Household Need for Liquidity and the Credit Card Debt Puzzle

Telyukova, Irina A.

University of California San Diego

26 December 2007

Household Need for Liquidity and the Credit Card Debt Puzzle

Irina A. Telyukova *
University of California, San Diego

First version: April 1, 2005

This version: December 26, 2007

Abstract

In the 2001 U.S. Survey of Consumer Finances (SCF), 27% of households report simultaneously revolving significant credit card debt and holding sizeable amounts of liquid assets. These consumers report paying, on average, a 14% interest rate on their debt, while earning only 1 or 2% on their liquid deposit accounts. This phenomenon is known in the literature as the "credit card debt puzzle." In this paper, I pose and quantitatively evaluate the following explanation 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 for goods for which credit cards cannot be used. Using both aggregate and survey data (SCF and CEX), I document that liquid assets are a substantial part of households' portfolios and that consumption in goods requiring liquid payments may have a sizeable unpredictable component. This would warrant holding precautionary balances in liquid accounts. I 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 calibration of the model parameters is based on the simulated method of moments. The calibrated model accounts for 73% of the households in the data who hold consumer debt and liquidity simultaneously, and for at least 55 cents of every dollar held by a median household in the puzzle group. I argue that these results are a lower bound, and that the liquidity-need hypothesis is thus successful in rendering most of the puzzle a rational phenomenon.

*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 invaluable input, I thank S. Borağan Aruoba, Andreas Lehnert, Ben Lester, Michael Palumbo, Shalini Roy, Gustavo Ventura, Ludo Visschers, and Neil Wallace, as well as the participants of seminars and conferences at the University of Pennsylvania, UCSD, Notre Dame, Washington University St. Louis, Queen's University, University of Western Ontario, University of British Columbia, the Board of Governors of the Federal Reserve, Federal Reserve Banks of New York, Cleveland, Philadelphia and Atlanta, USC, University of Maryland, European Central Bank, the Canadian Economic Association, the Midwest Macroeconomic Meetings, and the Society for Economic Dynamics. 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 since then. Lehnert & Maki (2001) study whether households may do this strategically, in preparation for a bankruptcy filing. Since in the U.S., each state offers some exemption level of assets in the event of a household bankruptcy filing, 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 in question, upon examination of the total portfolios of the puzzle households, it appears that most of them would be unlikely to file for bankruptcy. I will present the relevant evidence below.

Alternatively, Bertaut and Haliassos (2002), and Haliassos and Reiter (2003) have studied whether households may opt to hold both 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 small 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, which makes the behavior more puzzling still from this perspective.

Although the existing explanations for the credit card debt puzzle may have merit for some households, there are seemingly many households whose behavior they are not likely to capture, for reasons mentioned above. In this paper, I study in detail, for the first time, a hypothesis of why a household may choose rationally to hold liquid assets and revolve credit card debt simultaneously, and evaluate how much of the puzzle this hypothesis can account for. 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 auto dealerships accept credit cards, but smaller mechanics more trusted by households 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 is not 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 the expenditures requiring payment by a check, say, 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) briefly mention this idea as a possibility but do not analyze it, and indeed dismiss it as likely unimportant. A careful quantitative analysis of the hypothesis, however, is an involved exercise, from both theoretical and empirical perspectives, and it is

[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.

crucial, because it allows us to evaluate the possibility that the "puzzle" is in fact a rational phenomenon, at least for many households.

The goal of this paper is to measure how much of the puzzle the hypothesis presented here can account for. Specifically, I answer the following two questions: (1) Can the need for liquidity explain why *so many* households 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 characteristics of households who choose to borrow on credit cards and save in liquid accounts simultaneously. I will show that there is nothing inherent about them, from a demographic perspective, that would distinguish them from other households, so that the phenomenon may have economic causes. I will also show evidence that gives support to the importance of liquid assets in monthly household expenditures, and to the fact that uncertainty in these expenditures appears to play a significant role.

Next, I develop a dynamic stochastic partial-equilibrium model of household portfolio choice, in order to develop the intuition, and to study the hypothesis rigorously, both analytically and quantitatively. The basis is a standard incomplete-market heterogeneous-agent model with two types of idiosyncratic risk. The model's novel features are a two-market structure, where in one of the markets credit cannot be used, and the timing of the two risk realizations during the period such that portfolio decisions have to be made before spending decisions. There is also a restriction that while some of the spending decisions are made, no access to additional income or portfolio rebalancing is allowed. In its treatment of money, the model is consistent both with Lucas-Stokey-style cash-credit good models and with a more recent generation of monetary models that treat the reasons for why money is essential in trade explicitly.[3] As I will show, the model has all the analytical implications important for addressing the credit card debt puzzle.

