Household Need for Liquidity and the Credit Card Debt Puzzle

Irina A. Telyukova *
University of California, San Diego

December 20, 2011

Abstract

In the 2001 U.S. Survey of Consumer Finances (SCF), 27% of households report simultaneously revolving significant costly credit card debt and holding sizeable amounts of low-return liquid assets; this is known as the "credit card debt puzzle". In this paper, I quantitatively evaluate the role of predictable and precautionary liquidity demand in accounting for this puzzle: households that accumulate credit card debt may not pay it off using their money in the bank, because they expect to use that money in situations where credit cards cannot be used. Using both aggregate and survey data (SCF and CEX), I characterize the puzzle and document support in the data for the liquidity hypothesis. I then develop a dynamic heterogeneous-agent model of household portfolio choice, where households are subject to uninsurable income and preference uncertainty, and consumer credit and liquidity coexist as means of consumption and saving/borrowing. The calibrated model accounts for between 44% and 56% of the households in the data who hold consumer debt and liquidity simultaneously, and for 100% of the liquidity held by a median household in the puzzle group. I show that under reasonable calibration alternatives, the model can capture the entire puzzle group size as well. I also assess the quantitative contribution of each type of idiosyncratic risk and find that about one-half of the liquidity demand in the model is precautionary. The calibration also yields novel measurements of important parameters of preferences and idiosyncratic risk.

*I am indebted to Vìctor Rìos-Rull, Randall Wright, Jesùs Fernàndez-Villaverde and Dirk Krueger for their guidance in this project. For many helpful discussions and suggestions, I thank Kjetil Storesletten and three anonymous referees, as well as Orazio Attanasio, Marjorie Flavin, Andreas Lehnert, Ben Lester, Michael Palumbo, Shalini Roy, Gustavo Ventura, Ludo Visschers, and Neil Wallace, and the participants of seminars and conferences at the University of Pennsylvania, UCSD, USC, Notre Dame, Arizona State, Wash. U. St. Louis, Maryland, Queen's, Western Ontario, University of British Columbia, the Board of Governors of the Federal Reserve, Federal Reserve Banks of New York, Cleveland, Philadelphia and Atlanta, European Central Bank, the Canadian Economic Association, the Midwest Macroeconomic Meetings, and the SED. I am grateful to the Jacob K. Javits Graduate Research Fellowship Fund and the Federal Reserve Board of Governors Dissertation Internship program for research support.

1 Introduction

In the 2001 U.S. Survey of Consumer Finances, 27% of households reported revolving an average of \$5,766 in credit card debt, with an APR of 14%, and simultaneously, holding an average of \$7,338 in liquid assets, with a return rate of around 1%. In fact, 84% of households who revolved credit card debt had some liquid assets that could be, but were not, used for credit card debt repayment. This apparent violation of the no-arbitrage condition has been termed the "credit card debt puzzle".

Gross and Souleles (2002) were among the first to document this fact. They suggested several possible explanations for this behavior, two of which have been pursued in the literature. Lehnert & Maki (2001) study whether households may do this strategically, in preparation for a bankruptcy filing under Chapter 7. Since in the U.S., each state offers some exemption level of assets in the event of household bankruptcy, the authors argue that households may run up their credit card debt since it would be discharged during the filing, while keeping their assets in liquid form, in order to convert them to exemptible assets when filing. The authors examine exemption level by state, and find that in states where exemption levels are higher, the puzzle is more prevalent. While this may be a compelling idea to a small number of households, upon examination of the total portfolios of the puzzle households, it appears that most of them would be unlikely to file for bankruptcy, as they hold significant and positive financial and nonfinancial wealth.

Alternatively, Bertaut and Haliassos (2002), and Haliassos and Reiter (2003) have studied whether households may opt to hold liquidity and credit card debt simultaneously as a means of self-(or spouse) control. If one spouse in the household is the earner, and the other is the compulsive shopper, it is argued that the earner will choose not to pay off credit card debt in full in order to leave less of the credit line open for the shopper to spend. This again may apply to some share of households, but is unlikely to account for many of the households in the puzzle category, since it is a costly way of performing this kind of control. A household in the puzzle group loses, on average, \$734 per year, largely from the costs of debt revolving, which amounts to 1.5% of their total annual after-tax income. Less expensive control options are available, such as lowering the credit limit or holding fewer credit cards.

Laibson et al (2001) examine a related puzzle: the coexistence in household portfolios of credit card debt and retirement assets. The difference is key: retirement assets, such as IRA accounts, are nonliquid and involve a significant penalty for early withdrawal. The authors explain this behavior with time-inconsistent decision-making by households, which makes them patient in the long run, but impatient in the short run. The explanation cannot apply to the credit card debt puzzle, however, because the tradeoff here is between two short-run decisions, and because liquid asset withdrawal does not incur a penalty.

In this paper, I offer and evaluate an alternative hypothesis. I focus on the need for liquidity as the possible reason. The premise is that there are large parts of household monthly expenditures that cannot be paid for by credit card, so they must be paid by liquid instruments.[1] Such payments often are substantial in size, and include predicted expenses (such as mortgage and rent payments, utilities, babysitting and daycare services), as well as significant unpredictable ones (such as major household repairs, auto repairs and other types of emergencies).[2] Some of these are universally cash-only goods, while others may or may not be. For example, large contractors may accept credit cards for home repairs, while smaller outfits may not. All of these expenses warrant keeping money in the bank. Thus, even for a household that has accumulated credit card debt, drawing down its liquid assets below some threshold may not be an optimal choice, and the household may prioritize building its liquid asset holdings over debt repayment in the short to medium run. The unpredictable nature of some of these expenditures may warrant holding fairly large liquid balances for *precautionary* reasons, as inability to pay if emergency strikes may be very costly.

Gross and Souleles (2002) mention the idea of transaction demand for liquidity as a possible contributor to why people hold debt and money simultaneously, but dismiss it as insufficient for the purposes of explaining the puzzle. A careful quantitative analysis of the hypothesis presented here, however, is an involved exercise, from both theoretical and empirical perspectives, and it is crucial, because it allows us to evaluate the possibility that the puzzle may largely fit within the standard rational expectations framework, and hence appear much less puzzling in that context. I am particularly interested in understanding the nature and magnitude of not only transactions demand, but also precautionary demand for liquidity, which turns out to be quite nontrivial to

[1] I use the term "liquid assets" such as checks, debit cards and savings accounts, interchangeably with "money" and "cash", since their liquidity properties are the same for my purposes.

[2] Below, I discuss the survey evidence of the fact that such goods tend to be cash-only goods.

quantify, and in linking this demand to the puzzle in question.

The main goal of this paper is to measure how much of the puzzle the liquidity need hypothesis can account for. Specifically, I answer the following two questions: (1) Can the need for liquidity account for the number of households that revolve debt while having money in the bank?; and (2) *How much* liquidity is it optimal for a household to have, given the risk characteristics that it is exposed to, especially if it revolves credit card debt?

I use data from the Survey of Consumer Finances and the Consumer Expenditure Survey to study in detail the demographic and economic characteristics of households who choose to borrow on credit cards and save in liquid accounts simultaneously. I show that there is nothing inherent about them, from a demographic perspective, that would distinguish them from other households, but that they do have distinguishing characteristics in terms of their income and asset holdings. I also show evidence that gives support to the importance of liquid assets in monthly household expenditures, and to the presence of uncertainty in these expenses.

Next, I develop a dynamic stochastic partial-equilibrium model of household portfolio choice, in order to study the hypothesis quantitatively. The starting point is a standard infinite-horizon heterogeneous-agent model with exogenously incomplete markets. Households face two types of idiosyncratic risk, consume two goods, and allocate their portfolios between two assets. There is a two-market structure, where in one of the markets credit cannot be used. The two types of idiosyncratic risk are income shocks and shocks to expenses in liquid consumption, the latter of which is captured as preference shocks. In its treatment of money, the model can be thought of as a stochastic incomplete-market version of a Lucas-Stokey-style cash-credit good model. As I will show, the model has all the qualitative implications important for addressing the credit card debt puzzle.

I calibrate the model by matching it to properties of liquid-asset consumption in household-level data, as well as to distributional characteristics in the data. I pay particular attention to the division of consumption in the data into cash-only and cash-or-credit goods, and their relative properties, as well as to measuring the idiosyncratic expense uncertainty that households face. The calibration method is based on the simulated method of moments, and leaves the quantities crucial for answering the central questions untargeted. The benchmark calibrated model accounts for between 44 and 56% of the households who choose to revolve debt while holding money in the bank, and for a median such household, for 100% or more of their liquidity holdings. The ranges refer to two alternative calibrations, depending on the choice of the income process that households face. I also show that for a reasonable alternative for the borrowing limit, the model can account for 100% of the size of the puzzle group, without compromising liquidity demand predictions. I then use the model to disentangle the quantitative contribution of each of the shocks to the overall results. I find that the expense shocks are essential for generating the puzzle group in the data, and contribute about one-half of the household liquidity demand. That is, about one-half of the liquidity holdings that the model generates is driven by precautionary motive.

The main contributions of this paper are four. First, this paper is the first to evaluate the liquidity-need hypothesis in relation to the standing puzzle, and I find that it goes a long way toward matching the facts. Debt puzzles of this nature have led the literature to challenge the standard consumption and saving models as incapable of explaining them. This paper can be seen partly as a response to this challenge. Second, this paper is the first, to my knowledge, to attempt to separate consumption in the data into cash-only and cash-or-credit goods, and to measure a new type of idiosyncratic risk - namely, expense risk that leads to precautionary liquidity demand. These are nontrivial tasks. In light of the interest in the incomplete-market literature on idiosyncratic risk, this is an important step toward understanding the types and magnitude of risk that households face. Third, in the process of calibrating the model, I also obtain new estimates of the elasticity of substitution between cash and credit goods, that has previously only been estimated in deterministic representative-agent models. Modeling precautionary demand for money affects these parameter estimates considerably. My estimate of the elasticity of substitution is significantly lower than existing estimates, suggesting that cash and credit goods are complements rather than substitutes, as was previously believed. The intuition for the result is clear-cut in the model. Fourth, the need for liquidity arises in this paper because liquid assets are the most versatile and sometimes the unique payment option available. This mechanism then accounts for a much broader class of debt puzzles than just the one having to do with credit card debt. The co-existence of any kind of debt and liquid assets in a household portfolio could have the same explanation as the one presented here, and the model may be useful in accounting for such portfolio allocation puzzles.

In complementary work, Zinman(2006) uses survey data to demonstrate also, via data calculations, that "borrowing high and lending low" is largely not puzzling and could be accommodated within the standard theoretical framework. His claim is that once one accounts for the liquidity premium of checking and savings accounts, the return differential between the two assets is largely calculated away, and the puzzle stops being prevalent. Thus, Zinman's findings provide informal support for the formal treatment of the liquidity need hypothesis presented in this paper.

The paper is structured as follows. In section 2, I characterize the credit card debt puzzle in the data, by studying the Survey of Consumer Finances and Consumer Expenditure Survey. Section 3 lays out the model and briefly analyzes its properties. Section 4 presents detailed information on the calibration strategy. Section 5 shows the fit of the model and resulting calibration. Section 6 presents the results from the calibrated model, the shock decomposition and discusses the results. Section 7 concludes. Some details of the data are relegated to the appendix.

2 Data

In this section, I characterize the credit card debt puzzle in the data. The facts that I document here demonstrate that the data are consistent with my hypothesis, and inform the modeling choices that I make in the next section.

I use two U.S. household surveys in order to describe the puzzle in the data: the 2001 Survey of Consumer Finances (SCF) and the Consumer Expenditure Survey (CEX) from 2000-2002. The SCF is a triennial cross-sectional survey that has detailed information on household assets and liabilities. For example, it distinguishes revolving credit card debt from purchase balances that are immediately paid off, and despite its cross-sectional nature, allows to assert persistence of this revolving debt. The CEX is a rotating panel, where each household is interviewed for five consecutive quarters, four of which (second through fifth) are made into public data files. The advantage of the survey is detailed measurement of all aspects of household monthly consumption: in each interview, the household is asked to recall all of its expenditures in the preceding three months.[3] Although it is less careful about measuring assets and credit card

[3] 65% of the expenditure data are collected via direct questions about the month and amount of expenditure,

debt, there is sufficient information on credit card debt and liquid asset holdings. I use the CEX to study the properties of household consumption in goods paid by liquid assets versus other methods.

I focus on the post-college working-age population, studying all households with heads of age 25 to 64. I separate the samples in both surveys into three subgroups: those who have more than $500 in revolving credit card debt and less than $500 in liquid assets ("borrowers"), those who have more than $500 of both ("borrowers and savers", i.e. the puzzle group), and those who have liquid assets but less than $500 of revolving credit card debt ("savers").[4] To define the puzzle group, I take only the households that revolve debt *habitually*, that is, report repaying their balance off in full only sometimes or never. As credit card debt, I include only the balance due on the credit card left over after the last statement was paid - thus excluding recent purchases and balances that were paid off. The definition of liquid assets used in this paper includes checking accounts, savings accounts, and brokerage accounts (i.e. idle money in a brokerage house that is not being invested in stocks). As no data are collected on household cash holdings, I am not able to include currency holdings in the definition of liquidity. Under the premise that those with bank accounts do not hold much currency, the only households that are likely to be affected by this data restriction are those in the borrower category, some of whom may not have bank accounts and are thus forced to hold currency. It is difficult to know how much of it they are holding - there are no survey data on this in the U.S. According to Prescott and Tatar (1999), between 47% and 67% of those without bank accounts in different surveys report not to have enough money to make it worth opening an account, which could be interpreted as such households spending most of their available money during the month. Since this is a small group in the data (see below), I don't expect non-inclusion of cash to have a strong bearing on my results. Additional details of the surveys, the sample selection process, and the puzzle measurement methods are described in the data appendices A.1 and A.2.

Below, I first describe the credit card debt puzzle, and compare the households in the puzzle group to the rest of the population, in terms of demographics, income and assets. I then

while 35% of the expenditures are measured by questions on quarterly spending, and then divided into three average-monthly amounts. The latter procedure applies to food, for example. This procedure will understate volatility of consumption of such goods; see below.

