, and h_n is a bandwidth constant which is required to converge to 0 as sample size increases. The effect of this weighting scheme is that, asymptotically, only households with the same selection indices (for which the selection biases exactly offset each other) contribute to the estimation.

We now turn to the estimation of the loan demand equation (3.1). If we assume that the endogenous variables, r and m , enter linearly in that equation, as in (3.7), we may proceed by replacing the structural equations (3.3)-(3.4) to obtain the reduced form equation (3.8). Assuming that the unobservables u_r , u_m , and ε depend on the covariates X and Z only through the selection index $Z\beta_d$, the error term in this reduced form equation, $u_l \equiv \varepsilon + \gamma_1 u_r + \gamma_2 u_m$, also shares the same property. The estimation idea is then the same as above. Note that l^* is only observed for $d = 1$, i.e. for $\tilde{a}_1 - Z\beta_d < u_d < \tilde{a}_2 - Z\beta_d$. Thus, for a household i the magnitude of the selection term depends only on the magnitude of the scalar index $Z_i\beta_d$, and for two households that have approximately equal selection indices, i.e. $Z_i\beta_d \approx Z_j\beta_d$, the magnitude of the selection bias terms is also approximately equal. Hence pairwise differencing eliminates the sample selection bias. We may therefore estimate β_l by minimizing:

$$\sum_{i < j} 1 \{d_i = 1\} 1 \{d_j = 1\} K \left(\frac{(Z_i - Z_j) \hat{\beta}_d}{h_n} \right) [(l_i - l_j) - (X_i - X_j) \beta_l]^2$$

Having obtained $\hat{\beta}_d$, $\hat{\beta}_l$, $\hat{\beta}_r$, and $\hat{\beta}_m$ it is possible to obtain an analytic form for the asymptotic variance-covariance matrix of the estimators and proceed to apply minimum distance in order to estimate γ_1 and γ_2 . Let π be the vector that stacks the reduced form parameters and $\hat{\pi}$ its estimate, $h(\cdot)$ the function that maps the structural parameters into the parameters of the reduced form equations, and \hat{V} a consistent estimator of the variance-covariance matrix of the reduced form parameters. The structural parameters γ_1 and γ_2 are estimated by minimizing the quadratic form:

$$Q = (\hat{\pi} - h(\theta_l, \gamma_1, \gamma_2, \theta_r, \delta_r, \theta_m, \delta_m))' \hat{V}^{-1} (\hat{\pi} - h(\theta_l, \gamma_1, \gamma_2, \theta_r, \delta_r, \theta_m, \delta_m))$$

The estimator may be shown to be consistent and asymptotically normal with asymptotic variance equal to $\hat{H}'\hat{V}^{-1}\hat{H}$ where \hat{H} is an estimate of the gradient of h evaluated at the structural estimates. Denoting by $\theta (= \theta_l, \gamma_1, \gamma_2, \theta_r, \delta_r, \theta_m, \delta_m)$ the vector of structural parameters, a test of overidentifying restrictions is immediately available when $\dim(\pi) > \dim(\theta)$, exploiting the fact that under the null, $\pi_0 = h(\theta_0)$, the quadratic form above is distributed as $\chi^2(\dim(\pi) - \dim(\theta))$.

The method described above, despite its apparent advantages over traditional MLE, suffers from certain setbacks. The most important one is that it practically rules out conditional heteroskedasticity in the unobservables. In particular, having a higher than first order polynomial in r and m in the demand equation rules out the basic assumption on the distribution of the error in the reduced form loan demand equation, namely that it only depends on the selection index. As we would like in principle to be more flexible in the specification of the functional form via which r and m enter that equation, i.e. the form of f , we will next consider an alternative method for estimating sample selection models with endogenous regressors, which allows us to do so at the cost however of restricting the nature of endogeneity.

In order to apply the methodology described above, we need to choose the kernel function K and the bandwidth h_n . Although the theory requires that K is a 4th order bias-reducing kernel with bounded support, in the application we found it convenient to use the standard normal density function which is only a 2nd order kernel with unbounded support.¹⁴ The choice of the bandwidth is typically considered more crucial than the choice of the kernel function. We chose the bandwidth to be data dependent, of the form $h_n = c * \hat{\sigma} * n^{-1/7}$ where c is a constant, $\hat{\sigma}$ is the estimated standard deviation of the (estimated) selection index $Z\hat{\beta}_d$ and n is the sample size (see equation 6.7.1 in Powell (2001)). Although we experimented with the choice of c we did not find much sensitivity of the estimates to this choice and we only report the results for $c = 0.1$. Finally, in order to compute the minimum distance estimates of γ_1 and γ_2 we need to estimate the covariance matrix of the joint asymptotic distribution of $\hat{\beta}_d, \hat{\beta}_l, \hat{\beta}_r,$ and $\hat{\beta}_m$. Although it is possible to derive the latter and estimate the different components as in Powell (2001), in the application we chose to bootstrap the variance covariance matrix. The bootstrap standard errors are in general consistent as the number of bootstrap repetitions increases to infinity. In the application, the number of bootstraps was set equal to 100.

The second approach follows Das, Newey and Vella (2001).¹⁵ Consider first, a standard sample selection model such as (3.2)-(3.5) and (3.2)-(3.6). Here the covariates of the main equations (3.5) and (3.6) are assumed exogenous. The basic assumption is that the conditional mean of the error in these equations given the selection and all exogenous covariates depends only on the selection probability, or propensity score, $p \equiv \Pr(d > 0) = \Pr(u_d > \tilde{a}_1 - Z\beta_d)$, so that

$$E(r|X, Z, d > 0) = X\beta_r + E(u_r|X, Z, u_d > \tilde{a}_1 - Z\beta_d) = X\beta_r + \lambda_r(p) \quad (3.11)$$

$$E(m|X, Z, d > 0) = X\beta_m + E(u_m|X, Z, u_d > \tilde{a}_1 - Z\beta_d) = X\beta_m + \lambda_m(p) \quad (3.12)$$

¹⁴Using a quartic kernel of the form

$$K(v) = \left(3/(4 * \sqrt{5}) - 3/(20 * \sqrt{5}) * v^2\right) * 1 \left\{|v| \leq \sqrt{5}\right\}$$

hardly affected the estimates.

¹⁵We are grateful to Frank Vella for most useful conversations.

where $\lambda_r()$ and $\lambda_m()$ are unknown functions of the scalar propensity score p . It should be noted that the method allows, in principle, the functional form via which X enters the conditional means to be nonparametric. The idea then is to nonparametrically estimate the selection correction terms in (3.11) and (3.12), say via a series approximation, instead of differencing them out. Because p is unknown it needs to be estimated first, say by \hat{p} . Estimation of each of the main equations then proceeds by least squares regression of each of the dependent variables on the exogenous regressors and a polynomial in \hat{p} using only observations with $d > 0$.