[3] In a related theoretical exercise, Telyukova and Wright (2008) approach this puzzle as the rate-of-return dominance puzzle and develop a micro-founded monetary model to analyze it. In that paper, the model we develop treats explicitly the frictions that are needed to make both money and credit essential in an economy. In the current paper, in contrast, my focus is on quantitative analysis using a heterogeneous-agent model of realistic complexity - so I abstract from the reasons for why credit may be accepted in some markets, but not others, and simply take this fact as given.

I calibrate the model by matching it to properties of liquid-asset consumption and main distributional characteristics in the data. It is crucial that I leave all properties of household portfolio choice, as well as the numbers of people who choose different portfolios, untargeted in the calibration, in order to be able to judge in a disciplined way how well the hypothesis presented here does at explaining the puzzle. The calibration method is based on the simulated method of moments, where I minimize the weighted squared distance between relevant moments in the data and their simulated counterparts in the model. The calibrated model accounts for 73% of the households who choose to revolve debt while holding money in the bank, and for a median such household, for at least 55 cents of every dollar it holds in liquid accounts.

The main contributions of this paper are three. First, I propose and carefully evaluate a new answer to a still-unresolved puzzle, and find that it can account for the puzzle to a very significant degree. Debt puzzles of this nature frequently lead the observer to wonder whether households are capable of making rational decisions, while rationality is the most fundamental assumption of the majority of economic theory. The implication of this work is that we need not question rationality of at least most households in this context. Second, 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. Third, in the process, I obtain estimates of some parameters of interest. Based on survey data, I measure one type of unobservable idiosyncratic uncertainty that faces households in the context of a liquidity-based model. I also obtain estimates of some preference parameters of interest - especially the elasticity of substitution between cash and credit goods - that have only been estimated in deterministic models until now. My estimates suggest that idiosyncratic uncertainty affects these parameter estimates considerably.

In complimentary work, Zinman(2006) uses survey data to demonstrate, via some simple calculations, that “borrowing high and lending low” is largely not puzzling and can be seen as rational. 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 analyzes its properties. Section 4 discusses calibration and computation. Section 5 presents the results from the calibrated model, and section 6 discusses them. Section 9 concludes. Some details of the data and of the computational analysis are relegated to the appendices.

2 Data

I use two U.S. household surveys in order to describe the puzzle in the data. One data source is the Survey of Consumer Finances (SCF), a triennial cross-sectional survey that has detailed information on household assets and liabilities. In particular, it measures carefully both household liquid asset holdings and revolving credit card debt, and despite its cross-sectional nature, allows to assert persistence of this debt, by asking households about frequency of complete debt payoff (see appendix). I use the 2001 wave of the SCF in this analysis. I separate the SCF sample into three subgroups: those who have sizeable revolving credit card debt and no significant liquid assets ("borrowers"), those who have both in significant amounts ("borrowers and savers", i.e. the puzzle group), and those who have liquid assets but no revolving credit card debt ("savers"). Notice that the borrowing behavior here is defined solely by credit cards, and saving solely by liquid assets - which include checking, savings and brokerage accounts. I abstract, in choosing the terminology and focus, from the fact that these households may be, and usually are, borrowing or saving in other assets.

In addition, I use the 2000-2002 Consumer Expenditure Survey (CEX) to study consumption patterns of the households who revolved credit card debt in 2001, to match the SCF timeline. This survey is a quarterly rotating panel, where each household is interviewed for five consecutive quarters, four of which (second through fifth) are available in the public data set. 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.[4] Although it is less careful about measuring assets and credit card debt, there is enough information to subdivide this population into the same subgroups as in the SCF. I study the properties of household consumption in goods paid by liquid assets versus other methods.

In both surveys, I consider those who hold more than $500 in revolving credit card debt and more than $500 in liquid assets as the borrower-and-saver group.[5] I study all households with heads of age 25 to 64; thus, I exclude college students and retirees, whose saving and borrowing behavior may differ from the rest of the population (for example, borrowing behavior among college students may be hard to analyze, since the repayment of their debts is often undertaken by their parents, as is well documented). The details of the surveys, the sample selection process, and the puzzle measurement methods are described in the data appendix.