[4] I choose the $500 threshold to follow other literature on this subject. Higher thresholds still yield a significant puzzle in the data, and the subgroups' characteristics are robust to the threshold as well.

Table 1: The Credit Card Debt Puzzle in 2001

	Borrow	Borrow & Save	Save
Puzzle size:		*Percent distribution*	
SCF	5%	27%	68%
CEX	7%	29%	64%
Interest rates:			
Credit cards	14.8%	13.7%	9.8%
Checking accounts (avg. across groups)		0.7%	
Savings accounts (avg. across groups)		1.2%	

Notes: Credit cards are bank-type and store cards that allow revolving debt. Liquid assets are checking, savings, and brokerage accounts. Interest rates on checking and savings accounts are from a survey by bankrate.com, and represent national averages for the entire population. Credit card interest rates are self-reported in the SCF.

characterize in more detail the household use of liquid assets in consumption, and the extent of uncertainty that households may face in liquid consumption.

2.1 Demographics of the Puzzle

Table 1 measures the credit card debt puzzle in the data. I present the measurements from both data sets to demonstrate that they are close. (The groups are very similar in both surveys in terms of relevant characteristics; I omit further comparison here for space considerations.) Around 27% to 29% of the U.S. population were simultaneously borrowing and saving in 2001. Only between 5 and 7% of the population were credit card borrowers with little or no observed liquid assets, and the rest have no significant credit card debt. Notice that these numbers imply that of all habitual credit card debt revolvers, 80 to 84% have some liquid assets that they could in principle use to pay down their debt. The last three rows of the table give average interest rates that households report paying on their credit card debt versus national interest average rates on checking and savings accounts. Very few of the puzzle households report paying zero "teaser" interest rates on credit cards, and the average rate is around 14% for borrower-savers, around 15% for borrowers, and around 10% for savers. It is also clear that there is a significant premium in credit card rates, giving the appearance of a violation of the standard no-arbitrage condition.

Table 2 breaks down some of the demographic characteristics of the subgroups from the

Table 2: Demographics

	Borrow	Borrow & Save	Save	Share in Population
	Share of subgroup with characteristic			
Race: white	0.70	0.78	0.74	0.75
Marital status: married	0.48	0.62	0.58	0.59
Have dependent children	0.45	0.41	0.39	0.40
Head works full-time	0.76	0.85	0.80	0.81
Head white-collar/prof.	0.48	0.61	0.58	0.58
Education: less than HS	0.13	0.05	0.13	0.11
HS/some college	0.73	0.61	0.51	0.55
College degree or more	0.14	0.33	0.36	0.34

Source: 2001 SCF. Weighted averages within subgroups.

SCF. Each cell of the table shows a percentage of the subgroup that has the characteristic. For example, the first line shows that 70% of the borrower group, 74% of the saver group, and 78% of the borrower-saver group are white. Comparing the numbers for different characteristics to the overall sample average shown in the right column, it appears that the borrower-saver group is not demographically distinct relative to the overall population. The group is skewed very slightly toward white households (78% versus 75% overall average), toward married households (62% versus 59%), toward heads employed full-time (84% versus 81%) and in white-collar occupations (61% versus 58%). The share of households in this group with dependent children is on par with the overall average. They also tend to be slightly better educated: the group has the fewest households with education of less than high school (5% versus 11%), while the share of those with a college degree or above is the same as it is in the total sample. The saver group compares similarly to population averages, while the borrower group is the one that is least educated, comprises most unmarried households, and is skewed most toward nonwhite households.[5]

In addition, figure 1 gives the size of the borrower-saver group by age category. While there is a slight hump in this profile between ages 30 and 50, the size of the puzzle is significant in all age groups, giving the impression that life-cycle differences are not the first-order issue for this puzzle.

[5]These conclusions are confirmed in formal probit analysis, not presented here.

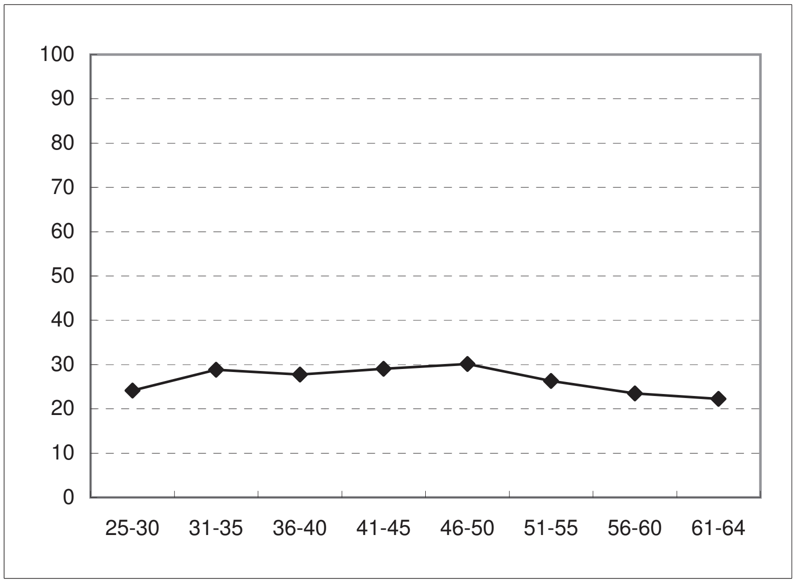

Figure 1: Size of the Puzzle Group by Age, Percent. (Source: 2001 SCF)

2.2 Asset Data

Table 3 presents income and asset information for each subgroup. The puzzle group is in the middle of the income distribution; their mean total after-tax annual income is $52,114, as compared to $64,331 for the saver group, and $28,032 for the borrower group. They hold, on average, about 1.7 times their monthly income in liquid assets (and 0.8 in the median), as compared to the liquidity holdings of the savers of 2.5 times monthly income (and equal to it in the median).[6] Several further insights are important. First, the median borrower-saver household has $3,000 in liquid assets. Another way to present this is in the last column of the table: a household with credit card debt in the 45th to 55th percentile in the borrower-saver category has median liquid assets of $4,000. Secondly, these households have significant amounts of nonliquid financial assets as well, which is inconsistent with the view that the puzzle may arise out of lack of financial sophistication.

In addition, a look at the significant and positive net worth of these households suggests

[6]A concern may arise that these numbers could be collected at the beginning of the month, say, when the paycheck has just arrived into the account. As per the Federal Reserve Board of Governors, which collects the data, SCF interviews are conducted throughout the month, and these asset numbers, averaged across households, thus represent a monthly average on the account. The Federal Reserve Board declined to release interview dates. Thus, I will treat liquidity measurements as monthly averages, and will carefully treat liquidity in the model to match the same average concept.

Table 3: Income and Asset Holding, U.S. Dollars

		Borrow	Borrow & Save	Save	B&S 45-55th debt pctile
Credit card debt:	*Mean*	5,172	5,766	317	3,622
	Median	3,340	3,800	0	3,800
Liquid assets:	*Mean*	227	7,237	17,386	7,652
	Median	200	3,000	3,200	4,000
Total after-tax income:	*Mean*	28,032	52,114	64,331	51,843
	Median	25,350	43,600	39,950	47,250
Other financial assets:	*Mean*	5,293	45,641	129,357	34,536
	Median	0	5,100	5,500	2,300
Net wealth:	*Mean*	36,231	187,912	466,462	181,413
	Median	9,450	84,640	104,830	104,450
Liquid assets as share of	*Mean*	*0.12*	*1.71*	*2.53*	*2.15*
monthly after-tax income	*Median*	*0.10*	*0.79*	*0.88*	*0.9*

Source: 2001 SCF. "Other financial assets" include IRA's, mutual funds, bond and equity holdings, annuities, life insurance. Net wealth is all financial and nonfinancial assets, net of liabilities.

that strategic bankruptcy behavior, as per Lehnert and Maki (2001), is unlikely for at least the majority of the puzzle households. Finally, note that the median household in the puzzle group, in either presentation, has credit card debt about equal to its liquidity holdings; if it were to use the liquidity to pay off debt, the household would be left with little or no money in the bank in most cases.

Table 4 shows that homeowners, especially those who still have a mortgage on their home, are more likely to be in the puzzle group. Homeowners with a mortgage constitute 59% of

Table 4: Home Ownership by Subgroup

	Borrow	Borrow & Save	Save	Share in Population
	% of subgroup with characteristic			
Own house with mortgage	0.41	0.59	0.47	0.50
Own house without mortgage	0.06	0.10	0.13	0.12
Rent	0.40	0.23	0.28	0.28

Source: 2001 SCF. Totals do not add up to one because some categories of homes are excluded.

Table 5: More Details on Financial Assets and Housing

		Borrow	Borrow & Save	Save	B&S 45-55th debt pctile
Money markets, CD's:	*% own*	3	16	23	22
	Mean	37	3,558	11,521	3,405
	Median	0	0	0	0
Mut. funds, bonds, stocks:	*% own*	19	46	44	51
	Mean	1,768	15,794	70,634	11,561
	Median	0	0	0	100
IRA's, annuities, life insurance:	*% own*	22	54	52	54
	Mean	3,287	24,670	51,735	20,757
	Median	0	1,000	500	400
Housing, net of debt:	*% own*	42	71	62	73
	Mean	11,016	50,612	82,088	50,475
	Median	0	27,000	28,000	37,000

Source: 2001 SCF.

the borrower-saver group, compared to only 50% of the overall population, while renters are underrepresented in this group. This is important because home owners are more likely to encounter significant, and potentially unexpected, expenses on homes and their maintenance.

One fact that comes out in the previous two tables is that the households in the puzzle category do have holdings of assets outside the liquid category – for example, households with the median amount of credit card debt in the borrower-saver category have a median of $2,300 in other financial assets (the corresponding mean is $34,536, table 3). The spread between the median and mean amounts is large, and the median holdings are small; nevertheless, perhaps such assets could be liquidated to pay down credit card debt, which might be a less costly option than holding on to the revolving debt. I investigate this possibility in tables 5 and 6, where I break down these asset holdings in order of decreasing liquidity, and document the frequency with which households transact in them.

Turning to table 5, it is first apparent that the majority of financial assets held by borrower-saver households are retirement accounts such as IRA's, which are subject to large penalties for early withdrawal, and thus illiquid - these constitute around 60% of financial assets. Secondly, a median household in the borrower-saver category has no other financial assets, aside from the liquid ones presented above; only 16% of borrower-saver households have CD's or money market

Table 6: Frequency of Transacting in Financial and Housing Assets

	Borrow	Borrow & Save	Save
% money mkt. holders who can write checks	0	47	66
% stock/bond holders with trading act.	7	38	48
% trading act. holders who traded[a]	100	67	78
Thus: % of stock/bond holders who did not transact[a]	93	74	63
% of home owners with cashout refinance[a]	4	3	3
% of home owners with HEL's[a]	5	1	2
% of home owners with HELOC's[b]	9	15	15

Source: 2001 SCF. (a) In the last year. (b) A current open line of credit, opened any time.

accounts, and 46% of borrower-saver households have stocks, bonds or mutual funds. Third, the majority of borrower-saver households have home equity, at about $25,000 in the median, and just over $50,000 in the mean.

Table 6 further looks at how often households with these assets report transacting in them. This information is available in some form for most of the assets in the table. First, only 47% of borrower-savers with a money market account have check-writing privileges on that account. Given that 8% of borrower-savers have money market accounts, only 4% of borrower-saver households would be able to access liquidity in a money market account easily. Second, CD's are subject to early-withdrawal penalties. Third, about 75% of those who held either stocks or bonds directly reported not transacting in them over the previous year. In addition, 43% of stock holders have stock of their own employer, and 20% of stock holders have stock *only* of their employer, further suggesting that liquidating this stock is not an option. Finally, few households report accessing their home equity. Only 4% of the borrower-saver *home owners* reported having either refinanced their mortgage with a cash-out option, or taken out a home equity loan (HEL). About 15% of the borrower-saver home owners (hence just 9.3% of all borrower-savers) reported having an open home equity line of credit (HELOC).

Thus, most households in the borrower-saver category either do not have less-liquid assets that they can liquidate to pay down their credit card debt, or even if they have them, do not exercise this option. This is consistent with these assets having high observed or unobserved transaction costs. For example, transacting in stocks and bonds requires payment of brokerage fees, but capital-gains tax considerations may add even more significant costs. Similarly, tapping home equity is typically quite costly, as appraisal and closing costs are usually at 1-2% of home equity. It may be cheaper even for those households that do have less-liquid financial assets not to exercise these options in order to repay credit card debt, which would be consistent with the behavior we observe in the data.

A related issue is the optimality of the portfolio choice in the first place where a household chooses to tie up assets in less liquid form (housing or stocks and bonds) while revolving credit card debt. From the data on holdings of less-liquid assets, it may appear puzzling that some borrower-saver households would not keep more liquidity in the bank and pay down their credit card debt, instead of allocating their portfolios to less liquid assets. However, this is not necessarily surprising. First, timing is key: these assets may have been locked up long before the household became a credit card debtor, and liquidating is then costly, as shown above. Second, thinking about housing purchase, the choice of how much money to use as downpayment is not just about the immediate tradeoff between debt and the amount of equity, but also about the terms of the mortgage for the following 30 years. For instance, putting an extra amount into a downpayment on a house may reduce the home owner's interest rate on the mortgage for the life of the loan; this benefit could easily outweigh the cost of carrying significant credit card debt at 14% for several years. Third, in relation to stocks and bonds, many of these assets are likely acquired passively by lower-income households, for example as compensation or through inheritance. The SCF partially reveals this to be true through data on holding of own employer's stock cited above.

Finally, to address the life-cycle angle one more time, figure 2 presents the breakdown of liquid assets and credit card debt for borrower-savers by age. These age profiles are fairly flat, confirming that life-cycle differences are not a key characteristic of the puzzle.