In the presence of endogenous regressors as in the loan demand equation (3.1), the crucial assumption in Das, Newey and Vella is that the conditional mean of the error in that equation, ε , given the selection ($d = 1$) and given the (exogenous and endogenous) covariates of the model, x , Z , r , and m , is only a function of the propensity score, say $p_1 \equiv \Pr(d = 1|Z)$, and the unobservable components of the endogenous variables, u_r and u_m ,¹⁶ which implies that

$$E(l|x, Z, r, m, d = 1) = x\theta_l + f(r, m) + \lambda_l(u_r, u_m, p_1) \quad (3.13)$$

Note that we do not restrict the functional form of $f()$ via which the endogenous variables enter the loan equation to be linear in r and m , as we have to do in Powell's approach. In fact, f need not be parametrically specified. As mentioned also above, in principle it is possible to allow the functional form via which x enters the conditional mean also to be nonparametric. The selection correction function $\lambda_l()$ is again unknown but may be approximated with a polynomial in its three scalar arguments. The latter, although not observed may be consistently estimated: u_r and u_m by the residuals from the regressions of r and m on X and a polynomial in \hat{p} using only those households that have bought and financed a car; and p_1 as the predicted probabilities for observing $d = 1$. The unknowns of the model (3.13), θ_l , $f()$ and $\lambda_l()$ may be estimated by least squares of l on x , a polynomial in r and m , and a polynomial in \hat{u}_r , \hat{u}_m , and \hat{p}_1 .

The methodology described above is implemented in the next section by imposing a linear structure in the exogenous variables of the main equations (X and x) in (3.11), (3.12), and (3.13). The selection correction terms are approximated via either a second or a third order polynomial in their arguments. Finally, we specify both a linear and a cubic polynomial for $f()$.¹⁷ Although it is possible to consistently estimate the asymptotic covariance matrix, in the application we bootstrap the standard errors of our estimates with the number of bootstraps set equal to 100. Increasing the number of bootstrap repetitions up to 500 did not alter the results.

¹⁶This would be implied if the errors were independent of all the exogenous variables.

¹⁷When we specify a linear polynomial in r and m , then we approximate the selection correction term λ_l with a quadratic polynomial in its arguments, while we specify the latter to be a cubic polynomial when a third order degree polynomial is specified for f .

Compared to Powell’s method, the Das, Newey, and Vella (2001) approach has the advantage that it allows for higher order terms in r and m in the loan demand equation, at the cost however of restricting the nature of endogeneity of r and m . In particular it needs to assume that the errors in the main equation of interest (3.1) are independent of the endogenous variables conditional on the unobservable components of the latter. An additional advantage is that it is easily implemented as it only involves applying least squares once the degree of approximation of the unknown functions f and λ is decided. The approach however illustrates a common drawback in all control function approaches in dealing with endogeneity in nonlinear simultaneous equations models, namely that, in contrast to linear simultaneous equations, a structural interpretation is given to the errors in the endogenous variable equations as they are included as regressors in the equation of interest.¹⁸

4. Data Description

4.1. The CES Data

The data used in the estimation are provided by the Consumer Expenditure Survey (CES), 1984-95. The CES is collected by the Bureau of Labor Statistics to compute the Consumer Price Index. It is a rotating panel in which each household is interviewed four consecutive times over a one year period. Each quarter 25% of the sample are replaced by new households.

The data provided by the CES include an extensive number of socioeconomic characteristics and information on the vehicle stock holdings of each household each quarter. In particular, for each vehicle the household owns, the BLS collects data on the purchase date and source, various vehicle characteristics, including whether the car was purchased as new or used, the purchase price, the trade-in allowance, and detailed information on the financing of the purchase. The latter includes the source of financing (dealer, bank, credit union, other financial institution, or other private source), the downpayment, the amount of the principal, the size of the monthly payments, the maturity of the loan, and the effective interest rate, computed by the direct ratio formula financial institutions use. A household enters our sample as one that bought a car only if the car purchase occurred during the interview period, or in the three months prior to the first interview. In principle we have information about the finance of the car whenever the purchase has occurred in the 12 months preceding the interview, *provided that the loan has not been fully repaid*. The reason we do not use the information on cars purchased between 12 and 3 months before the first interview is that in this case we would miss all households that took a loan with very short maturity, and repaid it before the first interview.

Our estimation approach uses three loan variables: the real interest rate, the maturity and the

¹⁸We thank a referee for pointing this out. See Blundell and Powell (2000) for a comprehensive discussion of methods that allow for continuous endogenous regressors in semiparametric and nonparametric models.

loan size. The maturity is directly given by the CES. To compute the real interest rate, we subtract the inflation rate in the consumer price index from the effective interest rate. For those households who bought more than one vehicle during the interview period we use the total auto debt taken by the household during the year as a measure of the loan size; the interest rate and maturity are computed in this case as weighted averages of the interest rates and maturities referring to the individual loans, with the loan amounts used as weights. This way, each household appears in our sample only once.

Tables 1 and 2 provide some descriptive statistics and Table 3 contains a list with the acronyms of the variables used in the estimation. From the 11666 households who bought at least one car during our sample period, 46% took auto loans; the average finance share for these households is 0.78. Among those who financed, approximately 18% financed 100% of the car value, while 33% financed more than 90% of the car price. These numbers suggest that a substantial portion of the households who take car loans may be at a corner of the IBC. The most popular finance source appears to be the banking sector with 44%, followed by dealers with 22%, and credit unions with 17%. The means of the interest rate and maturity are consistent with common wisdom, as the average interest rate for new cars is slightly lower than the one for used cars, while the opposite holds for the average maturity. As indicated by Table 2, both interest rate and maturity exhibit substantial variation in our sample (note that the interest rates are negative for some observations, because of inflation). But while most of the interest rate variation comes from the cross-sectional dimension of our data, maturities also exhibit substantial time series variation. Interest rates exhibit a slight downward trend during our sample period; maturities rise substantially from approximately 45 months for new, and 28 months for used cars at the beginning of 1984, to 55 (38 for used) months towards the end of the sample.

4.2. Preliminary Data Analysis

To summarize the main features of our data set, we related auto loan variables (such as interest rates, maturities, sources of financing, etc.) to various socioeconomic characteristics. This preliminary data analysis is purely descriptive and does not correct either for sample selection bias, or for the potential endogeneity of some of the variables appearing on the right hand side. Some interesting patterns, however, are evident. We use some of these facts to justify the modelling choices we make in the following section in which we estimate the structural parameters of the demand for loans.