Tables 1 through 4 describe the credit card debt puzzle, and compare the households in the puzzle group to the rest of the population. I show that these households appear to have the same demographic characteristics as everyone else, and they lie in the middle of the economic distribution. I also present evidence that the need for liquidity may be a good candidate for explaining the puzzle, because the liquid assets that these households have do not seem unreasonably large in amount relative to their income, spending and credit card debt. Tables 5, 6 and 7 then characterize in more detail household liquid asset holdings and their use, in order to show, in support of the central hypothesis here, that liquid assets do appear to have a significant and unique role in household finances that cannot be replaced by other instruments.

Table 1 gives the size of the credit card debt puzzle in the data. I present the measurements from both data sets, to demonstrate that they are close. Judging by descriptive statistics of both groups (not presented here for the CEX), it is clear that these groups are comparable in the two surveys, so that analyzing their consumption in the CEX and assets in the SCF is a valid exercise. To my knowledge, this kind of joint use of the two data sets is the first of its kind. As is clear, around 27% to 29% of the U.S. population were simultaneously borrowing and saving in 2001. Only between 5 and 7% of the population are credit card borrowers with little

[4]To be precise, 65% of the expenditure data are collected via direct questions about the month and amount of expenditure, while 35% of the expenditures are measured via questions on quarterly spending, which is then divided into three average-monthly amounts. The latter procedure applies to food, for example. This will affect some of the results discussed later, but favorably for my purposes (see below).

[5]I choose the $500 threshold mainly to follow other literature on this subject. Having studied alternatives, I came to the conclusion that the puzzle measured in different ways is still a significant one in the U.S., while the subgroups' characteristics remain stable regardless of the specification.

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: "Borrow" refers to revolving debt on credit cards; "save" to saving in liquid assets. 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 from the SCF question "What is the interest rate you pay on the credit card with the highest balance?" For the puzzle group ("borrow & save") measurement, I take everyone with liquid asset holdings above $500, and credit card debt above $500.

or no 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. It is clear that there is a significant difference in the rates, which gives the appearance of a violation of the standard no-arbitrage condition, and which originally gave rise to the term "credit card debt puzzle".

Table 2 breaks down some of the demographic characteristics of the subgroups from the SCF; the numbers are nearly replicated in the CEX, and not presented here. 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 puzzle group are white. Comparing the numbers for different characteristics to the overall sample average shown in the right column, we see that none of them seem particularly pronounced for the borrower-and-saver group. The borrower-and-saver group is skewed 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%) - perhaps contrary to

Table 2: Demographics

	Borrow	Borrow & Save	Save	Share in Population
	% 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.

what we might expect. 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 nationally. The saver group compares similarly to national averages, while the borrower group is the one that is least educated, comprises most unmarried households, and is skewed most toward nonwhite households. The main idea here is to show that there is nothing inherent in demographic or "socioeconomic" terms about the borrowing and saving group that might lead them to behave differently from others.[6]

Table 3 presents income and asset information for each subgroup. The puzzle group clearly lies in the middle of the economic 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 only 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).[7] Several further insights are important. First, while liquidity holdings of the borrower-and-saver group are not negligible, at \$3,000 in the median, they are

[6]This is confirmed in formal probit analysis, not presented here.

[7]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 thus represent a monthly average on the account.

Table 3: Income and Asset Holding

		Borrow	Borrow & Save	Save
		U.S. Dollars		
Credit card debt:	*Mean*	5,172	5,766	317
	Median	3,340	3,800	0
Liquid assets:	*Mean*	227	7,237	17,386
	Median	200	3,000	3,200
Total after-tax income:	*Mean*	28,032	52,114	64,331
	Median	25,350	43,600	39,950
Other financial assets:	*Mean*	4,424	40,545	102,558
	Median	0	4,400	4,100
Net wealth:	*Mean*	36,231	187,508	466,462
	Median	9,450	84,640	104,500
Liquid assets as share of monthly after-tax income	*Mean*	*0.12*	*1.71*	*2.53*
	Median	*0.10*	*0.79*	*0.88*

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

not unreasonable either, relative to their income. Secondly, these households have significant amounts of nonliquid financial assets as well, so there is no evidence that they are unaware of more lucrative saving opportunities. These facts suggest that the liquidity holdings of these households may, in fact, be geared toward some well thought-out purpose in any given month. Compare this to the savers, who evidently have enough liquidity both to cover their credit card expenses, so they need not revolve debt (the majority of them do have and use credit cards), and to cover any monthly liquid expenditure needs as well. Insofar as we may think of the savers as the least constrained group - i.e. most able to achieve their first-best allocation - these data suggest that the borrower-and-saver group might like to hold even more liquidity than they are able to.