2.3 Liquidity and Consumption

Having characterized the borrower-saver group, I now turn to characterizing the role that liquid assets play in the portfolios and consumption of these households. The evidence presented below is consistent with the hypothesis that households have liquid assets to self-insure against expense shocks in goods that cannot be paid using credit cards, which may lead them to hold liquidity

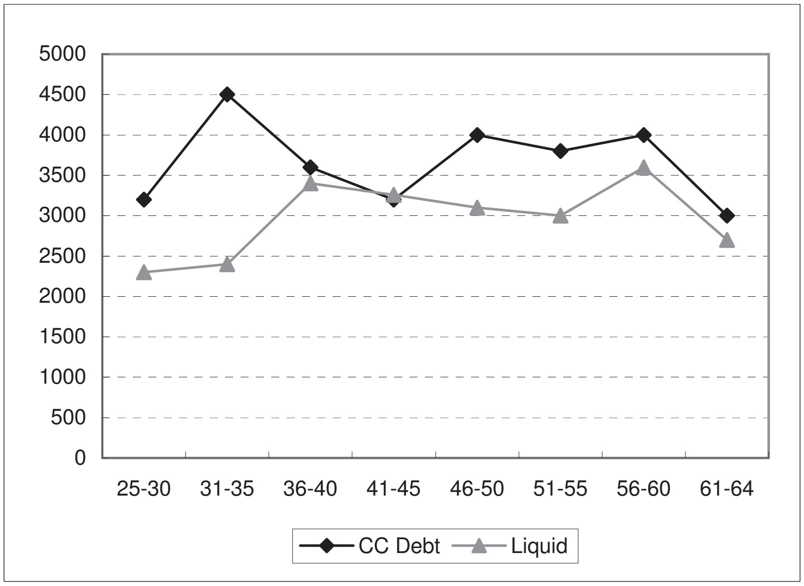

Figure 2: Median Credit Card Debt and Liquid Assets, Borrower-Saver Group by Age

simultaneously with the debt.

Table 7 shows that liquid assets retain a dominant role in consumer transactions, even though credit card usage grew noticeably between 1990 and 2002. In 2002, liquid payment methods – cash, checks, and debit cards – accounted for 77% of total consumer transactions, or 65% of their total value. Adding in electronic payments which are often backed by a checking account directly, the numbers go up to 79% and 71%, respectively. In contrast, credit cards accounted for only 24% of the value of all consumer purchases in 2002.

Table 7: Aggregate Consumer Transactions, Shares by Method of Payment

	Transaction number			Transaction value			
	1999	2000	2002	1990	1999	2000	2002
Liquid	78.2	77.8	76.7	81.2	70.3	68.8	64.9
Checks	27.9	26.9	24.4	61.3	46.2	43.9	39.0
Cash	44.2	43.5	41.3	19.6	19.4	18.9	19.5
Debit	6.1	7.4	11.0	0.3	4.7	6.0	8.4
Electronic	1.5	1.8	2.4	0.7	3.4	4.2	5.6
Credit Cards	17.4	17.7	17.6	14.5	22.5	23.9	24.0

Source: Statistical Abstract of the U.S. 2003

I next characterize, using the CEX, household-level expenditures using liquid assets. I am interested in their magnitude, as well as their volatility, as a gauge of uncertainty against which households may have to insure using liquid assets. First, I separate out the group of goods that can be viewed as payable predominantly by liquid assets. Direct information on how people pay for a given good is not collected in the CEX. I rely on the 2004 American Bankers Association (ABA) survey of payment methods to extract the relevant goods, making some conservative assumptions along the way to overcome data limitations. In this survey, consumers were asked how they normally pay at different types of stores and for different types of bills. The details of the 2004 wave of this survey are in appendix A.3. Tables A.3.1 and A.3.2 summarize the relevant information.

The ABA survey confirms the aggregate consumer transaction picture: liquid payment methods dominate household expenditures. Consumers report paying house-related types of bills, such as rents, mortgages, insurance, and utilities, by check or direct debit from the account. They also tend to pay for child care and tuition with liquid instruments, though I do not include intermittent expenses such as tuition in the cash-only group, as they are likely to skew upwards the perception of volatility. Home repairs are not asked about in the survey; however, in the SCF, households name emergencies as their number two reason for saving, preceded only by retirement planning.[7] Judging by the SCF data, households save for retirement in nonliquid retirement accounts, while emergencies, including home-related ones, by definition are likely to require liquid savings. In terms of payment methods in stores, the evidence suggests that while credit cards are dominant in department stores, gas stations and convenience stores, liquid payment methods dominate in supermarkets, drug stores, restaurants and transit systems. Backed by this information, I choose the group of cash-only goods that consists of rents, mortgages, utilities, household maintenance and repairs, household operations, property taxes, public transportation, health insurance, cash contributions, food, alcohol and tobacco. For most of these goods, a liquid payment method is required. This is not true for food, alcohol and tobacco, where consumers often have the credit option. I include these as cash goods since consumers still predominantly choose to pay for them using liquid methods; this issue is discussed in more detail in the appendix.

[7] The question reads “What are your most important reasons for saving?” Respondents get to choose as many as they want in the order of declining importance.

Table 8: Household Liquidity Holding and Consumption Patterns

		Borrow	Borrow & Save	Save
		U.S. Dollars		
Liquid assets:	*Mean*	227	7,237	17,386
	Median	200	3,000	3,200
Monthly cash-only good cons:	*Mean*	1,659	2,223	1,763
	Median	1,464	1,979	1,512
Liquid assets/cons:	*Mean*	*0.1*	*3.4*	*10.0*
	Median	*0.1*	*1.5*	*2.0*

Source: SCF, CEX. Household levels, weighted averages.

The cash good group selection is designed to be conservative. First, no durable goods, such as appliance or auto purchases, are included. Thus, for example, the cash-good category excludes many situations that may be reflections of emergencies that require liquid payment - such as an emergency purchase of (or downpayment on) a durable to replace - rather than repair - a broken one. Similarly, medical payments, which include co-pays or other out-of-pocket expenses, many of which can be unpredictable and may require a liquid payment - are not included either, because some medical expenses may be payable by credit card and I do not have sufficient information to discern the liquid portion of these payments. Instead, many of the categories that are included - such as food, property insurance, etc., - are paid on monthly basis and are predictable. Thus, in measuring the volatility of cash-good consumption using a lot of the "smooth" good categories, while excluding many that may reflect other types of emergencies, will tend to understate my measurements of the uncertainty in liquid expenditures that households face. Third, auto insurance payments and auto repairs, education expenses, and pension and insurance payments are not included as cash goods, even though many of these expenses may be liquid. I show below robustness of volatility measures to alternative definitions of the cash-only good group.

Table 8 presents household liquid asset holdings relative to average monthly consumption of cash-only goods. In the borrower-saver group, the median household has 1.5 times its average monthly liquid consumption in liquid assets, while the mean household has 3.4 times the amount. Compare these with the holdings of the savers, who have on average 10 times their mean monthly liquid spending, or twice the monthly spending amount in the median. This evidence is consistent with precautionary demand for money: households have liquid asset amounts that are in excess of what they spend on average per month, and those who are sufficiently well-off are holding much more liquidity than those in the middle. That is, richer households choose to buffer themselves more fully, while borrower-saver households may become constrained from doing so completely, but still choose not to use all of their liquid assets to pay down debt, instead keeping a precautionary balance.

To see further whether the notion of precautionary demand for liquidity is supported by the data, I look at volatility of cash-only consumption at household level as a reflection of possible uncertainty in liquid expenses that households face. Measuring raw volatility of consumption may not be fully informative about uncertainty, as it may also reflect seasonal volatility, for example, as well as other factors that may be predictable to the household. In my measurement, I first exclude from the expenditures all purchases made as gifts, to remove some of the seasonality in the consumption series. Second, I filter out the predictable component of expenditures, by estimating the following model:[8]

$$\begin{aligned} \log(c_{it}^{liq}) &= \beta\mathbf{X_{it}} + u_i + \varepsilon_{it} \qquad (1) \\ \varepsilon_{it} &= \rho\varepsilon_{i,t-1} + \eta_{it}. \end{aligned}$$

The vector $\mathbf{X}$ includes, depending on specification, household observables, such as age (a cubic), education, marital status, race, earnings, family size, home ownership status, as well as seasonal effects (a set of month dummies). Several such specifications all produced nearly identical results. u_i is the household fixed effect. The residual ε_{it} is the idiosyncratic component of liquid consumption, which I model as an AR(1) process with a normally-distributed disturbance η_{it}.

Table 9, rows 1 and 4, show variance over time of log-consumption in the cash-only good category. To construct this measure, I first take the variance of the residuals ε and η over the 12 months of observation for each household, and then average this variance across households. Thus, I get the average measure of household-level consumption volatility, which I take to

[8] One important distinction between measuring income versus consumption uncertainty is that the measures of income volatility are often translated directly into measures of income shocks, while consumption volatility reflects only the *endogenous response* of the household to its idiosyncratic shocks, which may be larger than the response. E.g. after a breakdown in the home, one may choose to make fewer repairs than is recommended, to conserve the expense. This mapping between volatility and uncertainty will be discussed further in the Calibration section.

Table 9: Average Variance of Household Cash-Good Log-Consumption, **Monthly** Data

	Borrow	Borrow & Save	Save
Liquid consumption (residual), ε_{it}			
Benchmark	0.056	0.058	0.065
Excluding food[a]	0.084	0.082	0.096
Excluding food and property taxes	0.096	0.099	0.113
Unpredictable liquid consumption, η_{it}			
Benchmark	0.053	0.057	0.064
Excluding food[a]	0.079	0.080	0.096
Excluding food and property taxes	0.092	0.098	0.113

Source: CEX. Measures variance over time of the unpredictable component of household log-consumption in cash-only goods, averaged across households. (a) "Food" includes food, alcohol and tobacco.

capture household response to idiosyncratic expense risk. First, variance of liquid consumption is significant in the benchmark measure, ranging between 0.056 and 0.065 for the total residual ε, and between 0.053 and 0.064 for η. Liquid consumption volatility is slightly higher for savers, and lowest for borrowers, which is consistent with differing ability of these groups, given their asset positions, to insure against shocks in consumption. Again, housing-related expenditures constitute the bulk of the cash-only good group and a sizeable portion of them is likely to be unpredictable. Indeed, expenses that pertain to home maintenance are the most volatile in the cash-only category, while expenses such as food are the least volatile.

To show robustness of these measures to the inclusion of food, alcohol and tobacco, as well as to the inclusion of predictable but more "lumpy" expenditures, like property taxes, I examined many different permutations of cash-good group measurements, taking out from the benchmark measure above food/alcohol/tobacco, insurance payments, property tax payments, and other predictable expenses. I present results for two such permutations: (a) the benchmark minus food/alcohol/tobacco, and (b) group (a) minus property taxes. As is evident, the more I exclude such predictable expenses, the more volatility of the remaining group increases. The benchmark cash-good category gives by far the most dampened measure of consumption volatility. To be conservative, this is the measure I will use to calibrate the model, with the understanding that it gives a lower bound on liquid expenditure volatility.

2.3.1 Consumption Volatility and Measurement Error in the CEX

When measuring idiosyncratic volatility in expenses in the CEX, an important concern that can come up is that this volatility, as measured by variance of ϵ above, or at least the transitory component η, is created not by underlying expense shocks, but by measurement error. This issue is worth examining further since the standard deviation of ϵ will be an important calibration target in the model. In this section, I use information from Attanasio et al (2011, 2004) and the BLS (Garner et al, 2006) to investigate possible nature and impact of measurement error on cash-good consumption.

For each commodity in the cash-good category, table 10 presents three pieces of information. The first shows whether the BLS considers the good to be better measured in the diary survey (DS) or the interview survey (IS), as cited by Attanasio et al (2011), Garner et al (2006), and Bee et al (2011). According to the table, the only goods in the cash-good group that have been shown to be better measured by the diary than the interview are food away from home, alcohol and tobacco. This suggests that the interview survey, which is the one I use, is best for the majority of the cash-good group; this is encouraging, since the interview covers a household for twelve months, while the diary – only for two weeks, which would make it impossible to measure household-level time variation in expenses.

The second column in the table reports the criterion that the BLS uses to study reliability of CEX measurement. This is the CEX/PCE ratio, which calculates how closely a particular commodity in the CEX, when aggregated using household weights, approximates total consumption of the same commodity in the NIPA Personal Consumption Expenditures measure. Garner et al (2006), the source of the numbers in this table, point out that many commodities in the CEX are measured differently than in the PCE, so that the aggregated CEX commodities are not always directly comparable and the ratio is often not equal to 1 for that reason. In table 10, the comparably-measured categories are marked by an asterisk(*).

As is clear, for most comparably-measured goods that are in the cash-good category, the ratio of CEX to PCE is high, approaching 1, which suggests relatively accurate measurement. For example, one key category among cash goods is the household operations category, which constitutes 56% of cash-good expenses on average, and which replicates the aggregate NIPA measure very well. The most salient exception for cash goods overall are again alcohol and

Table 10: Quality of CEX Data: Cash-Good Categories, 1997 CEX

Good	Best survey component	CEX/PCE ratio	share directly reported
Food at home*	IS	0.86	0.78
Food away *	DS	0.74	0.69
Alcohol*	DS	≈0.35	0.88
Tobacco*	DS	0.51	0.96
Household operations*	IS	1.09	0.96
owner-occupied*	IS	1.26	
rent + utilities*	IS	0.98	0.84
tenant-occupied	IS	1.05	0.85
electricity	IS	1.02	
gas	IS	0.86	
water	IS	0.69	
telephone	IS	0.82	0.99
other household ops[a]	IS	1.03	
Mass transit	IS	0.98	0.59
Taxi	IS	0.67	
Health insurance	IS	1.84	0.90
Auto repairs[b]	IS	0.67	0.91
Medical expenses, ex. insurance premia[b]	IS	0.17	0.91

(*)Category comparably measured between CEX and PCE (BLS). ([a])Other household operations include household insurance, furnishings, repairs. ([b])Not part of cash-good definition. DS and IS are diary and interview surveys, respectively. Sources: Attanasio et al (2011), Bee et al (2011), Garner et al (2006).

tobacco, and to a lesser extent, food away from home; their total share in the cash-good category is 11%. The BLS cites alcohol and tobacco as categories that are systematically underreported, due to the sensitive nature of the goods, and the CEX/PCE ratios clearly demonstrate that.