To draw the profile of the car buyers who finance their purchases, we started by estimating probit equations relating the existence of auto loans to a set of demographics. But the results (omitted here for brevity), while intuitive, are not informative regarding the question of credit constraints. This is hardly surprising; the absence of auto financing may indicate the presence of credit constraints, or alternatively, reflect the car buyer's ability to pay in cash.

A perhaps more informative exercise is to characterize the individuals who have credit contracts

with long maturities and the individuals with finance shares higher than 90% of the car value. Note that two of the empirical implications of the theoretical discussion were that liquidity constrained consumers would prefer longer maturities, and that some liquidity constrained consumers will be financing the entire value of the car purchased. The maximum maturity term was 60 months for the early years of our sample; in later years the maximum was increased to 72 months. A few observations in our sample have loans with maturities in excess of 100 months; such cases are, however, very rare and probably reflect loans obtained from special sources. In Table 4 we report the results from the estimation of probit equations, in which the dependent variable is 1 for observations taking loans with maturities greater than or equal to 60 months. While the sign of some of the coefficients is hard to interpret (e.g. education dummies, income), other parameter estimates are consistent with the presence of credit constraints; for example, young households are more likely to finance at long maturities. Note also that long maturities seem to be used more by households financing through dealers.

Table 5 characterizes the households with large finance shares (in the reported probits the dependent variable is 1 if the household financed more than 90% of the car value). The results here are very intuitive. Households with large finance shares are young, have little education, low income and low financial assets. To the extent that one interprets large finance shares as an indicator of credit constraints, one can use these results to identify consumer groups that are likely to be liquidity constrained. We exploit this idea in the next section, where we split our sample into various subgroups according to the criteria of age and income, and test for the existence of liquidity constraints separately in each case.

As mentioned in the previous section, our approach for identifying liquidity constraints exploits exogenous variation in interest rates and maturities. Before we discuss the estimation of the structural model, it is therefore useful to take a look at the main determinants of the variability (both cross sectional and over time) of interest rate and maturity. To this end, we regressed these variables on various household characteristics, credit source dummies, and dummies for new vs. used car. We conducted this exercise without correcting for sample selection, or potential endogeneity of some of the regressors, but experimented with a variety of specifications. Two patterns clearly emerge out of these regressions: First, a large fraction of both interest rate and maturity variability is accounted for by the new vs. used dummy. Second, both interest rates and maturities are highly correlated with credit source dummies (e.g., credit unions are associated with lower interest rates and shorter maturities). The choice of credit source is itself highly correlated with socioeconomic characteristics; estimation of a simple multinomial logit on the choice of finance source points to age, education, race and gender as the main determinants of this choice.

We use these descriptive results to justify two important choices we make in the specification of

the empirical model. First, we do not use the source of financing as an instrument, even though this variable captures a substantial amount of the cross sectional variability of interest rates. The reason is that it is likely to be correlated with unobserved heterogeneity. Perhaps a preferable treatment of the credit source would explicitly model the decision where to obtain credit from; this, however, would complicate our model considerably. Second, for the same reason, that is the likely endogeneity of such a variable, we do not use the choice between new vs. used as an instrument. Instead we rely on the identification assumptions outlined in subsection 2.2.

5. Results from the Estimation of the Empirical Model

In section 3 we laid out two estimation approaches, both of which involve estimation of the selection equation, and the reduced form equations for the interest rate and maturity as a first step. To recap, these are specified as follows. The vector Z of variables that enter the selection equation includes: *all* exogenous variables included in the vectors x and W that enter the reduced form interest rate and maturity equations (see below for details); quarter and regional dummies; the population size of the town of residence; and the following vehicle stock variables: a dummy for not owning a car prior to the current auto purchase, number of cars in the stock, the mean age of the stock, the median age of the stock, the age of the oldest car, and the age of the newest car. The population size and vehicle stock variables represent our first set of exclusion restrictions; it is assumed that while these variables affect the decision to purchase a car (so that they must be included in the ordered probit estimation), they have no effect on the loan size.¹⁹

The vector x includes the exogenous variables that enter both the interest rate and maturity equations, as well as the loan demand equation. These are: a quadratic in the age of the household head, three educational dummies (one for household heads who did not graduate from high school, one for high school graduates, and one for individuals who have received some college education but without receiving a college degree), the number of income earners in the household, the number of adults, the number of children, a dummy for African Americans, a dummy for males, the household's after-tax income, regional unemployment rate, and regional personal disposable income. The regional unemployment rate and personal disposable income are supposed to capture macroeconomic factors, as well as expectations about future macroeconomic developments that may affect loan markets. In the actual estimation, we always found that the regional personal disposable income had no explanatory power once the unemployment rate was controlled for. We therefore dropped it from the reported specifications. The other variables capture life cycle effects that are likely to affect the demand for

¹⁹The total number of reduced form parameters in the loan demand, interest rate and maturity equations (vector X) is 25; the total number of reduced form parameters in the first stage (vector Z) is 86 (the latter include a set of population dummies, vehicle stock variables and quarter and regional dummies).

loans; in addition, they can be thought of as the determinants of the value of the car purchased by the consumer in a reduced form specification. The vector W consisting of the exogenous variables that enter the equations for r and m , but *not* the loan demand equation, includes the average (by region) tax price of the debt ($avtaxp$), which is discussed in the Appendix; the inflation rate; and a set of four depreciation rates: $depr01$ (depreciation rate between the year of car production and 1st year), $depr02$ (depreciation rate between the year of production and 2nd year), $depr14$ (depreciation rate between the 1st and 4th year), and $depr24$ (depreciation rate between the 2nd and 4th year), which are also discussed in the Appendix. This vector represents our second set of exclusion restrictions that were discussed extensively in section 2.2.

The inclusion of a household’s income in the estimation is potentially problematic, as income could itself be endogenous; this would be for example the case if consumers took a second job or worked overtime in order to purchase a car; or, alternatively, if they felt they could afford to reduce their work hours after obtaining a sizeable loan. We therefore also experimented with an alternative specification, in which income was excluded from all equations. The results were almost unchanged; we therefore report only the specifications with income, since these are more complete.

Our sample includes 70184 households, of which 16.6% bought a car (new or used) during the completed interview period. The maximum number of completed interviews in the CES is four. Because households who have completed four interviews are naturally more likely to have bought a car than households who have completed a smaller number of interviews, we also included the number of completed interviews as an explanatory variable in the estimation of the sample selection equation. Among those who bought a car, about 47% used auto financing. In the following discussion we focus on the coefficients of the finance equation; we do not report the results of the reduced form equations, since these are not interesting per se.