In addition, the presence of significant nonliquid financial assets in all but the borrowers' portfolios, as well as a look at the net worth of these households, suggest that strategic bankruptcy behavior, as per Lehnert and Maki (2001), is highly unlikely for at least the majority of the puzzle households. Finally, note that on average, the amount of debt these households have is approximately equal (higher in the median, at $3,800, but lower in the mean) to their liquid

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 (such as townhouse/condo association) are excluded.

holdings; if they were to use their liquidity to pay off debt, they would be left with little or no money in the bank in most cases.

Table 4 presents a further aspect of household asset holdings: homeowners (especially those who pay mortgage) are more likely to be in the puzzle subgroup. They are overrepresented in this group compared to the overall average: homeowners with mortgage constitute 59% of this group, relative to only 50% of the population.

The evidence presented so far would suggest that there is no apparent reason to assume anything different about the preferences of these households, and it seems likely that the motivation for this observed behavior is economic in nature. Moreover, households appear to diversify their portfolios, as they tend to have investments in real estate and significant holdings of nonliquid financial assets. In other words, it appears that the liquid holdings that households have may be designated for a specific purpose which may have priority over credit card repayment up to a certain level of liquid assets. Those households that are not overly cash-rich (see table 3) may have liquid assets under that level, so it may be optimal for them to delay debt repayment in favor of keeping the liquid assets available in the bank. In addition, as discussed, homeowners are more likely to be in the puzzle group than non-homeowners. This makes sense once we consider that the expenditures for which credit is not accepted in payment have most to do with home ownership - examples are mortgage payments and especially household operations and repairs, for which the owner of the house, rather than a renter, would be responsible, and which also are often unexpected and large in magnitude. The next three tables demonstrate in more detail that liquid assets appear to have an important autonomous role in household finances that

Table 5: Aggregate Consumer Transactions, Shares by Method of Payment

	Transaction number			Transaction volume			
	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

cannot be replaced by other assets, which would support the hypothesis under investigation.

In aggregate, it is clear that liquid assets have retained an obviously dominant role in consumer transactions, even though credit card usage has been growing somewhat. Table 5 gives aggregate consumer transactions by payment method for selected years from 1990 to 2002. In 2002, liquid payment methods, such as cash, checks, and debit cards, accounted for 77% of total consumer transactions, or 65% of their total value. If we include electronic payments in this category, since they are most 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.

I turn to the CEX to study household liquid asset holdings relative to their consumption patterns in goods that require the use of liquid assets. Many subtle issues are involved in separating out the group of goods that may be considered cash-only goods. Although survey data on consumer payment method choice are scant to nonexistent, one such survey was conducted in 2004 by the American Bankers Association. In it, consumers were asked questions about their perceptions and usage of payment methods; in particular, they were asked how they normally pay at different types of stores and for different types of bills. I present the details of this survey in appendix A.3. Tables A.3.1 and A.3.2 summarize the relevant information. It is clear from the survey that liquid payment methods dominate household expenditures. Consumers overwhelmingly pay all house-related types of bills that are asked about in the survey, such as rents, mortgages, insurance, and utilities, by check or related liquid instruments (e.g. direct debit from the account). They also tend to pay for child care and tuition with liquid instruments, but I do not include intermittent expenses such as tuition in the cash-only group, as they are likely to skew the perception of volatility (see appendix). Payments for 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 asset investment for retirement.[8] While we see evidence that they save for retirement in retirement accounts, emergencies, including home-related ones, by their definition are likely to require liquid savings. In terms of payment methods in stores, the evidence suggests that while credit cards are predominant 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, repairs, household operations, property taxes, insurance, public transportation, health insurance, and also food, alcohol and tobacco. For most of these goods, it is largely a requirement that a liquid payment method be used, but this is not true for the last three good groups. For these goods, we see that consumers do pay for them predominantly in liquid instruments, but they frequently are likely to have the option to use a credit card as well (as evidenced by the survey, where some 20-30% of such expenditures are made by credit). The justification for including these three good groups as cash-only goods are also in the appendix. As a sensitivity check of the results, I will experiment with exclusion of these goods, but I take inclusion of food as the main case to study. In any event, food, alcohol and tobacco are a minority of this expenditure category.

