Habit Persistence and Keeping Up with the Joneses: Evidence from Micro Data

Enrichetta Ravina*
New York University

November 2007

Abstract

This paper provides evidence of habit persistence in household consumption choices. I find that the strength of external habit, captured by the fraction of the consumption of the reference group that enters the utility function, is 0.290, and the strength of internal habit, represented by household past consumption, is 0.503. The results are robust to controlling for various measures of economic activity, tests for the presence of aggregate shocks, liquidity constraints, precautionary saving motives, and learning. Aggregation of the Euler equations as a weighted average of individual marginal rates of substitution accounts for heterogeneity and market incompleteness and preserves the results.

Keywords: Habit Persistence; Micro data; Household Finance; Social Interactions.

*I would like to thank Debbie Lucas, Kent Daniel, Bill Rogerson and Paola Sapienza for long discussions, support and guidance. Special thanks go to Donald P. Jacobs and the U.S. credit card company that provided the data. I am also grateful to Roc Armenter, Paco Buera, Francesco Caselli, John Campbell, Janice Eberly, Martin Eichenbaum, Mike Fishman, Joel Hasbrouck, Johannes Horner, Ming Hsu, Arvind Krishnamurthy, David Laibson, Martin Lettau, Annamaria Lusardi, Stefan Nagel, John Panzar, Marcin Peski, Karl Schmedders, Christopher Polk, Matthew Rabin, Marciano Siniscalchi, Mark Seasholes, Costis Skiadas, Annette Vissing-Jorgensen, partecipants at the Russell Sage Foundation Behavioral Camp 2004, and seminar participants at Kellogg School of Management (Finance), Northwestern University, Cornell, NYU-Stern, Harvard Business School NOM and Finance groups, Brandeis, UNC at Chapell Hill, London Business School, Stockholm School of Economics, the FED Board of Governors, the Western Finance Association 2005 Meeting, the Society of Economic Dynamics 2005 Meeting, the 2005 Econometric Society World Congress, and the NYU Stern Macro Lunch for valuable discussions and suggestions. Financial support from the Northwestern University Dissertation Year Fellowship, Banca d'Italia and Fondazione Einaudi is gratefully acknowledged. All remaining errors are mine. Comments are welcome: eravina@stern.nyu.edu.

The ideas of habit formation and social comparisons in consumption choices have a long history in economics, dating back to Thorstein Veblen's 1899 "The Theory of the Leisure Class," and James Duesenberry's 1949 "Income, Saving, and the Theory of Consumer Behavior". These studies, and the many studies written thereafter, postulate that individuals derive utility not only from the level of their current consumption, but also from how their consumption compares to their own past consumption (internal habit) and the consumption of the people around them ("Keeping Up with the Joneses," or external habit).

Habit formation models have proven very successful in providing theoretical explanations of a variety of dynamic asset pricing phenomena and macroeconomic facts. In the asset pricing literature, they have been employed to explain the equity premium puzzle (Constantinides (1990), Abel (1990, 1999), Campbell and Cochrane (1999)), the procyclical variation of stock prices (Campbell and Shiller (1988)), the countercyclical variation of stock market volatility (Harvey (1989)), and the term structure of interest rates (Buraschi and Jiltsov (2007)). In the macroeconomics literature, habit persistence frameworks explain business cycle facts (Boldrin, Christiano and Fisher (2001)), savings and growth (Carroll, Overland and Weil (2000)), and consumption's response to monetary and other shocks (Fuhrer (2000)). However, despite these models' impressive track record in simulations with aggregate data, there is only mixed evidence on whether they reflect actual preferences. The empirical studies that have addressed this question so far have mostly followed the macroeconomists' approach to aggregate consumption, leaving the micro foundations of the phenomenon largely unexplored.[1]

In this paper, I test the micro story behind habit formation models. I look into actual household consumption decisions and estimate a log-linearized Euler equation for a representative sample of U.S. credit-card holders. The estimation equation incorporates internal and external habit motives, using a setting characterized by uninsurable income shocks and household-specific interest rates.

The empirical strategy builds on the micro consumption literature, and shows two novel features. First, I exploit the detailed information on the evolution of household financial variables. I use this information to build a more powerful instrument set, and to directly control for the effect of household-specific interest rates, debt burden, and credit availability. These controls are valuable in disentangling the effect of internal habit from that of liquidity constraints. They also have the

[1] An exception is Dynan (2000), who investigates habit formation in annual food consumption data from the PSID and reaches negative conclusions. I illustrate in detail the differences between my approach and hers, and the possible reasons for the different findings, later in this section. I also defer a discussion of a few other papers that investigate the effects of habit formation on stock investments at the micro level to Section 2.

advantage of allowing me to estimate the reaction of consumption to household-specific interest rates, rather than the risk-free rate most of the previous literature has had to use. Second, I define and test a more intuitive measure of the external reference point of each household (HH), the consumption level of the city in which the household lives, measured by city-level sales, rather than the consumption of the entire nation, as in previous studies. To address the potential endogeneity of city-level consumption in the Euler equation, in addition to using an instrumental variables approach that exploits lagged information, I show that my results are robust to using an alternative, strictly exogenous measure of changes in the external reference point, whether somebody in the city of the household wins the lottery.

To measure consumption, I use a novel panel data set, the Credit Card Panel (CCP), which consists of 2,674 U.S. credit-card accounts located in California, for the period between the third quarter of 1999 and the third quarter of 2002. I construct the consumption measure as the sum of all the credit-card purchases over the quarter for the households that are active credit-card users. The comparison with the U.S. Census and the Survey of Consumer Finances described in detail below show that the sample is representative of the U.S. population for both demographic characteristics and borrowing behavior. Moreover, this variable exhibits the characteristics we expect to see in household consumption: a hump-shaped path over the life cycle, and an increasing relationship with income. Although it is far from perfect, this measure of consumption allows me to overcome some of the drawbacks of previously used data sets, such as the Panel Study of Income Dynamics (PSID), which contains only information on food consumed both at home and at restaurants, and the Consumer Expenditure Survey (CEX), which provides a very detailed measure of consumption, but only follows households up to five quarters, and does not contain any household-specific geographic or detailed financial information.

I find that the strength of external habit, captured by the consumption of the reference group, is 0.29 (significant at the 5% level), and that the strength of internal habit, represented by household past consumption, is 0.503 (significant at the 1% level). These coefficients represent the fraction of city-level aggregate consumption (or household's own past consumption) that enters the utility function as the reference level to which the household compares itself. A coefficient of zero would imply that a household is not influenced by the consumption of its neighbors (its own past consumption), and the model would then collapse to the standard one used in the literature. On the contrary, a coefficient of one would mean that the household only cares about the way its consumption compares to the neighbors' (its past own consumption), and not about the absolute level.

These findings provide micro evidence that supports the theories that explain macroeconomic and asset pricing phenomena by introducing habit persistence in the utility function. The magnitude of the coefficients is in line with the conclusions of these theories, and points to a large effect of habit formation.

The results are robust to a variety of checks. To control for the possibility that city and lagged household-level consumption growth capture the effect of some omitted variable, I repeat the estimations adding to the regression the change in city-level unemployment rate, current and future state-wide income growth rates, and measures of city housing market conditions. I also use lottery winnings as an exogenous change in external reference points. To address the concern that unobservable aggregate shocks might influence the results, I add time dummies to the regression and repeat the estimation separately for each major occupation in the data set to address the concern that such shocks might affect different HHs in a different way. I also directly test for the presence of aggregate shocks following the methods in Runkle (1991), and I detect no aggregate shock. These results reduce the concern that aggregate shocks influence the findings. All the tests confirm the economic and statistical significance of the habit persistence coefficients.

Alternative explanations of the findings include the presence of liquidity constraints, precautionary saving motives, and learning about the household's income profile. These phenomena, like internal habit formation, cause consumption to adjust slowly to changes in income, and therefore induce a positive correlation between current and lagged consumption growth rates. Following Zeldes (1989), I test the habit formation hypothesis against a liquidity constraints model by re-estimating the Euler equation on two subsamples of unconstrained and credit-constrained HHs. I perform further tests by including the lagged growth rate of income and a credit constrained indicator into the regression. The tests show some evidence of liquidity constraints in addition to those accounted for by the household-specific borrowing rate, but indicate that these liquidity constraints are not the cause of the results. The precautionary motive and learning stories have similar implications. I test them further by adding a measure of consumption uncertainty to the regression. Again, the tests confirm the validity of the habit persistence interpretation of the evidence.

My results differ from those in Dynan (2000), who investigates internal habit formation in annual food consumption using the PSID and reaches negative conclusions. A comparison of the two methodologies indicates that the main reason for these discrepancies is the instrument set used in this study, due to the availability of household-specific financial information.[2] I find that once I drop the financial variables from the instrument set I cannot capture the endogenous variables as well as before and, most important, that the internal habit coefficient drops from 0.60 to 0.13.

Finally, I examine the aggregate implications of the micro findings. Following Attanasio and Weber (1993b) and Brav, Constantinides and Geczy (2002), I estimate an Euler equation on aggregate data, which I obtain by taking the weighted average of HHs' marginal rates of substitution (MRS), $\Delta\left(\frac{1}{N}\sum_{i=1}^{N} lnc_{i,t}\right)$. This procedure takes into account that due to market incompleteness and liquidity constraints, at any point in time the MRS of different HHs are not necessarily the same. I find a habit persistence coefficient of 0.515 (significant at the 5% level). On the contrary, an econometrician who uses the same data to calculate the MRS of a representative agent that consumes per capita consumption, $\Delta ln\left(\frac{1}{N}\sum_{i=1}^{N} c_{i,t}\right)$, will find no evidence of habit persistence.[3] These findings suggest the inadequacy of the representative agent as a description of real households consumption choices and of their aggregate implications.

The paper is also related to the literature that studies the effect of social interactions on stock market participation and investments. Hong, Kubik and Stein (2004) show that controlling for wealth, race, education and risk tolerance investors that have more social interactions participate more in the stock market, especially when their neighbors' participation is higher. Hong, Kubik and Stein (2005) document the importance of word of mouth transmission of information about stocks among mutual fund managers. This paper analyzes the effect of social interactions on consumption decisions. In addition, the paper is related to the household finance literature that studies individual financial decisions. A large part of this literature focuses on stock and mutual fund investments (Calvet, Campbell and Sodini (2007a,b), Odean (1998, 1999), Barber and Odean (2001)), or retirement decisions (Choi et al. (2005a,b)). This paper studies consumption choices to better understand investors' preferences and the drivers of their financial decisions.

The remainder of the paper is organized as follows. Section 1 describes the data and compares them to those traditionally used in the literature. Section 2 presents the Euler equation that guides the estimation, describes the empirical strategy, and illustrates the results. Section 3 contains various robustness checks, and Section 4 examines alternative explanations for the results.

[2]The inclusion of financial variables in the instrument set generates both higher partial R^2's in the first stage regressions and F statistics robust to the Stock and Yogo (2002) critique.

[3]The two measures differ if the cross sectional distribution of HHs' consumption growth rates changes over time. The first measure, the weighted average of HHs' MRS, will pick up changes in such distribution, while the representative agent MRS will change only when the growth rate of aggregate consumption changes. The results in Section 6 show that indeed the distribution of HHs consumption growth rates evolves from fairly flat to quite peaked over the time period analyzed.

Section 5 examines the aggregate implications of the micro findings. Section 6 concludes. The appendix contains detailed information on the sample selection, the features of the data, and the HH maximization problem that generates the Euler equation I estimate.

1 Data Description

To measure household consumption, I use the CCP panel data set, which comprises 2,674 U.S. credit-card accounts located in California, for the period between the third quarter of 1999 and the third quarter of 2002. The data provide information on spending and borrowing patterns, the evolution of interest rates, and credit availability. The data also give a snapshot of the main economic and demographic characteristics of the account holder and the zip code of the area in which he or she lives.[4] To measure the external reference point, I link these data to city-level quarterly retail sales data. To perform robustness checks I augment the data with Census, BLS and ACCRA city-level information on median house values and rent, median income, unemployment, inflation data, and city-level mortgage rates.

Table I provides the definition and sources of the variables, while Table II contains summary statistics. In constructing the sample, to obtain a more meaningful measure of consumption I exclude people whose accounts are inactive and those who seldom use their credit cards. Following earlier studies, I also exclude retired account holders, because modeling their consumption and borrowing decisions is very complex and requires the consideration of issues such as bequests, failing health and other factors that are not specific to the scope of this study. Section A of Appendix I provides a detailed description of the sample selection procedure.

1.1 Demographic and Financial Characteristics

The Credit Card Panel (CCP) is representative of the U.S. population in terms of both demographic characteristics and borrowing behavior. Section B of Appendix I compares the demographics of the account holders to those of the U.S. population reported in the 2000 U.S. Census. The Appendix shows that the distributions of income and age are very alike. The main discrepancy is due to the fact that HHs in the lower income range, and individuals who are either very young or old, are under-represented in the credit-card data set. The Survey of Consumer Finances (SCF), which is the main source of information for assets and liabilities of U.S. households, also shares this feature,

[4] One of the major U.S. credit card issuers, which wishes to remain anonymous, provided the data set.

because people in those age categories are less likely to own financial instruments. Finally, my data set contains information on the occupation of the account holders. However, I note that it is harder to make a comparison with the Census Bureau figures, given the different classification criteria.

My data set also compares well with the U.S. data on household borrowing behavior. The sample is similar to the SCF on the percentage of people who do not pay their credit-card balance in full at the end of the month. The SCF estimates this percentage at 44.4%. In my sample it is 45.75%. However, studies such as those by Laibson et al. (2000) and Gross and Souleles (2002) show that the SCF survey suffers from under-reporting of debt. One of the specific advantages of my data set is that it is not a survey. Therefore under-reporting and measurement error in the financial variables are not an issue. Comparisons with a large multi-issuer credit card data set covering the period between 1995 and 1998 and used by Gross and Souleles (2002) further confirms that the CCP correctly represents the level of indebtedness and interest rates faced by the households.

1.2 Credit Card Expenditures as a Measure of Consumption

I measure household consumption as the sum of the purchases and cash advances charged on the credit card each quarter. The data set also contains information on the size and timing of balance transfers, and such quantities are excluded from the consumption measure.

The previous section shows that the sample is representative of the U.S. credit-card accounts. An important question is whether credit-card expenditures are a good measure of household consumption. Panels A and B of Figure I show that this measure of consumption exhibits the characteristics we expect to see in household consumption, a hump-shaped path over the life cycle and an increasing relation with income.

Credit-card purchases represent an increasing fraction of U.S. consumer spending. Roughly 24% of people purchases in the United States are made by credit cards, retail cards, and debit cards. Statistics from the Survey of Consumer Finances (SCF) and other sources indicate that in the year 2000, the average American had five or six credit cards and used them to charge over $1 trillion in purchases, more than he or she spent in cash.[5] While these statistics provide support for using credit-card spending as a measure of consumption, they also raise the issue of whether having information about one credit card is enough. To address this issue, I build quarterly nondurables and services expenditures from the CEX, following the criteria in Attanasio and Weber (1995).

[5] Lim, Paul J., and Matthew Benjamin. "Digging Your Way Out of Debt", U.S. News and World Report, (3/19/01); Gerdes and Walton (2002) and Zinman (2004).

Using CEX data, I estimate the amount that the average household spends on credit cards is between \$1,143 and \$1,212. In comparison, the average credit-card expenditure in my sample is approximately \$700, which shows that the HHs in my data set make heavy use of this credit card and therefore provide a good measure of their expenditures. Section C of the Appendix provides details on this calculations.

The main advantage of my data in comparison to the PSID is that they provide a more comprehensive measure of consumption than food consumed at home or at restaurants, which has been proven inadequate for many reasons.[6] The CEX solves some of the drawbacks of the PSID by providing a very detailed measure of overall consumption. However, the CEX only interviews families for up to five quarters, and the survey does not contain any household-specific geographic or detailed financial information.

One feature of the data set is that we can observe expenditures fluctuate on this card even though they stay relatively stable overall, if the appeal of using this card versus another varies over time. Information on interest rates and credit line variations, unused portion of the credit line and balance transfers in and out of the card helps to control for such effects. Also, low activity accounts are excluded from the sample, and time dummies and account characteristics are included in the estimation to capture aggregate changes and individual time invariant effects, respectively. In addition, there is no particular pattern in the way the appeal of this card should vary over time, economic conditions, or across different HHs. This phenomenon will therefore add noise to the data and make the estimation harder, but after controlling for the variables listed above, will not biased the results in a particular direction.

Another feature of the data is that some of the expenditures on the credit card might be related to durable goods. I note that, despite I don't have information on the fraction of credit-card purchases related to durables goods, the presence of such goods would bias the results against finding the positive coefficient implied by internal habit.

Table III compares the mean and standard deviation of my consumption measure to those of both micro-level and aggregate data, and shows that these measures are comparable. The statistics

[6] First, food is a necessity, its share of expenditure falls with wealth, and it might not represent well the overall consumption basket. Moreover, using food as a proxy for total consumption implicitly assumes separability between food and other commodities, and this has been rejected by numerous studies. In particular, Attanasio and Weber (1995) analyze the bias induced by using PSID food consumption rather than a more general measure and find it sizeable. Another problem with the PSID measure of consumption is that the way the question about food expenditure is posed leaves a lot of space to the interpretation of the timing of the variable making the choice of the correct timing for the instrumental variables very hard. Finally, Runkle (1991) estimates that up to 76% of the variability of food expenditure is noise.

confirm that household consumption is very noisy, with a standard deviation of 2.73 if I use all the observations down to 0.37 if, like Zeldes (1989), I consider only those cases in which the growth rate of consumption is between -1.1 and 1.1. This statistics is similar to the one in Zeldes (1989), who uses annual data on food consumption from the PSID. I also provide a comparison with the CEX data used in Brav et al. (2002). Since they analyze cross-sectional averages, I build the same quantity in my data set. The volatility of my individual consumption measure is 0.33, compared to 0.06 obtained by Brav et al. (2002) for the 1982-1996 period.

Finally, I note that the assumption underlying my estimation is the separability between credit-card expenditures and the rest of the consumption basket. Unfortunately, detailed statistics on the type of goods people buy on the credit card is not available and this assumption cannot be directly tested. Macroeconomic models of cash and credit goods can shed light on the issue of separability.

1.3 City-level Consumption

As the reference group of the household I use the city in which the household lives. My measure of the consumption of the reference group is city-level quarterly per capita taxable sales, which I obtain from the California Board of Equalization (BOE). I use this variable because it is a very comprehensive, natural measure of city-level consumption and it is a good aggregate equivalent to credit card expenditures. The main categories it excludes are necessities (such as food consumed at home, prescription medicines, etc.) and sales of goods intended for resale.

Also, retail sales constitute an important component of personal consumption expenditures at the national level. The National Income and Product Accounts (NIPA) constructs this quantity from a variety of sources, among which a monthly sample of national retail sales plays a central role. The measure I use here is not a sample, but the total of all reported taxable sales. Therefore, it is not influenced by sampling procedures and errors.[7] Table IV presents a comparison of the summary statistics of aggregate consumption constructed with city-level sales data and personal consumption expenditures on nondurables and services from NIPA. The two measures compare very well in their variability and are highly correlated. The main difference between them is due to the fact that NIPA includes housing services.

[7] See Wilcox (1992) for a thorough description of the way NIPA personal consumption expenditures are constructed and the implications of these imperfections for empirical work.