Whether one considers a CEX/PCE ratio close to 1 to be sufficient evidence of lack of measurement error depends on the model of measurement error that one has in mind. In particular, if one believes that error is mean-zero and independent across time and households, then a CEX/PCE ratio of 1 would only mean that the error nets out across households, but that at household level, there is still fluctuation created by measurement error. However, in that case we would expect the CEX/PCE ratio to be 1 for all commodities, while even for the subset of the *comparable* ones presented in the table, it is clear that the ratios are heterogeneous. If we examine comparable goods outside of the cash-good group, such as apparel, for example, we find further heterogeneity (for example, the ratio is 0.8 for shoes, 0.63 for women's apparel, etc.)

In addition, for nearly all comparably-measured goods and services, we find the CEX/PCE ratio to be either at or below 1, and very rarely above (see Garner et al (2006), table 2). By introspection, we might expect that measurement error could be the result of survey respondents forgetting some of their expenses. If so, when compared to NIPA data, aggregated CEX expenses should understate consumption, and the CEX/PCE ratios below 1 confirm that. Further, we might expect that households are more likely to forget smaller expenses on the goods that they purchase often, such as groceries or clothing, and remember well expenses that are caused by more major events, such as a repair. Examining the ratios broadly confirms this view; while household operations have CEX/PCE ratios of about 1, food and clothing have ratios of around 0.7-0.8. Under this view of measurement error caused by forgetting incidental expenses in more minor goods, a high CEX/PCE ratio is an indication of relatively less error in reporting.

This evidence suggests that if measurement error is present in the CEX, as is likely, it does not just consist of a classical error, but also has a significant, possibly dominant, memory error component. It is then useful to examine what this memory error would do not only for the mean, but also for the variance, of measured consumption, relative to true consumption. To formalize the argument, in Appendix B I present a simple example model of measurement error that includes both classical error and memory error of the sort just discussed.[9] This model demonstrates that if memory error is present, the mean of measured consumption will be smaller than the mean of true consumption, resulting in the CEX/PCE evidence that we observe in the data. In addition, I show that memory error would also reduce the *variance* of measured consumption relative to true consumption variance. Thus, if both classical and memory error are present, it is possible that measured consumption volatility could be understated by measurement error, rather than exaggerated, if memory error is sufficiently large.

Finally, the third column of the table includes the share of observations on the given good that are directly reported, as opposed to allocated or imputed, which could be an additional source of error. Again, we see that the shares of directly-reported expenses are lowest for food and alcohol, while for household operations, for instance, they are directly reported. Incidentally, this is also consistent with the presence of memory error: households may be less prone to directly report

[9] I thank Marjorie Flavin for suggesting the idea and a setup of the formal argument.

expenses that they are more likely to forget.

Table 10 includes two expense categories that are not in my cash-good definition: medical expenses, net of insurance premia, and auto repairs. I do not include them because they may be payable by credit card; however, these categories are also likely to include expenses that result from unexpected, potentially major, events and may require a liquid payment. While these categories are not defined comparably in the CEX and PCE, because in the case of both, PCE, unlike CEX, includes payments from insurance companies and nonprofits on behalf of households, they are categories that are measured well by the interview survey and are relatively free of imputation. I include them in the discussion to compare measured volatility of these expenses to that of the goods included in the cash-good group.

To sum up, although measurement error requires some speculation, evidence seems to suggest that most cash goods are measured fairly reliably nevertheless, that memory error appears to play a role, and that as a result, overall measurement error need not overstate volatility of consumption. While food away from home, alcohol and tobacco could be the least well-measured, I showed in table 9 that excluding these would make volatility of cash-good consumption even higher. I complete the discussion by measuring variability of reported consumption of each of the cash goods. Table 11 shows the results, together with average share of each good in total cash-good consumption. Note that I report coefficients of variation (CV) of consumption, rather than standard deviation or variance of logs, because for some goods, households report zero expenses in some months, which lead to a missing value when the log is taken.

It is clear that most categories are much more volatile than food, and certainly significantly more than the cash-good category overall. The CV of housing overall is 0.39, compared to 0.32 for total cash goods; the CV of household operations (which includes repairs) is 1.5. For auto repairs, the coefficient is 2.3, and for medical expenses it is 2. The orders of magnitude are such that if I were to remove the more stable items from the cash-good definition, the target volatility of consumption would increase, thus increasing the degree of uncertainty that drives precautionary demand for liquidity in the model.

It is also worth pointing out that in the data over time, there is variation in relative prices of cash and credit goods that I do not observe in my short time sample, nor model explicitly. The possibility of such price variation likely affects precautionary motive of households as well,

Table 11: Variances of Individual Cash-Good Categories

Good	Avg share of cash-goods	Avg household coef. of variation
Food at home	0.23	0.33
Food away	0.08	0.87
Alcohol	0.02	1.45
Tobacco	0.02	0.92
Housing expenses	0.56	0.39
Housing excl. mortgage, rent, tax	0.23	0.63
mortgage	0.17	0.36
property tax	0.06	0.35
repairs/maintenance	0.03	2.26
rent	0.11	0.39
utilities	0.16	0.39
other household ops.	0.03	1.52
Public transportation	0.02	2.65
Health insurance	0.05	0.71
Auto repairs[b]	0.02[c]	2.32
Medical expenses, ex. insurance premia[b]	0.02[c]	1.97
All cash goods	0.68[c]	0.32

([a])Other household operations include household insurance, furnishings, repairs not elsewhere classified. ([b])Not part of cash-good definition. ([c]) Share given is of total expenses.

and would act similarly to the idiosyncratic preference shocks in the model. Insofar as my measurements do not capture this source of variation, the measured volatility of cash-good consumption likely understates true volatility for this reason as well.

The facts that I documented in this section motivate my model and inform many of the modeling and calibration choices that I make in the next section. First, demographic analysis suggests that households in the borrower-saver category are not distinguishable from other households; in particular, there is not a pronounced life-cycle component to the puzzle, and no obvious reason to suspect that something inherent in household preferences leads them to accumulate liquid assets but not use them to repay debt. Thus, the model is an infinite-horizon one where all agents are ex-ante identical and have identical preferences; heterogeneity results ex-post from idiosyncratic risk only. Second, the majority of such households have only liquid financial assets, nonliquid retirement ones, and a house with a mortgage. Those who do have home equity or other less-liquid financial assets do not frequently transact in those assets. Based on this, the model will have two assets: a riskless bond and a liquid asset, which the agents use to buffer against income risk and preference risk. I do not model the nonliquid part of the portfolio; modeling the portfolio choice problem as a more broad one would require that this infrequency of holding and transacting is captured, which in turn would require modeling high transaction costs on an asset that is a composite of primarily a house and a retirement account, and in small part a less-liquid financial asset. This would add greatly to the complexity of the exercise, while possibly detracting from the main focus of this paper. Third, cash-only good consumption is a significant portion of household expenses, so all households in the data appear to require liquid assets for transaction purposes, including to buffer against sizeable uncertainty in consumption of such goods. The model will reflect this in the fact that households partly consume in a market where only liquid assets are accepted, and this consumption is subject to idiosyncratic preference shocks. I next develop the model, and then use it to evaluate whether precautionary demand for liquidity may lead households to keep liquid assets simultaneously with credit card debt.

3 Model

The model captures the decision problem of a household faced with a portfolio choice each period. Time is discrete. There is a [0,1] continuum of infinitely-lived agents. Each period is divided into two subperiods that differ by their market arrangements. There are two consumption goods: one consumed in subperiod 1, the other in subperiod 2. There are also two assets available to agents in each period. One is money, denoted m_{jt}, which represents all liquid assets, including checks and debit cards. Its essential feature is that it is an instant form of payment, rather than a form of credit. The subscript j stands for the subperiod, while t is for the period. The other instrument is a noncontingent bond, b_{jt}, borrowing through which at a rate r_t captures consumer credit (which can be interpreted as a credit card in the current context); saving in it is also allowed.

In the goods market in the first subperiod, either money or credit can be used in trade. In contrast, during the second subperiod, consumer credit is not allowed in trade.[10] Although mar-

[10] The question of why credit cannot be used is beyond the scope of this paper. There are several approaches

kets are competitive, they are incomplete: insurance markets are closed during both subperiods, so that the only way to insure against uncertainty is by saving.

During each period, households are subject to idiosyncratic income and preference uncertainty. There is no aggregate uncertainty. The shocks on income and preferences are independent of each other, and do not realize simultaneously. At the beginning of the first subperiod, the household's income shock s_t realizes. I model $s_t \in S$ as a discrete Markov process, with $S = \{s_1, s_2, ..., s_n\}$. The transition matrix is given by $\Gamma(s_t, s_{t+1})$, with each entry denoting probability of entering state s_{t+1} given that the currently realized state is s_t. Agents then supply labor inelastically and earn their income, consume with either credit or money, and allocate their resources between the two instruments in a household portfolio.

At the start of the second subperiod, the consumer's preference shock z_t realizes, also assumed to be a discrete Markov process with $z \in Z = \{z_1, z_2, ..., z_k\}$, and transition matrix $\Pi(z_t, z_{t+1})$. The shocks on income and preferences are assumed to be independent of each other. After the realization of z, the second subperiod's market opens. Here, households choose consumption conditional on their preference shock realization; crucially, they cannot produce or borrow in this market, so they do not have access to additional income when they need to consume. Note that the sequential timing structure is not critical for the results. The model could have the two markets co-existing in time, for example; the important feature is only that a household makes its portfolio decisions for the entire period at its start - which is realistic, given that liquid spending opportunities can arrive continually and randomly throughout the month, while additional income does not.

In the first subperiod, the household's state variables are its current knowledge of the idiosyncratic shock processes and its current portfolio: $x_{1t} \equiv (s_t, z_{t-1}, m_{1t}, b_{1t})$. The state in the second subperiod also incorporates previous subperiod's consumption: $x_{2t} \equiv (s_t, z_t, m_{2t}, b_{2t}, c_{1t})$. Agents take prices as given, so prices, or alternatively the distribution of agents, are aggregate state variables, which I make implicit in the notation. Due to the absence of aggregate uncertainty, the environment is stationary. Consequently, I will treat prices as time-invariant below.

to it in the macro literature in similar contexts: one is to assume spatial separation between the earner and the shopper, as in Stokey-Lucas-style cash-credit good models; another is to assume that agents are anonymous, as in money search models following Kiyotaki and Wright (1989). See Telyukova and Wright (2008) for a theoretical treatment of a monetary search model of money and credit that addresses the issue in more detail.

Lifetime utility, nonseparable in the two consumption goods, is given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_{1t}, z_t c_{2t}),$$

where it is assumed that $\forall\, j = \{1,2\}$, where j denotes the subperiod, $u \in C^3$, $u_j(\cdot,\cdot) > 0$, $u_{jj}(\cdot,\cdot) < 0$, and the functions satisfy Inada conditions, $\lim_{c_j \to 0} u_j(c_j, \cdot) = \infty$ and $\lim_{c_j \to \infty} u_j(c_j, \cdot) = 0$. I assume that the preference shock is multiplicative on the utility of second-subperiod consumption.

I formulate the household problem recursively.[11] In the first subperiod, a household solves the following problem:

$$\begin{aligned} V_1(s_t, z_{t-1}, m_{1t}, b_{1t}) &= \max_{c_{1t}, m_{2t}, b_{2t}} \mathbb{E}_{z_t|z_{t-1}} V_2(s_t, z_t, m_{2t}, b_{2t}, c_{1t}) \\ s.t.\ \ c_{1t} + m_{2t} &= s_t + m_{1t} + b_{2t} - b_{1t}(1+r) \\ b_{2t} &\leq \bar{B} \\ c_{1t} &\geq 0, m_{2t} \geq 0 \end{aligned} \tag{2}$$

The second constraint imposes an exogenous credit limit on the household. There is no non-negativity constraint on debt: agents can save in b_{2t}, here denoted as $b_{2t} < 0$. r is the interest rate that is charged on debt (or paid on savings) at the beginning of subperiod 1. Further, r incorporates an interest spread, as the rate of return on saving is below the borrowing rate:

$$\begin{aligned} r &= r^b \text{ if } b_{1t} > 0 \\ r &= r^s < r^b \text{ if } b_{1t} < 0 \end{aligned}$$

In the second subperiod, households choose cash-only consumption, once the preference shock realizes:

$$\begin{aligned} V_2(s_t, z_t, m_{2t}, b_{2t}, c_{1t}) &= \max_{c_{2t}} u(c_{1t}, z_t c_{2t}) + \beta \mathbb{E}_{s_{t+1}|s_t} V_1(s_{t+1}, z_t, m_{1,t+1}, b_{1,t+1}) \\ s.t.\ \ c_{2t} &\leq m_{2t} \\ m_{1,t+1} &= m_{2t} - c_{2t} \\ b_{1,t+1} &= b_{2t} \end{aligned} \tag{3}$$

[11] The Principle of Optimality applies here as is standard. In addition, existence and uniqueness are guaranteed as long as standard assumptions are made on the utility function and the constraint space to make the problem bounded.

Notice from the third constraint that no interest on consumer debt is accumulated in the second subperiod - this captures the grace period typical of a credit card billing cycle. Note also that in this subperiod, no portfolio rebalancing can take place if a household experiences a low shock and has money left over at the end of the period. This restriction is meant to capture the continual nature of the unpredictable expenses in the data: since in reality, expense shocks could hit any time throughout the month, experiencing a low expense shock at any point would not cause the household to spend the remainder of its precautionary liquid balances to pay off debt before the month is over.