Table 6 presents household liquid asset holdings relative to average monthly consumption of cash-only goods. In the borrower-and-saver group, the median household has 1.5 times its average monthly liquid consumption in the bank accounts, while the mean household has 3.4 times the amount. Again, these are numbers that are significant but seemingly not unreasonable. Compare these with the holdings of the saver group, who have on average 10 times their mean monthly liquid spending, or twice the monthly spending amount in the median. Again we see that the savers are better equipped to handle *both* their liquid spending needs and credit card bills, rather than having to prioritize one over the other due to scarce liquid resources.

The evidence in table 6 points to precautionary demand for money: households have liquid asset amounts that are in excess of what they spend on average per month, and those who

[8] 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 6: Household Liquidity Holding and Consumption Patterns

		Borrow	Borrow & Save	Save
		U.S. Dollars		
Liquid assets:	*Mean*	227	7,338	17,435
	Median	200	3,000	3,200
Monthly cash-only good cons:	*Mean*	1,561	2,106	1,665
	Median	1,369	1,890	1,433
Liquid assets/cons:	*Mean*	*0.1*	*3.4*	*10.0*
	Median	*0.1*	*1.5*	*2.0*

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

Table 7: Unpredictable Volatility of Average Monthly Household Cash-Good Consumption

	Borrow	Borrow & Save	Save
	Avg. conditional standard deviation		
Cash-only goods, including food	20.1%	21.0%	21.6%
Cash-only goods, excluding food	23.3%	23.9%	25.3%

Source: CEX. Conditional standard deviation: population average of individual conditional standard deviation of month-to-month liquid consumption, taken across a 12-month period in which the household appears in the survey. Measured by regressing log liquid household consumption on a set of month and year dummies and household fixed effect. The residual is taken as the idiosyncratic unpredictable component, with its conditional standard deviation used here. Cash goods: see appendix A.3.

are sufficiently well-off are holding much more liquidity than those in the middle, suggesting that richer households choose to buffer themselves more fully, and that some households become constrained from doing so completely, which may lead to borrowing-and-saving behavior on their part.

I now turn to characterizing one of the likely causes of this precautionary behavior. Table 7 shows volatility of consumption in the cash-only good category, measured as average monthly conditional standard deviation of household liquid expenditures. There are several issues that arise in constructing this measure of volatility. First, measuring raw volatility of consumption may not be fully informative about unpredictable volatility, since it may also reflect seasonal volatility, for example, as well as other factors that may be predictable to the household. Second, users of the CEX data frequently use quarterly averages of consumption rather than the monthly measure because some questions are asked only as averages over three months, as mentioned before. To answer in part the first concern, I exclude from the expenditures all purchases made as gifts; this information is explicitly collected in the CEX for each purchase reported. This should help remove some of the seasonality in the consumption series, since much of seasonal purchasing is done in holiday gifts. In addition, following literature on idiosyncratic income and consumption uncertainty (see, e.g., Storesletten, Telmer and Yaron 2004b), I filter out the predictable component of expenditures, by regressing the log of cash-only consumption for each household on a household fixed effect, to control for household characteristics known to the household, as well as a full set of month and year dummies, to control for any seasonal effects. I treat the volatility of the residual as a measure of unpredictable consumption volatility; the average conditional standard deviation of this process across households is taken as the final measure of this volatility. For the group of goods selected in the cash-only good category, the household and seasonal factors reduce, but not substantially, the volatility of the consumption series. To answer the second concern, I also measured volatility based on quarterly aggregates of the monthly expenditure responses. Although this leaves only three observations per household, the measures of volatility remain robust to this specification - they go down, but only slightly. Thus, insofar as it is possible with such short panels, I can be fairly confident that I have an accurate measure of consumption volatility.

In addition, as discussed in the data appendix, including food in the cash-good category downplays the estimate of volatility, because it is measured as an average across three months in each household interview; thus, I show volatility measures both including and excluding food. Volatility appears quite significant at 20-22% of the average when food is included, and 23-25% when it is not. Volatility is slightly higher for savers, and lowest for borrowers, which may reflect 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. The volatility we observe in cash-only good consumption may be a reflection of such unexpected, and possibly large, spending shocks; households try to insure against them by holding extra liquidity in the bank.