2 Estimation and Empirical Evidence on the Micro Foundations of Habit Persistence

Households face uninsurable income shocks and borrowing rates that depend on their asset position and credit history. Davis, Kubler and Willen (2004) show that disregarding borrowing frictions, and especially the wedge between borrowing and lending rates, can lead to unrealistic predictions in terms of the amounts borrowed and the portfolio allocations of the households. Despite this evidence, empirical analyses of consumption decisions usually assume that HHs can borrow and save at the same rate, due to lack of data. In this paper I depart from these assumptions. In Appendix II, I present a standard model of intertemporal choice that incorporates uninsurable income risk and household-specific borrowing rates. The Euler equation I derive from this model guides my choice of the variables to consider in the empirical analysis and offers a framework for interpreting the results:

$$u^c_{i,t} = \beta E_t \left[[u^c_{i,t+1} + \beta\zeta E_Y u^h_{i,t+2}](1 + (R^f_{i,t+1} - 1)1[Y^H_{i,t+1}])(1 + (R^C_{i,t} - 1)1[B]) - \zeta u^h_{i,t+1} \right] \tag{2.1}$$

where $R^C_{i,t}$ and $R^f_{i,t}$ are the gross credit card and the gross risk free rate, respectively. $u^c_{i,t}$ and $u^h_{i,t}$ are the marginal utility w.r.t. consumption and the internal habit level, and ζ is the internal habit coefficient. $1[Y^H_{t+1}]$ is an indicator function that equals one for high realizations of income that will allow the HH to repay the credit card in full next period. $1[B]$ is an indicator function that equals one if the HH has an unpaid balance in the current period.[8]

The household decides how much to consume today versus tomorrow by weighting future utility and different interest rates against the probability that it will actually face them. If we disregard for a moment the effect of the habit stock, we can see that if the HH consumes \$1 less today it loses $u^c(c_t)$ and gains the following: in the next period the HH's credit-card balance will be \$1 lower, making one more dollar available for consumption, and yielding a utility of $u^c(c_{t+1})$. In addition to this gain, if the income realization is high enough that the HH is able to repay the balance in full, the HH will earn the gross risk-free rate on the dollar moved through time and the utility will be $u^c(c_{t+1})R^f_{t+1}$. If the HH carries a balance, consuming one dollar less today means that the credit-card balance next period will be R^C_t dollars less. The utility deriving from this intertemporal transfer will be $u^c(c_{t+1})R^C_t$ if the HH does not have enough resources to pay the balance in full

[8] See Appendix II for more details.

in period *t+1*, and $u^c(c_{t+1})R_t^C R_{t+1}^f$ if it does and can invest the dollar charged on the credit card at the risk-free rate. The presence of the habit stock generates an additional effect because when the HH consumes one dollar less today it increases tomorrow's utility not only directly, but also indirectly by decreasing its habit level.

I note that HH income does not enter the Euler equation, since HHs optimize the consumption path and offset predicted changes in income with asset transactions, while unanticipated changes in income are reflected in the error term. This result does not hold for those HHs that are liquidity constrained. A big advantage of the data set is that I can take the degree of liquidity constraints into account in my empirical analysis.

The assumptions about HHs' expectations on the evolution of income that are consistent with (2.1) are quite flexible, and include income processes following ARMA models and both trend and difference stationary income processes.

2.1 Specification and Estimation Strategy

Following Deaton (1992), I can express equation (2.1) as a second-order difference equation in $u^c_{i,t}$, whose solution is given by:

$$u^c_{i,t} = \beta E_t \left[u^c_{i,t+1}(1 + (R^f_{i,t+1} - 1)1[Y^H_{i,t+1}])(1 + (R^C_{i,t} - 1)1[B]) \right] \tag{2.2}$$

Equation (2.2) holds approximately if the number of lags of consumption entering the habit stock is small relative to the HH lifetime horizon, and if the HH has static expectations about future interest rates.[9] The equation holds exactly if the interest rates are constant.

Following the consumption literature, I consider a log-linear version of (2.2):

$$\ln u^c_{i,t} = \ln \beta + k + \ln u^c_{i,t+1} + \ln(1 + (R^f_{i,t+1} - 1)1[Y^H_{i,t+1}]) + \ln(1 + (R^C_{i,t} - 1)1[B]) + \varepsilon_{i,t+1} \tag{2.3}$$

where $\varepsilon_{i,t+1}$ contains an expectation error, a multiplicative measurement error in consumption and preference shocks, and k contains second and possibly higher moments of the variables. As is traditional in the literature, I assume that these variables are either constant or uncorrelated with the instruments used in the estimation.[10]

[9] Hayashi (1985a) obtains a similar result and provides a proof of this statement.

[10] In the robustness Section I provide evidence in support of this assumption by showing that including a measure of the variance of consumption in the regression doesn't change the results.

The utility function depends not only on the level of current consumption, but also on own past consumption, the consumption of the reference group and demographic characteristics:

$$u(c_{i,t}, H_{i,t}, \Theta_i, h_{i,t}) = u(c_{i,t} - h_{i,t} - H_{i,t}) \exp(\theta' \Theta_{i,t}) \tag{2.4}$$

The variable $h_{i,t}$ represents the internal habit stock and depends on the household's past consumption:

$$h_{i,t} = \zeta c_{i,t-1}$$

In this framework, to maintain the marginal utility constant, an increase in past consumption must be followed by an increase in present consumption, because the habit stock to which consumption is compared is higher. Therefore the HH will try to smooth not only consumption levels, but also changes. The HH's consumption will react slowly to changes in permanent income to avoid the risk of building a habit too quickly. Theoretical papers usually model internal habit formation by including a large number of consumption lags in the utility function. However, empirical investigations are usually limited to one lag as adding extra ones would require very long lags of the instruments to insure exogeneity.

The variable $H_{i,t}$ represents the external habit level. I model it as a function of the consumption of the city in which the HH lives:

$$H_{i,t} = \alpha_0 C_{i,t} + \alpha_{-1} C_{i,t-1} \tag{2.5}$$

$\Theta_{i,t}$ represents household demographic characteristics. Attanasio and Weber (1993 and 1995) show that age, family characteristics and labor supply choices are very important explanatory factors for individual consumption. I condition on the optimal value of these variables by incorporating them into the utility function in the multiplicative way shown in (2.4):

$$\Theta_{i,t} = \theta_1 age_{i,t} + \theta_2 age_{i,t}^2 + a_i + t_t + e_{i,t} \tag{2.6}$$

where a_i is an unobservable HH-specific effect, t_t is a time-varying effect that is constant across HHs, and $e_{i,t}$ an idiosyncratic component orthogonal to the previous two.

I also add to some of the estimations individual characteristics such as the marital status of the HH's head, homeownership, income bracket, occupation and, in some specifications, the median income, house value and unemployment rate in the zip code area of the HH at the end of 1999.

To control for cyclical fluctuations in consumption, I include seasonal dummies in the estimation. This specification is equivalent to modeling the discount factor as dependent on HH socioeconomic characteristics and the time and seasonal dummies.

The log-linear Euler equation (2.3) can be re-written more extensively as:

$$\begin{aligned}\Delta \ln c_{i,t} = k_1 + \alpha_0 \Delta \ln C_{i,t} + \alpha_{-1} \Delta \ln C_{i,t-1} + \zeta \Delta \ln c_{i,t-1} + \gamma \ln(1 + (R^f_{i,t} - 1) \Pr[Y^H_{i,t}]) \\ + \eta \ln(1 + (R^C_{i,t-1} - 1) 1[B]) + \theta_1 \Delta age_{i,t} + \theta_2 \Delta age^2_{i,t} + Seas.Dummies + \varepsilon_{i,t}\end{aligned} \tag{2.7}$$

where, following Muellbauer (1988) and Dynan (2000), I approximate the expression $\ln u(c_{i,t} - H_{i,t} - h_{i,t})$ with $\ln u(c_{i,t}) - \ln u(H_{i,t}) - \ln u(h_{i,t})$.[11] For simplicity and flexibility, I specify the utility function to depend linearly on the growth rate of HH and aggregate consumption. This specification subsumes the isoelastic functional form widely used in the literature, but also the ratio specification in Abel (1990). Since this equation reflects the fundamental dynamics of habit persistence, the results of the estimation do not constitute only a test of the specific model presented, but also of the general idea of habit formation. I also note that the estimation takes into account unobserved household heterogeneity in consumption levels, since the equation is in first differences and the HH fixed effect a_i contained in (2.3) drops out.

The **estimation strategy** exploits the panel aspect of the data set and the wide cross-sectional variation in the external reference point faced by households living in different cities. The identification of the parameters is achieved in the cross sectional dimension. I estimate the equation by using a GMM procedure with robust standard errors. I treat as endogenous both household and city-level consumption and the interest rates either because they are simultaneous or uncertain at the time the HH makes the consumption decision, or because they are affected by measurement error. Table V shows that the autocorrelation in HH consumption growth rates is negative and consistent with a MA(1) structure that is induced by the presence of measurement error and taste shocks when true consumption changes are not serially correlated.[12] In my specification I follow the previous literature that shows that consistent estimates of the parameters can be obtained if the error is multiplicative in levels and independent of true consumption, returns and instruments.[13]

The instruments I use in the estimation are the exogenous variables and second and previous lags

[11] Muellbauer (1988) and Dynan (2000) show that the correlation between this approximation and the exact expression is higher than 0.90 for reasonable values of the coefficients.

[12] Similar values of the autocorrelation coefficients are obtained by Hayashi (1985b) who analyzes a panel of Japanese household expenditures.

[13] See Browning Lusardi (1996) for an excellent survey of the literature. In the next Section, I provide tests for the validity of these assumptions.

of the marginal tax rate for the HH income bracket, the city-level unemployment rate, inflation rate, mortgage rate, the state-level disposable income growth rate, and some individual variables such as lags of the growth rate of debt, amount charged off, automatic credit line changes, and a credit constrained indicator. The macroeconomic literature shows that aggregate and city-level indicators of economic activity are good predictors of the risk-free rate and city-level consumption growth. The individual variables, like the marginal tax rate and the financial and demographic variables are used as instruments for the household level variables. Given that household consumption is affected by measurement error, I choose not to use past lags of the consumption growth rate as an instrument. The basic assumption required for identification is that the instruments are uncorrelated with the preference shocks, measurement error and expectation errors contained in $\varepsilon_{i,t}$. I test this assumption by performing various tests of overidentifying restrictions, such as the Hansen J test for the entire instrument set and the difference-in-Sargan statistic on the subset of household-specific instruments. Both tests support the validity of the instruments set.

Finally, the standard errors are robust to arbitrary correlation and heteroskedasticity within households. However, the errors are assumed to be uncorrelated across households, once the time dummies account for a common aggregate component.

2.2 Empirical Evidence

Table VI reports the main results of my estimation of (2.7). I regress the household consumption growth on past HH consumption growth, as a measure of internal habit stock, city-level per-capita consumption and its first lag, as a measure of the external habit level, the household-specific interest rate, HH demographic characteristics and seasonal dummies. Column I of Table VI estimates the basic model, and columns II to IV progressively add to the specification family composition, home ownership and occupation. I correct the standard errors for the non-independence of the observations within the same household. I also include controls for the evolution of city-level prices and seasonal dummies in all the regressions.

All versions of the estimates are consistent with the presence of habit formation in household consumption decisions. The specification reported in column IV of Table VI shows that after controlling for demographic and socioeconomic characteristics, the **strength of internal habit** is very high: the coefficient on past household consumption growth is 0.503, which is significant at the 1% level. A coefficient of zero would imply that the HH is not influenced by its past consumption level, and the model would collapse to the standard model used in the literature. A coefficient of one would mean that the HH cares only about the way its consumption compares to the past and not about the absolute level. The **effect of the external habit** is significant at the 5% level. This coefficient captures complementarities in consumption and represents the fraction of city-level consumption that enters the utility function as the reference level to which the household compares itself. The strength of the external habit is 0.29 in the specification reported in Column IV and ranges between 0.258 and 0.295 in the other columns of the table. These findings provide empirical micro foundations to the theories that explain macroeconomic facts and the equity premium puzzle by introducing habit persistence in the utility function. Although, market incompleteness and heterogeneity don't permit to compare this coefficient directly to those in the representative agent models, it is interesting that the coefficients are of the same order of magnitude required in Constantinides (1990) and Campbell and Cochrane (1999) to explain the equity premium puzzle, and in Boldrin et al. (2001) to explain output persistency.

The papers studying habit formation at the aggregate level vary in the instruments used, the horizon analyzed and the specifics of the estimation equation. These papers reach mixed conclusions. Eichenbaum, Hansen, and Singleton (1988) find evidence of habit persistence in leisure choices, but not in consumption, while Heaton (1995) finds evidence of durability at very short horizons and of habit persistence at quarterly frequencies. Conversely, Constantinides and Ferson (1991) find support for habit formation in monthly, quarterly and annual data. More recently, Chen and Ludvigson (2003) finds support for a non-linear form of internal habit. My paper differs from these studies in that it does not use the abstraction of the representative agent, with all the necessary assumptions related to market completeness and the types of heterogeneity that such an abstraction involves. Instead, I test the micro story behind habit formation models. A few papers investigate the implications of habit formation at the micro level. In addition to Dynan (2000), the other such studies are Lupton (2003) and Brunnermeier and Nagel (2007). Rather than directly examining consumption, these papers look at the implications of habit formation for households' stock market investment. In particular, Lupton (2003) estimates a proxy for the habit level and shows that, consistent with the theory, the habit level is negatively related to the share of the household portfolio that is invested in stocks. On the contrary, Brunnermeier and Nagel (2007) analyze the link between idiosyncratic wealth changes and portfolio allocations and conclude against habit persistence.

A natural question is what drives the discrepancy between my findings and **Dynan's (2000)**. The **differences between the two studies** are many: the measure of consumption used, food vs. credit-card expenditures, the annual versus quarterly frequency of observation, the different estimation equation, which in her case doesn't take the variation of the interest rates into account, and the instrument set. In particular, Dynan uses the second lag of income, hours worked and job loss as instruments for HH consumption, while I use individual financial and demographic variables. Table VII presents an attempt to compare the two studies. Columns I and II show that neither the annual frequency nor the presence of external habit seem to cause the difference, since simply re-estimating (2.7) with my data annualized or without external habit still generates coefficients of the same magnitude on the internal habit parameter, although not always statistically significant. On the contrary, Columns III and IV show that excluding the household-specific financial variables from the instrument set, or dropping the HH-specific interest rate generates a drop from 0.60 to 0.10 in the internal habit coefficient, although the coefficient is still statistically significant. Panel B of Table VII shows that both the partial R^2s and the F tests of the first-stage regressions are higher and safe from the Stock and Yogo (2002)'s critique when financial variables are among the instruments. These results indicate that the different findings stem from the availability of HH-specific financial information.[14]

Another result of the paper is my finding that households respond to the price of consumption. The availability of household-specific borrowing rates and financial information allows to study the **sensitivity of consumption to individual-specific interest rates.** I estimate that the short run elasticity of consumption to the borrowing rate, R^C, is -1.876, which is statistically significant at the 1% level. Previous studies had difficulties in getting precise estimates of this parameter, as they were forced to use the after tax risk-free rate, which by nature displays very limited cross sectional variability. The magnitude of my finding is consistent with the results of Gross and Souleles (2002), who estimate that the elasticity of debt, and therefore consumption, to the borrowing rate is -1.3.

The coefficient on the risk-free rate represents the Elasticity of Intertemporal Substitution (EIS). The estimates indicate that an increase of one percentage point in the risk free rate leads on average to an increase of 0.824% in the consumption growth rate. As in most of the literature, due to the small amount of cross sectional variation exhibited by this variable, I cannot precisely measure the effect. Nevertheless, the value of the coefficient is very similar to those obtained in previous studies of individual consumption choices. For some very common specifications of the

[14] Unfortunately, it is not possible to test directly whether the use of food vs. credit card consumption is another reason for the differences, as information credit card expenditures on food is not available. An indirect indication that this might be part of the story comes from the findings in Lupton (2003) who using the PSID finds negative evidence for habit formation in food consumption, but positive evidence in portfolio decisions.

utility function the inverse of the EIS constitutes the coefficient of relative risk aversion. This variable measures the curvature of the value function, which can be interpreted as the willingness to substitute consumption across different states of nature. When this is the case the results imply a coefficient of risk aversion of 1.21, which is reasonable based on the range of values found by other micro-level studies and surveys. I note that the sensitivity of consumption to the borrowing rate and to the risk free rate need not be the same, as the first quantity represents the interest rate at which the HH can borrow and also reflects the effect of liquidity constraints, while the second is a proxy of the interest rate at which it can save.

Consistent with the predictions of economic theory, the coefficients on age and age squared indicate that consumption exhibits a hump-shaped path over the life cycle. I include occupation dummies in the regression, although they are imprecisely measured. Among them, particularly interesting is the positive, although not significant, effect on consumption growth of being self-employed. Since this group can be using the credit card for business related expenses, as a robustness check I re-estimate all the regressions discarding households headed by a self-employed individual and find that the results do not change. Column IV also shows that the effect of home ownership and income bracket on consumption growth is negative, although it is so small as to be almost indistinguishable from zero. I can interpret these findings as mild evidence of precautionary savings or liquidity constraints: households that own a house or which are in a higher income bracket have less need to save for a rainy day and they face fewer limitations on the amount of funds they can borrow. Consequently, they are better able to smooth consumption and, on average, exhibit a lower consumption growth rate.

The **Hansen J statistics** of overidentifying restrictions confirms the validity of the instruments in all the specifications. Further, to make sure that the household-specific instruments are orthogonal to the error term, I compute an additional difference-in-Sargan statistic for this group of instruments. Again, the hypothesis of orthogonality is not rejected. Nonetheless, a too large instrument set could decrease the power of the overidentifying restrictions tests and also bias the coefficients toward the inconsistent ordinary least square estimates. I note that this effect would bias the results against finding habit persistence, because due to measurement error, the OLS coefficients are biased toward zero. In Table VIII, I re-estimate the equation by progressively reducing the number of instruments, to the point in which the system is just identified. For comparison, column I reproduces the first column of Table VI. In Column II, I eliminate the marginal tax rate from the instrument set illustrated above. In Column III I further eliminate the unemployment rate; in Column IV the inflation rate; in Column V all but one lag of aggregate income; in column VI the amount charged off. Finally, in Column VII, I eliminate the automatic credit-line increases. The resulting instrument set exactly identifies the system and is comprised of the first available lag of the mortgage rate, the state-level income growth, the household debt growth rate and the credit-constrained indicator. The results of this test are very encouraging, because the coefficients on ΔC_t, Δc_{t-1} and R^C either remain approximately the same or increase slightly, and continue to be statistically significant.[15] The coefficient that is most sensitive to the reduction of the instrument set is that on the risk-free interest rate, which is sometimes negative, even though very close to zero. However, this coefficient is very imprecisely measured. Also, I report the $R^{2'}$**s from the first stage regressions** for Table VI at the bottom. The value of the adjusted R^2s are very good. The F statistic shows that the coefficients on the instruments are statistically different from zero and way outside the critical values indicated by Stock and Yogo (2002) in relation to concerns of weak instruments. The external habit variable in some of the regressions presents an exception. The value of the F tests are way outside of the critical interval required by traditional theory, but they are sometimes too low to dismiss the possibility that the variable is weakly instrumented.

3 Robustness Checks

In this section, I investigate the possibility that the findings are affected by an omitted variable problem which causes city-level and past household consumption growth rates to enter the regression with an economically and statistically significant coefficient even if per se they are irrelevant for the optimization problem of the household.