5.1. The loan demand equation

5.1.1. Results Based on the Full Sample

The results from the estimation of the loan demand equation are reported in Tables 6 to 10. We start by focusing on the full sample. We estimate various specifications, and examine the robustness of our results to the functional form of the loan demand equation, and to the estimation approach. Then we estimate a subset of specifications for various subgroups in the population, which differ in their likelihood of being liquidity constrained.

The main results for the full sample are reported in Table 6. The total number of households who finance is 4323.²⁰ The table reports only the two derivatives of interest: the interest rate, and

²⁰Note that this number is significantly lower than the number of observations associated with a positive finance share

maturity derivatives. The standard errors are based on bootstrapping, where the sample drawn in each replication is a bootstrap sample of clusters defined on the basis of the quarter of the car purchase (there are 48 quarters in our data); this way we allow for correlation in the error terms across households who bought a car in the same quarter. The coefficients on the variables included in the vector x (not reported here) were jointly significant in every specification; this indicates that life cycle effects are indeed important. The individual coefficients are intuitive and consistent with the results reported earlier in the descriptive analysis.

We start by estimating various specifications based on the DNV approach since this gives us more flexibility regarding the functional form. The first reported specification includes (in addition to the variables in the x vector) the interest rate, but not maturity, on the right hand side of the loan demand equation. The coefficient is negative and significant suggesting that loan demand is, as expected, negatively related to the interest rate. This result is however reversed once the maturity is included in the specification. In specification [2] we estimate a specification linear in r and m , using again the DNV approach. The point estimate for the interest rate drops to -0.09 and becomes statistically insignificant. In contrast, the maturity coefficient is large (7.38) and highly significant. To interpret these numbers, note that maturity is measured in months divided by 100. Our results accordingly suggest that increasing the maturity of a loan by 1 month, increases loan demand by 7.38%; extending maturities by 1 year, increases loan demand by 88.56% (7.38 x 12). These are economically significant effects.

In specification [3] we experiment with a more flexible functional form (3rd degree polynomial in r and m) and estimate it again using the DNV approach. The results are similar to the ones reported for the linear specification. The interest rate derivative is now larger in absolute value, but remains statistically insignificant. The maturity derivative on the other hand is almost unchanged compared to the linear specification. In specification [4], we also estimate the loan demand equation using Powell's method outlined in section 3. We only experimented with one functional form for r and m in this case, namely a linear specification, since the single index assumption on the distribution of the error terms prevents us from including higher order terms, or interactions between r and m in the loan demand equation. The results are similar to the ones reported in row [2] of the table: the interest rate derivative is insignificant, while the maturity derivative is positive, large in magnitude, and highly significant.²¹ Based on the results of Table 6 one would conclude that consumers are highly sensitive to changes in

in Tables 1-5 (5407). The reason is that we now include regional dummies in the specification, and are therefore forced to drop all observations with missing values for region. Unfortunately, the BLS blanks out information on region in certain cases, in order to preserve confidentiality. Given the information on population size, the CES user would be able to infer the exact residence of households in particular areas if the regional information were available (e.g., if the population size is greater than 4 million, and the region is West, one knows that the household lives in Los Angeles).

²¹ A test of the (four) overidentifying restrictions discussed in Section 3 yielded a statistic equal to 1.69 with a p -value of 0.79 which suggests that we do not reject the null that the model is well specified.

the loan maturity, but unresponsive to changes in the interest rate. This seems to confirm the common wisdom that consumer credit is driven primarily by the size of monthly payments, and it is strongly suggestive of the existence of liquidity constraints in the population.

In the previous sections we emphasized the importance of accounting for selection bias and the endogeneity of interest rate and maturity. To examine how our approach of dealing with these issues impacts the results, we report in Table 6a the interest rate and maturity derivatives based on specifications that ignore selection bias, simultaneity bias, or both. Row [1] of the table repeats the earlier results from Table 6, row [2] for comparison purposes; that is, [1] corrects for selection and simultaneity bias. In row [2], we estimate the loan demand equation on the subsample of consumers who borrowed for a car purchase using OLS, thus ignoring all of the above issues. The results are remarkably different from the ones reported in [1]: the interest rate derivative is negative, large in magnitude, and statistically significant; in contrast, the maturity derivative is about 2.5 times smaller than the one reported for the specification with selection bias and endogeneity correction. Estimating loan demand using OLS would thus lead to the erroneous conclusion that consumer borrowing is highly responsive to interest rate changes. In the following two rows of the Table, we sequentially control for simultaneity and selection bias to examine where the major differences between OLS and our estimation results stem from. As evident from a comparison between rows [2] and [4], the selection bias correction based on the propensity score obtained from the ordered probit estimation²² has (by itself) a negligible effect on the results (the differences are visible only at the 2-digit level of the coefficients).²³ In contrast, the treatment of interest rate and maturity as endogenous variables substantially affects the estimated derivatives. A comparison between rows [3] and [2] suggests that the simultaneity bias correction increases the maturity derivative almost threefold, while it substantially reduces the magnitude of the interest rate derivative and renders it insignificant. Results based on a specification using a 3rd degree polynomial in r and m produced a similar pattern. These results once again indicate the importance of unobserved heterogeneity in the choice of loan size and terms. In particular, the change in the derivatives between rows [2] and [3] suggests that unobserved factors affecting loan demand induce a negative correlation between loan size, interest rate, and maturity. Consumers who borrow a lot, tend to have loans with low rates and low maturities (this could be for example consumers who choose to finance from a credit union), so that the interest rate and maturity

²²We should point out that we obtained very similar results when we estimated the model using a probit rather than an ordered probit in the first stage, thus ignoring the fact that the loan amount cannot exceed the car value (but still accounting for selection into our sample). Overall, we found that the inclusion of the propensity score from the first stage of the estimation had little impact on our results. In principle, we could have used a different, less restrictive approach than the ordered probit to deal with selection in the first stage, but based on the results of Table 6a we do not believe that this would have changed our results.

²³The results in Row 4 were obtained by first estimating the ordered probit, and then using the propensity score to control for selection in the estimation of the loan demand equation. We do not instrument for the interest rate and maturity in this specification.

derivatives based on OLS are biased downwards.

Finally, correcting for both selection and endogeneity bias has a small effect on the point estimate of the maturity derivative compared to the specification that corrects for endogeneity only, but a larger effect on the interest rate derivative, which becomes now negative (but remains insignificant). Despite the fact that the selection bias correction does not seem to make a big difference in the results for the whole sample, we correct for both simultaneity and selection in our subsequent estimation, as the change in the point estimate of the interest rate derivative between rows [1] and [3] indicates that selection bias could potentially be important when the model is estimated for various subsamples.