To sum up, data suggest that the credit card debt puzzle is significant in magnitude, but it may not be as puzzling as it appears initially. There are situations where liquidity is a non-substitutable resource, and the resulting demand for liquidity may be significant enough to account for households who choose to hold on to their liquid assets instead of paying down credit card debt. The rest of the paper is devoted to evaluating formally whether this hypothesis can account for the data. First, I lay down a model that can address this question in a disciplined way. Then, I calibrate this model and use it to measure the ability of the need for liquidity to account for the credit card debt puzzle.

3 Model

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 instruments available to agents in each period. One is money, denoted m_{jt} - a storable, perfectly divisible, intrinsically worthless object, potentially useful only as a medium of exchange. This instrument 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.[9] In both subperiods, there are competitive firms producing the consumption good in the background. In the first subperiod, they take labor supplied by households as input, while in the second, households do not provide any inputs into production, and simply buy consumption goods from the firms at prices they take parametrically. Although markets are competitive, they are incomplete: insurance markets are closed during both subperiods.

[9] The question of why credit cannot be used is beyond the scope of this paper, since it is not pertinent to the empirical problem at hand. There are several approaches to it in the literature: one is to assume spatial separation between the earner and the shopper, as in 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 related model of money and credit that addresses the issue in more detail in a similar context.

During each period, households are subject to idiosyncratic income and preference uncertainty. There is no aggregate uncertainty. The shocks on income and preferences do not realize simultaneously: income shocks realize at the beginning of the first subperiod, while preference shocks realize at the beginning of the second. Since there are no insurance markets for these shocks, the only way to insure is by accumulating one or both of the assets m and b.

At the beginning of the first subperiod, the household's income shock s_t realizes. Agents then supply labor inelastically (that is, there is no labor choice) and earn their income, consume with either credit or money, and allocate their resources between the two instruments in a household portfolio. Let us assume that $s_t \in S$ is a discrete Markov process, with $S = \{\underline{s}, s_2, ..., \bar{s}\}$, $\underline{s} > 0$. 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.

At the start of the second subperiod, the consumer's preference shock z_t realizes, also assumed to be a discrete Markov process with $z \in Z = \{\underline{z}, z_2, ..., \bar{z}\}$, and transition matrix $\Pi(z_t, z_{t+1})$. Note that the shocks on income and preferences, and their transitions, are assumed to be independent of each other. After the realization of z, the subperiod's market opens. Here, households choose consumption conditional on their preference shock realization, but it is crucial to note that they cannot produce or borrow in this market, so they do not have access to additional income when they need to consume. This assumption is meant to capture the fact that in any given month, a household is likely to encounter liquid-asset spending opportunities continually and randomly, without simultaneous opportunities to rebalance their portfolios or get additional income.

In each subperiod, the household's state variables are its current knowledge of the idiosyncratic shock processes s and z, and its current portfolio (m, b). Since the income shock s_t realizes at the beginning of the first subperiod, while the preference shock z_t does not realize until the second, in the first subperiod the state is $(s_t, z_{t-1}, m_{1t}, b_{1t})$. Correspondingly, the state in the second subperiod is $(s_t, z_t, m_{2t}, b_{2t})$. Agents take prices as given, so prices, or alternatively the distribution of agents, are aggregate state variables, which I make implicit in the notation.

Lifetime utility is given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [u_1(c_{1t}) + z_t u_2(c_{2t})],$$

where it is assumed that $\forall j = \{1, 2\}$, where j denotes the subperiod, $u_j \in C^3$, $u_j'(\cdot) > 0$, $u_j''(\cdot) < 0$, $u_j'''(\cdot) > 0$ and the functions satisfy Inada conditions, $\lim_{c_{jt}\to 0} u_j'(c_{jt}) = \infty$ and $\lim_{c_{jt}\to\infty} u_j'(c_{jt}) = 0$. I assume that the preference shock is multiplicative on the utility of consumption in the second subperiod. Note that in this formulation of the problem, the utility function is assumed to be separable in first- and second-subperiod consumption. This is not necessary for any of the results that I want to emphasize, but does make analysis more transparent. For computation, I will make the utility function nonseparable, as it is more realistic from the data point of view, and adds interesting empirical insights.[10]