Despite the use of lagged information to instrument for the contemporaneous city-level consumption growth rate, there might still be questions about the endogeneity of this variable. A central focus of the social interaction literature is finding out why people who belong to the same group behave similarly. Manski (1993) identifies three distinct reasons why it could happen: the people in the group face the same shocks, they have similar characteristics, which lead them to behave similarly, and social interactions. To further analyze this issue, I use a strictly exogenous random change in the external reference point of the household: the event that somebody in the city wins the **lottery**. I merge the data set with information about lottery winnings from the California lottery and repeat the estimates reported in Table VI but now I drop the city-level information

[15] These results are not due to a particular choice of which instruments to drop. Results available upon request show the robustness of the findings to changes in the order the variables are dropped from the instrument set.

from the instrument set, since the city consumption does not enter the estimation equation anymore and use a lottery dummy instead. Table IX shows that the coefficient on the lottery dummy is economically and statistically significant across all the specifications. The other coefficients are similar to those in Table VI. If we interpret the lottery dummy as a proxy for the change in the reference point, then the coefficient indicates that the fraction of city-level consumption that enters the utility function of the household is approximately 0.22. The estimation equation also includes the lag of the lottery dummy, which displays a negative, although not statistically significant, coefficient. To investigate this aspect further, I re-estimate all the regressions in Table IX without the lagged dummy. I obtain a coefficient of 0.13 for the contemporaneous dummy, which is statistically significant at the 5% level. These findings provide complementary evidence on the effect of social comparisons on consumption decisions. I note that these two measures of the consumption of the external reference group most likely have effect through different mechanisms. In the case of city-level consumption what matters for the HH is the average of the reference group, while in the case of the lottery, it is the consumption of one outlier, like the effect that the lifestyle of celebrities can have on common people's consumption.

Another concern is that unobservable **aggregate shocks** might affect the results. In column VI of Table XI, I add time dummies to the estimation equation to address the case in which the shocks can be decomposed into an economy-wide shock and in an idiosyncratic one. However, if the shocks affect different households to a different extent, time dummies do not address the issue. The reason is that when there is an unexpected shock, HHs make errors in predicting future income. As a consequence, the magnitude of the surprise change in income, and therefore today's consumption, is correlated with lagged changes in income (and lagged consumption). For example, suppose everybody expects a recession that does not happen. It is likely that the lower the predicted income change, the higher will be the surprise. This mechanism generates a negative correlation between prediction and surprise that goes against finding habit formation. Nevertheless, I check for the stability of the coefficients across the major occupations in the data set, as occupation is one of the main reasons why different economic agents might be affected differently by macroeconomic shocks. Table X reports the results. The coefficient on the internal habit parameter is statistically significant in five out of six cases and is also very stable, ranging between 0.425 and 0.546. Following Runkle (1991), I also directly test for the presence of these aggregate shocks by including year dummies in the instrument set. I check whether they are valid instruments by looking at the difference in the J statistics from the estimation with and without these dummies. The difference is distributed as a χ^2_3 and is equal to 7.817. From this evidence I conclude I cannot reject the null hypothesis of no aggregate shocks. This evidence reduces concerns that aggregate shocks are driving the results.

Another concern is that city-level consumption growth influences households' behavior simply because it captures some measure of economic activity that helps them predict their income. In other words, the results could constitute another facet of the "excess sensitivity" of consumption to income documented by Flavin (1981), Campbell and Mankiw (1989) and the many studies thereafter. To exclude this possibility, I add both the **current and future growth rate of income** to the Euler equation and perform a horse race between city-level consumption and income. Since my data doesn't contain household-specific income evolution, I use California state income instead. Below I provide evidence that this choice doesn't influence the conclusions. Column I and II of Table XI show that the coefficient on the external habit variable is unchanged and significant at the 5% level, while income is not statistically significant (p-value=0.319 and 0.434, respectively). The first stage regressions exclude the possibility that this result is due to income being poorly instrumented, since the adjusted R^2 for the income regression is 0.698, higher than that for city-level consumption, past HH consumption and borrowing rate. Although it seems reasonable that aggregate income is a better proxy for individual income than are city sales, perhaps the smaller information contained in the sales is orthogonal to that contained in aggregate income and therefore relevant in forming expectations on future household income. Encouragingly, Attanasio and Weber (1995) find that once demographic variables and labor choice are accounted for, household income doesn't enter the Euler equation in a significant way. While indicative, these arguments alone are unconvincing: for example, Lusardi (1996) using micro data finds evidence of excess sensitivity, so the debate is still open. For this reason, using the PSID, I directly investigate whether city sales provide any information about household income once state income is available:

$$\begin{aligned} \Delta lnincome_{individual,it} &= \alpha + \beta_1 \Delta lnINCOME_{state-level,it} \\ &+ \beta_2 \Delta lnC_{city-level,it} + \beta_3 age_i + \beta_4 age_i^2 + \varepsilon_{it} \end{aligned} \tag{3.1}$$

I also include in the regression dummy variables that capture marital status, occupational choice, seasonal fluctuations and homeownership. The null hypothesis is that once I control for state income, the coefficient on city sales is small and not statistically significant. Table XII confirms this hypothesis and suggests that the reason why income is not significant in Table IX cannot be ascribed to the fact that I use state rather than individual income. Some caution should be exercised in interpreting these results since the time period analyzed spans only from 1997 to 2001. Nevertheless, the various pieces of evidence presented above, taken together, seem to indicate that city-level consumption doesn't proxy for individual income.[16] To further control for local economic activity, in column III and IV of Table XI, I add to the regression the **variation in city-level unemployment rate** and **median zip code house value and rent**: these variables are neither economically nor statistically significant, while the habit persistence coefficients are stable and statistically significant.

In column V of Table XI, I investigate the effect of changing the specification and **add an extra lag of household own consumption** to the regression. Since the theoretical models of habit formation usually involve many lags of HH consumption, while the empirical analyses are limited to a few to insure the availability of enough instruments, it is interesting to see the effect of this simplification on the estimates, especially since HH consumption exhibits autocorrelation over time. The table shows that the coefficient on the second lag of HH consumption growth is statistically significant, although not very big. Most important, the coefficient on the first lag of consumption growth decreases only slightly (from 0.503 to 0.453), while the one on city-level consumption actually increases, and the statistical significance is the same.

Results available upon request also indicate that the findings are robust to changes in the instrument set beyond those reported in Table VIII, and to the inclusion in the regression of various measures of household financial conditions, as the growth rate of debt, and balance transfers indicators. I note that the measure of consumption is not affected by the presence of balance transfers on the credit card, because I separately identify and exclude these sums. However, it could be that HHs that make balance transfers behave differently from others in ways that the variables in the regression do not account for. This robustness tests controls for that possibility. Column VII of Table XI contains, as an illustrative example, the estimates obtained by **adding the household-specific debt growth rate** to the regression. The results on the balance transfer indicators are similar.

4 Alternative Explanations of the Results

Alternative explanations for my findings include the presence of liquidity constraints, precautionary saving motives, and learning about one's own income profile. These phenomena, like internal habit

[16] Similar regressions using the growth rate of individual wealth as the dependent variable provide evidence that city-level consumption doesn't proxy for this variable either.

formation, cause consumption to adjust slowly to predictable changes in income, and therefore induce positive correlation between current and lagged consumption growth rates.

Liquidity constraints consist in borrowing interest rates that are higher than the lending rates and quantity constraints on the amount of funds that an HH can borrow. I can control for both of these aspects in the current setting. One way to test for the presence of liquidity constraints consists in splitting the sample between unconstrained and credit constrained households, in the spirit of Zeldes (1989). For this purpose, I build a measure of the tightness of liquidity constraints based on the ratio of unpaid balance to the credit limit. Figure II plots the frequency distribution of this indicator and shows that the majority of the observations display quite low credit usage, mostly due to the very high credit limit they enjoy. Nevertheless, many cases of very high usage and therefore binding credit constraints are present in the data. Column I and II of Table XIII contain the results of estimating the baseline regression on subsamples constituted by the lowest and highest quartile of the credit constrained indicator distribution. The lowest quartile subsample contains the unconstrained HHs for whom the credit constraint measure is zero; while the highest quartile subsample contains the liquidity constrained households, for which the measure is above 0.76. If liquidity constraints are the cause of the correlation between current and lagged household consumption growth rate, I would expect to see that the past consumption growth rate matters only in the highly constrained subsample. The estimates indicate that this is not the case. The magnitude of the internal habit coefficient is similar and highly statistically significant in both subsamples: 0.565 in the unconstrained sample, and 0.557 in the liquidity constrained sample. The results are robust to the choice of other cutoff points. Moreover, liquidity constraints are more likely to be binding for HHs that experience a low level of cash on hand. Therefore, I test the habit formation Euler equation against a liquidity constraints model by including the lagged growth rate of income in the estimation equation for each subsample, and find that income is not economically nor statistically significant (columns III and IV of Table XIII, respectively). Nevertheless, it slightly decreases the coefficient on the internal habit parameter in the unconstrained case and increases it in the constrained one. This result indicates some sign of liquidity constraints, although they are not strong enough to dismiss the effect of habit persistence. Another interesting result from column I is that *keeping up with the Joneses* is a phenomenon typical of higher-income HHs, who are over-represented in the unconstrained group, while internal habit persistence characterizes all income levels.

Another consequence of binding liquidity constraints is that the HH cannot set the consumption at the optimal level. Therefore I will observe that consumption is too low today relative to tomorrow, and that the lagged income growth rate is negatively related to current consumption growth (Zeldes (1989)). Column V of Table XIII shows that when I include the lagged state income growth rate in the regression for the entire sample, the coefficient of the internal habit parameter equals 0.502, similar to 0.503 in the baseline regression, and is statistically significant at the one percent level. On the contrary, the coefficient on the income growth rate doesn't have the expected negative sign and cannot be statistically distinguished from zero.

I present a further robustness test in Column VI. This test consists of adding the credit-constrained indicator directly in the regression, where it acts as a proxy of the Lagrange multiplier on the borrowing constraint. This indicator is increasing in the level of credit line usage and therefore should display an increasing relationship with the unobserved shadow price of resources. The coefficient on this variable has the expected sign, although it is not significantly different than zero.

All the tests discussed above suggest that the presence of liquidity constraints does not explain away the results. The **precautionary motives story** has implications that are very similar to the liquidity constraints model, and the tests of the significance of lagged income growth and splitting of the sample contained in column I to V of Table XIII speak to this alternative as well. HHs that have low resources available and uncertain future prospects save more and therefore display a higher growth rate of consumption from one period to the next, much like liquidity constrained HHs. Following Dynan (1993) and Carroll (1997), I perform a further test of the precautionary savings alternative by adding the square of household consumption growth rate to the regression, to capture future household-specific uncertainty. To address Carroll's concerns about the validity of group-specific instruments, I include in the instrument set individual variables proxying for household resources, such as lags of the growth rate of debt, amount charged off, automatic credit line changes, and a credit constrained indicator. Column VII of Table XIII shows that the coefficient on the square of consumption innovations is neither economically nor statistically significant and its presence in the regression doesn't affect the estimates of the habit persistence parameters or the interest rates sensitivities.

The latter test also addresses the hypothesis that consumption reacts slowly to expected income changes because the HH learns whether the change is permanent or temporary. To the extent that more volatile consumption indicates more uncertain income and thus more scope for **learning**, we should expect higher consumption growth rates for HHs with more volatile Δc_{it}. Column VII shows that the data do not bear out this hypothesis.

5 Aggregate Implications of the Micro Findings

In addition to characterizing household preferences, does habit persistence matter in the aggregate? And can the representative agent model capture this phenomenon? I investigate this issue by aggregating the micro data, with the caveat that, given the relatively short time dimension of the panel, the results are only suggestive.

Leaving aside the statistical problems of the per-capita-consumption series (Wilcox (1992), Dynan (2000)), the representative agent framework does not capture heterogenous information sets, finite lives, and liquidity constraints. These factors cause the marginal rate of substitution (MRS) to differ across agents and, unless very restrictive conditions are met, the estimates of the preference parameters obtained with such models can be severely biased.[17] While in the theory, there are contrasting conclusions on whether these biases arise, the empirical work on aggregation issues unanimously indicates the crucial importance of accounting for heterogeneity. Brav et al. (2002) show that in order to obtain plausible asset pricing implications the stochastic discount factor needs to be constructed as a weighted average of individual MRS. Attanasio and Weber (1993) illustrate how correct aggregation matters for estimating the EIS and the excess sensitivity of consumption to income.[18]

Exploiting the household-level data, I can estimate the Euler equations using a weighted average of the individual MRS that takes heterogeneity and non linearities into account:

$$\Delta \frac{1}{N} \sum_{i=1}^{N} (\ln c_{i,t}) = \alpha + \beta \Delta \frac{1}{N} \sum_{i=1}^{N} (\ln c_{i,t-1}) + R_t^f + \varepsilon_t \tag{5.1}$$

but also using the per capita MRS, as in the representative agent framework:

$$\Delta \ln \left(\frac{1}{N} \sum_{i=1}^{N} c_{i,t} \right) = \alpha + \beta \Delta \ln \left(\frac{1}{N} \sum_{i=1}^{N} c_{i,t-1} \right) + R_t^f + \varepsilon_t \tag{5.2}$$

[17] Grossman and Shiller (1982) show the conditions under which individual choices can be aggregated to a representative agent with the same type of preferences: quadratic utility, infinitely lived consumers (or dynasties) and homogeneous individual information sets containing all the macro variables. If these conditions are not met, only in the case of complete markets heterogenous consumers are able to pool risks and equate their marginal rates of substitution in every state, despite having different marginal utilities of consumption (Constantinides (1982)). These assumptions are hard to meet in practice.

[18] In the theoretical literature, Mankiw (1986) and Constantinides and Duffie (1996) show that in the presence of idiosyncratic income shocks the Euler equations don't depend only on the growth rate of consumption, but also on its cross sectional variability. On the contrary, Krusell and Smith (1998), show, using calibrations, that in a setting with transitory idiosyncratic income shocks the mean of the wealth distribution is enough to describe the macroeconomic aggregates.

where $c_{i,t}$ is household consumption, and R_t^f is the risk-free rate.[19]

Table XIV contains the estimates from (5.1) and (5.2). As in the micro-level regressions, household consumption and the risk free rate are treated as endogenous and instrumented with the second lag of income growth rate, average mortgage rate, and unemployment rate. Column I shows that when estimating the per capita regression (5.2) there is no evidence of habit persistence: the coefficient on past consumption growth rate is actually negative, albeit not statistically different from zero. Columns II and III illustrate that the estimates display signs of excess sensitivity of consumption to income. This finding is often interpreted as evidence that some households are liquidity constrained or rule-of-thumb consumers, and that per capita regressions are not able to capture the effect of this phenomenon on the cross sectional distribution of consumption growth and therefore the Euler equations. A different picture arises when the properly aggregated quantities are used. Column IV reports the results of the estimation of (5.1). The habit persistence coefficient is now equal to 0.515 and it is statistically significant at the 5% level. This number is quite sizeable, and suggests that, once the correct measure of consumption is used, there is evidence of habit formation also at the aggregate level. Tests of excess sensitivity, reported in Columns V and VI, display no sign of it, using either the contemporaneous or lagged income growth rate. The Hansen J statistic indicates that the overidentifying restrictions don't reject the model, but has low power, due to the relatively big number of instruments in comparison to the number of observations available. I have therefore performed some robustness checks by changing the instrument set and found that the habit persistence coefficient is relatively stable when the properly aggregated measure is used, while it varies from -0.8 to 0.23, and usually not statistically significant, when per capita consumption is analyzed.

Although the short time period analyzed makes these findings only indicative, it is striking that they show a big discrepancy between the two procedures. What are the reasons for this difference? Attanasio and Weber (1993) show that, for any distribution of consumption growth, the difference between $\Delta \frac{1}{N}\sum_{i=1}^{N}(\ln c_{i,t})$ and $\Delta \ln\left(\frac{1}{N}\sum_{i=1}^{N} c_{i,t}\right)$ represents the change in the Theil's measure of entropy and can be approximated by the first four central moments of the cross sectional distribution:

$$\Delta Theil = \Delta E_i \ln c_{i,t} - \Delta ln(E_i c_{i,t}) = \Delta \ln\left(1 + \frac{1}{2}\mu_{2,t} + \frac{1}{6}\mu_{3,t} + \frac{1}{24}\mu_{4,t}\right) + \widetilde{R} \qquad (5.3)$$

[19] Since, by definition, the representative agent framework is populated only by one consumer, the difference between internal and external reference point dissipates and only the dependence of current consumption growth on last period growth can be tested.

where $\mu_{k,t}$ represents the kth central moment and $\widetilde{R}$ is an approximation residual. If the cross sectional moments of consumption growth rates change over time, the two measures will move in an asynchronous way and the per capita regression will be biased. Figure III, panel A, shows the evolution of $E_i \ln c_{i,t}$ and $ln(E_i c_{i,t})$ over time. As the economy slows down, both measures of aggregate consumption fall, but the difference between them widens. Fig. III, panel B, confirms that the two series not only diverge in levels, but have also different growth rates, which causes marginal utilities to differ. To further investigate the origins of the differences, Figure IV plots each of the higher moments of the cross sectional distribution of household consumption included in (5.3): the cross sectional standard deviation of consumption growth displays a strong seasonal pattern (Panel A), the skewness is always positive and tends to increase as time passes and the economy gets deeper into recession (Panel B), and the distribution evolves from fairly flat to quite peaked (Panel C). Following Attanasio and Weber (1993b), Abel and Eberly (2002), columns VII to IX of Table XIV report the results of adding to the regression cross sectional standard deviation, the skewness and the kurtosis, respectively, to see whether this reduces the bias. The estimates indicate that adding these moments strongly improves the significance of the EIS, but has no effect on uncovering the habit persistence revealed by the micro data. These results suggest that omitted demographic variables and other measures of heterogeneity also play a very important role. Unfortunately, it is not possible to check directly the effects of omitting these variables in the aggregate regression.

6 Conclusions

In this paper I use a novel data set on the credit-card expenditures of a representative sample of U.S. households to investigate whether household preferences exhibit habit formation.

The estimation results provide support for this hypothesis: I find that the strength of the external habit, captured by the fraction of the consumption of the reference group that enters the utility function, is 0.290; while the strength of internal habit, represented by household past consumption, is 0.503. My results are robust to the inclusion of various measures of economic activity in the regression, tests for the presence of aggregate shocks, liquidity constraints, precautionary saving motives, and learning. The results are also confirmed by the use of lottery winnings in the city as an independent exogenous measure of the change in the reference point. Another result is the finding that households respond to the price of consumption. I estimate the short run elasticity of consumption to the borrowing rate at -1.876, which is statistically significant at the 1% level.

An open issue is whether the results generalize to a more comprehensive measure of consumption. Unfortunately, there are no statistics available on the type of goods bought with credit cards, so I am unable to check whether these HHs are more likely to exhibit habit formation or to display status seeking behavior. Yet, the analysis in Section 2 shows that this measure of consumption displays the same features that we would expect to see in overall consumption. Further, the statistics on payment methods illustrate the exponential pervasiveness of credit-card payments.

Another question is whether we can distinguish internal habit formation from adjustment costs. Gabaix and Laibson (2002) build a model based on adjustment costs. Their model shows aggregate implications that are observationally equivalent to habit persistence. The individual behavior behind these dynamics is however one of infrequent adjustments that in the micro data would generate a negative, rather than a positive, coefficient on the lagged consumption growth. Nevertheless, it is always possible to model the adjustment costs in a different way and develop a setting that is observationally equivalent to internal habit. For example, quadratic adjustment costs would generate these results. However, while this type of adjustment costs is intuitively appealing for firms making capital investments, it is not clear why a consumer should face psychological or physical costs that would induce her to buy something in installments of increasing amounts over time. On the contrary, search costs and the cost of processing information or paying attention would lead to a framework similar to that of Gabaix and Laibson (2002) and induce negative rather than positive coefficients on past consumption growth.[20] While these remarks are suggestive, one would like to directly test this hypothesis; unfortunately, very detailed data on good purchased and possibly survey evidence would be necessary to address this question.