Using the state-variable notation defined above, the stationary decision rules from the first-subperiod problem are $c_1(x_{1t})$, $m_2(x_{1t})$, and $b_2(x_{1t})$; the decision rule of the second-subperiod problem is $c_2(x_{2t})$. In addition, let $\lambda(x_{1t})$ and $\mu(x_{2t})$ be the Lagrange multipliers associated with the credit constraint and the money constraint, respectively. The first-order conditions that characterize the solution to the first-subperiod problem are, $\forall\, x_{1t}$:

$$-\mathbb{E}_{z_t|z_{t-1}} V_{2c}(x_{2t}) + \mathbb{E}_{z_t|z_{t-1}} V_{2m}(x_{2t}) = 0 \tag{4}$$

$$\mathbb{E}_{z_t|z_{t-1}} V_{2c}(x_{2t}) + \mathbb{E}_{z_t|z_{t-1}} V_{2b}(x_{2t}) - \lambda_t = 0 \tag{5}$$

The envelope conditions of the first subperiod are:

$$V_{1m}(x_{1t}) = \mathbb{E}_{z_t|z_{t-1}} V_{2c}(x_{2t}) \tag{6}$$

$$V_{1b}(x_{1t}) = -(1+r)\mathbb{E}_{z_t|z_{t-1}} V_{2c}(x_{2t}) \tag{7}$$

The first-order condition of the second-subperiod problem is:

$$z_t u_2(c_{1t}, z_t c_{2t}) = \mu_t + \beta \mathbb{E}_{s_{t+1}|s_t} V_{1m}(x_{1,t+1})) \tag{8}$$

The envelope conditions are :

$$V_{2c}(x_{2t}) = u_1(c_{1t}, z_t c_{2t}) \tag{9}$$

$$V_{2m}(x_{2t}) = \beta \mathbb{E}_{s_{t+1}|s_t} V_{1m}(x_{1,t+1}) + \mu_t \tag{10}$$

$$V_{2b}(x_{2t}) = \beta \mathbb{E}_{s_{t+1}|s_t} V_{1b}(x_{1,t+1}) \tag{11}$$

Combining the first-order conditions with the envelope conditions, we get the following Euler equations, which, along with the budget constraint and the Kuhn-Tucker conditions on the multipliers characterize the solution to the household decision problem:

$$\mathbb{E}_{z_t|z_{t-1}} u_1(c_{1t}, z_t c_{2t}) = \mathbb{E}_{z_t|z_{t-1}} \{\beta \mathbb{E}_{s_{t+1}|s_t} \mathbb{E}_{z_{t+1}|z_t} u_1(c_{1,t+1}, z_{t+1} c_{2,t+1}) + \mu_t\} \quad (12)$$

$$\mathbb{E}_{z_t|z_{t-1}} u_1(c_{1t}, z_t c_{2t}) - \lambda(x_{1t}) = \mathbb{E}_{z_t|z_{t-1}} \{\beta \mathbb{E}_{s_{t+1}|s_t} (1+r) \mathbb{E}_{z_{t+1}|z_t} u_1(c_{1,t+1}, z_{t+1} c_{2,t+1})\} \quad (13)$$

$$z_t u_2(c_{1t}, z_t c_{2t}) = \beta \mathbb{E}_{s_{t+1}|s_t} \mathbb{E}_{z_{t+1}|z_t} u_1(c_{1,t+1}, z_{t+1} c_{2,t+1}) + \mu_t \quad (14)$$

In this model, households will insure against income shocks by saving in the bond b (or, on the flip side, by taking out consumer loans against this asset for consumption in the first subperiod), and against preference shocks by setting aside a part of their assets in the liquid asset m each period. That is, income uncertainty does not affect precautionary motive for holding liquidity, as the portfolio can be rebalanced every period; similarly, preference uncertainty does not lead to precautionary bond demand. Note also that the bond here is modeled as one asset: households can borrow against it or save in it, but never do both. One could model this asset as two separate ones, but as long as the bond remained as liquid as it currently is in the model (i.e. no transactions costs for liquidation), and the interest rate on saving is below the interest rate on borrowing, no household would ever not deplete its holdings of the bond to repay all of the consumer debt. Hence, capturing these as one asset delivers the same result.

The population subgroups in the model are generated endogenously by idiosyncratic shock dynamics. Some agents may, after a series of low income shocks, find themselves depleting their savings in the bond and going into debt, but even when they accumulate debt, they will always choose to hold positive amounts of the liquid asset. These will be the "borrower-savers" in the model, and the point of the computational exercise is to evaluate how large this group can be in the model, and how much money they will choose to hold optimally even when it co-exists with consumer credit. The "saver" and "borrower" groups also come naturally out of the model. If a household borrows against b and then is hit by a high preference shock, and as a result spends all of its money holdings at the end of the period, it will be classified as a "borrower" for that period, as we will observe it at one point during the period as having no liquidity. If a household has savings in the asset b, then regardless of its liquidity position, it will be classified as a "saver". Clearly, households in the model will move in and out of these subgroups depending on their income and preference shock histories, so that no households would be in any subgroup permanently. This feature is mirrored in the data.

In order to compute the model, I merge the assets m_1 and b_1 into a "cash-on-hand" asset in the first subperiod. I solve the problem of the household in two stages: the first-subperiod problem (the outer maximization) is solved by value function iteration with piecewise linear interpolation, while the second-subperiod problem (the inner maximization) is solved directly from the first-order condition, by approximating the derivative of the value function. The inner maximization can, alternatively, be solved by value function iteration as well - results are robust to the choice of method.

4 Calibration

I choose model period to be a month, which is a natural frequency for studying household decisions that involve credit card statements and paychecks. The functional form for the household utility function is of the standard CRRA form, which incorporates a CES aggregator between the two consumption goods:

$$u(c_{1t}, z_t c_{2t}) = \frac{((1-\alpha)c_{1t}^{\nu} + z_t \alpha c_{2t}^{\nu})^{\frac{1-\gamma}{\nu}}}{1-\gamma} \text{ with } \gamma > 1.$$

The utility function gives three parameters to calibrate: α, ν and γ. β, the discount factor, is the fourth. The other parameters have to do with the shock processes on income and preferences, as well as prices (the interest rates). I calibrate the parameters of the income process outside the model, set γ to be in the standard range in the literature, set the interest rates to those reported in the SCF, and calibrate the remaining parameters within the model. I perform this within-model calibration by a minimum distance estimator based on the simulated method of moments. As is standard, I select the target moments so that they cover the relevant properties of data and provide discipline for the calibration, but the moments are all unrelated to the main data observations that I am trying to explain - the size of the credit card debt puzzle in the data, as well as the magnitude of household liquidity demand. Thus, these key quantities are left free to speak for the performance of the liquidity need hypothesis in accounting for the puzzle.

I set the risk aversion parameter $\gamma = 2$, chosen in the lower part of the standard range in the literature, in order to give a lower bound of the possible results. I set the debt limit B to be equal to one-half the largest annual income in the economy for all households; I discuss the sensitivity of the results to this choice below. The monthly interest rate on saving in nonliquid financial assets is set to match the annual rate of 4%, so that $r_s = 0.0033$. I set $r_b = 0.011$, which corresponds to the annual rate of 14%, the average interest rate paid on revolving credit card debt as reported by the debtors I observe in the SCF.

4.1 Income Process

The calibration of the income process at monthly frequency is a non-trivial task, as the longitudinal income data that are available and normally used for measuring idiosyncratic risk are annual. In order to calibrate the income process for the model, I use estimates of income uncertainty after observables have been controlled for from two different studies: Guvenen and Smith (2010), hereafter GS, and Heathcote, Storesletten and Violante (2010), hereafter HSV. The two estimates differ by how they control for observables; in GS, the estimation is based on a heterogeneous-income-profile specification, while in HSV, the specification is a homogeneous-profile one. From each, I use the estimates of the residual income process consisting of an AR(1) component with persistence parameter ρ_s and standard deviation of innovation σ_η, and a transitory component with the standard deviation σ_ε. The GS parameters are given at annual frequency by $\rho_s = 0.75$, $\sigma_\eta = 0.19$, $\sigma_\varepsilon = 0.15$. The HSV estimates are $\rho = 0.97$, $\sigma_\eta = 0.10$, $\sigma_\varepsilon = 0.25$. I explore these two alternatives in my model because they have very different implications in terms of persistence and overall variance of the income process, and hence different implied amounts of risk that households face.

I convert these parameters into a monthly income process by simulating, at monthly frequency, a log-income process that also has the AR(1) + transitory component structure. The simulated monthly observations are then aggregated into annual ones, on which I estimate the annual process as described above. This estimation is done by a minimum-distance estimator on the variance-covariance matrix as described in Guvenen (2009). The monthly parameters ρ_s^m, σ_η^m and σ_ε^m are estimated by recursion, such that the postulated monthly process aggregates to one with the annual parameters listed above. For the GS process, I get $\rho_s^m = 0.975$, $\sigma_\eta^m = 0.076$ and $\sigma_\varepsilon^m = 0.576$; for the HSV process, I get $\rho_s^m = 0.997$, $\sigma_\eta^m = 0.04$ and $\sigma_\varepsilon^m = 0.75$. I discretize these processes into a six-state Markov chain using the Rouwenhorst (1995) discretization method.

4.2 Idiosyncratic Preference Risk

The remaining parameters – the discount rate β, the parameters of the consumption aggregator α and ν, and the preference process parameters – are calibrated together within the model. In this subsection, I describe in some detail the calibration of preference shocks. In the subsection that follows, the remaining parameters are described.

For the preference shock parameters, I assume that the log of the preference shock, $\log(z_t)$, follows an AR(1) process with a Gaussian disturbance, so the parameters to calibrate will be a persistence parameter ρ_z and standard deviation σ_z of this process. I then discretize this AR(1) into a five-state Markov chain. The choice of an AR(1) is motivated by the idea that households have both constant pre-committed expenditures, and some additional expenditure shocks (extreme events), both of which have to be captured in the shock process. In terms of data already described, the shock's AR(1) is meant to mirror the AR(1) in the residual of liquid consumption ε_t, as described in (1).

The preference shock process is clearly not observed in the data, but the way households respond to these shocks is, through their liquid consumption. Thus, the preference shock process has to match properties of consumption of cash-only goods in the data, namely its persistence (measured as autocorrelation) and volatility (standard deviation). For the calibration targets, the standard deviation is computed by subgroup of households, so in total, I get four calibration targets for the shock process. As described extensively in the data section (table 9), the volatility of liquid consumption is sensitive to how the cash-good group is computed. Of all the measures that I have examined, the benchmark measure (the most inclusive) produces the smallest volatility of consumption; in the estimation below, I use this benchmark measure for maximum discipline on the model, but this will also make the results on the puzzle a lower bound.

In order to convince the reader that normal disturbances are a reasonable assumption for the shocks, and also that the calibration does not overstate the tail shocks through such a representation, figure 3 plots the consumption residual ε_{it} for the benchmark measure of liquid consumption, together with a nonparametric kernel estimator of its density (thick red line) and the corresponding normal approximation (thin green line).

According to the graph, normal distribution approximates the actual distribution well. While

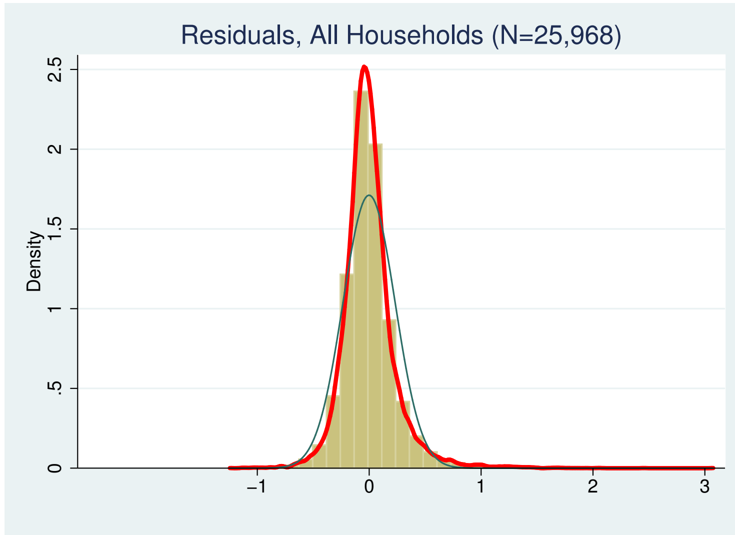

Figure 3: Residuals of Log-Consumption of the Cash Good, Benchmark Measure, All Households

it understates the density of consumption at the mean (which will be corrected by the fact that I will match the autocorrelation of consumption as a targeted moment), it overstates consumption very near the mean (within one standard deviation), and understates the tails of the density. Recall that the right tail of the distribution is key for determining money holdings, as it maps to the highest preference shock, which causes the money constraint of the household to bind, which in turn determines liquidity demand. I treat this right tail conservatively in the calibration. First, in the discretization of the shocks by the Tauchen method, I restrict the five discrete states to fall within just two standard deviations of the mean. This approximates the AR(1) well, but is clearly an understatement of the tails in the data. Also, I only target up to the second moment of the residual distribution of liquid consumption in the data, without attempting to match the tails.

4.3 Remaining Parameters and Mapping to Targets

The remaining parameters are β, α and ν. To calibrate these, I add three more targets: the mean revolving debt-to-income ratio in the population, the share of liquid consumption in total household consumption, and a measure of the elasticity of substitution between cash and credit goods. Although all seven (together with properties of liquid consumption above) calibration targets interact and jointly determine all five parameters, it is clear that β is pinned down primarily by the debt-to-income ratio, that the time series properties of liquid consumption help determine the shock process, and that α and ν will be pinned down by the cash-good share of consumption and the elasticity measure.

The estimation of the parameter ν as part of the SMM procedure deserves some attention, as it is nonstandard. Previous estimates of this parameter come from deterministic cash-credit good models, in which the cash-in-advance constraint always binds, so that aggregate cash-good consumption equals aggregate money demand. A direct implication of this class of models is that ν can be measured in closed form from the regression coefficient that gauges sensitivity of aggregate money demand to the gross nominal interest rate. In contrast, in my model cash-good consumption and money demand are distinct, due to the presence of idiosyncratic risk, so that the cash-in-advance constraint is often not binding. As a result, the model no longer has a closed-form implication for the parameter ν, and the implications for the elasticity of substitution between the two goods will be dramatically different as well. This is discussed in more detail in the next section.