I formulate the household problem recursively.[11] The nature of the question makes it sufficient to study the partial equilibrium of this problem: that is, I will set prices exogenously and study the resulting decision rules. In the first subperiod, a household solves the following problem:

$$
\begin{aligned}
V_1(s_t, z_{t-1}, m_{1t}, b_{1t}) &= \max_{c_{1t}, m_{2t}, b_{2t}} u_1(c_{1t}) + \mathbb{E}_{z_t|z_{t-1}} V_2(s_t, z_t, m_{2t}, b_{2t}) \\
s.t.\ c_{1t} + \phi_{1t} m_{2t} &= s_t + \phi_{1t} m_{1t} + b_{2t} - b_{1t}(1 + r_t) \\
b_{2t} &\leq \bar{B} \\
c_{1t} &\geq 0, m_{2t} \geq 0
\end{aligned}
\tag{1}
$$

Here, ϕ_{1t} is the real value of money, that is, the inverse of the price on the consumption good. r_t is the interest rate that is charged on debt at the beginning of subperiod 1. I assume, as is necessary for existence of a stationary equilibrium, that $\beta < 1/(1+r_t)\,\forall\, t$ (Aiyagari, 1994). The expectation term is written conditional on only the previous realization of the shock, reflecting the assumption above that the shock has a Markov form. The second constraint imposes a credit limit on the household, here taken to be exogenous. Notice that there is no nonnegativity constraint on debt: agents can save in b_{2t}.[12]

[10] To be precise, if in this study I were concerned with the interaction, for example, of monetary policy with the credit market, the separability assumption would be restrictive in the analytical context, since it severs an interaction channel between the two markets. In the current context, however, the analytical results I emphasize do not hinge on the separability assumption. Empirically, the interaction of the two consumption goods may play a part in the magnitude of the results, and it seems natural to assume that it is non-trivial in reality; I will take up this issue in the computational part of the paper.

[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.

[12] In computation, I will allow for an interest spread: interest rate on borrowing, $b_{2t} > 0$, will be higher than

In the second subperiod, households solve the following problem, once the preference shock realizes:

$$\begin{aligned} V_2(s_t, z_t, m_{2t}, b_{2t}) &= \max_{c_{2t}} z_t u_2(c_{2t}) + \beta \mathbb{E}_{s_{t+1}|s_t} V_1(s_{t+1}, z_t, m_{1,t+1}, b_{1,t+1}) \\ s.t.\ c_{2t} &\leq \phi_{2t} m_{2t} \\ m_{1,t+1} &= m_{2t} - \frac{c_{2t}}{\phi_{2t}} \\ b_{1,t+1} &= b_{2t} \end{aligned} \tag{2}$$

ϕ_{2t} again denotes the subperiod's real value of money. Notice from the third equality that no interest on consumer debt is accumulated in the second subperiod - this captures the grace period typical of a credit card billing cycle.

Because in this problem the timing of the decisions between the two subperiods affects the state variables on which these decisions depend, it helps to keep track of the states explicitly while discussing the solution. Denote the state variables of the first subperiod as $x_{1t} = (s_t, z_{t-1}, m_{1t}, b_{1t})$. Then, the decision rules from the first-subperiod problem are $c_{1t}(x_{1t})$, $m_{2t}(x_{1t})$, and $b_{2t}(x_{1t})$. In addition, let $\lambda(x_{1t})$ be the Lagrange multiplier associated with the credit constraint. The first-order conditions that characterize the solution to this problem are, $\forall\, x_{1t}$:

$$-u_1'(c_{1t}(x_{1t}))\phi_{1t} + \mathbb{E}_{z_t|z_{t-1}} V_{2m}(s_t, z_t, m_{2t}(x_{1t}), b_{2t}(x_{1t})) = 0 \tag{3}$$

$$u_1'(c_{1t}(x_{1t})) + \mathbb{E}_{z_t|z_{t-1}} V_{2b}(s_t, z_t, m_{2t}(x_{1t}), b_{2t}(x_{1t})) - \lambda(x_{1t}) = 0 \tag{4}$$

The envelope conditions of the first subperiod are:

$$V_{1m}(x_{1t}) = \phi_{1t} u_1'(c_{1t}(x_{1t})) \tag{5}$$

$$V_{1b}(x_{1t}) = -(1 + r_t) u_1'(c_{1t}(x_{1t})) \tag{6}$$

Denote by $x_{2t} = (s_t, z_t, m_{2t}, b_{2t})$ the state of the agent in subperiod 2; note again that it is different from the state in subperiod 1. Then the decision rule of this subperiod is $c_{2t}(x_{2t})$, and I denote the Lagrange multiplier on the money constraint $\mu(x_{2t})$. The first-order condition of this problem is:

$$z_t u_2'(c_{2t}(x_{2t})) - \mu(x_{2t}) - \frac{\beta}{\phi_{2t}} \mathbb{E}_{s_{t+1}|s_t} V_{1m}(s_{t+1}, z_t, m_{2t} - \frac{c_{2t}(x_{2t})}{\phi_{2t}}, b_{2t}) = 0. \tag{7}$$

that on saving, $b_{2t} < 0$. This does not change the nature of the problem, but would require additional notation. In the analytical discussion, I abstract from this.