Finally and most important, these findings provide micro evidence in support of the theories that explain macroeconomic and asset pricing phenomena by introducing habit persistence in the utility function. Most of this work is based on a representative agent framework, and although a direct parallel between the magnitude of the coefficients cannot be made, my findings are line with those the theoretical models require for habit persistence to explain the macroeconomic and asset pricing phenomena in question. Most important for policy making, and to understand what drives consumption, saving and financial investments decisions, is that these preferences are not the artifact of aggregation, but rather a feature of the decision making by the households in the economy.

[20] If we consider liquidity constraints as *sui generis* adjustment costs, we would indeed get positive autocorrelation in consumption growth rates. Section 6 indicates that despite liquidity constraints are important for some householdes they are not teh source of the results.

The results in this paper suggest that this is the case. Section 6 indicates that these phenomena matter at the macro level as well and, once I aggregate individual consumption choices properly, and take into account heterogeneity and nonlinearity of marginal utility, habit persistence is present at the aggregate level as well. Theoretical investigation of the aggregate implications of heterogenous households exhibiting habit persistence and facing liquidity constraints is an important task for future research.

References

[1] Abel, Andrew B. (1990), "Asset Prices under Habit Formation and Catching Up with the Joneses." *American Economic Review Papers and Proceedings*, pp. 38-42.

[2] Abel, Andrew B. (1999), "Risk Premia and Term Premia in General Equilibrium." *Journal of Monetary Economics*, Vol. 43, No. 1, p. 3-33.

[3] Abel, Andrew B., and Eberly, Janice C. (2002), "Investment and q with Fixed Costs: An Empirical Analysis," Working Paper.

[4] Attanasio, Orazio, and Browning, Martin. (1993a) "Consumption over the Life Cycle and over the Business Cycle." *The American Economic Review*, Vol. 85, No. 5, pp. 1118-1137.

[5] Attanasio, Orazio, and Weber, Guglielmo. (1993b), "Consumption Growth, the Interest Rate and Aggregation" *The Review of Economic Studies*, Vol. 60, No. 3, pp. 631-649.

[6] Attanasio, Orazio, and Weber, Guglielmo. (1995), "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *The Journal of Political Economy*, Vol. 103, No. 6, pp. 1121-1157.

[7] Barber, Brad, and Odean, Terrance. (2001) "Boys will be Boys: Gender, Overconfidence, and Common Stock Investment." *The Quarterly Journal of Economics*, 116, 1, pp. 261-292.

[8] Boldrin, Michele, Christiano, Lawrence J., and Fisher, Jonas D.M. (2001) "Habit Persistence, Asset Returns and the Business Cycle." *American Economic Review*, Vol. 91, pp. 149-166.

[9] Brav, Alon, Constantinides, George M., and Geczy, Christopher C. (2002), "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence." *The Journal of Political Economy*, Vol. 110, No. 4, pp. 793-824.

[10] Browning, Martin, and Lusardi, Annamaria. (1996), "Household Saving: Micro Theories and Micro Facts." *Journal of Economic Literature*, Vol. 34, pp. 1797-1855.

[11] Brunnermeier, Markus, and Nagel, Stefan. (2007), "Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation," *American Economic Review*, forthcoming.

[12] Buraschi, A., and Jiltsov, A. (2007), "Habit Formation and Macroeconomi Models of the Term Structure of Interest Rates." Journal of Finance, Vol. 62, pp. 3009-3063.

[13] Calvet, Laurent, Campbell, John Y., and Sodini, Paolo. 2007, "Down or Out: Assessing the Welfare Costs of Household Investment Mistakes", *Journal of Political Economy*, forthcoming.

[14] Calvet, Laurent, Campbell, John Y., and Sodini, Paolo. 2007, "Fight or Flight? Portfolio Rebalancing by Individual Investors" Working Paper.

[15] Campbell, John Y., and Cochrane, John H. (1999), "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *The Journal of Political Economy*, Vol. 107, No. 2, pp. 205-251.

[16] Campbell, John Y., and Mankiw, Gregory N. (1989), "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence." in O.J. Blanchard and S. Fischer (eds.), *NBER Macroeconomics Annual 1989*, Cambridge, Mass. MIT Press, pp. 185-216.

[17] Carroll, Christopher D. (1997), "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, Vol. 112, no. 1, pp. 1-56.

[18] Carroll, Christopher D., Overland, Jody R., and Weil, David N. (2000) "Saving and Growth with Habit Formation," *American Economic Review*, 90(3), 341-355.

[19] Chen, Xiaohong, and Ludvigson, Sydney C. (2003), "Land of Addicts? An Empirical Investigation of Habit-Based Asset Pricing Models," Working Paper.

[20] Choi, James, Laibson, David, and Madrian, Brigitte. (2005a) "Are Empowerment and Education Enough? Under-Diversification in 401(k) Plans" *Brookings Papers on Economic Activity*, 2, pp. 151-198.

[21] Choi, James, Laibson, David, Madrian, , Brigitte, and Metrick, Andrew (2005b), "Passive Decisions and Potent Defaults" In David Wise, ed., *Analyses in the Economics of Aging.* University of Chicago Press, pp. 59-78.

[22] Constantinides, George M. (1982), "Intertemporal Asset Pricing with Heterogeneous Consumers and Without Demand Aggregation." *The Journal of Business*, Vol. 55, No. 2, pp. 253-267.

[23] Constantinides, George M. (1990), "Habit Formation: A Resolution of the Equity Premium Puzzle." *The Journal of Political Economy*, Vol. 98, No. 3, pp. 519-543.

[24] Constantinides, George M., and Ferson, Wayne E. (1991), "Habit Persistence and Durability in Aggregate Consumption: Empirical Tests." *Journal of Financial Economics*, Vol. 29, pp.199-240.

[25] Constantinides, George M., and Duffie, Darrell. (1996), "Asset Pricing with Heterogeneous Consumers." *The Journal of Political Economy*, Vol. 104, No. 2, pp. 219-240.

[26] Davis, Steven J., Kubler, Felix, and Willen, Paul. (2004), "Borrowing Costs and the Demand for Equity over the Lifecycle," Working Paper.

[27] Deaton, Angus. (1992), *Understanding Consumption.* Oxford: Clarendon Press ; New York : Oxford University Press.

[28] Duesenberry, James S. (1949), *Income, Saving, and the Theory of Consumer Behavior*, Harvard University Press, Cambridge, Massachusetts.

[29] Dynan, Karen E. (1993), "How Prudent are Consumers?" *The Journal of Political Economy*, Vol. 101, No. 6, pp. 1104-1113.

[30] Dynan, Karen E. (2000), "Habit Formation in Consumer Preferences: Evidence from Panel Data." *The American Economic Review*, Vol. 90, No. 3, pp. 391-406.

[31] Eichenbaum, Martin S., Hansen, Lars P., and Singleton, Kenneth J. (1988), "A Time Series Analysis of Representative Agent Models of Consumption and Leisure Choice Under Uncertainty." *The Quarterly Journal of Economics*, Vol. 103, pp. 51-78.

[32] Flavin, Marjorie. (1981), "The Adjustment of Consumption to Changing Expectations About Future Income." T*he Journal of Political Economy*, Vol. 89, No. 5, pp. 974-1009.

[33] Fuhrer, Jeffrey C., (2000), "Optimal Monetary Policy in a Model with Habit Formation." FRB Boston Series, paper no. 00-5.

[34] Gabaix, Xavier, and Laibson, David. (2002) "The 6D Bias and the Equity Premium Puzzle" in Benjamin Bernanke and Kenneth Rogoff ed., *NBER Macroeconomics Annual*, vol.16, p. 257-312.

[35] Gerdes, Geoffrey R., and Walton, Jack K. II, (2002), "The Use of Checks and Other Noncash Payment Instrument in the United States." *Federal Reserve Bulletin*, pp. 360-374.

[36] Gross, David B., and Souleles, Nicholas S. (2002), "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data" *The Quarterly Journal of Economics*, Vol. 117, No. 1, pp. 149-186.

[37] Grossman, Sanford J., and Shiller, Robert J. (1982), "Consumption Correlatedness and Risk Measurement in Economies with Nontraded Assets and Heterogenous Information." *The Journal of Financial Economics*, Vol. 10, No. 3, pp. 195-210.

[38] Harvey, Campbell R. (1989), "Time-Varying Conditional Covariances in Tests of Asset Pricing Models." *The Journal of Financial Economics*, Vol. 24, pp. 289-317.

[39] Hayashi, Fumio. (1985*a*), "The Effect of Liquidity Constraints on Consumption: A Cross-Sectional Analysis." *The Quarterly Journal of Economics*, Vol. 100, No. 1, pp. 183-206.

[40] Hayashi, Fumio (1985*b*). "The Permanent Income Hypothesis and Consumption Durability: Analysis Based on Japanese Panel Data." *The Quarterly Journal of Economics*, Vol. 100, No. 4, pp. 1083-1113.

[41] Heaton, John. (1995), "An Empirical Investigation of Asset Pricing with Temporally Dependent Preference Specifications." *Econometrica*, Vol. 63, No. 3, pp. 681-717.

[42] Hong, Harrison, Kubik, Jeffrey D., and Stein, Jeremy C. (2004), "Social Interactions and Stock-Market Participation." *The Journal of Finance*, Vol. 49, no. 1, pp. 137-163.

[43] Hong, Harrison, Kubik, Jeffrey D., and Stein, Jeremy C. (2005), "Thy's Neighbor Portfolio: Word-of-Mouth Effects in the Holdings and Trades of Money Managers." *The Journal of Finance*, Vol. 60, no. 6, pp. 2801-2824.

[44] Krusell, Per, and Smith, Anthony. (1998) "Income and Wealth Heterogeneity in the Macroeconomy" *The Journal of Political Economy*, Vol. 106, No. 5, pp. 867-896.

[45] Laibson, David, Repetto, Andrea, and Tobacman, Jeremy. (2000), "A Debt Puzzle," *Working Paper, Harvard University*.

[46] Lim, Paul J., and Benjamin, Matthew. (2001) "Digging Your Way Out of Debt," *U.S. News and World Report*.

[47] Lupton, Joseph P. (2003), "Household Portfolio Choice and the Habit Liability: Evidence from Panel Data." *Working Paper, Federal Reserve Board*.

[48] Lusardi, Annamaria. (1996), "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets." *Journal of Business Economics and Statistics*, Vol. 14, no. 1, pp. 81-90.

[49] Mankiw, Gregory N. (1986), "The Equity Premium Puzzle and the Concentration of Aggregate Shocks." *Journal of Financial Economics*, Vol. 17, pp. 211-219.

[50] Manski, Charles F. (1993), "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies*, Vol. 60, No. 3, pp. 531-542.

[51] Muellbauer, John N.J. (1988), "Habits, Rationality and Myopia in the Life Cycle Consumption Function" *Annales d'Economie et de Statistique*, Vol. 9, pp. 47-70.

[52] Odean, Terrance. (1998), "Are Investors Reluctant to Realize Their Losses?", *Journal of Finance*, 53, 5, pp. 1775-1798.

[53] Odean, Terrance. (1999), "Do Investors Trade Too Much?", *American Economic Review*, 89, pp. 1279-1298.

[54] Runkle, David E. (1991), "Liquidity Constraints and the Permanent-Income Hypothesis." *Journal of Monetary Economics*, Vol. 27, pp. 73-98.

[55] Stock, James H., and Yogo, Motohiro, (2002), "Testing for Weak Instruments in Linear IV Regression." *NBER Working Paper*, No. T0284.

[56] Veblen, Thorstein. (1899), *The Theory of the Leisure Class. An Economic Study of Institutions*, Random House.

[57] Wilcox, David W. (1992), "The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work." *The American Economic Review*, Vol. 82, No. 4, pp. 922-941.

[58] Zeldes, Stephen P. (1989), "Consumption and Liquidity Constraints: An Empirical Investigation" *The Journal of Political Economy*, Vol. 97, No. 2, pp. 305-346.

[59] Zinman, Jonathan. (2004) "Why Use Debit Instead of Credit? Consumer Choice in a Trillion Dollar Market." Working Paper.

Appendix I

IA Sample Selection

The data set used in the analysis is a random sample of credit card accounts, active and not delinquent as of July 1999. Each account is followed for the period between the third quarter of 1999 and the third quarter of 2002. The unit of observation is the account in a given quarter. The card could be used by more than an individual and therefore I will refer to the decision maker behind it as household (HH).

Although the original sample covers the entire U.S. territory, I restrict the analysis to California, the only state for which retail sales are available at the city level and quarterly frequency. For the same reason, although the credit card data are available at monthly frequencies, I construct quarterly variables, a choice that has also the benefit of reducing the noise that plagues individual monthly consumption data. In constructing the sample, I exclude people whose accounts are inactive and those that don't use the card very often, in order to obtain a more meaningful measure of consumption. For an account to be in the sample the expenditure can never fall below \$50 in any given quarter. The choice of this cutoff is meant to compromise between sample size and representativeness. According to Lim and Benjamin (2001), the average transaction amount on a credit card is \$87, 112% higher than that made in cash. Following the literature, I also exclude retired account holders and people living in military areas, because their expenditures are influenced by special conditions and required specific modelling that is outside the scope of this paper. After this selection procedure, the sample contains 2,674 accounts.

Extreme outliers. Some of the series present a lot of variability. In order to address this issue, part of the previous literature excludes observations in which the growth rate of consumption is too large. Zeldes (1989) for example excludes observations in which the growth rate of expenditures is bigger than 1.1; Brav et al. (2002) exclude observations for which the growth rate goes from less than 1/2 to more than 2 or it is bigger than 5; Vissing-Jorgensen (2002) excludes observations for which the growth rate goes from less than 0.2 or more than 5. I keep all the observation because the distributions of expenditure growth rates is very symmetric and the cutoff points seem somewhat arbitrary. I have also tried to exclude observations whose growth rate is bigger than 5 and I get very similar results. Similarly, I have tried winsorizing the city retail sales data at the 5% and 95% cutoff points and obtained similar estimates.

Missing data. I exclude from the sample, for a given quarter, any observation with missing data on any of the variables included in the basic regression (including income). One of the reasons for missing observations is bankruptcy. In particular, 0.92% of the accounts ends up bankrupt at the end of the sample period. This makes the panel unbalanced, but doesn't biased the results, since in the estimation I control the financial health of the account through debt outstanding, amount charged off and credit constrained indicators.

IB Comparison with the U.S. Census 2000: Analysis of Demographic Characteristics and Borrowing Behavior

Table AI and Figure AI compare the distributions of income and age in the data set to those in the U.S. Census 2000. Both the distributions are very alike. The main discrepancy is due to the fact that HHs in the lower range of income and individuals in very young or old age are under-represented in the credit card data set. However, this phenomenon is to be expected given the nature of the variables analyzed: as documented in the Survey of Consumer Finances, HHs with lower income and whose head is younger than 35 or older than 65 are less likely to hold bank-type credit cards. For the same reason, the proportion of HHs that own the house in which they live is 74.88% in my data set, compared to the 66.2% in the Census. Finally, the percentage of married people is slightly less than the national average, 44.76% versus 52.4%, even though the sizeable amount of missing data, 35.27%, could be the cause of the difference. Table AII illustrates the breakdown of the data set by occupation: the "professional/technical" category is the most widely represented, accounting for 13.76% of the observations, followed by the "administrative/managerial" with 7.59%. An interesting feature related to this variable is that 2.51% of the HHs in the sample are headed by self-employed individuals.

The sample compares well to the U.S. data with regard to HHs borrowing behavior as well. Panel A of Table AIII contains a comparison between my data set and a large multi-issuer credit card data set covering the period between 1995 and 1998 and used by Gross and Souleles (2002). Indebtedness is highly skew in both data sets; while the median debt outstanding is $0 and $70 respectively, the average debt outstanding is $1,486 for my data set and similar for Gross and Souleles. The figure more than doubles if we consider only those HHs that have debt outstanding, reaching $3,400. Both the mean and median credit limits are higher in my data set, probably reflecting the increase in credit availability over the period. Similarly, average and median interest rates are lower in the latter period, due to a trend in the reduction of interest rates that is reflected in the statistics on the rate changes as well.

Table AIII Panel B provides a comparison with the Survey of Consumer Finances. The samples are similar with regard to the percentage of people not paying the balance in full at the end of the month, which is estimated to be 44.4% in the SCF and 45.75% in my sample. Unfortunately, the different classification criteria cause some difficulties in comparing the statistics, as the SCF reports the total debt outstanding on all the credit cards available for the HH. Moreover, this survey has been proven to suffer from under-reporting of debt.

IC Credit Cards Expenditures as a Measure of Consumption

Data from the Consumer Expenditure Survey (CEX), show that in 2001 the average annual expenditure was $38,045. Following the selection criteria of Attanasio and Weber (1995), I construct an estimate of consumption of non-durable goods and services by excluding from the aggregate consumer durables, housing, health and education expenditures. This gives a quarterly expenditure between $4,765 and $5,053, depending on the classification criteria. According to the statistics presented above, 24% of this amount, between $1,143 and $1,212, is paid by credit card. The average expenditure on the accounts I analyze is $701.19. This means that on average the expenditures analyzed cover between 13.88 and 14.72 percent of non-durable goods and services consumption. This figure represents good news, because it indicates that on average the HHs in my data set use this credit card conspicuously and are therefore offering a good measure of their expenditures.

Appendix II

In this Section I model the consumption decisions of a household that borrows through a credit card and saves in a savings account. The specification proposed is very flexible and subsumes both the specification used in Constantinides (1990) and Campbell and Cochrane (1999), where the habit variables enter the utility function as a difference, i.e. $U(c_{i,t}, h_{i,t}, H_{i,t}) = U(c_{i,t} + \zeta h_{i,t} + \alpha H_{i,t})$, and the specification used in Abel (1990, 1999), where the habit variables enter the utility function as a ratio, i.e. $U(c_{i,t}, h_{i,t}, H_{i,t}) = U(\frac{c_{i,t}}{h_{i,t}H_{i,t}})$.