Instead, I run a similar regression on household-level data, using the interest rate that the household currently faces as the measure of the opportunity cost of holding cash, and thus of the cost of the cash good. The regression is: $\ln(c_{2i}/c_{1i}) = \kappa_0 + \kappa_1 \ln(1+r_i) + \omega_i$. Those who are debtors face a different opportunity cost of the cash good ($r_i = r^b$) than those who are savers ($r_i = r^s$); I measure the sensitivity of the cash-credit good ratio to the cross-sectional variation in this cost. Since this regression parameter does not directly translate into the parameter ν in closed form, I will run the same regression on simulated model data and estimate ν such that the two regression coefficients match.

The CEX does not provide information on the interest rate that the households are paying on their credit card. To measure r^b in the CEX, I use the fifth-interview question on the finance charges that the household reports paying on the credit card, dividing that amount by the average of the two credit card balances reported by the households in the second and fifth interviews. For households who are not debtors, I assign the current Federal Funds rate (based on month and year) as the bond interest rate.[12] The resulting regression gives the coefficient of about -0.29 on the interest rate (with a standard error of 0.037), suggesting a

[12] The results are robust to using the 3-month Treasury Bill rate instead.

Table 12: Calibration Targets - Data and Model

Target	Data	Model (GS inc.)	Model (HSV inc.)
Autocorrelation (annual) of log liquid consumption:	0.226	0.226	0.225
St. dev. of log liquid consumption: Borrowers	0.203	0.158	0.151
Borrowers & savers	0.211	0.212	0.213
Savers	0.217	0.248	0.235
Share of liquid cons in total	0.683	0.628	0.672
Regression coefficient: $log(c_2/c_1)$ *on* r	-0.285	-0.257	-0.273
Mean debt/income ratio	0.070	0.071	0.070

negative relationship, as would be expected, but little sensitivity of the consumption ratio to the interest rate.

I measure the cash-good share of household consumption directly in the CEX at household level, then average across households. The last target, the debt-to-income ratio, is meant to gauge how well the model does in reproducing the most relevant dimension of the aggregate economy, given the paper's focus. As the debt measure, I choose total revolving debt, computed in the SCF. This measure includes all unsecured debt as well as home equity lines of credit, with credit card debt in the vast majority, since the uptake of home equity lines is very low in the 2001 data.

In sum, I estimate the five parameters within the model based on seven moments. For each set of parameters in the minimization process, the procedure solves the model, simulates a 502-month panel of 100,000 households, computes the moments from it, and compares them with the moments in the data. I use the simplex method of Nelder and Mead (1965), parallelized at parameter level as suggested by Lee and Wiswall (2007). The weighting matrix is the identity matrix in the first step, subsequently adjusted to correct for moments computed with highest variance (those moments that concern the borrower group, which is smallest in the data). Data covariances of the moments in question are not possible to compute in this exercise, since the moments come from two different data sets.

Table 13: Calibration

	Parameter	Model (GS)	Model (HSV)
Interest rates	r_s	0.0033 (annual $r_s = 0.04$)	0.0033
	r_b	0.0107	0.0107
Risk aversion/IES	γ	2.0	2.0
Income process parameters	ρ_s^m	0.975	0.997
	σ_η^m	0.076	0.040
	σ_ε^m	0.576	0.750
Discount rate	β	0.992	0.992
Consumption aggregator parameters	α	0.934	0.956
	ν	-3.482	-3.260
Preference shock process: AR(1) with discretization	ρ_z	0.536	0.504
	σ_z	0.947	0.677

5 Model Fit and Resulting Parameters

From this point forward, I will discuss the results of two calibrations, one with the GS income process, and one with the HSV process. In each case, all the parameters are re-estimated to the same targets. In order to assess the fit of the calibrated model, table 12 presents the target moments in the data and the model.

The calibrated model fits most targets closely. The model with the GS income calibration understates the share of liquid consumption in total consumption slightly (model's 0.63 versus the data's 0.68). It matches nearly perfectly the standard deviation of log-liquid consumption for the borrower-saver group, although it overpredicts somewhat the dispersion of that moment across groups. All other targets are matched very well. Under the HSV calibration, this pattern is repeated, although the share of liquid consumption in total is better matched, and the dispersion of standard deviation of liquid consumption for savers is closer to the data as well.

Table 13 presents the two resulting calibrations; the parameters are similar in both, with the exception of σ_z. The discount factor in the two calibrations is equivalent to 0.91 in annual terms. The parameters of the CES aggregator are in themselves of interest and a contribution of this paper: to date, these parameters for aggregators of cash and credit goods were estimated in deterministic cash-credit models only. In order to match the high cash-good share of consumption, α, the weight on cash-only goods in the CES utility function has to be high, at 0.93-0.96. The parameter that measures elasticity of substitution between cash and credit goods is between -3.26 and -3.48, that is, cash and credit goods are complements rather than substitutes, with the elasticity of substitution of 0.23-0.25.

It is worthwhile to compare these estimates to previous ones from the deterministic cash-in-advance literature. For example, Chari et al (1991) and others after them find a lower estimate for α of around 0.62. Their ν tends to be on the order of 0.79-0.84, producing the elasticity of substitution on the order of 4.76 to 6. These results come from the model where, as I discussed above, (aggregate) money demand and liquid consumption are equal to each other. Once idiosyncratic preference uncertainty is introduced into the model, aggregate liquid consumption becomes much less sensitive than money demand to nominal interest rates, because all the households who do not hit their liquidity constraint have liquid consumption that is not sensitive to the movement in interest rates. Insofar as the elasticity of substitution can still be linked to this sensitivity, the elasticity should be expected to be much lower in this model than in the deterministic version, which is confirmed in my results using micro data.

Finally, the estimates of the preference process are of importance, since this study presents a new, and to my knowledge, first, effort to quantify unobservable idiosyncratic uncertainty to preferences from micro data specific to liquidity needs. The estimated monthly AR(1) parameter on $\log(z)$ is around 0.50-0.54. The AR(1) specification is flexible, encompassing anything from a very persistent shock process to an i.i.d. one; these estimates suggest that the extreme realizations of the shock are not persistent, which is to be expected given the properties of liquid consumption discussed in the data section. The standard deviation of the shock process is estimated at 0.947 in the GS calibration, but only 0.677 in the HSV calibration. Under HSV, the agents face more persistent income risk and higher income dispersion, so the preference shocks need not be as disperse to match household-level standard deviation of liquid consumption over time. Moreover, as partial insurance is always optimal and agents prefer to smooth consumption, it is intuitive that the observed consumption process "mutes" the variability of the underlying shock process.

Based on the analysis of the calibration targets, the parameterization described above produces a realistic economy in terms of the mapping to the relevant data dimensions. Before I

Table 14: Cross-Sectional Dispersion of Annual Income and Consumption

Target	Data	Model (GS)	Model (HSV)
Variance of income (GS)	0.11	0.13	
Variance of income (HSV)	0.25		0.31
Variance of cash-only consumption	0.17	0.08	0.31
Variance of cash-or-credit consumption	0.54	0.07	0.24
Variance of total nondurable consumption	0.20	0.07	0.29

Notes: The measures of variance here refer to the residual variance of the log-variable, once observables have been controlled for. See text for details.

turn to describing the model's implications for the credit card debt puzzle itself, I examine an additional set of the model's implications: namely those for cross-sectional dispersion of income and consumption. Table 14 assesses the performance of the model in this respect.

Overall cross-sectional dispersion in the data results from many types of heterogeneity across households, which the model abstracts from.[13] The income calibration takes this into account by putting in only the residual income dispersion, estimated after observables, in the form of factors such as labor experience, have been controlled for in panel data. I construct a comparable measure of residual variation of consumption in the data, but I do not have information on labor market experience in the CEX annual cross-section. The consumption variance measures used in table 14 are derived using the HSV methodology, by first regressing log-annual consumption (cash, credit and total nondurable consumption) on a cubic in potential labor market experience (age - years of education - 5) and an education dummy, and extracting the residual.

The variance of income in the model is by construction close to that in the data (the discrepancy comes from the fact that I decompose the income process into a monthly one, discretize, and then re-aggregate - with resulting approximation issues). The resulting variance of total (nondurable) consumption varies, as is to be expected for a model of this kind, based on the properties of the stochastic process on income. The model with the GS income calibration, which puts more emphasis on transitory variation in income, produces variance of total consumption that is under one-half of that in the data. In contrast, the model with HSV calibration produces

[13] In my sample, overall cross-sectional variance of log-income is 0.45 and of log-nondurable consumption is 0.24, which is in line with the numbers reported, e.g., by Krueger and Perri (2005).

variance of total consumption that is only slightly lower than variance of income, and so overstates dispersion by about one-third. This is not surprising, since in the HSV calibration, the income shocks have significantly higher persistence than the GS case. It is well-known that in models of this type, agents are unable to self-insure against very persistent shocks well, and do much better with transitory shocks, and this model confirms that result.

In addition, it is worthwhile to look at the components of consumption. In the GS calibration, cash-good consumption variance is nearly equal to credit-good consumption variance, and both are at about half of income variance. In the HSV calibration, cash-good consumption variance is higher than credit-good consumption variance; there is noticeable consumption-smoothing in credit goods, but not in cash goods. The intuition comes from the two sources of idiosyncratic uncertainty which both contribute in different ways to creating fluctuations in consumption. Preference shocks create individual-level volatility in cash-good consumption. The cross-sectional dispersion in both credit-good and cash-good consumption is driven primarily by income shocks, since these create long-run dispersion in wealth, which in turn translates into differing abilities of households to self-insure against the expense shocks as well. In addition, the dispersion is reinforced by complementarity between the two consumption goods in either calibration. Thus, the logic for smoothing in each component of consumption is still linked primarily to the properties of the income process; the more persistent this process, the harder it is for agents to self-insure against the fluctuations in either component of consumption.

Finally, notice that both models produce the counterfactual result that credit-good consumption variance is equal to or below that of cash goods, and under both calibrations, the model underpredicts volatility of credit-good consumption. This is likely the result of the fact that in the data, many credit goods are semi-durable (e.g. clothing, vehicle maintenance expenses, etc.), and are possibly also affected by expense shocks that I do not model.

6 Results

6.1 The Credit Card Debt Puzzle

To measure how much of the puzzle is accounted for by the liquidity need hypothesis, I focus on the size of the subgroups (especially the borrower-saver group), as well as liquidity holdings of the borrower-saver group.

Table 15: Results - Subgroup Size (Percent)

	Data	Model (GS)	Model (HSV)
Borrowers	5.2	2.2	8.1
Borrowers & savers	27.1	11.8	15.1
Savers	67.7	86.0	76.8

Table 16: Results - Liquid-Asset-to-Income Ratio, Median Household

	Data	Model (GS)	Model/Data	Model (HSV)	Model/Data
Borrowers	0.10	0.73	7.3	0.99	9.9
Borrowers & savers	0.79	0.82	1.04	1.07	1.35
Savers	0.88	0.73	0.83	0.49	0.56

Table 15 gives the size of the three subgroups in the data and the model. In the model, the size of the borrower-and-saver group is between 12% of the population in the GS calibration, and 15% in the HSV calibration, while in the data, it is 27%. Thus, the model, depending on the calibration, accounts for between 44% and 56% of the puzzle group size. Further, the model overstates the size of the saver group. Under the GS calibration, the size of the borrower group is understated. The reasons for why the borrower group is underpredicted is that the model's borrower group consists only of those who are constrained at the end of the month, while in the data, there may be some households who have very few liquid assets throughout the month and year; these households are not captured by the model, since it is never optimal to hold zero liquidity. Under the HSV calibration, the size of the borrower group is larger than in the data: more households under this calibration who are constrained at the end of the period are also debtors; this is likely driven by a higher dispersion of the low income states.

In order to measure liquid assets, I have to map money holdings in the data to those in the model. As I discussed, a cross-sectional average of money holdings in the SCF reflects an average monthly amount of money in the bank accounts. Since in the model I observe money holdings at two points during the month, I study average monthly money holdings for all households to map to the observed amount in the data.

In table 16, I present average monthly liquid asset holdings relative to income for a median household in each subgroup, in the data and in the model. The "Model/Data" columns translate the model's results into per-dollar amounts relative to the data. For the median household in the borrower-saver group, the model matches its liquidity holdings perfectly; in the HSV calibration, it even overstates them. The model accounts for 83% of the liquidity holdings of the saver group under the GS calibration. Under both calibrations, and particularly under the HSV one, the model counterfactually predicts savers' liquidity-to-income ratio to be below that of borrowers-savers. This is because money holdings in the model are not as disperse as in the data (in particular, the upper tail is not matched), and income dispersion is higher under the HSV calibration. Looking separately at distribution of money holdings and income in the model by subgroup (not presented) reveals the expected increasing pattern of both from borrowers to savers. [14]

For the borrowers, the model generates over 700% of the money holdings in the data, based on the average monthly liquidity holdings; even higher under the HSV calibration. This comparison between the model and data is the most tenuous; it is not likely that *all* borrower households in the data never hold liquid assets during the month, given their average liquid spending documented above. Many of these may instead be households observed at the end of the month, when they have drawn down all of their liquid assets, most likely due to binding resource constraints. Thus, it may be more appropriate, in the borrower case, to compare the data number to the end-of-month liquid holdings in the model, which are 0, in which case the model and data once again match.

As one sensitivity check, I re-compute the model with the GS calibration and a lower borrowing limit, assuming that all households can borrow only up to one-half of the annual *average* income in the economy. If I do not re-estimate any of the parameters of the model, the central implications of the model are all robust to the re-specification; in particular, the size of the borrower-saver group, as well as the liquidity holdings of each group do not change. However, there is one important and predictable change: the debt-to-income ratio in this economy falls

[14] I use the median household as the comparison point in the data because the model is not calibrated to account for the upper tail of the income or wealth distributions, and thus has a hard time accounting for the dispersion between the median and the mean in the data: the liquidity/earnings ratio in the data for the borrower-saver group is 0.8 in the median, but 1.7 for the mean, while in the model they are much closer to each other. Still, the model generates about 46-53 cents of the dollar held by the *mean* borrower-saver household in the data, accounting for nearly one-half of the average puzzle household's liquidity holdings.