5.1.2. Results Based on Subgroups

The results for the whole sample suggest that consumers are not responsive to interest rates, while they are very sensitive to maturity changes. To the extent that this pattern is taken as evidence for the significance of liquidity constraints, one would expect the interest rate and maturity responsiveness to vary across groups of the population that have a different likelihood of being credit constrained. To get a better idea as to who is most affected by the presence of liquidity constraints, we repeat the estimation in this subsection focusing on subsamples that are drawn on the basis of alternative criteria.²⁴

We start by using age as a criterion for dividing consumers into subgroups. Specifically, the sample is split into three age groups. Age group 1 consists of households with head 35 years and younger, age group 2 includes household heads in the age range of 35-55, and age group 3 consists of all households with heads 55 years and older. A-priori we would expect the younger households to be the ones most constrained by the existence of liquidity constraints. Such households face a steeper income profile and are thus most likely to be affected by constraints that prevent them from transferring expected future wealth to the present.

Table 7 reports results for each age group. Given the robustness of the results for the whole sample to alternative functional forms and estimation approaches, we only estimated two specifications here,

²⁴In addition to identifying the groups that are most likely to face binding liquidity constraints, the focus on subgroups has the econometric advantage that it allows for heterogeneous responses of loan demand to changes in r and m across the different subgroups. Since ignoring the presence of such heterogeneity would introduce heteroskedasticity with respect to r and m which might invalidate our single index assumption on the distribution of errors, we test for the presence of general heteroskedasticity via a White test. In particular, to conduct this test we ignored selection bias (which as indicated by Table 6a does not seem to affect our results in a significant manner), and estimated the loan demand equation using 2SLS. Due to the large number of variables included in the estimation (and the even larger number of cross-products), we concentrated on the linear specification in r and m . The resulting test-statistic was 311 for the full sample, which leads us to unequivocally reject homoskedasticity. When we performed the same test for the subgroups, the test statistics were significantly lower, between 121 and 120. With 164 degrees of freedom we still reject homoskedasticity in these cases, but the test statistics indicate an improvement. Note that these tests ignored selection issues, and did not employ higher order polynomials. It is possible that the correction for selection bias and the use of higher order polynomials in r and m , might further eliminate sources of heteroskedasticity.

one linear in r and m , and one that uses a 3rd degree polynomial in r and m . As evident in Table 7, the results across the two specifications are very similar, so we concentrate our discussion on the first one.

Contrary to our expectations, the first group consisting of the younger consumers does not appear to be more constrained than the second group of middle-aged consumers. Both groups exhibit negative and significant interest rate derivatives, and positive and significant maturity derivatives. Though based on the point estimates, the interest rate responsiveness of the middle-aged group seems slightly higher than the one of the younger consumers, the difference is not statistically significant. In contrast, the interest rate derivative of the third group of older consumers is not statistically different from zero. But for this group the maturity derivative, though positive and large in magnitude, is also not statistically different from zero, suggesting that our results could be due to the small number of observations in this subgroup. The similarity of the results between Groups 1 and 2 does not offer any support to our prior that younger consumers are more liquidity constrained than middle-aged consumers, however this lack of evidence could be due to the relatively high age cutoff point: 35 years. Because the first group spans ages between 21 and 35 years, one could argue that the older consumers in that group are not that different from the younger ones in the second group. Ideally we would like to use more narrowly defined age groupings, but then we are faced with the problem that the cells become too small, and we have difficulty obtaining bootstrap standard errors.

We encountered a similar problem when we tried to define groups based on both age and education. Within the set of young consumers, one would expect the more educated ones to be the most liquidity constrained, as such consumers have a steep income profile. However, splitting the sample into groups based on both age and education produces very small cells, and we have difficulty obtaining standard errors for many groups, while for the ones we do obtain standard errors, the results are not statistically significant.

Table 5 of the descriptive section suggests that households with low income are more likely to have finance shares in excess of 90% and hence be at a corner. This empirical finding motivates our next experiment, in which we group households based on income-related criteria. We start by using the per-capita income (after-tax household income divided by household size) to split the sample into three groups: the first one consists of households with annual after-tax per capita income below \$10,000; the second group includes all households with per-capita income between \$10,000 and \$20,000 per year; and the third group is the ‘high income’ group, including households with \$20,000 and above per-capita income.

The results displayed in Table 8 provide strong support for the hypothesis that the interest rate and maturity sensitivity of borrowers depends on their likelihood of being liquidity constrained, in a pattern that is consistent with our intuition. Note first that the interest rate derivatives have the

expected negative sign, and are fairly precisely estimated for all three groups. But the interest rate sensitivity of the third (high income) group is (focusing on the results for the 3rd degree polynomial in r and m) about 4.5 times the interest rate sensitivity of the first (low income) group, and 1.8 times the interest rate sensitivity of the second (middle-income) group. Put differently, our results imply that decreasing the interest rate by 1% would increase the loan demand of high income consumers by 14%, versus a 3.1% increase for low income borrowers, and an 8% increase for middle-income borrowers.

Moreover, both lower income groups exhibit high and statistically significant derivatives with respect to maturity; for the third group, on the other hand, there is no evidence that its loan demand responds to maturity increases. For the low income group, our estimates imply that increasing maturity by 1 year increases loan demand by 74%; the effect of a similar increase in maturity for the middle-income group drops to 55%, while the effect for high income households is not statistically significant from zero. Overall, the picture that emerges from these results strongly suggests that the lower income groups are highly sensitive to maturity extensions, while the high income group is not. In contrast, the high income group exhibits the highest sensitivity to interest rate changes.

A potential problem with the specifications in Table 8, where we divide the sample on the basis of per-capita income, is that by using household size to compute per-capita income we assign the same weight to the spouse and children as to the household head. This is arguably problematic as family expenses are not proportional to the number of household members (especially children). To address this concern, we recomputed per-capita income using OECD's formula for the adult equivalence of the spouse and children to derive an alternative (expense-adjusted) measure of family size; this measure attaches a weight of 0.67 to the spouse, and a weight of 0.4 to each of the children in the household (see also notes in Table 8). We then reestimated the specifications reported earlier. The results are reported in Table 9, and clearly display the same pattern discussed earlier. The high income group is highly interest rate sensitive, while its maturity elasticity is not significantly different from zero. The other two groups are both interest rate and maturity sensitive, with the interest rate sensitivity increasing, and the maturity sensitivity decreasing in the higher income group.

Overall, the results suggest that interest rate responsiveness is substantially more pronounced among high income households, and less so among lower income households. In contrast, high income households are not responsive to maturity changes, while lower income households are.