In each period t, household i chooses consumption $c_{i,t}$ and the amount $P_{i,t}$ of the credit card balance to pay back, with the objective of maximizing the expected value of a lifetime utility function:

$$\underset{\{c_{i,t};P_{i,t}\}_0^{T-1}}{Max} E_t \sum_{t=0}^{T-1} \beta^t U(c_{i,t}, h_{i,t}, H_{i,t}, \Theta_{i,t}) \tag{0.1}$$

subject to:

$$A_{i,t+1} = (A_{i,t} + Y_{i,t} - P_{i,t})R^f_{i,t} \tag{0.2}$$

$$B_{i,t+1} = (B_{i,t} - P_{i,t} + c_{i,t})\widetilde{R}^C_{i,t} \tag{0.3}$$

$$c_{i,t} \leq A_{i,t} + Y_{i,t} + \overline{B}_i - B_{i,t} \tag{0.4}$$

$$h_{i,t+1} = \zeta c_{i,t} \tag{0.5}$$

$$A_{i,t} \geq 0, \quad B_{it} \leq \overline{B_i}, \quad c_{i,t} > 0 \tag{0.6}$$

where $h_{i,t}$ is the level of the habit stock that HH i derives from its own past consumption, $H_{i,t}$ is the consumption of the reference group and $\Theta_{i,t}$ are demographic and socioeconomic characteristics; $A_{i,t}$ is the amount of money in the savings account at the beginning of period t, which earns the risk-free rate, $R^f_{i,t}$; $Y_{i,t}$ is the income realization; $B_{i,t}$ is the credit card balance at the beginning of period t, $\overline{B}_i$ the credit limit on the card, and $\widetilde{R}^C_{i,t}$ the gross interest rate the HH is charged on the balance outstanding. The value of $\widetilde{R}^C_{i,t}$ depends on whether or not the HH pays the balance in full and is given by the following expression:

$$\widetilde{R}^C_{i,t} = \begin{array}{ll} R^C_{i,t} > R^f_{i,t} > 1 & \text{if } P_{i,t} < B_{i,t} \\ 1 < R^f_t & \text{if } P_{i,t} \geq B_{i,t} \end{array}$$

Equation (0.2) describes the evolution of the savings account balance: at the end of period t, the account contains the initial funds, plus that period income, minus the credit card payment. Analogously, (0.3) shows that the credit card balance at the beginning of period t+1 consists of the unpaid balance from the previous period, $B_{i,t} - P_{i,t}$, and any new expenditure charged on the card. Equation (0.4) states that consumption cannot exceed the sum of the resources on hand and the unused part of the credit line. Equation (0.5) describes the evolution of the habit stock of the HH, which is assumed to depend on last period consumption only in order to make the empirical analysis more tractable. The specification of the evolution of local aggregate consumption and the demographic characteristics is not necessary for the derivation of the optimal consumption rule and is left for later.[1] Finally, equations (0.6) represents the no short sales constraint, the borrowing limit, and the condition that consumption must be strictly positive, respectively.

I solve the maximization problem delineated above by expressing it in recursive form. To make the model more tractable and help focusing on the key elements of the problem, I make the following simplifying assumptions.

Assumption 1 *Household income lies between* $\underline{Y_i}$ *and* $\overline{Y}_i$ *and evolves according to* $Y_{i,t} = Y_i + \epsilon_{i,t}$*, where* Y_i *is HH's i mean income and* $\epsilon_{i,t}$ *is i.i.d. and has mean 0 and variance* σ^2*.*

This assumption implies that in any given period the income realization doesn't depend on previous history. Despite being restrictive, it is aimed at capturing the presence of uninsurable income shocks that causing fluctuations around the income level expected by the household or, alternatively, random unexpected expenses. Exploring the effect of more general income processes constitutes material for future research.

Assumption 2 *The credit limit is set to be* $\overline{B_i} = \sum_0^T \frac{\overline{Y_i}}{R^f_{i,t}}$*, the present value of an income stream in which the highest possible income* $\overline{Y_i}$ *is realized in every period.*

Assumption 3 *There is no default.*

Notice that the optimal payment rule for the HH is to employ all the resources on hand to pay back as much as possible of the outstanding credit card balance. The intuition is that since the household will be charged for $c_{i,t}$ only in period t+1, paying the balance doesn't subtract any resource from current consumption. On the contrary, since there is a wedge between borrowing and lending rates, a bigger payment increases the wealth of the HH because it reduces the amount on which it is charged the high interest rate. Therefore, the only choice variable left for the household is $c_{i,t}$.

The solution of the problem delineated in (0.1)-(0.6) is complicated by the presence of a discontinuity in the value function at the point in which the HH switches from lending at the risk-free rate to borrowing at a higher interest rate. To solve this problem I use as state variable to be the amount of resources on hand, $Z_{i,t} = A_{i,t} + Y_{i,t} - B_{i,t}$. The problem can be then defined over two different regions: the first, region B, in which the HH has negative resources on hand and thus borrows at a high interest rate; and the second,

[1] $H_{i,t}$ and $\Theta_{i,t}$ are therefore suppressed in the rest of the section for notational simplicity.

region $\overline{B}$, in which it has enough resources to pay the balance in full, put some money into the savings account and earn the risk-free rate. Since the point of discontinuity coincides with the point in which the HH switches regions, within each region continuity and differentiability are satisfied and standard solution techniques can be applied.

More precisely, the problem of a HH that in period t is not borrowing can be expressed as follows:[2]

$$V_t^{\overline{B}}(Z_t,h_t) = \max_{\{c_t\}} u(c_t,h_t) + \beta\left[\int_{-Z_tR_t^f+c_t}^{\overline{Y}} V_{t+1}^{\overline{B}}(Z_{t+1},h_{t+1})f(Y)dY + \int_{\underline{Y}}^{-Z_tR_t^f+c_t} V_{t+1}^{B}(Z_{t+1},h_{t+1})f(Y)dY\right] \\ +\lambda_t(Z_t+\overline{B}-c_t) \tag{0.7}$$

$$\text{subject to:} \qquad Z_{t+1} = Z_tR_t^f + Y_{t+1} - c_t \tag{0.8}$$

$$h_{t+1} = \zeta c_t \tag{0.9}$$

where $V_t^{\overline{B}}(Z_t,h_t)$ is the value function; the expression in brackets is the expectation of the future value function, taken with respect to next period income realization; and the last expression is the product of the Lagrange multiplier and the resource constraint. Equations (0.8) and (0.9) describes the evolution of the state variables. The cutoffs of the two integrals show that whether next period the HH falls in one region rather than the other is determined by the resources accumulated from the previous period, $A_{i,t}$, and the realization of the income shock: if $Y_{i,t+1}$ is high enough that the HH is able to repay the credit card balance in full, then it keeps staying in the non-borrowing region; otherwise, it will start borrowing. When the HH makes its consumption decision in period t, it knows that this choice will determine the credit card balance and therefore the likelihood of staying in the non-borrowing region. As time passes and the HH accumulates or decumulates resources, the probability of falling into a certain region gets smaller and smaller, independently from the income realization and the consumption choice.

The formulation of the problem of a HH that enters period t with an unpaid balance is very similar to (0.7).The only difference is the law of motion of the state variable:

$$V_t^{B}(Z_t,h_t) = \max_{\{c_t\}} u(c_t,h_t) + \beta\left[\int_{(-Z_t+c_t)R_t^C}^{\overline{Y}} V_{t+1}^{\overline{B}}(Z_{t+1},h_{t+1})f(Y)dY + \int_{\underline{Y}}^{(-Z_t+c_t)R_t^C} V_{t+1}^{B}(Z_{t+1},h_{t+1})f(Y)dY\right] \\ +\lambda_t(Z_t+\overline{B}-c_t) \tag{0.10}$$

$$\text{subject to:} \qquad Z_{t+1} = Y_{t+1} - (-Z_t+c_t)R_t^C \tag{0.11}$$

$$h_{t+1} = \zeta c_t \tag{0.12}$$

The maximization problem delineated above generates the following Euler equations for a HH that at period t is in the non-borrowing region:

$$u^c(c_t,h_t) = \beta\int_{-Z_tR_t^f+c_t}^{\overline{Y}}[u_{t+1}^cR_{t+1}^f + \beta R_{t+1}^f\zeta E_Y(u_{t+2}^h) - \zeta u_{t+1}^h]f(Y)dY \\ +\beta\int_{\underline{Y}}^{-Z_tR_t^f+c_t}[u_{t+1}^c + \beta\zeta E_Y(u_{t+2}^h) - \zeta u_{t+1}^h]f(Y)dY \tag{0.13}$$

and for one that is in the borrowing region:

$$u^c(c_t,h_t) = \beta R_t^C\int_{(-Z_t+c_t)R_t^C}^{\overline{Y}}\left[[u_{t+1}^cR_{t+1}^f + \beta R_{t+1}^f\zeta E_Y(u_{t+2}^h)]R_t^C - \zeta u_{t+1}^h\right]f(Y)dY \\ +\int_{\underline{Y}}^{(-Z_t+c_t)R_t^C}\left[[u_{t+1}^c + \beta\zeta E_Y(u_{t+2}^h)]R_t^C - \zeta u_{t+1}^h\right]f(Y)dY \tag{0.14}$$

These results are very intuitive. The household decides how much to consume today versus tomorrow by weighting future utility and different interest rates by the probability that it will actually face them. If

[2]From now on I suppress the subscript i, for notational simplicity.

the HH consumes \$1 less today it looses $u^c(c_t)$ and gains the following: next period the credit card balance will be \$1 lower and so one more dollar will be available for consumption, yielding a utility of $u^c(c_{t+1})$; if, in addition to this gain, the income realization is high enough that the HH is able to repay the balance in full, it will earn the gross risk-free rate on the dollar moved through time and the utility will be $u^c(c_{t+1})R^f_{t+1}$. These events realize with probabilities $\int_{\underline{Y}}^{(-Z_t+c_t)R^C_t} f(Y)dY$ and $\int_{(-Z_t+c_t)R^C_t}^{\overline{Y}} f(Y)dY$, respectively.

Analogously, in the case of (0.14), consuming one dollar less today means that the credit card balance next period will be R^C_t dollars less.[3] The utility deriving from this intertemporal transfer will be $u^c(c_{t+1})R^C_t$ if the HH doesn't have enough resources to pay the balance in full in period t+1, and $u^c(c_{t+1})R^C_t R^f_{t+1}$ in case it does and can invest the dollar charged on the credit card at the risk-free rate.

The presence of the habit stock generates an additional effect due to the fact that when the HH consumes one dollar less today it increases tomorrow's utility not only directly, but also by decreasing the habit level. This effect is given by $-\zeta u^h_{t+1} > 0$. [4] The extra dollar consumed in period t+1 will increase the habit stock of period t+2 at a cost in term of utility given by $\beta\zeta E_Y(u^h_{t+2})$ or $\beta\zeta R^f_{t+1} E_Y(u^h_{t+2})$, depending on the HH asset position.

It is also possible to draw a parallel with the Euler equation traditionally obtained in the literature, where the HHs are assumed to be able to borrow and lend at the risk-free rate. In particular, if in case $\overline{B}$ we assume that the resources on hand are high enough that the HH will always pay the balance in full, then (0.13) collapses to the usual Euler equation:

$$u^c(c_t) = \beta \left\{ \int_{\underline{Y}}^{\overline{Y}} u^c(c_{t+1}) R^f_{t+1} f(Y)dY \right\} \tag{0.15}$$

References

[1] Abel, Andrew B. (1990), "Asset Prices under Habit Formation and Catching Up with the Joneses." *American Economic Review Papers and Proceedings*, pp. 38-42.

[2] Abel, Andrew B. (1999), "Risk Premia and Term Premia in General Equilibrium." *Journal of Monetary Economics*, Vol. 43, No. 1, p. 3-33.

[3] Attanasio, Orazio, and Weber, Guglielmo. (1995), "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *The Journal of Political Economy*, Vol. 103, No. 6, pp. 1121-1157.

[4] Brav, Alon, Constantinides, George M., and Geczy, Christopher C. (2002), "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence." *The Journal of Political Economy*, Vol. 110, No. 4, pp. 793-824.

[5] Campbell, John Y., and Cochrane, John H. (1999), "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *The Journal of Political Economy*, Vol. 107, No. 2, pp. 205-251.

[6] Constantinides, George M. (1990), "Habit Formation: A Resolution of the Equity Premium Puzzle." *The Journal of Political Economy*, Vol. 98, No. 3, pp. 519-543.

[7] Gross, David B., and Souleles, Nicholas S. (2002), "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data" *The Quarterly Journal of Economics*, Vol. 117, No. 1, pp. 149-186.

[8] Lim, Paul J., and Benjamin, Matthew. (2001) "Digging Your Way Out of Debt," *U.S. News and World Report.*

[3] Since in period t the HH was not able to pay the balance completely it was charged interest on both the unpaid portion of the balance and any new purchases.

[4] Recall that an increase in the habit stock decreases utility and therefore $u^h < 0$.

[9] State Board of Equalization. *Taxable Sales in California,* various issues.

[10] State Board of Equalization. (1999) *Sales and Use Taxes: Exemptions and Exclusions*, No. 61.

[11] Vissing-Jorgensen, Annette. (2002), "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution." *The Journal of Political Economy*, Vol. 110, No. 4, pp. 825-853.

[12] Zeldes, Stephen P. (1989), "Consumption and Liquidity Constraints: An Empirical Investigation" *The Journal of Political Economy*, Vol. 97, No. 2, pp. 305-346.

Fig. I
Panel A

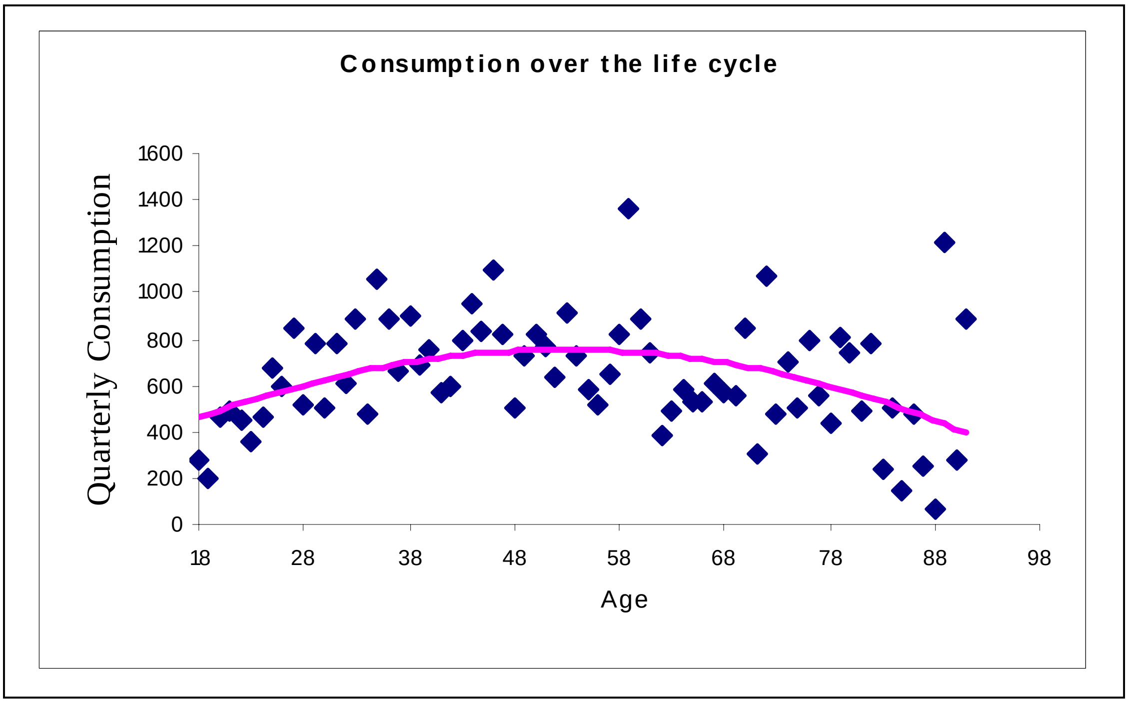

Panel B

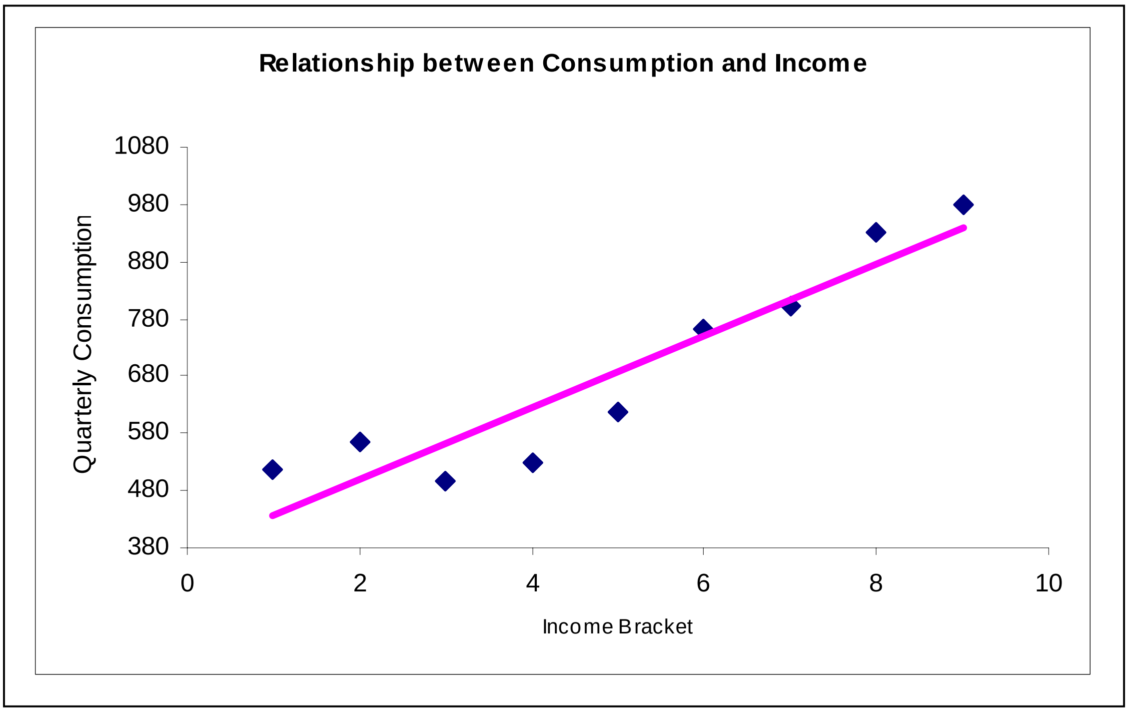

Fig. I – Plot of average quarterly credit-card expenditures against age (Panel A) and income bracket (Panel B).

Fig II

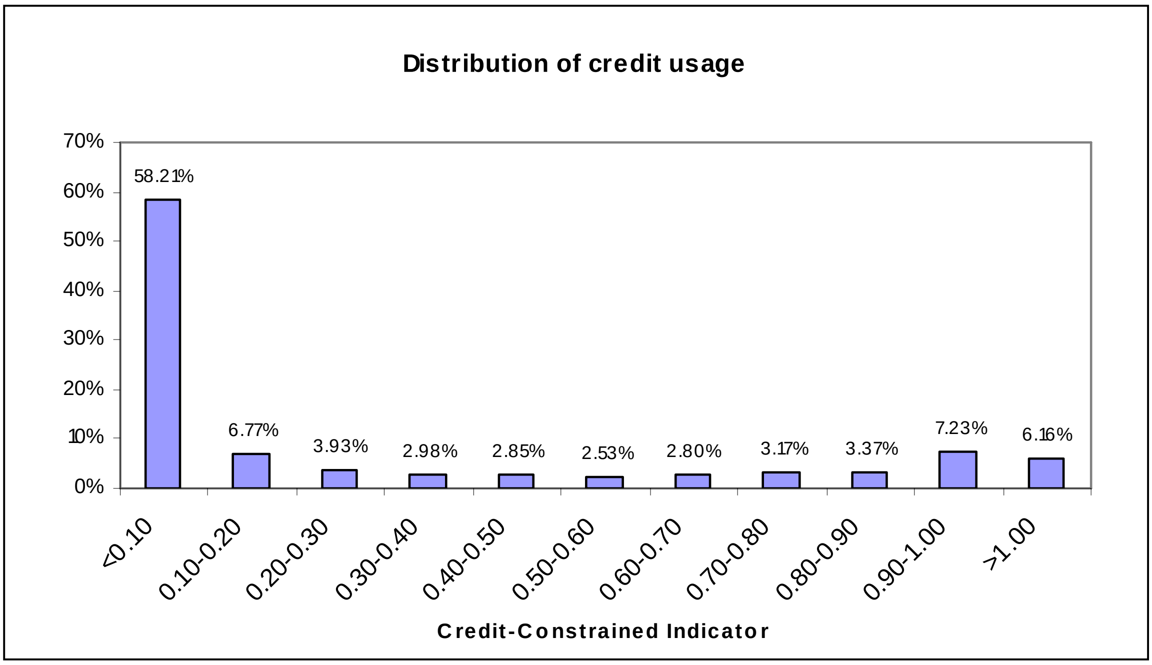

Fig . II – Frequency distribution of the credit-constrained indicator. This variable is calculated as the ratio between unpaid balance and credit line.