Table 17: Shock Decomposition - Subgroup Size (Percent)

	Data	Model (GS, no shocks)	Model (GS, no pref. shock)
Borrowers	5.0	0.0	15.0
Borrowers & savers	27.1	0.0	0.0
Savers	67.7	100.0	85.0

Table 18: Shock Decomposition - Liquid-Asset-to-Income Ratio, Median Household

	Data	Model (GS, no shocks)	Model/Data	Model (GS, no pref. shock)	Model/Data
Borrowers	0.10	–	–	0.50	5
Borrowers & savers	0.79	–	–	–	–
Savers	0.88	0.34	0.39	0.36	0.41

by more than half. If I instead re-estimate this economy to match the debt-to-income ratio in the data, the parameters change; notably, the discount factor β has to be lowered somewhat (to about $\beta = 0.987$). In this circumstance, the size of the borrower-saver group in the economy comes up to 28%, thus matching the size of the puzzle in the data, while optimal liquidity demand remains robust.

6.2 Model Decomposition: Role of the Shocks

In order to understand the role of the idiosyncratic income and preference shocks in generating the credit card debt puzzle, I re-compute two versions of the model with the GS calibration, first shutting down all shocks, and then shutting down only preference shocks. All the parameters are kept the same as in the baseline calibration. Tables 17 and 18 present the results for the credit card debt puzzle. It is immediately clear that unless both shocks are present, the model cannot account for the puzzle, since it does not generate a borrower-saver group; moreover, the model in that case severely underpredicts liquidity demand.

The model with both the income and preference shocks shut down (referred to as "GS, no shocks" in both tables) reduces to a deterministic representative-agent model. In this model, there is no dispersion of consumption, and in steady state, the model would generate constant consumption in both cash and credit goods. The computation of the model, given the benchmark calibration, produces a solution that is just out of the steady state (as the condition $\beta = 1 + r$ is violated), so there is a very slight decline in household consumption over time. However, the solution (as well as the steady state) features positive savings, and thus, the household is a saver. The size of the borrower-saver group is 0. The liquidity holdings of the household are at 34% of its income, which is about 39% of the data. Thus, the model with all the shocks shut off generates average monthly liquidity holdings for savers of just under 47% of the full model.

If only the preference shocks are shut off ("GS, no pref. shock"), there is still no borrower-saver group in the model. This model is a heterogeneous-agent model due to idiosyncratic histories of income shocks, but there is no precautionary liquidity demand here, since cash-good consumption is deterministic. The dispersion properties of this model (not in the tables) are nearly identical to those of the GS benchmark model, confirming the intuition that it is the income uncertainty that drives cross-sectional dispersion of consumption in both goods. Because in this model everyone's liquidity constraint binds at the end of the period, anyone with debt here is a borrower. This group is 15% of the population in the stationary distribution. A saver household's liquidity holdings are essentially the same as in the no-shock model, confirming the intuition that income shocks alone do not alter the size of liquidity demand. At 36% of household income, this version of the model generates 41 cents of each dollar held by the median saver household in the data, and accounts for less than 50% of the savers' liquidity holdings in the full model. In sum, preference shocks are necessary for generating the borrower-saver group in the model, and for matching the magnitude of liquidity holdings of this group in the data; when present, these shocks at least double the liquidity demand relative to the model without such shocks.

6.3 Discussion of the Results

The results presented above lead to the conclusion that precautionary demand for liquidity is a key factor in accounting for the credit card debt puzzle. In addition, there are many reasons why the results presented above may be seen as a lower bound on both liquidity demand and the size of the borrower-saver group. For the former, money demand can be directly affected by aspects not captured by the model; one that comes to mind is the minimum balance requirement on checking accounts. Many checking accounts allow their holders to avoid sizeable fees by maintaining a minimum balance in the account at all times. Anecdotally, this minimum balance requirement can go as high as $1,000 or more. I do not account for such a requirement in the model, in large part because I do not have data on what these requirements might be and what the share of the population is that has them. If, however, it is assumed that many or all checking account balances have some minimum positive amount that they need to exceed, then the total amount of liquidity that the model can account for will rise, possibly substantially. The argument would, of course, be more nuanced given that one would have to consider when it may be optimal to dip below the minimum balance for a household that finds itself in the borrowing-and-saving situation. But if this situation is temporary, this channel may still increase the puzzle household's liquidity demand in the model, and it will certainly increase the demand of saver households; this may be one way to help account for average liquidity holdings as well.

In terms of both liquidity demand and puzzle size, the model currently captures only one channel that gives rise to precautionary liquidity demand, namely, preference uncertainty. There is, of course, a second source of uncertainty in the model - income uncertainty - but it plays a role only in generating disperse nonliquid asset holdings, as households insure against this shock by saving or borrowing in the asset b. The reason for the lack of a link between income uncertainty and liquidity demand is that it is costless in the first subperiod to acquire additional liquidity from a credit card in the event of a low income shock. Yet this link may be important: even predictable expenses may require precautionary money holdings in the face of income risk. For example, if one should lose one's job and paycheck, one still needs to pay the mortgage each month. The idea, then, is that both preference (expense) and income uncertainty, both of which are present in the data, may provide a precautionary motive for holding liquidity for most households, and if this channel is modeled by adding a direct cost of transfers from consumer credit to liquidity, both the size of the borrower-saver group and liquidity holdings in the model would likely increase.[15]

To summarize, while the discussion of these additional channels of precautionary liquidity demand would have to be more rigorous and nuanced, as many tradeoffs need to be evaluated, it

[15]This intuition was confirmed in two-period examples, since a full extension of the model becomes significantly harder computationally. For example, liquidity demand in an extended example increases by 30-50% relative to the benchmark.

can be reasonably concluded that a model that abstracts from behavioral, self-control, or strategic bankruptcy considerations, when carefully calibrated, goes a long way toward accounting for the facts of the puzzle in the data, and that the results presented above can, for several reasons, be seen as a lower bound on accounting for the credit card debt puzzle.

7 Conclusion

This paper presents the first examination of liquidity demand as an explanation for the credit card debt puzzle. I examine the hypothesis that there is a significant share of household expenditures each month that cannot be paid by credit card, so that households need to keep liquidity in the bank at all times to pay for these expenditures. Thus, if a household accumulates credit card debt, but does not have enough money both for its needed precautionary amount and for debt repayment, it will optimally choose to revolve the debt in favor of keeping a sufficient supply of liquidity.

The central contribution of the paper is a detailed measurement of how much of the puzzle this hypothesis can account for. First, I document the puzzle carefully in the data, which requires a novel split of household consumption into cash-only goods and cash-or-credit goods, based on survey evidence. I then pose a dynamic stochastic model of household portfolio choice with two consumption goods and types of idiosyncratic uncertainty timed so that portfolio decisions have to be made before spending needs are known. This model successfully accounts, qualitatively, for the salient empirical features of the credit card debt puzzle. The model is then calibrated via a disciplined match of moments in the data to model moments. The parameter estimates are in themselves of interest, providing new or even first measurements of magnitudes of unobservable idiosyncratic expense risk and the elasticity of substitution between cash and credit goods in micro data. For example, whereas in deterministic cash-credit good models cash and credit goods were estimated to be substitutes, my model implies that they are complements.

In terms of the puzzle, I find that, depending on the calibration, the hypothesis successfully accounts for at least one-half of the households who revolve debt while having money in the bank, and under some calibrations can account for the entire group. For a median such borrower-saver household, the model accounts for their entire liquid asset holdings. I also decompose the model in terms of the role of the two types of idiosyncratic shocks, and find, for example, that one-half of the liquidity demand in the model is precautionary. Thus, even though there are likely households for which alternative explanations along the lines of time inconsistency or strategic bankruptcy behavior are valid, or even dominant, it should be apparent that liquidity demand - including precautionary demand for liquidity - is a factor that goes a long way toward accounting for the puzzle.

References

[1] Aiyagari, S. R. and M. Gertler. 1991. "Asset Returns with Transactions Costs and Uninsurable Individual Risk." *Journal of Monetary Economics*, 27(3), 311-331.

[2] American Bankers Association / Dove Consulting. "2004 Study of Consumer Payment Preferences".

[3] Attanasio, O., E. Battistin and M Padula. 2011. *"Inequality in Living Standards since 1980"*, The AEI Press.

[4] Attanasio, O., E. Battistin and H. Ichimura. 2004. "What really happened to consumption inequality in the US?". Mimeo, UCL.

[5] Ausubel, L. 1991. "The Failure of Competition in the Credit Card Market." *The American Economic Review*, 81, 50-81.

[6] Bee, A., B. Meyer, and J. Sullivan. 2011. "Micro and Macro Validation of the Consumer Expenditure Survey". NBER conference slides.

[7] Bertaut, C. and M. Haliassos, 2002. "Debt Revolvers for Self Control." University of Cyprus.

[8] Baumol, W.J. 1952. "The Transactions Demand for Cash: an Inventory Theoretic Approach." *Quarterly Journal of Economics*, 66, 545-556.

[9] Blundell, R., L. Pistaferri, I. Preston. 2008. "Consumption Inequality and Partial Insurance". *American Economic Review*, 98:5, 1887-1921.

[10] Brito, D. and P. Hartley, 1995. "Consumer Rationality and Credit Cards". *Journal of Political Economy*, 103(2), 400-433.

[11] Bruce, L. 2002. "A Graphic Look at Deposit Rates, 1998-2002." bankrate.com.

[12] Cargill, T. and J. Wendel, 1996. "Bank Credit Cards: Consumer Irrationality versus Market Forces?" *The Journal of Consumer Affairs*, 30, 373-389.

[13] Chari, V. V., L. Christiano and P. J. Kehoe, 1991. "Optimal Fiscal and Monetary Policy: Some Recent Results". *Journal of Money, Credit and Banking*, 23, 519-539.

[14] Chatterjee, S., D. Corbae, M. Nakajima, J.V. Rìos-Rull, 2007. "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default." Forthcoming, *Econometrica*.

[15] Duca, J. V., and W. C. Whitesell, 1995. "Credit Cards and Money Demand: A Cross-Sectional Study." *Journal of Money, Credit and Banking*, 27, 604-623.

[16] Durkin, T., 2000. "Credit Cards: Use and Consumer Attitudes, 1970-2000." *Federal Reserve Bulletin,* September.

[17] Durkin, T., 2002. "Consumers and Credit Disclosures: Credit Cards and Credit Insurance." *Federal Reserve Bulletin,* April.

[18] Eisfeldt, A., 2003. "Smoothing with Liquid and Illiquid Assets". Kellogg School of Management.

[19] Erosa, A. and G. Ventura, 2002. "On Inflation as a Regressive Consumption Tax." *Journal of Monetary Economics*, May.

[20] Evans, D. and R. Schmalensee, 1999. *Paying with Plastic: The Digital Revolution in Buying and Selling.* MIT Press.

[21] Federal Reserve Board. Statistical Release G.19, Consumer Credit. Various issues.

[22] Fernàndez-Villaverde, J. and D. Krueger, 2004. "Consumption over the Life Cycle: Facts from the Consumer Expenditure Survey." University of Pennsylvania.

[23] Garner, T., G. Janini, W. Passero, L. Paszkiewicz, and M. Vendemia. 2006. "The CE and the PCE: A Comparison". *Monthly Labor Review*, September, p. 20-46.

[24] Gourieroux, C. and A. Monfort, 1996. *Simulation-Based Econometric Methods*, Oxford University Press.

[25] Gross, D. and N. Souleles, 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics,* February, 149-185.

[26] Gross, D. and N. Souleles, 2001. "An Empirical Analysis of Personal Bankruptcy and Delinquency." Wharton School, University of Pennsylvania.

[27] Guvenen, F. and A. Smith. 2010. "Inferring Labor Income Risk from Economic Choices: an Indirect Inference Approach". Mimeo, University of Minnesota.

[28] Haliassos, M. and M. Reiter, 2003. "Credit Card Debt Puzzles." University of Cyprus.

[29] Hayashi, F. and E. Klee, 2003. "Technology Adoption and Consumer Payments: Evidence from Survey Data." *Review of Network Economics,* 2, 175-190.

[30] Heathcote, J., K. Storesletten, G. Violante. 2010. "The Maroeconomic Implications of Rising Wage Inequality in the United States". *Journal of Political Economy*, 118, 681722.

[31] Hopenhayn, H. and E. Prescott. 1992. "Stochastic Monotonicity and Stationary Distributions for Dynamic Economies," *Econometrica*, 60(6), 1387-1406.

[32] Jasper, M. 2000. *Legal Almanac: Credit Cards and the Law.* Oceana Publications, Inc.

[33] Johnson, K.W., 2004. "Convenience or Necessity? Understanding the Recent Rise in Credit Card Debt." *Finance and Economics Discussion Series*, Federal Reserve Board.

[34] Krueger, D. and F. Perri, 2003. "On the Welfare Consequences of the Increase in Inequality in the United States", in *NBER Macroeconomics Annual 2003*, M. Gertler and K. Rogoff, eds. The MIT Press.

[35] Krueger, D. and F. Perri. 2005. "Does Income Inequality Lead to Consumption Inequality?" *Review of Economic Studies*, 73(1), 163-193.

[36] Lacker, J. and S. Schreft. 1996. "Money and credit as means of payment". *Journal of Monetary Economics*, 38(1), 3-23.

[37] Laibson, D., A. Repetto and J. Tobacman, 2002. "A Debt Puzzle." Harvard University.