6. Conclusion

To summarize our results, we find evidence that liquidity constraints exist, and have important implications for the borrowing behavior of lower income households. This conclusion is based on an empirical approach that makes direct use of loan data for new car purchases. The approach is general as it does not rely on specific functional forms of the utility function. The basic idea is simply that

for consumers who are not liquidity constrained, loan demand should be a function of the price of the loan (the interest rate), while liquidity constrained consumers should be more responsive to maturity changes. To implement this idea we exploited exogenous variation in interest rates and maturities and employed an estimation approach that deals with the main challenges posed by our data, that is, selection bias and endogeneity of interest rates and maturities faced by car buyers.

The specific empirical results are consistent with the existence of liquidity constraints in that we find: (1) Consumer groups that we would a-priori consider more likely to be liquidity constrained (e.g., low income households) are highly sensitive to maturity changes, but less sensitive to interest rate changes; (2) consumer groups that do not seem likely candidates for liquidity constraints (e.g., high income consumers) exhibit significant interest rate, but no maturity sensitivity. We do not find any evidence that young consumers are more liquidity constrained than older ones.

A drawback of the approach we propose is that it only works for interior values of financing, but fails to identify liquidity constrained consumers whose situation corresponds to corner solutions; in other words, we can identify the intensive, but not extensive margin. In this sense, we can think of our approach as potentially underestimating the importance of liquidity constraints. However, this problem would have been more severe if we had failed to find any evidence in favor of liquidity constraints. Given that despite the above limitation we still find strong evidence for liquidity constraints, we believe that our conclusion that such constraints exist, in particular among low income households, is fairly robust.

BIBLIOGRAPHY

- Aizcorbe, A. and M. Starr-McCluer (1997): “Vehicle Ownership, Purchases, and Leasing: Consumer Survey Data”, *Monthly Labor Review*, 120, 34-40.
- Alessie, R., M. Devereux, and G. Weber (1997): “Intertemporal Consumption, Durables and Liquidity Constraints: A Cohort Analysis,” *European Economic Review*, 41, n1, 37-59.
- Altonji, J. G., and A. Siow (1987): “Testing the Response of Consumption to Income Changes with (Noisy) Panel Data,” *Quarterly Journal of Economics*, 102, 293-328.
- Amemiya, T. (1985): *Advanced Econometrics*, Harvard University Press.
- Andrews, D. W. K., and M. Buchinsky (2000): “A Three-Step Method for Choosing the Number of Bootstrap Repetitions,” *Econometrica*, 68, 23-51.
- Attanasio, O.P. and S.J. Davis (1996): “Relative Wage Movements and the Distribution of Consumption”, *Journal of Political Economy*, 104, 1227-1262.
- Attanasio, O.P. , P. K. Goldberg and E. Kyriazidou (2000): “Credit Constraints in the Market for Consumer Durables: Evidence from Micro Data on Car Loans”, *NBER Working Paper*, No. 7694, May 2000.
- Browning, M. and C. Meghir (1991), “The Effects of Male and Female Labor Supply on Commodity Demands”, *Econometrica*, 59, n4, 925-951.
- Brugiavini, A. and G. Weber (1992): “Durable and Nondurable Consumption: Evidence from Italian Household Data, ” IFS Working Paper 92/13.
- Chah, E. Y., V. A. Ramey, and R. M. Starr (1995): “Liquidity Constraints and Intertemporal Consumer Optimization: Theory and Evidence from Durable Goods,” *Journal of Money, Credit, and Banking*, 27, n1, 272-287.
- Das, M., W. Newey, and F. Vella (2001): “Nonparametric Estimation of Sample Selection Models,” *Review of Economic Studies*, 70, 33-58.
- Flemming, J. S. (1973): “The Consumption Function when Capital Markets are Imperfect, the Permanent Income Hypothesis Reconsidered,” *Oxford Economic Papers*, 25, 160-172.
- Hall, R. E., and F. Mishkin (1982): “The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data of Households,” *Econometrica*, 50, 461-481.

- Hansen, L., and K. Singleton (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectation Models," *Econometrica*, 50, 1269-1286.
- Hansen, L., and K. Singleton (1983): "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns", *Journal of Political Economy*, 91, n2, 249-265.
- Hendel, I., and A. Lizzeri (2002): "The Role of Leasing under Adverse Selection", *Journal of Political Economy*, 110, n1, 113-143.
- Honore, B. E. and J. L. Powell (2002): "Pairwise Difference Estimation of Nonlinear Models", working paper.
- Jaffee, D. and J. Stiglitz (1990): "Credit Rationing" in Friedman, B. M. and F. H. Hahn (eds.) *Handbook of Monetary Economics*, Vol. 2. Amsterdam: North Holland.
- Jappelli, T. (1990): "Who is Credit Constrained in the US Economy?" *Quarterly Journal of Economics*, 105, 219-234.
- Juster, F. T., and R. P. Shay (1964): "Consumer Sensitivity to Finance Rates: An Empirical and Analytical Investigation," NBER Occasional Paper No. 88.
- Maki, D. (2001), "Household Debt and the Tax Reform Act of 1986", *American Economic Review*, 91, n1, 305-319.
- Pissarides, C. (1978): "Liquidity Considerations in the Theory of Consumption," *Quarterly Journal of Economics*, 92, 279-296.
- Powell, J. L. (2001): "Semiparametric Estimation of Censored Selection Models," in *Nonlinear Statistical Modeling: Proceedings of the Thirteenth International Symposium in Economic Theory and Econometrics: Essays in honor of Takeshi Amemiya*, edited by C. Hsiao, K. Morimunem and J.L. Powell. Cambridge University Press.
- Runkle, D. (1991): "Liquidity Constraints and the Permanent Income Hypothesis," *Journal of Monetary Economics*, 27, 73-98.
- Stiglitz, J. and A. Weiss (1981), "Credit Rationing in Markets with Imperfect Information", *American Economic Review*, 71, n3, 393-410.
- Sullivan, C: "Consumer Credit: Are There Limits?," *Journal of Retail Banking*, VIII, n4, 5-18.
- Zeldes, S. (1989): "Consumption and Liquidity Constraints: An Empirical Analysis," *Journal of Political Economy*, 97, 305-346.