Fig. III
Panel A
Correctly aggregated and per-capita aggregate consumption

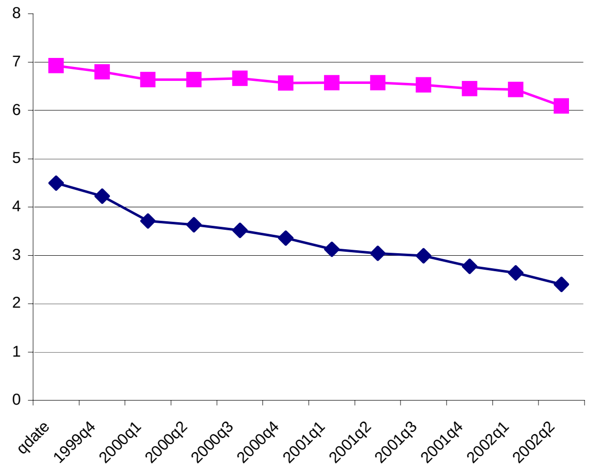

Panel B
Growth rates of correctly aggregated and per-capita aggregate consumption

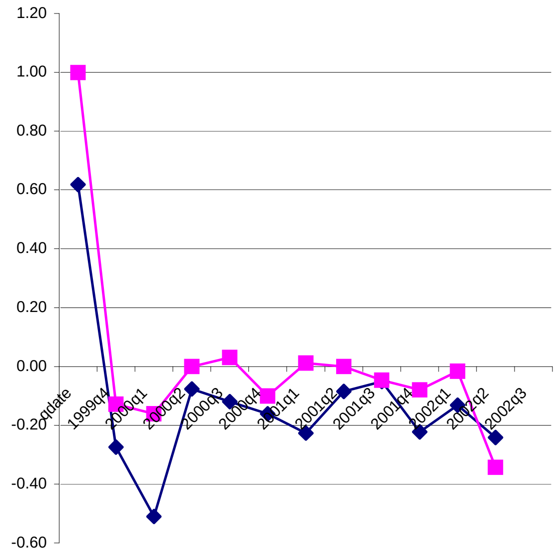

Fig. III – Aggregate consumption series, constructed from the credit card data. The blue line represents the correctly aggregated variable, as in (7.1), and the pink line represents per-capita consumption, as in (7.2).

Fig. IV
Evolution of standard deviation, skewness and kurtosis of the cross sectional distribution of household consumption growth rates

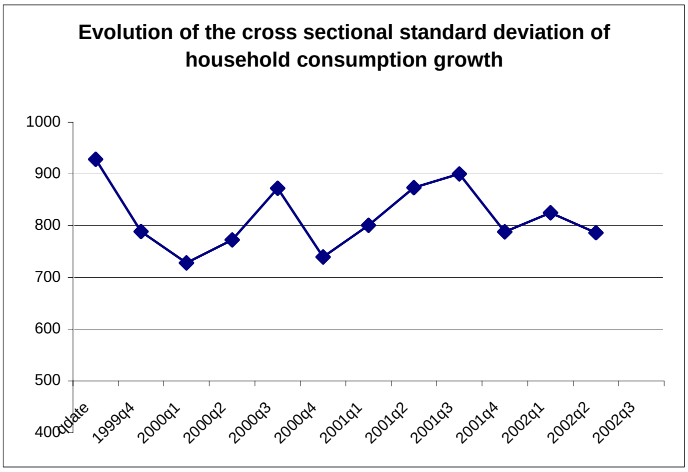

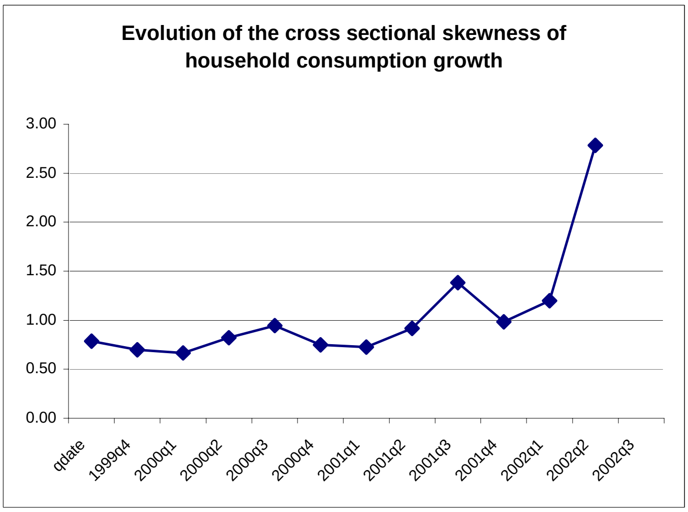

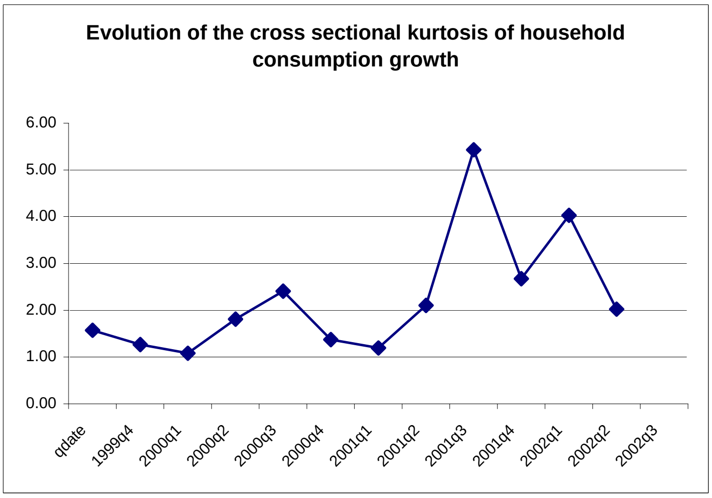

Table I
Data Sources and Variables Description

Variable	Description	Source
Δc_{it}	**Household consumption growth rate** Amount charged on the credit cardduring the quarter for purchases and cash advances. I construct growth rates by taking the difference in the logarithms of this variable between time t and (t-1)	Credit card data set
ΔC_{jt}	**External Reference Point: Citi-level consumption growth rate** city level quarterly taxable sales, defined as "the dollar amount of California retail transactions, excluding those transactions specifically exempt from the California Sales and Use tax". This measure excludes prescription medicines, sales of nontaxable items such as some food products consumed at home and prescription medicines, and taxable transactions disclosed by BOE audits. A detail description is available in "Publication Number 61, Sales and Use Taxes: Exemptions and Exclusions" (March 2003). People not living in a city are associated to the nearest one by geographic matching based on distance between zip code's centroids. The average distance between the centroids of the originating and matching point for HHs not living in a city is 1.1 miles, while the median is 0.4 miles. I construct per capita sales by dividing this quantity by the city population available from Current Population Survey at annual intervals.	State Board of Equalization (BOE)
	Financial Variables	
R^C_{it}	**Household-specific Borrowing Rate** Interest charged on the debt outstanding on the credit card, calculated by taking the ratio of the total charges incurred in a period (finance charges, late charges and over the limit charges) to the balance outstanding.. It is different from the stated APR, because it takes into account compounding and the effective period of time over which the money is borrowed. This results in a more accurate measure of the cost of borrowing	Credit card data set
R^f_t	**Risk Free Rate** 3-month T-bill rate. It represents the risk free at which the HHs are supposed to invest the funds that are left after paying the balance on the credit card.	Federal Reserve Bank of St. Louis (FRED)
Credit constrained indicator_{it}	Ratio of debt outstanding to credit limit.It captures credit availability	Credit card data set
Δdebt_{it}	**Household-specific Debt Growth Rate** Amount of the credit card balance unpaid and on which the HH is charged interest. Two measures of debt are considered: debt, which represents the overall debt on the card, including balance transfers from other cards. And debt2, which excludes balance transfers and assumes that of any debt outstanding the HH first repays the newly generated debt and only after that the one transferred from other cards.	Credit card data set
$\Delta \text{credit\_line}_{it}$	Household total credit limit on the card	Credit card data set
$\Delta \text{charged\_off}_{it}$	**Amount charged off.** Amount of debt outstanding that the credit card issuer will not be able to recoup and thus writes down as a loss on the account. The reason is bankruptcy, both formally recognized by a court or informal. According to this measure 2.68% of the observations have a positive amount charged off, corresponding to 0.92% of the HHs.	Credit card data set

Balance Transfer$_{it}$	Amount of debt outstanding on another credit card and transferred to this one, or transferred from this card to another. The total number of balance transfers is 562, equal to 1.6% of the observations. The measure of consumption doesn't include balance transfers.	Credit card data set
	Demographic Variables	
Age	Age of the main account holder as of July 1999.	Credit card data set
Marit_status	Dummy variable equal to one if the main account holder is married as of July 1999	Credit card data set
Homeowner	Dummy variable equal to one if the HH owns the house it lives in.	Credit card data set
Income bracket	Income category the HH belongs to	Credit card data set
Occupation dummies	dummy variables indicating the occupation of the primary card holder	Credit card data set
Self_empl	Dummy variable equal to one if the HH head is self-employed.	Credit card data set
Marginal Tax Rate	Marginal tax rate faced by a family or single individual in a given income bracket	Internal Revenue Service (IRS)
	Zip Code-Level Economic Variables	
Median House Value	Median value of the house for a specific sample of owner-occupied houses in the 2000 U.S. Census. The value is the respondent's estimate of how much the property would sell for if it were for sale.	2000 U.S. Census
Median Rent	Median rent asked in the zip code area. No adjustment is made for the inclusion of utilities and fuel.	2000 U.S. Census
	City-Level Economic Variables	
Mortgage Rate$_t$	Average mortgage rate faced by people living in the city in a given quarter.	American Chamber of Commerce Research Association (ACCRA)
Δln U_rate$_t$	Quarterly average of the monthly MSA level unemployment rate	Bureau of Labor Statistics (BLS)
Inflation Rate$_t$	**All Urban Consumers Price Index, base 182:84, not seasonally adjusted.** The quantities are deflated using the BLS Consumer Price Index of the MSA to which they belong (Los Angeles, San Francisco-San Jose or the index for the West Region).	Bureau of Labor Statistics (BLS)
Lottery	Dummy variable equal to one if the HH lives in a city in which there has been a lottery winning in that quarter	California Lottery
Prop. annuity	fraction of the winners who opt to receive the prize in installments, rather than in a lump-sum payment	California Lottery
Total winnings	Total $ winnings in the city where the HH lives, in a given quarter	California Lottery
	State-Level Economic Variables	
Δlnincome$_t$	Growth rate of quarterly per-capita disposable income in current dollars	California Department of Finance

Table II
Summary Statistics

	Mean	Median	Std. Dev.
Individual Consumption	701.19	114.92	1596.92
growth rate	-0.12	0.00	2.73
Aggregate Consumption	2367.22	2239.58	975.53
growth rate	0.01	0.03	0.19
R^{C}_{it}	16.95%	0.00%	114.65%
R^{f}	3.88%	4.42%	1.71%
Inflation rate	0.79%	0.87%	0.46%
DEMOGRAPHIC VARIABLES			
Age	46	45	15
Marital Status (% married)	44.76%		
Home Owner (%)	74.88%		
Marginal Tax Rate	28.44%	28.00%	7.37%
CHARACTERISTICS of the CONTRACT			
Unpaid balance	$1,060.05	$0.00	$2,048.60
Credit Line	$8,823.51	$10,000.00	$3,733.95
Charged Off	$152.33	$0.00	$1,058.07
dummy	0.027	0.000	0.160
Credit-Constrained Indicator	0.33	0.07	0.41
Balance Transfers	$132.64	$0.00	$1,004.97
dummy	0.016	0	0.126
ZIP CODE-LEVEL DATA			
Median House Value	$258,980.80	$222,700.00	$146,479.10
Median Income	$54,618.55	$50,825.00	$19,602.71
Rent	$791.10	$752.00	$245.95
Unemployment Rate	6.33%	5.49%	3.62%
CITY-LEVEL DATA			
Mortgage Rate	7.32%	7.21%	0.59%
Unemployment rate	5.07%	4.43%	3.22%
Lottery	0.0501906	0	0.2183412
Prop. Annuity	0.0021785	0	0.0466238
Total winnings	156133.4	0	3095676
STATE-LEVEL DATA			
Personal Income	$26,792.17	$26,906.59	$895.74

Table III
Mean and Std. Dev. of Individual Consumption

In this Table I compare the mean and standard deviation of my measure of consumption to data from the publicly available data sets traditionally used in the literature: the Panel Study of Income Dynamics (PSID), the Consumer Expenditure Survey (CEX) and aggregate data.

		Credit Card Panel		Comparison data set	
		Mean	Std.Dev.	Mean	Std.Dev.
Individual Data					
	All	-0.120	2.730		
	-3.3/3.3	0.005	0.870		
	-1.1/1.1	-0.002	0.368	0.000	0.32*
Cross Sectional Data					
	All	0.014	0.330	-0.01*	0.06*
Aggregate Data					
	All	0.000	0.013	0.003***	0.009***

* Comparison data set: PSID, data from Zeldes (1989)

* Comparison data set: CEX, data from Brav et al. (2002)

* Comparison data set: aggregate data, from Constantinides et al. (1991)

Table IV
Aggregate Local Consumption
Comparison with NIPA Aggregate Consumption

In this table I compare the city-level sales data that I use to construct the measure of the external reference point to the aggregate data traditionally used in the literature: the aggregate consumption from the National Income and Product Accounts (NIPA). The main difference between the two measures is the lack of housing services in the California aggregate sales.

	California Aggregate Sales	Non-durables and Services (NIPA)
Mean	3674.394	5659.72
Median	3669.729	5737
Std. Dev.	402.2181	404.1349
Correlation Coefficient	0.8583	

For a definition of the California aggregate sales see Table I

Table V
Autocorrelations of Household Consumption

	Δc_t	Δc_{t-1}	Δc_{t-2}	Δc_{t-3}
Δc_t	1			
Δc_{t-1}	-0.3863	1		
Δc_{t-2}	-0.0281	-0.3866	1	
Δc_{t-3}	-0.0212	-0.0317	-0.3914	1

Table VI
Basic Estimation

In this Table I present the results of the estimation of the Euler equation:

$$\Delta lnc_{i,t}=k1+\alpha_0\Delta lnC_{i,t}+\alpha_{-1}\Delta lnC_{i,t-1}+\zeta\Delta lnc_{i,t-1}+\gamma ln(1+(R_{i,t}^{f}-1)Pr[Y_{i,t}^{H}]) +\eta ln(1+(R_{i,t}^{C}\ 1)1[B])+\theta_1\Delta age_{i,t}+\theta_2\Delta age_{i,t}+\varepsilon_{i,t}$$

Column I estimates the basic model, while columns II to IV progressively add to the specification family composition, home ownership and occupation.

The standard errors are corrected for the non-independence of the observations within the same household. In addition, controls for the evolution of city-level prices and seasonal dummies are included in all the regressions.

	(I)	(II)	(III)	(IV)
R^C_t	-1.727***	-1.845***	-1.842***	-1.876***
	(0.000)	(0.000)	(0.000)	(0.000)
R^f_t	0.448	0.770	0.765	0.824
	(0.537)	(0.381)	(0.384)	(0.349)
ΔC_t	0.258***	0.295**	0.294**	0.290**
	(0.005)	(0.024)	(0.024)	(0.026)
Δc_{t-1}	0.530***	0.501***	0.502***	0.503***
	(0.000)	(0.000)	(0.000)	(0.000)
ΔC_{t-1}	0.007	0.014	0.014	0.013
	(0.195)	(0.176)	(0.176)	(0.189)
Age		0.004	0.004	0.004
		(0.334)	(0.363)	(0.305)
Age^2		-0.000	-0.000	-0.000
		(0.141)	(0.152)	(0.127)
marital status		0.022	0.023	0.023
		(0.399)	(0.375)	(0.386)
homeowner		-0.018	-0.017	-0.017
		(0.656)	(0.676)	(0.668)
income bracket		-0.003	-0.004	-0.002
		(0.524)	(0.470)	(0.705)
self_empl			0.085	0.071
			(0.158)	(0.252)
Occupation dummies				Yes
Seasonal dummies	Yes	Yes	Yes	Yes
# Households	2220	1432	1432	1432
Hansen J statistic	8.725	7.712	7.730	7.853
(pvalue)	0.463	0.441	0.562	0.441
C statistic[+]	1.021	3.748	3.761	3.747
(pvalue)	0.907	0.563	0.439	0.549
Adj. R-squared	0.214	0.226	0.226	0.225

Robust p values in parentheses

* significant at 10%; ** significant at 5%; *** significant at 1%

[+] Instruments tested: lags of debt outstanding, amount charged off, change in credit line and credit constraints measure.

A description of the variables is reported in Table I.

Instrument set: marginal tax rate, the local unemployment rate, the inflation rate, aggregate disposable income growth rate, mortgage rate, and some individual variables such as lags of the growth rate of debt, amount charged off, automatic credit line changes, and a credit constrained indicator.

First Stage Regressions

	(I)	(II)	(III)	(IV)
	Adj. R^2	Adj. R^2	Adj. R^2	Adj. R^2
R^C_t	0.0635	0.0602	0.0605	0.0618
R^f_t	0.9543	0.9543	0.9543	0.9543
ΔC_t	0.2351	0.3666	0.3666	0.3661
Δc_{t-1}	0.0281	0.0282	0.0281	0.028
	p-value	p-value	p-value	p-value
R^C_t	0.0000	0.0000	0.0000	0.0000
R^f_t	0.0000	0.0000	0.0000	0.0000
ΔC_t	0.0000	0.0000	0.0000	0.0000
Δc_{t-1}	0.0000	0.0000	0.0000	0.0000

Table VII
Comparison with Dynan (2000)

This Table illustrates the comparison between my results and Dynan (2000). In particular, I re-estimate the same regression as in Col (V) of Table VI using an annualized measure of credit card consumption and of the other variables:

$$\Delta \ln c_{i,t} = k1 + \alpha_0 \Delta \ln C_{i,t} + \alpha_{-1} \Delta \ln C_{i,t-1} + \zeta \Delta \ln c_{i,t-1} + \gamma \ln(1+(R_{i,t}^{f}-1)\Pr[Y_{i,t}^{H}]) + \eta \ln(1+(R_{i,t}^{C}-1)1[B]) + \theta_1 \Delta age_{i,t} + \theta_2 \Delta age_{i,t} + \varepsilon_{i,t}$$

Col (I) presents the baseline estimation; Col (II) investigates the effect of estimating the above regression without the external habit, as Dynan's estimation doesn't contain this variable. Col (III) drops the HH-specific financial variables from the instrument set; while Col (IV) drops the interest rate from the estimation equation, as Dynan considers an Euler equation with constant interest rates. The standard errors are corrected for the non-independence of the observations within the same household. In addition, controls for the evolution of city-level prices and seasonal dummies are included in all the regressions.