[38] Lee, D. and M. Wiswald, 2007. "A Parallel Implementation of the Simplex Function Minimization Routine". Forthcoming, *Computational Economics.*

[39] Lehnert, A. and D.M. Maki, 2002. "Consumption, Debt and Portfolio Choice: Testing the Effect of Bankruptcy Law." Board of Governors of the Federal Reserve Bank.

[40] Lucas, R., E. Prescott and N. Stokey, 1989. "Recursive Methods in Economic Dynamics". Harvard University Press.

[41] Nelder, J. A., and R. Mead. 1965. "A Simplex Method for Function Minimization". *The Computer Journal.*

[42] Prescott, E.S. and D.D. Tatar, 1999. "Means of Payment, the Unbanked, and EFT'99." *Federal Reserve Bank of Richmond Economic Quarterly,* 85, 49-70.

[43] Quadrini, V. and J.-V. Ríos-Rull, 1997. "Understanding the U.S. Distribution of Wealth." *Quarterly Review,* Federal Reserve Bank of Minneapolis, spring, 22-36.

[44] Rouwenhorst, K. 1995. "Asset Pricing Implications of Equilibrium Business Cycle Models". In "Frontiers of Business Cycle Research", T.F. Cooley, ed. Princeton University Press.

[45] Schreft, S. 1992. "Transaction Costs and the Use of Cash and Credit". *Economic Theory*, 2(2), 283-96.

[46] Stavins, J., 2000. "Credit Card Borrowing, Delinquency, and Personal Bankruptcy." *New England Economic Review,* July/August.

[47] Storesletten, K., C. I. Telmer and A. Yaron, 2004a. "Consumption and Risk Sharing Over the Life Cycle". *Journal of Monetary Economics*, 51, 609-633.

[48] Storesletten, K., C. I. Telmer and A. Yaron, 2004b. "Cyclical Dynamics in Idiosyncratic Labor-Market Risk". *Journal of Political Economy*, 112, 695-717.

[49] Tauchen, G., 1986. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions". *Economics Letters,* 20, 177-181.

[50] Telyukova, I.A. and L. Visschers, 2011. "Precautionary Demand for Money in a Monetary Business Cycle Model". Mimeo, UCSD.

[51] Telyukova, I. A. and R. Wright, 2008. "A Model of Money and Credit, with Application to the Credit Card Debt Puzzle." *Review of Economic Studies*, 75(2), 629-647.

[52] U.S. Census Bureau, 2003. *Statistical Abstract of the United States.*

[53] Yoo, P.S., 1998. “Still Charging: The Growth of Credit Card Debt Between 1992 and 1995.” *Federal Reserve Bank of St. Louis Review,* , January/February.

[54] Zinman, J., 2006. “Household Borrowing High and Lending Low under No-Arbitrage”. Dartmouth University.

Appendix A Data

A.1 Sample Selection

I use the 2001 wave of the SCF, and the Q2 2000 - Q1 2001 of the CEX, to capture all households who were interviewed in 2001. In both surveys, I restrict the sample to people of ages between 25 and 64. I drop low-income outliers below a threshold of $200 per month, and also those who are incomplete income reporters in either survey. Further, I drop those who fail to report valid asset and credit card debt information (if a CEX household has no such information in its fifth interview, then I drop it for all the quarters in which it is present). This leaves me with 2,878 households in the SCF, and 2,743 households in the CEX, with 2,164 of them present for the entire 12 months of the survey.

A.2 Household Assets and Subdivision of Population into Subgroups

I select the subgroups with the intention of matching their characteristics as closely as possible in the two data sets. In the SCF, liquid asset holdings are measured in detail, as are credit card debt data. The SCF asks the following questions about credit card balances that I use here:

- "After the last payment [on your credit card accounts], roughly what was the balance still owed on these accounts?"

- "How often do you pay off your credit card balance in full?" Answer choices are: Always or almost always, Sometimes, Almost never.

From the first question, I can clearly distinguish revolving balance from the new purchases that appear before the bill is paid. I use the second question to select only habitual credit card debtors to be in the puzzle group, that is, those who answer "Sometimes" or "Almost never"; of all households who report to have positive credit card debt at the time of the interview, 77% are in this group.

Liquid assets are defined as all household checking and savings account balances, and I also include brokerage accounts, because in the CEX there is no way to separate them out. Credit cards that I consider are bank-type and store credit cards, that is, those that allow to revolve debt.

In the CEX, credit card balance information is collected in the second and fifth interviews, and in the fifth interview, households are also asked the amount they paid in the last year in finance charges on credit cards (distinct from late fees). The relevant questions in the CEX are:

- "On the first of this month, what was the balance on your credit card account(s)?"
- "What was the amount paid in finance charges on all credit card accounts over the last 12 months?"

As is clear from the first question, it is harder to distinguish revolving debt from new purchases in the CEX, but I can do so using the finance charge question. In the CEX, credit cards are defined similarly to the SCF, as store and bank-type cards that allow debt to be revolved. Selecting a threshold of \$500 for revolving debt, and assuming it is revolved for a year, I take all households who paid an average of 14% APR on this balance as credit card revolvers. (The 14% interest rate is the SCF-reported interest rate paid on average on credit cards, shown in the text). Again, liquid assets are savings, checking and brokerage accounts.

In both surveys, those who report credit card debt above \$500 and liquid assets below \$500 (and those who are habitual debtors in the SCF, or paid positive finance charges in the CEX) are then put in the subgroup "debtors". The remaining subgroup - those with little non-habitual debt or no credit card debt - are "savers".

A.3 Separating Consumption Goods into Groups by Payment Method; ABA Survey of Consumer Payment Preferences

In looking at household consumption in the CEX, it was necessary to separate consumption into goods that people have to pay for with liquid instruments (cash, check, debit card) and goods that can be paid by either credit or liquidity. I separate household expenditures in the CEX into "cash-only goods", "cash-or-credit goods", education and durables. I separate out education and durables because expenditures for these goods occur rarely, while consumption is continuous but not measured through expenditure alone (see Krueger and Perri, 2003). Thus, studying volatility of expenditure on these goods is uninformative. This is true of cash-or-credit goods to some extent also, since they include many semi-durable items, such as clothing; it is important that the point of this exercise is not to compare volatilities across good groups.

Table A.3.1: ABA Survey: Most Used Payment Method by Bill Type

Bill type	Check, cash, direct debit	Debit Card	Credit Card
Rent or mortgage	99.4	0.3	0.4
Loan or lease	98.2	1.0	0.8
Insurance	96.2	1.2	2.6
Child care, tuition	91.8	2.2	6.0
Utilities	95.0	2.5	2.5
Charity contributions	96.0	1.3	2.7
Memberships, subscriptions	85.2	3.1	11.7

Table A.3.2: ABA Survey: Most Used Payment Method by Store

Store	Cash or check	Debit Card	Credit Card
Grocery store	45.4	35.7	18.9
Gas station/convenience store	34.1	26.8	39.1
Department store	27.6	26.4	46.0
Discount store/warehouse club	43.4	27.2	29.4
Drug store	47.3	29.7	23.0
Restaurants	42.3	23.4	34.3
Fast food	85.6	7.8	6.6
Transit system	81.4	8.6	10.0

To accomplish the separation, I relied on the 2004 Survey of Consumer Payment Preferences conducted by the American Bankers Association and Dove Consulting. This survey is not representative of all U.S. households, but is the only up-to-date survey that studies consumer payment methods. The sample that it does study consists of people with access to internet, so arguably, these are households who have the broadest payment options, and thus it should give a fairly accurate idea of payment methods used for most common good groups. In the survey, consumers are asked how they pay for certain types of goods and services, as well as at certain types of stores. Tables A.3.1 and A.3.2 present a summary of all results from the survey that pertain to consumer choice of payment methods. The questions were all phrased in the same way: “When you make purchases at [type of store], which method of payment do you use most often?”, and “When you pay for [type of bill], which payment method do you use most often?”

The resulting categories are presented in table A.3.3.

Table A.3.3: Goods Categories for CEX Analysis

Good group	Components
Cash-only goods (paid by check, debit, cash)	Rent, mortgage, utilities, property taxes, insurance, household operations/repairs/maintenance, child care, public transportation, health insurance, cash contributions; food in and out, alcohol, tobacco.
Cash-or-credit goods	Apparel, entertainment, gasoline, medical services, medical equipment, prescription drugs, reading, personal care, membership fees, funeral expenses, legal fees, auto insurance, auto repairs, other vehicle expenses.
Durables	Households furnishings and major appliances, vehicle purchases
Education	Tuition and fee expenses, textbook purchases

Appendix B An Example Model of Measurement Error

Suppose that expenses on a generic good in the CEX have the following expenditure structure: for good i, household j at time t, true consumption satisfies

$$c_{ijt} = x_{ijt} + \sum_{k=1}^{K} z_{ijt,k}, \tag{15}$$

where x_{ijt} represents a regular expense (like a weekly grocery trip, monthly utility payment, monthly transit pass, a shoe purchase), and $z_{ijt,k}$ are incidental expenses, possibly irregular, or very frequent, of any magnitude (i.e. a sandwich purchase, a taxi ride, purchase of new running shoes). Let $x_{ijt} \sim N(\mu_x^{ij}, \sigma_x^{ij})$,[16] and $z \sim U[0, \lambda^i \sigma_x^{ij}]$. For certain goods, like utilities, we'd expect $K = 0$, for most others, $K > 0$. For this formulation we have the mean and variance of true consumption at the household level, assuming independence of individual expenses and a long enough time horizon, as

$$\mu_C^{ij} = \mu_x^{ij} + K\frac{\lambda^i \sigma_x^{ij}}{2} \tag{16}$$

$$(\sigma_C^{ij})^2 = (\sigma_x^{ij})^2 + K\frac{(\lambda^{ij}\sigma_x^{ij})^2}{12}. \tag{17}$$

[16]The assumption of normality on the level of consumption is unusual, and would be more standard on the log of consumption. Here, I make this assumption purely for algebraic convenience, supposing the parameters of the distribution such that the probability of a negative event is negligible.

Now suppose that there are two types of measurement error. First, households may forget the exact size of their regular purchase X. This error $\varepsilon_{ijt} \sim iid(0, \sigma_\varepsilon^i)$ is the classical measurement error. Second, the memory error is a function of the size of the expense, and is given by $R(z_{ijt,k})$. This function captures that households are more likely to forget small expenses than large ones, and is given by

$$R(z_{ijt,k}) = \begin{cases} \gamma_h > \frac{1}{2} & \text{if } z_{ijt,k} \geq \frac{\lambda \sigma_x^{ijt}}{2} \\ \gamma_l < \frac{1}{2} & \text{if } z_{ijt,k} < \frac{\lambda \sigma_x^{ijt}}{2} \\ 0 & \text{otherwise} \end{cases} \tag{18}$$

Assume for simplicity that $\gamma_h = 1$ and $\gamma_l = 0$, so households remember all incidental expenses higher than the mean of the uniform distribution, and forget all such expenses smaller than the mean. With this error structure, suppose that measured consumption looks as follows:

$$\tilde{C}_{ijt} = (X_{ijt} + \varepsilon_{ijt}) + \sum_{k=1}^{K} R(z_{ijt,k}) z_{ijt,k} = (X_{ijt} + \varepsilon_{ijt}) + \sum_{k=1}^{K/2} \tilde{z}_{ijt,k}, \tag{19}$$

where $\tilde{z}_{ijt,k} \sim U[\frac{\lambda \sigma_x^{ijt}}{2}, \lambda \sigma_x^{ijt}]$.[17] Under the assumption of a long enough time horizon for a given household, the mean of measured consumption is given by

$$\mu_{\tilde{C}}^{ij} = \mu_x^{ij} + \frac{K}{2} \frac{3\lambda^i \sigma_x^{ij}}{4} < \mu_C^{ij}. \tag{20}$$

Notice that for a good that has no incidental expenses, like a monthly utility payment, $K = 0$, so the mean of measured consumption will equal the mean of true consumption, but for any good that has incidental expenses, mean measured consumption will be lower than mean true consumption. Thus, if only classical error were present, we would expect CEX/PCE ratios of 1 for all goods; if memory error is also present, at least for goods with incidental purchases, we would expect the CEX/PCE ratios below 1, which is consistent with what we see in the data.

What about the variances? For measured consumption, we get

$$(\sigma_{\tilde{C}}^{ij})^2 = (\sigma_x^{ij})^2 + (\sigma_\varepsilon^i)^2 + \frac{K}{2} \frac{(\lambda^{ij} \sigma_x^{ij})^2}{48}. \tag{21}$$

For a good that has no incidental expenses, like utilities, classical error would exaggerate household-level variance since $K = 0$. The size of the distortion would depend on the variance of the error. However, if most goods in the CEX feature incidental payments, then it is

[17] Obviously, the model could also incorporate classical error on the reporting of incidental expenses z_{ijt}. Here I omit this for simplicity, but the spirit of the argument would still go through.

quite clear that the memory error would reduce the variance of these incidental payments, and possibly undo the exaggeration of variance of measured consumption created by the classical error. It is not hard to derive the values of the parameters of this model under which measured volatility of consumption would be *lower* than true volatility. In this example,

$$(\sigma_{\tilde{C}}^{ij})^2 < (\sigma_C^{ij})^2 \text{ whenever } (\sigma_\varepsilon^i)^2 < \frac{7}{96} K (\lambda^{ij} \sigma_x^{ij})^2. \tag{22}$$

To tie this model back to the discussion of expense emergencies, those kinds of events are not well-represented by a uniform distribution like that on z_{ijt} here. One way to model this would be with a third component in (15) described by a Poisson distribution, for example. If we assumed that those events are remembered accurately, it would not change any of the discussion above: memory error would affect more minor expenses, classical error could affect those as well as the regular expenses X, but the combination of the memory and classical error would still understate the measured mean consumption, and may understate the measured variance, relative to the true consumption moments, depending on the parameters of the two errors. In any case, if we apply this model, with or without Poisson events, to the cash-good category as a whole, we might expect that what is left in the uniform distribution of z_{ijt} are relatively large or unusual expenses, while all the small food expenses are forgotten; in that case, both K and λ would be large, and the above condition (22) would be more likely to be satisfied.