APPENDIX

Construction of the Instruments

1. The Tax Reform of 1986

As noted in the main text the nominal after-tax rate is given by the formula:

$$r_t = r_p * (1 - t_f * \alpha_t * I - t_s * \alpha_t * I + t_f * t_s * \alpha_t * I)$$

where r_p denotes the pre-tax interest rate, and the term in the parenthesis denotes the tax price of the debt. In order to both compute the after-tax rate, and construct the instruments related to the tax reform of 1986 described in the main text, we need to identify whether a household itemizes the deductions (the dummy I), and compute the federal and state marginal tax rates (t_f and t_s). To this end, we ran the NBER tax simulation program that uses information on the demographic structure of the household (and some consumption items), as well as information on whether the household has a mortgage, to “guess” whether the household is an itemizer or not. The program then uses state and time specific information on tax codes, income, and deductible items to compute marginal tax rates. The NBER web site provides the details of the program.

As noted in the text, both the marginal tax rates and the dummy I should be considered endogenous to the borrowing decision. However, we believe that endogeneity is unlikely to be a serious concern in our case because of the particular way these variables are constructed; in each case, we use the information the CEX household provided in the year prior to the year of the car purchase in order to compute the marginal tax rates and assign a value to the dummy I . Nevertheless, to be sure that our results are not driven by simultaneity bias, we computed the average (across regions) tax price of consumer debt, and its individual components, and used those as potential instruments instead of the household specific variables.

In the estimation we experimented with different subsets of components of the average tax price of the debt: $\alpha_t * I$, $t_f * \alpha_t * I$, $(t_f + t_s) * \alpha_t * I$, etc., all averaged across all households within a region in every year. The results we report are based on using just the average (by region) tax price of consumer debt $(1 - t_f * \alpha_t * I - t_s * \alpha_t * I + t_f * t_s * \alpha_t * I)$ in each year.

2. Depreciation Rates

The construction of depreciation rates is best explained using an example. Suppose that some households interviewed in the Consumer Expenditure Survey (CEX) in 1990 report that they own an Acura Integra that is 1 year old (this means the car is model-year 1989). The households also report the value of the car. Going back in the CEX to year 1989, we can identify households who purchased a new Acura Integra (model-year 1989) in that year. These households also report the value of their car (in this case the purchase price) in that year. Using this information we can construct an estimate of the 1-year depreciation rate of the Acura Integra 1989.

While this procedure can in principle be applied to every model separately, a practical problem is that for models that are not very popular we may have some years missing, especially as we go back in the sample. In order to avoid this problem, we first deflated all car prices using the CPI deflator, and then computed the *mean* and the *median* price for all cars of a particular cohort (that is model-year) in each year in our sample. Then we constructed two sets of depreciation rates, one based on the means, and one based on the medians of the car prices.

Note that since this method constructs depreciation rates going forward, and our sample ends in 1995, we need to utilize data from more recent surveys in order to construct depreciation rates in 1995. To this end, we use data from the CEX up to 1999. This allows us to compute depreciation rates up to the fourth year; for example, for households who bought new cars in 1995 we know the value of these cars in 1995; from the 1999 CEX we can compute the value of these cars in 1999, and then construct the 4-year depreciation rate of cars of model-year 1995.

The actual depreciation rates we use as instruments are:

depr01: depreciation rate between the year of production and 1st year.

depr02: depreciation rate between the year of production and 2nd year.

depr14: depreciation rate between the 1st and 4th year.

depr24: depreciation rate between the 2nd and 4th year.

These depreciation rates are based on means of car prices in each year. Depreciation rates based on medians were similar, so we ended up not using up in the estimation.

Table 1: Descriptive statistics of households

Variable	Number of Obs.	Mean	Std. Dev.	Min	Max
Age of household head	70184	47.0	17.8	16	90
Family size	70184	2.6	1.5	1	18
Adults	70184	1.96	0.94	1	14
Number of interviews	70184	3.4	0.81	2	4
Avg number of cars	70184	1.96	0.94	0	22
% of households with 0 cars	70184	14.7	-	-	-
Mean age of the stock of cars	69829	6.76	5.03	0	27
Financial assets	57057	11763	29836	0.	364500
% of households buying a car	16.6	-	-	-	-

Table 2: Descriptive statistics of recently bought cars

Variable	Number of Obs.	Mean	Std. Dev.	Min	Max
Fraction financing	11666 ¹	0.464	0.500	0	1
Fraction financing 100%	5409	0.180	0.384	0	1
Fraction financing >90%	5409	0.326	0.469	0	1
Finance share	5407	0.783	0.188	.02	1
Interest rate	5409	0.089	0.044	-0.053	0.229
Int. rate on new cars	2460	0.076	0.040	-0.053	0.224
Int. rate on used cars	2949	0.101	0.050	-0.053	0.229
Int.rate on finan >90%	1764	0.090	0.045	-0.053	0.214
Maturity (in months)	5409	41.4	17.7	2	252
Mat. on new cars	2460	49.6	15.8	2	252
Mat. on used cars	2949	34.6	16.3	2	180
Mat. on finan. >90%	1764	40.5	18.2	2	252

Finance Source	Frequency	Percent
Dealer	1178	22.3
Bank	2348	44.4
Credit union	914	17.30
Finance company	477	9.03
Other (including missing)	492	6.70

¹The number 11666 is the number of households who bought a car. Out of them 5409 (0.464 x 11666) financed the car purchase.

Table 3: List of Variables Used in the Estimation

AGE2	Dummy, 1 if age of household head is between 35 and 55
AGE3	Dummy, 1 if age of household head is greater than 55
EDUC3	Dummy, 1 if household head is a high school graduate
EDUC4	Dummy, 1 if household head attended (but not completed) college
EDUC56	Dummy, 1 if household head is college graduate or has higher education

Table 4: Choice of Long Maturity

Dependent variable: 1: If long maturity is used; 0: Otherwise

Method of estimation: Probit

Number of observations: 5409

Number of positive observations: 1305

Parameter (Std. Error)	Spec. 1	Spec. 2	Spec. 3
C	-0.82 (0.06)	-0.91 (0.06)	-0.73 (0.06)
AGE2	-0.10 (0.04)	-0.13 (0.04)	-0.10 (0.04)
AGE3	-0.12 (0.06)	-0.12 (0.06)	-0.15 (0.06)
EDUC3	0.16 (0.06)	0.12 (0.06)	0.16 (0.06)
EDUC4	0.29 (0.06)	0.23 (0.07)	0.29 (0.06)
EDUC56	0.16 (0.06)	0.05 (0.07)	0.16 (0.06)
FEMALE	0.00 (0.05)	0.03 (0.05)	0.01 (0.05)
MINOR	-0.02 (0.06)	-0.02 (0.07)	-0.04 (0.06)
INCOME	— (NA)	.46 E-05 (.82 E-06)	— (NA)
ASSET	— (NA)	-.93 E-06 (.94 E-06)	— (NA)
BANK	— (NA)	— (NA)	-0.05 (0.04)
CREDIT UNION	— (NA)	— (NA)	-0.19 (0.06)
FINANCIAL COMPANY	— (NA)	— (NA)	-0.02 (0.07)
OTHER	— (NA)	— (NA)	-0.88 (0.14)