	(I)	(II)	(III)	(IV)
R^C_t	-0.076	-0.116	-2.100	
	(0.487)	(0.246)	(0.352)	
ΔC_t	0.110		1.055***	0.985***
	(0.931)		(0.000)	(0.000)
Δc_{t-1}	0.602	0.662***	0.109	0.139**
	(0.422)	(0.000)	(0.179)	(0.048)
Age	-0.001	-0.001	0.007	0.000
	(0.837)	(0.622)	(0.367)	(0.910)
Age^2	0.000	0.000	-0.000	-0.000
	(0.829)	(0.584)	(0.346)	(0.872)
marital status	0.006	-0.000	0.045	0.010
	(0.740)	(0.974)	(0.273)	(0.529)
self_empl	-0.031	-0.031	-0.014	-0.038
	(0.519)	(0.151)	(0.839)	(0.290)
homeowner	-0.001	-0.005	-0.035	0.006
	(0.964)	(0.790)	(0.539)	(0.815)
income bracket	-0.002	-0.001	-0.008	-0.003
	(0.559)	(0.539)	(0.287)	(0.439)
Occupation dummies	Yes	Yes	Yes	Yes
Seasonal dummies	Yes	Yes	Yes	Yes
# Households	1330	1450	1336	1336
Hansen J statistic	29.145	29.067		1.616
(p-value)		0.060		0.204
Adj. R-squared	-0.490	0.712	0.386	0.441

Robust p values in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%
A description of the variables is reported in Table I.
Instrument Set: second lag of local unemployment rate, the inflation rate, aggregate income growth rate, mortgage rate, household-specific debt growth rate and credit-constrained indicator.

First stage Regressions Regressions

Variable	Column	Partial R^2	F test	P-value
R^C_t	(I)	0.0846	130.9	0.0000
ΔC_t		0.14	230.61	0.0000
Δc_{t-1}		0.1587	267.17	0.0000
R^C_t	(II)	0.0768	126.08	0.0000
Δc_{t-1}		0.162	292.84	0.0000
R^C_t	(III)	0.0014	3.89	0.0086
ΔC_t		0.0744	226.34	0.0000
Δc_{t-1}		0.0814	249.38	0.0000
ΔC_t	(IV)	0.0744	226.34	0.0000
Δc_{t-1}		0.0814	249.38	0.0000

Table VIII
Effect of a smaller IV set

In this table I perform the same regressions as in Column I of Table VI with a progressively smaller instrument set. The instruments used in the estimation in the paper are: marginal tax rate, various lags of local unemployment rate, inflation rate, aggregate disposable income growth rate, mortgage rate, and some individual variables such as lags of the growth rate of debt, amount charged off, automatic credit line changes, and credit constrained indicator.
In Column II, I eliminate the marginal tax rate from the IV set illustrated above. In Column III I further eliminate the unemployment rate; in Column IV the inflation rate; in Column V all but one lags of aggregate income; in column VI the amount charged off; finally, in Column VII, I eliminate the credit line increases. The IV set I am left with exactly identifies the system and is composed by the first available lag of mortgage rate, aggregate income growth, household debt growth rate and credit constrained indicator.
The results show that restricting the IV set doesn't change the coefficients on ΔC_t and Δc_{t-1}, or their significance. The coefficient that is most sensitive to the shrinking of the IV set is that on the risk free interest rate, which happens sometimes to be negative, even though very close to zero and very imprecisely measured.

	(I)	(II)	(III)	(IV)	(V)	(VI)	(VII)
R^C_t	-1.727***	-1.743***	-1.742***	-1.750***	-1.738***	-1.752***	-1.568***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
R^f_t	0.448	0.408	-0.002	0.362	-0.009	-0.024	0.322
	(0.537)	(0.576)	(0.998)	(0.670)	(0.992)	(0.981)	(0.762)
ΔC_t	0.258***	0.266***	0.370***	0.283**	0.354**	0.352**	0.365**
	(0.005)	(0.004)	(0.001)	(0.036)	(0.031)	(0.034)	(0.029)
Δc_{t-1}	0.530***	0.527***	0.530***	0.529***	0.526***	0.523***	0.575***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
ΔC_{t-1}	0.007	0.007	0.010*	0.007	0.009	0.009	0.010
	(0.195)	(0.202)	(0.083)	(0.247)	(0.190)	(0.215)	(0.188)
# of Households	2220	2220	2220	2220	2220	2220	2220
Hansen J statistic	8.725	8.292	3.322	2.286	0.819	0.816	-
(p-value)	0.463	0.405	0.650	0.683	0.664	0.366	-
Adj. R-squared	0.214	0.211	0.158	0.203	0.168	0.169	0.156

Robust p values in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table IX
External Reference Shifts measured by Lottery Winnings

In Column (I) to (IV) of this Table, I reproduce the estimations in Table VI, using lottery winnings dummies instead than the city consumption. I estimate the following Euler equation:

$$\Delta lnc_{i,t}=k1+\alpha_0 Lottery+\alpha_{-1} Lottery\_lagged+\zeta \Delta lnc_{i,t-1}+\gamma ln(1+(R_{i,t}^{f}-1)Pr[Y_{i,t}^{H}]) +\eta ln(1+(R_{i,t}^{C}\ 1)1[B])+\theta_1 \Delta age_{i,t}+\theta_2 \Delta age_{i,t}+\varepsilon_{i,t}$$

In Column (V), I repeat the regression in Column (I), adding some additional controls about the lottery event: the fraction of winners opting for annuity payments and total winnings in the city in the given quarter.
The standard errors are corrected for the non-independence of the observations within the same household. In addition, controls for the evolution of city-level prices and seasonal dummies are included in all the regressions.

	(I)	(II)	(III)	(IV)	(V)
R^C_t	-1.664***	-1.854***	-1.856***	-1.867***	-1.863***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
R^f_t	1.508***	1.929***	1.934***	1.955***	1.949***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Lottery	0.136*	0.229**	0.228**	0.229**	0.210**
	(0.087)	(0.025)	(0.026)	(0.025)	(0.040)
Δc_{t-1}	0.548***	0.537***	0.537***	0.538***	0.539***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Lottery_lagged	-0.093	-0.164	-0.165	-0.165	-0.156
	(0.259)	(0.128)	(0.125)	(0.124)	(0.146)
Age		0.004*	0.004*	0.005*	0.005*
		(0.070)	(0.084)	(0.054)	(0.054)
Age^2		-0.000**	-0.000**	-0.000***	-0.000***
		(0.012)	(0.014)	(0.009)	(0.009)
marital status		0.020	0.021	0.021	0.022
		(0.188)	(0.173)	(0.167)	(0.160)
homeowner		-0.005	-0.003	-0.004	-0.004
		(0.852)	(0.908)	(0.896)	(0.893)
income bracket		-0.006	-0.006	-0.006	-0.006
		(0.120)	(0.101)	(0.122)	(0.114)
self_empl			0.078**	0.066*	0.066*
			(0.043)	(0.093)	(0.098)
Total winnings					0.000
					(0.497)
Prop. annuity					-0.245
					(0.598)
Occupation dummies				Yes	Yes
Seasonal dummies	Yes	Yes	Yes	Yes	Yes
# of Households	24066	15435	15435	15435	15435
Hansen J statistic	9.650	6.737	6.647	6.929	6.918
(p-value)	0.140	0.241	0.248	0.924	0.912
C statistic*	3.214	0.565	0.628	0.476	0.532
(p-value)	0.360	0.904	0.890	0.226	0.227
Adj. R-squared	0.270	0.268	0.268	0.267	0.267

Robust p values in parentheses
 * significant at 10%; ** significant at 5%; *** significant at 1%
[+] Instruments tested: lags of debt outstanding, amount charged off, change in credit line and credit constraints measure.
A description of the variables is reported in Table I.
Instrument set: marginal tax rate, aggregate disposable income growth rate and some individual variables such as lags of the growth rate of debt, amount charged off, automatic credit line changes, and a credit constrained indicator.

First Stage Regressions

	(I)	(II)	(III)	(IV)
	Adj. R^2	Adj. R^2	Adj. R^2	Adj. R^2
R^C_t	0.0651	0.0538	0.0566	0.056
R^f_t	0.9537	0.9536	0.9536	0.9536
Δc_{t-1}	0.0247	0.0220	0.0220	0.0221
	p-value	p-value	p-value	p-value
R^C_t	0.0000	0.0000	0.0000	0.0000
R^f_t	0.0000	0.0000	0.0000	0.0000
Δc_{t-1}	0.0000	0.0000	0.0000	0.0000

Table X
Estimations by Occupation

In this Table I reproduce the regression in Column (IV) of Table VI, separately for each of the main occupations in the data set. The purpose is to control for the stability of the internal habit coefficient across occupations. This can help address the concern that unobservable aggregate shocks that affect different HHs differently might cause the findings. The estimation equation is:

$$\Delta \ln c_{i,t}=k1+\alpha_0\Delta \ln C_{i,t}+\alpha_{-1}\Delta \ln C_{i,t-1}+\zeta\Delta \ln c_{i,t-1}+\gamma \ln(1+(R_{i,t}{}^{f}-1)\Pr[Y_{i,t}{}^{H}])+\eta \ln(1+(R_{i,t}{}^{C}\ 1)1[B])+\theta_1\Delta age_{i,t}+\theta_2\Delta age_{i,t}+\varepsilon_{i,t}$$

Column I estimates the model for a generic group that didn't specify the occupation, Col. (II) for administrative and staff, Col. (III) for white collars, Col (IV) for blue collars, Col. (V) for people holding technical jobs, and Col. (VI) for self-employed. The Appendix contains a more detailed description of these occupations.

The standard errors are corrected for the non-independence of the observations within the same household. In addition, controls for the evolution of city-level prices and seasonal dummies are included in all the regressions.

	(I)	(II)	(III)	(IV)	(V)	(VI)
R^C_t	-1.906***	-2.532**	-1.379	-2.474	-1.191**	-1.139
	(0.002)	(0.022)	(0.414)	(0.198)	(0.037)	(0.462)
R^f_t	-0.547	2.600	5.209	1.443	0.491	-6.713
	(0.701)	(0.341)	(0.204)	(0.829)	(0.830)	(0.371)
ΔC_t	0.434**	0.677	-0.100	0.121	0.302	-0.005
	(0.044)	(0.131)	(0.822)	(0.912)	(0.125)	(0.995)
Δc_{t-1}	0.497***	0.546***	0.197	0.446*	0.502***	0.425**
	(0.000)	(0.000)	(0.318)	(0.086)	(0.000)	(0.016)
ΔC_{t-1}	0.151	0.136	0.436	3.013	0.405***	0.965
	(0.468)	(0.944)	(0.871)	(0.447)	(0.004)	(0.784)
Age	0.008	0.009	0.008	-0.038	-0.017	0.023
	(0.196)	(0.606)	(0.767)	(0.370)	(0.130)	(0.384)
Age^2	-0.000*	-0.000	-0.000	0.000	0.000	-0.000
	(0.090)	(0.530)	(0.710)	(0.487)	(0.111)	(0.273)
marital status	0.016	-0.148	0.024	0.214	0.095	0.059
	(0.652)	(0.145)	(0.864)	(0.384)	(0.267)	(0.565)
homeowner	-0.079	-0.232	0.125	0.222	0.008	-0.001
	(0.142)	(0.250)	(0.483)	(0.424)	(0.927)	(0.996)
#of Households	722	160	66	80	265	50
Hansen J statistic	2.414	3.344	6.694	4.873	3.341	2.638
(p-value)	0.660	0.502	0.375	0.301	0.638	0.942
C statistic	1.040	1.427	1.961	3.095	0.898	0.120
(p-value)	0.594	0.490	0.153	0.213	0.502	0.620
Adj. R-squared	0.156	0.259	0.175	0.169	0.121	0.243

Robust p values in parentheses

* significant at 10%; ** significant at 5%; *** significant at 1%

Table XI
Robustness Checks

In this Table I analyze the robustness of the results to the inclusion of various measures of economic activity to the regression: the state-level income growth rate (Col (I)), the lead of the state-level income growth rate (Col (II)), change in city-level unemployment rates (Col (III)), housing market conditions (Col. (IV)). I also check the effect of adding an extra lag of the consumption growth rate (Col. (V)), year dummies to control for aggregate shocks (Col. VI)), and the growth of household debt (Col. (VII)).
The standard errors are corrected for the non-independence of the observations within the same household. In addition, controls for the evolution of city-level prices and seasonal dummies are included in all the regressions.

	(I)	(II)	(III)	(IV)	(V)	(VI)	(VII)
R^C_t	-1.878***	-1.862***	-1.871***	-1.833***	-2.054***	-1.849***	-1.322***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.005)
R^f_t	2.116	1.213	0.550	0.711	0.266	3.113	0.189
	(0.174)	(0.230)	(0.527)	(0.432)	(0.777)	(0.548)	(0.835)
ΔC_t	0.271**	0.268**	0.442**	0.290**	0.311**	0.266**	0.295**
	(0.037)	(0.045)	(0.016)	(0.030)	(0.020)	(0.039)	(0.024)
Δc_{t-1}	0.501***	0.507***	0.509***	0.503***	0.453***	0.508***	0.445***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
ΔC_{t-1}	0.014	0.012	0.021	0.014	0.016	0.013	0.014
	(0.164)	(0.251)	(0.107)	(0.184)	(0.105)	(0.234)	(0.161)
Age	-0.000	-0.000	-0.000	-0.004	-0.000	-0.000	0.004
	(0.109)	(0.135)	(0.188)	(0.324)	(0.261)	(0.150)	(0.419)
Age^2	0.023	0.025	0.021	0.000	0.027	0.025	-0.000
	(0.375)	(0.360)	(0.466)	(0.138)	(0.348)	(0.366)	(0.199)
Marit_status	-0.014	-0.015	-0.028	0.026	-0.039	-0.014	0.023
	(0.724)	(0.708)	(0.498)	(0.336)	(0.361)	(0.722)	(0.401)
Homeowner	-0.002	-0.002	-0.003	-0.012	-0.000	-0.002	-0.025
	(0.692)	(0.712)	(0.650)	(0.767)	(0.956)	(0.701)	(0.538)
Income bracket	0.068	0.072	0.080	-0.004	0.092	0.071	-0.002
	(0.262)	(0.248)	(0.201)	(0.466)	(0.158)	(0.248)	(0.764)
Self_empl	-1.878***	-1.862***	-1.871***	0.065	-2.054***	-1.849***	0.057
	(0.000)	(0.000)	(0.000)	(0.288)	(0.000)	(0.000)	(0.331)
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Δlnincome	0.020						
	(0.319)						
Δlnincome _lead		0.012					
		(0.434)					
Δln U_rate			-0.003				
			(0.465)				
Median House Value				0.000			
				(0.705)			
Rent				0.000			
				(0.544)			
$(\Delta c_t)^2$				-0.013			
				(0.527)			
Δc_{t-2}					0.170***		
					(0.000)		
Δlndebt_t							-0.157*
							(0.060)
Year dummies						Yes	
Seasonal dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# Households	1432	1432	1432	1432	1432	1432	1432
Hansen J statistic	6.986	7.245	4.546	7.890	12.705	5.397	9.695
(p-value)	0.538	0.524	0.576	0.545	0.176	0.524	0.287
C statistic$^+$	5.224	3.204	2.892	3.933	5.665	3.204	2.487
(p-value)	0.265	0.510	0.715	0.415	0.226	0.798	0.647
Adj. R-squared	0.230	0.231	0.176	0.219	0.249	0.232	0.264

Robust p values in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%
$^+$ Instruments tested: lags of debt outstanding, amount charged off, change in credit line and credit constraints measure.
A description of the variables is reported in Table I.
Instrument set: marginal tax rate, the local unemployment rate, the inflation rate, aggregate disposable income growth rate, mortgage rate, and some individual variables such as lags of the growth rate of debt, amount charged off, automatic credit line changes, and a credit constrained indicator.

Table XII
Does Local Aggregate Consumption Proxy for Individual Income?

In order to answer this question I use household-level data from the PSID, in which individual income is provided.[1] I investigate whether local aggregate consumption provides any information about household income once aggregate income is available. The regression is the following:

$$\Delta \mathbf{lnincome}_{\text{individual, it}} = \alpha + \beta_1 \Delta \mathbf{lnINCOME}_{\text{aggregate, it}} + \beta_2 \Delta \mathbf{lnC}_{\text{citylevel, it}} + \beta_3 \mathbf{age}_i + \beta_4 \mathbf{age}^2_i + \varepsilon_{it}$$

Dummy variables capturing marital status, occupational choice, seasonal fluctuations and whether the individual owns the house in which he lives are included in some of the regressions as well.

The null hypothesis is that once we control for aggregate income, the coefficient on aggregate consumption is small and statistically insignificant. The results below confirm this hypothesis and suggest that aggregate consumption doesn't proxy for individual income.

I have also tried the above regression using the growth rate of future HH income as the dependent variable. The results are presented in columns IV to VI and show that neither aggregate income nor consumption are good predictors of future individual income. The regression is the following:

$$\Delta \mathbf{lnincome\_lead}_{\text{individual, it+1}} = \alpha + \beta_1 \Delta \mathbf{lnINCOME}_{\text{aggregate, it}} + \beta_2 \Delta \mathbf{lnC}_{\text{aggregate, it}} + \beta_3 \mathbf{age}_i + \beta_4 \mathbf{age}^2_i + \varepsilon_{it}$$

Notice should be given to the fact that the time period analyzed is not very long, spanning from 1997 to 2001. Unfortunately a longer time series of city-level consumption is not available.

Dependent Variable	Δ**lnincome**$_{\text{individual, it}}$			Δ**lnincome_lead**$_{\text{individual, it+1}}$		
	(I)	**(II)**	**(III)**	**(IV)**	**(V)**	**(VI)**
ΔlnC	0.045	0.066	0.051	0.093	0.057	0.053
	(0.550)	(0.377)	(0.492)	(0.290)	(0.474)	(0.508)
ΔlnINCOME	2.697***	2.660***	2.887***	-0.674	-0.334	-0.525
	(0.009)	(0.009)	(0.005)	(0.726)	(0.849)	(0.768)
Age	-0.018**	-0.023***	-0.026***	0.001	-0.000	-0.002
	(0.041)	(0.010)	(0.005)	(0.940)	(0.967)	(0.886)
Age^2	0.000*	0.000**	0.000***	-0.000	0.000	0.000
	(0.075)	(0.022)	(0.006)	(0.989)	(0.915)	(0.818)
Marital status	-0.057	-0.076	-0.072	-0.077	-0.094	-0.112*
	(0.241)	(0.144)	(0.172)	(0.243)	(0.153)	(0.098)
Homeowner dummy		Yes	Yes		Yes	Yes
Occupation dummies			Yes			Yes
Observations	921	891	891	426	413	413
R-squared	0.017	0.022	0.036	0.019	0.017	0.027

p values in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

<u>Sample:</u> households in the PSID living in California. Time span: 1997 to 2001.[2]

<u>Variables description and definitions:</u> $\Delta lnincome_{individual,\ it}$ is the quarterly growth rate of household total income; *ΔlnINC* is the quarterly growth rate of aggregate California per capita income (Source: California DOF); *ΔlnC* is the quarterly growth rate of city-level aggregate taxable sales and constitutes my measure of aggregate consumption (Source: California DOF); *Age* is the age of the head of the household; *Marital status* is a dummy equal to 1 if the head of the HH is married and 0 otherwise; the *Homeowner dummy* equals 1 if the HH owns the house in which it lives and 0 otherwise; the *Occupational dummies* are dummies that categorize HH heads in main occupational areas and are built to be as similar as possible to those used in the rest of the paper.

[1]Unfortunately, this dataset doesn't contain a good measure of consumption; for other drawbacks of the PSID see Section 3.2.
[2] This is the period for which data on taxable sales are available on the California Department of Finance website.

Table XIII
Alternative Explanations: Liquidity Constraints and Precautionary Saving Motives

This Table contains tests of the habit persistence hypothesis against the liquidity constraint and precautionary saving motive alternatives. The baseline regression to which the results are compared is:

$$\Delta lnc_{i,t}=k1+\alpha_0\Delta lnC_{i,t}+\alpha_{-1}\Delta lnC_{i,t-1}+\zeta\Delta lnc_{i,t-1}+\gamma ln(1+(R_{i,t}^{f}-1)Pr[Y_{i,t}^{H}]) +\eta ln(1+(R_{i,t}^{C} -1)1[B])+\theta_1\Delta age_{i,t}+\theta_2\Delta age_{i,t}+\varepsilon_{i,t}$$

Col. (I) and (II) re-estimate the regression on two sub-samples of unconstrained and credit constrained HHs; col (III) and (IV) perform the same regression as (I) and (II) adding the lag of income growth rate. Col (V) adds the lagged growth rate of income to the regression; Col (VI) adds a credit constrained indicator. Col (VII) adds the square of consumption growth to the regression.

The standard errors are corrected for the non-independence of the observations within the same household. In addition, controls for the evolution of city-level prices and seasonal dummies are included in all the regressions.

	(I)	(II)	(III)	(IV)	(V)	(VI)	(VII)
R^C_t	1.788	-1.109	-0.141	-2.395	-1.879***	-0.603	-1.695***
	(0.390)	(0.829)	(0.939)	(0.638)	(0.000)	(0.423)	(0.000)
R^f_t	-2.313	3.906	-6.201*	7.733	0.414	0.161	1.543
	(0.556)	(0.350)	(0.099)	(0.124)	(0.694)	(0.870)	(0.287)
ΔC_t	0.644**	0.186	0.267	0.233	0.274**	0.240*	0.290**
	(0.035)	(0.674)	(0.392)	(0.598)	(0.035)	(0.063)	(0.024)
Δc_{t-1}	0.565***	0.557***	0.416***	0.693***	0.502***	0.498***	0.504***
	(0.000)	(0.007)	(0.005)	(0.002)	(0.000)	(0.000)	(0.000)
$\Delta lnincome_{t-1}$			0.081**	-0.076	0.012		
			(0.048)	(0.124)	(0.473)		
ΔC_{t-1}	0.023	-0.362***	0.016	-0.396***	0.014	0.011	0.014
	(0.307)	(0.005)	(0.269)	(0.003)	(0.165)	(0.253)	(0.184)
Age	-0.015	0.001	-0.006	0.006	0.005	0.001	0.004
	(0.241)	(0.963)	(0.596)	(0.708)	(0.276)	(0.900)	(0.339)
Age^2	0.000	0.000	0.000	-0.000	-0.000	-0.000	-0.000
	(0.384)	(0.968)	(0.842)	(0.736)	(0.111)	(0.775)	(0.149)
Marit_status	0.026	-0.048	0.009	-0.038	0.023	0.012	0.026
	(0.759)	(0.490)	(0.902)	(0.590)	(0.372)	(0.623)	(0.372)
Homeowner	-0.072	-0.148	-0.021	-0.159	-0.015	0.007	-0.017
	(0.529)	(0.247)	(0.840)	(0.236)	(0.713)	(0.839)	(0.691)
Income bracket	0.015	0.027*	0.008	0.024	-0.002	0.001	-0.003
	(0.437)	(0.094)	(0.585)	(0.173)	(0.698)	(0.822)	(0.653)
Self_empl	-0.188	0.243	-0.031	0.287	0.068	0.063	0.067
	(0.473)	(0.347)	(0.885)	(0.271)	(0.266)	(0.177)	(0.296)
Credit constr. indic.						0.062	
						(0.327)	
$(\Delta c_t)^2$							-0.013
							(0.527)
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Seasonal dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# Obs	773	347	773	347	1432	1432	1432
Hansen J statistic	5.400	3.298	3.780	1.028	7.491	7.426	7.567
(p-value)	0.483	0.654	0.779	0.906	0.485	0.406	0.472
C statistic	0.493	0.000	0.079	0.002	5.174	2.911	3.540
(p-value)	0.249	0.990	0.286	0.966	0.270	0.491	0.477
Adj. R-squared	0.229	0.282	0.165	0.240	0.230	0.255	0.233

Robust p values in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

A description of the variables is reported in Table I.
Instrument set: marginal tax rate, the local unemployment rate, the inflation rate, aggregate disposable income growth rate, mortgage rate, and some individual variables such as lags of the growth rate of debt, amount charged off, automatic credit line changes, and a credit constrained indicator.

Table XIV
Aggregate Consumption Regressions

This Table investigates the aggregate implications of household level consumption choices. Columns (I), (II), and (III) present the results of estimating an aggregate Euler equation based on per capita consumption: (a) $\mathbf{\Delta ln\Sigma c_{i,t} = \alpha + \beta \Delta ln\Sigma c_{i,t-}1 + R^f_t + \varepsilon_t}$. Columns (IV), (V), and (VI) presents the results of estimating the same Euler equation using correctly aggregated data: (b) $\mathbf{\Delta\Sigma lnc_{i,t} = \alpha + \beta \Delta\Sigma lnc_{i,t-}1 + R^f_t + \varepsilon_t}$. Finally, Columns (VII), (VIII), and (IX) illustrate the effect of adding moments of the cross sectional distribution of consumption growth rates to regression.

The standard errors are corrected for the non-independence of the observations within the same household. In addition, controls for the evolution of city-level prices and seasonal dummies are included in all the regressions.

	Per-capita Consumption Aggregation Method			**Correct Aggregation Method**			**Per-capita Consumption Aggregation Method plus Moments**		
Dependent variable	$\Delta ln\Sigma c_t$			$\Delta\Sigma lnc_t$			$\Delta ln\Sigma c_t$		
	(I)	(II)	(III)	(IV)	(V)	(VI)	(VII)	(VIII)	(IX)
R^f_t	1.361	1.468	0.872	-0.112	-0.018	-0.042	2.401*	1.454**	1.444**
	(0.186)	(0.144)	(0.190)	(0.818)	(0.974)	(0.926)	(0.064)	(0.038)	(0.040)
$\Delta ln\Sigma c_{i,t-1}$	-0.685	-0.739	-0.494				0.139	-0.791**	0.233
	(0.205)	(0.191)	(0.164)				(0.942)	(0.033)	(0.704)
$\Delta\Sigma lnc_{i,t-1}$				0.515**	0.727**	0.538*			
				(0.042)	(0.039)	(0.089)			
$\Delta lnincome_t$		0.002			-0.131				
		(0.996)			(0.404)				
$\Delta lnincome_{t-1}$			0.471*			0.147			
			(0.061)			(0.344)			
Δstddev/2							36.669		
							(0.489)		
Δskewness/6								2.278	
								(0.571)	
Δkurtosis/24									-4.954
									(0.116)
Time dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# Obs	10	10	10	10	10	10	10	10	10
Hansen J statistic	2.244	2.061	4.108	2.286	1.213	1.210	0.002	3.532	0.370
(p-value)	0.326	0.151	0.043	0.319	0.271	0.271	0.965	0.060	0.543
Adj. R-squared	0.380	0.363	0.685	0.285	0.009	0.170	-.241	0.479	0.174

Robust p values in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

A description of the variables is reported in Table I. Δstddev, Δskewness, and Δkurtosis are the standard deviation, skewness and kurtosis of the cross sectional distribution of consumption.
Instrument set: second lag of aggregated city-level sales, income growth rate, average mortgage rate, and unemployment rate.

Table AI
Demographic Characteristics
Comparison with U.S. Census data

This Table contains a comparison between the breakdown by age and income of my data set versus the U.S. Census.

	Age	
	My dataset CA	*U.S. Census*
20 to 24 years	7.11%	9.44%
25 to 34 years	17.61%	19.85%
35 to 44 years	24.91%	22.47%
45 to 54 years	24.01%	18.75%
55 to 59 years	8.12%	6.70%
60 to 64 years	5.12%	5.38%
65 to 74 years	9.16%	9.15%
75 to 84 years	3.48%	6.15%
85 years and over	0.49%	2.11%
Total	**100%**	**100%**
	Income	
	My dataset CA	*U.S. Census*
Less than $15,000	7.21%	15.85%
$15,000 to $19,999	5.33%	6.25%
$20,000 to $29,999	8.04%	13.02%
$30,000 to $39,999	13.32%	12.27%
$40,000 to $49,999	11.04%	10.62%
$50,000 to $74,999	22.12%	19.46%
$75,000 to $99,999	19.74%	10.23%
$100,000 to $124,999	4.44%	5.20%
$125,000+	8.75%	7.09%
Total	100%	100%
	Other Demographic Characteristics	
	My dataset CA	*U.S. Census*
Home owner	74.68%	66.20%
Renter	9.69%	33.80%
Missing	15.63%	0.00%
Married	44.76%	52.40%
Single	19.97%	47.60%
Missing	35.27%	0.00%

Figure AI

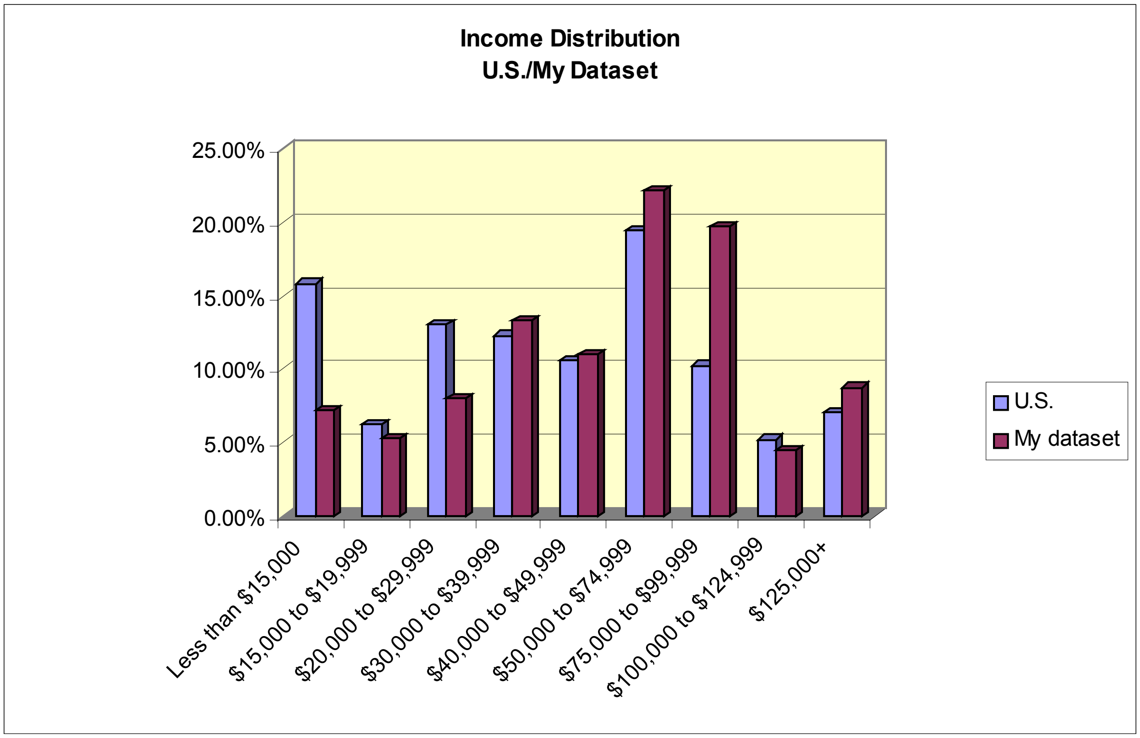

Source: U.S. Census and credit card data set.

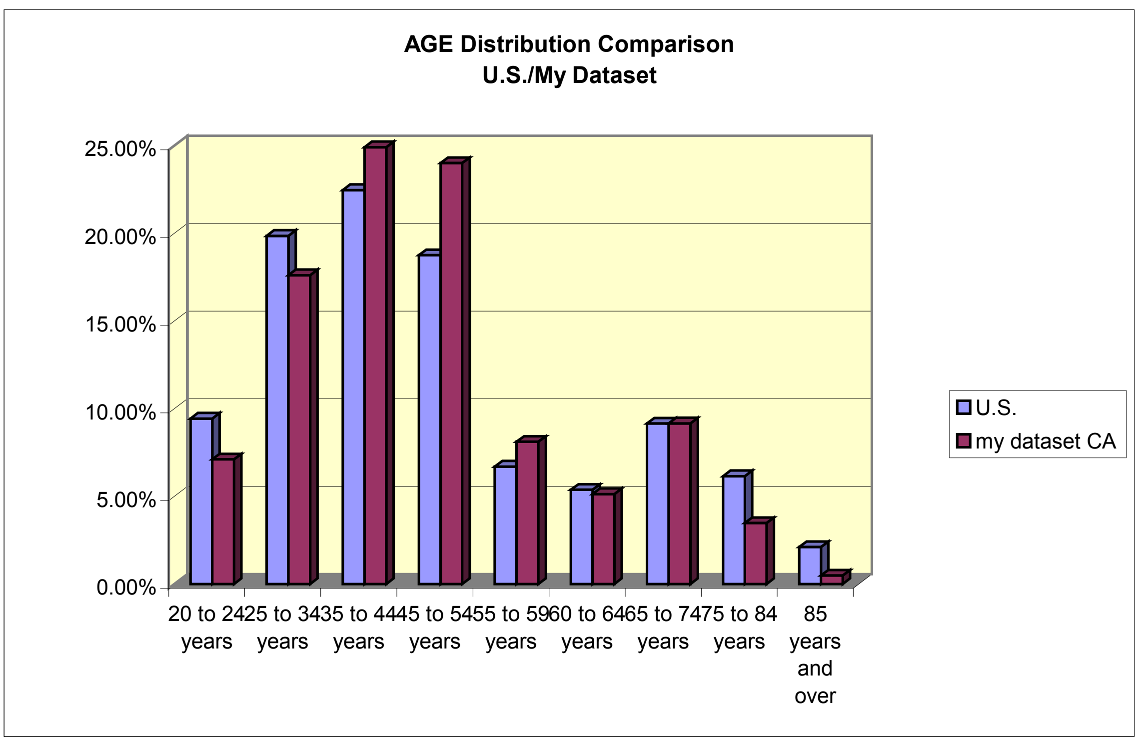

Table AII
Occupation:
Comparison with U.S. Census 2000

My Dataset	
Occupation	*Percent*
Administrative/Managerial	7.59%
Clerical/White_Collar	3.25%
Craftsman/Blue_Collar	3.93%
Farmer	0.11%
Housewife	0.67%
Military	0.19%
Professional/Technical	13.76%
Sales/Service	1.91%
Self_Employed	2.51%
Student	1.65%
Other	64.44%
Total	100%

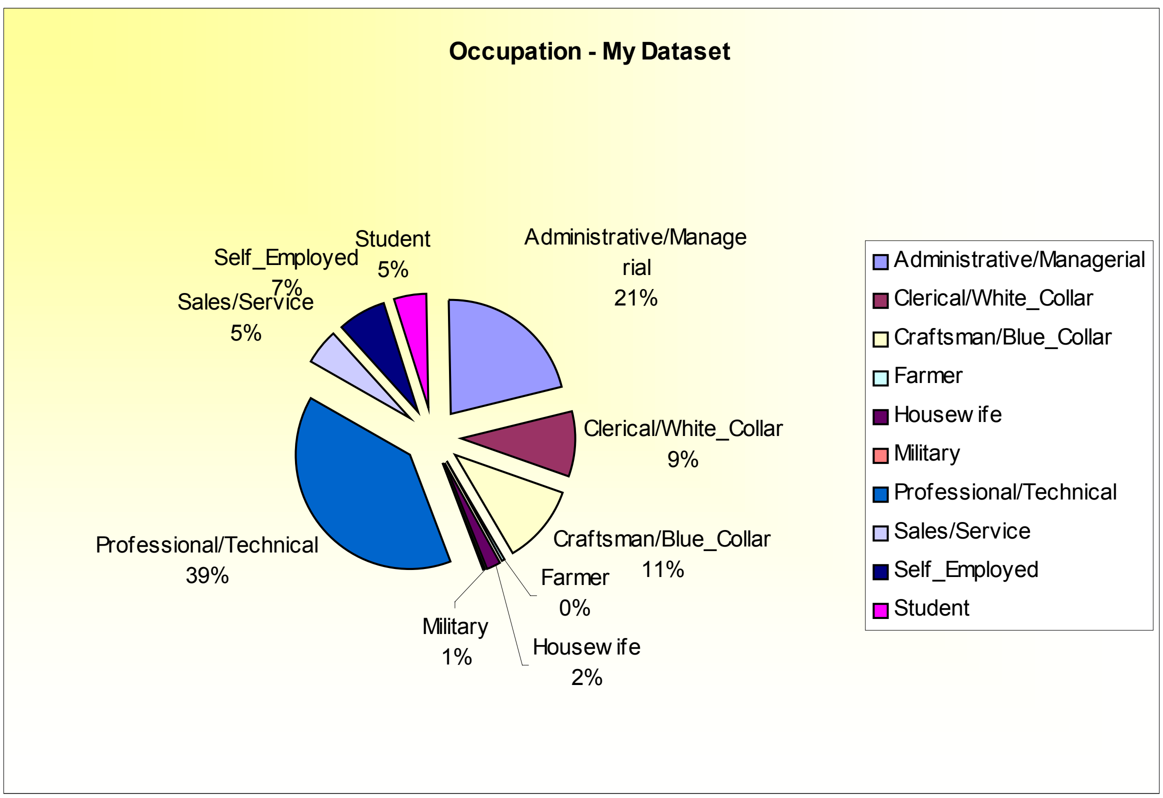

Table A III
Panel A
Financial Characteristics
Comparison with multi-issuer credit card dataset

	My dataset*		Gross and Souleles (2002)**	
	Mean	Median	Mean	Median
debt	$1,486.10	$0	$1,349.00	$70
debt|debt>0	$3,408.36	$2,821	$2,809.00	$2,120
credit limit	$8,823.51	$10,000	$6,207.00	$5,000
D credit limit	$154.15	$0	$76.80	$0
if D credit limit~=0	$1,181.65	$1,000	$1,985.00	$1,000
interest rate	16.13	14.99	16.60	17.20
D interest rate	-0.109	0	0.036	0
if D interest rate~=0	-1.012	-2.25	0.914	0.25

* The period analyzed is Aug.1998-Jul.2002
** Source: Gross and Souleles (2002), Table I. The period analyzed is Jan. 1995-Jan.1998

Panel B
Financial Characteristics
Comparison with the Survey of Consumer Finances

	My dataset		SCF	
	Mean	Median	Mean	Median
All	$3,408.36	$2,821	$4,100	$1,900
by Age				
Less than 35	$2,952	$2,312	$4,000	$2,000
35-44	$3,578	$2,967	$4,300	$2,000
45-54	$3,880	$3,202	$4,200	$2,300
55-64	$3,092	$2,463	$4,100	$1,900
65-74	$3,276	$2,810	$5,200	$1,000
older than 75	$3,720	$3,233	$1,900	$700
by Income Percentiles				
Less than 20	$3,039	$2,516	$2,100	$1,000
20-39.9	$3,285	$2,814	$2,800	$1,200
40-59.9	$3,820	$3,207	$3,700	$2,000
60-79.9	$3,551	$2,850	$4,700	$2,300
80-89.9	$3,215	$2,643	$7,200	$3,800
90-100	$3,836	$3,401	$6,600	$2,800
by Housing Status				
Home Owner	$3,559	$3,033	$4,500	$2,100
Renter	$3,405	$2,826	$3,400	$1,200