Table 5: Who Finances more than 90%

Dependent variable: 1: If fin. share>0.9; 0: Otherwise

Method of estimation: Probit

Number of observations: 5409

Number of positive observations: 1894

Parameter (Std. Error)	Spec. 1	Spec. 2
C	-0.18 (0.05)	-0.14 (0.06)
AGE2	0.01 (0.04)	0.04 (0.04)
AGE3	-0.28 (0.05)	-0.23 (0.06)
EDUC3	-0.16 (0.05)	-0.12 (0.06)
EDUC4	-0.21 (0.06)	-0.16 (0.06)
EDUC56	-0.37 (0.06)	-0.29 (0.06)
FEMALE	0.12 (0.04)	0.09 (0.05)
MINOR	0.05 (0.06)	0.04 (0.06)
INCOME	— (NA)	-.16 E-05 (.80 E-06)
ASSET	— (NA)	-.38 E-05 (.96 E-06)

Table 6: The demand for car loans

Full Sample

Dependent variable: Log of Loan Size

Number of Obs. Financing: 4323

Specification	r -Derivative		m -derivative	
[1] <i>DNV</i> , Only r	-0.82	(0.25)	—	—
[2] <i>DNV</i> , r and m	-0.09	(0.25)	7.38	(1.33)
[3] <i>DNV</i> , cubic polynomial in r and m	-0.27	(0.26)	7.45	(1.39)
[4] <i>P</i> , r and m	-0.05	(0.37)	8.77	(2.69)

Notes:

1) Bootstrap standard errors in parentheses. Number of bootstrap repetitions=100.

2) *P*: Powell (2001); *DNV*: Das, Newey, Vella (2001)

3) In specifications [2] and [4] the function $f(r, m)$ takes the form: $f(r, m) = \gamma_1 r + \gamma_2 m$.

4) In specification [3] the function $f(r, m)$ takes the form:

$$f(r, m) = \gamma_1 r + \gamma_2 m + \gamma_3 r^2 + \gamma_4 m^2 + \gamma_5 r m + \gamma_6 r^3 + \gamma_7 m^3 + \gamma_8 r^2 m + \gamma_9 r m^2$$

5) r is measured in percentage points, divided by 10; m is measured in months, divided by 100.

**Table 6a: The demand for Car Loans
Robustness Analysis**

Full Sample

Dependent variable: Log of Loan Size

Number of Obs. Financing: 4323

<i>DNV, r</i> and <i>m</i>	Specification	<i>r</i> -Derivative		<i>m</i> -derivative	
[1]	Correction for Selection and Endogeneity	-0.09	(0.25)	7.38	(1.33)
[2]	No Correction (OLS)	-0.20	(0.02)	2.94	(0.06)
[3]	Correction for Endogeneity only	0.08	(0.30)	7.79	(1.52)
[4]	Correction for Selection only	-0.20	(0.03)	2.94	(0.10)

Table 7: The demand for car loans

By Age Group

Dependent variable: Log of Loan Size

Specification	Age Gr. 1		Age Gr. 2		Age Gr. 3	
	<i>r</i>	<i>m</i>	<i>r</i>	<i>m</i>	<i>r</i>	<i>m</i>
<i>DNV</i> , <i>r</i> and <i>m</i>	-0.49 (0.24)	4.03 (0.73)	-0.61 (0.32)	4.26 (2.11)	0.27 (1.00)	7.46 (4.32)
<i>DNV</i> , cubic polynomial in <i>r</i> and <i>m</i>	-0.59 (0.26)	4.31 (0.87)	-0.67 (0.34)	4.79 (2.32)	0.13 (1.01)	7.81 (4.47)
Number of obs. financing	1534	1534	2053	2053	736	736

Notes:

Age groups are defined as follows:

Age Gr. 1: Age of Household Head ≤ 35

Age Gr. 2: $35 < \text{Age of Household Head} < 55$

Age Gr. 3: Age of Household Head ≥ 55

Table 8: The demand for car loans

Groups based on Per-Capita Income

Dependent variable: Log of Loan Size

Specification	Income Gr. 1		Income Gr. 2		Income Gr. 3	
	<i>r</i>	<i>m</i>	<i>r</i>	<i>m</i>	<i>r</i>	<i>m</i>
<i>DNV</i> , <i>r</i> and <i>m</i>	-0.24 (0.17)	6.04 (1.48)	-0.68 (0.37)	3.62 (1.49)	-1.34 (0.79)	-0.04 (1.84)
<i>DNV</i> , cubic polynomial in <i>r</i> and <i>m</i>	-0.31 (0.17)	6.17 (1.72)	-0.80 (0.42)	4.57 (1.62)	-1.40 (0.82)	0.21 (2.10)
Number of obs. financing	1752	1752	1525	1525	814	814

Notes:

Groups are defined as follows:

Inc. Gr. 1: $0 < \text{Per-Capita Income} \leq 10000$

Inc. Gr. 2: $10000 < \text{Per-Capita Income} < 20000$

Inc. Gr. 3: $\text{Per-Capita Income} \geq 20000$

Table 9: The demand for car loans

**Groups based on Per-Capita Income
with family size adjusted according to the
OECD adult-equivalence formula for children**

Dependent variable: Log of Loan Size

Specification	Income Gr. 1		Income Gr. 2		Income Gr. 3	
	<i>r</i>	<i>m</i>	<i>r</i>	<i>m</i>	<i>r</i>	<i>m</i>
<i>DNV</i> , <i>r</i> and <i>m</i>	-0.26 (0.23)	4.98 (1.65)	-0.58 (0.38)	3.45 (1.72)	-1.24 (0.66)	0.96 (1.49)
<i>DNV</i> , cubic polynomial in <i>r</i> and <i>m</i>	-0.35 (0.24)	5.22 (1.67)	-0.73 (0.44)	3.91 (1.93)	-1.21 (0.68)	1.28 (1.68)
Number of obs. financing	1934	1934	1618	1618	539	539

Notes:

The OECD formula is:

$$famsize = 1 + (adult - 1) * 0.67 + kids * 0.4$$

Groups are defined as follows:

Inc. Gr. 1: $0 < \text{Per-Capita Income} \leq 15000$

Inc. Gr. 2: $15000 < \text{Per-Capita Income} < 30000$

Inc. Gr. 3: $\text{Per-Capita Income} \geq 30000$




Figure 1: Average Depreciation Rate between 2nd and 4th year