The envelope conditions are, after substituting in (5) and (6),

$$V_{2m}(x_{2t}) = \beta \mathbb{E}_{s_{t+1}|s_t} \phi_{1,t+1} u_1'(c_{1,t+1}(x_{1,t+1})) + \phi_{2t}\mu(x_{2t}) \tag{8}$$

$$V_{2b}(x_{2t}) = -\beta \mathbb{E}_{s_{t+1}|s_t} (1 + r_{t+1}) u_1'(c_{1,t+1}(x_{1,t+1})). \tag{9}$$

Combining the first-order conditions with the envelope conditions, we get the following characterization. In any equilibrium, the solution to the household problem in this economy (a partial equilibrium) is given by the set of decision rules $\{c_{1t}(x_{1t}), m_{2t}(x_{1t}), b_{2t}(x_{1t}), c_{2t}(x_{2t})\}$ that satisfy the following Euler equations (along with the budget constraint and the Kuhn-Tucker conditions for the multipliers), $\forall x_{1t}, x_{2t}$:

$$\phi_{1t} u_1'(c_{1t}(x_{1t})) = \mathbb{E}_{z_t|z_{t-1}} \{\beta \mathbb{E}_{s_{t+1}|s_t} \phi_{1,t+1} u_1'(c_{1,t+1}(x_{1,t+1})) + \phi_{2t}\mu(x_{2t})\} \tag{10}$$

$$u_1'(c_{1t}(x_{1t})) - \lambda(x_{1t}) = \mathbb{E}_{z_t|z_{t-1}} \{\beta \mathbb{E}_{s_{t+1}|s_t} (1 + r_{t+1}) u_1'(c_{1,t+1}(x_{1,t+1}))\} \tag{11}$$

$$z_t u_2'(c_{2t}(x_{2t})) = \beta \mathbb{E}_{s_{t+1}|s_t} \frac{\phi_{1,t+1}}{\phi_{2t}} u_1'(c_{1,t+1}(x_{1,t+1})) + \mu(x_{2t}) \tag{12}$$

In a *stationary* equilibrium, the solution to the household problem is characterized by the above equations, with $r_t = r \forall t$, and $\phi_{1t} = \phi_1$, $\phi_{2t} = \phi_2 \forall t$. In addition, as long as the Markov transition matrices for the shocks satisfy monotone mixing condition (Hopenhayn and Prescott, 1992) and given the assumption on r_t relative to β, associated with the solution is a stationary distribution of agents, which does not change period to period in aggregate, although individual agents change states due to the idiosyncratic shocks.

In what follows, I describe the properties of the model related to the credit card debt puzzle. Some of these properties are quite standard, and are presented for completeness, and in order to highlight the features of the model that relate to the credit card debt puzzle.

Property 1. ***Nontrivial distribution of assets***. *Given the assumptions on the utility functions, the equilibrium distribution of households across money and debt holdings is nondegenerate. That is, $m_{2t}(x_{1t})$ and $b_{2t}(x_{1t})$ are nontrivial functions of their states.*

As is standard, the distribution of agents is driven by heterogeneous histories of idiosyncratic shocks here, which reflect in the asset decisions and states m and b. This is obvious from the Euler Equations (10) and (11), which equate the marginal utility of first-subperiod consumption with the marginal value of carrying a dollar in cash or of "saving" a dollar by repaying debt, and from the budget constraint. It clearly follows from this property that $c_{1t}(x_{1t})$ and $c_{2t}(x_{2t})$

are also nontrivial functions of their states. Having established that there is a distribution of agents across states, I will from now on make the dependence of the decision rules on the states implicit in the notation. I next show that it is always optimal to partially insure against the preference shocks, and that the level of insurance will depend on the cost of insurance as well as the individual state.

Property 2. ***Optimally Incomplete Insurance.*** *In any equilibrium,*

1. *Optimal decisions involve partial insurance against preference shocks. That is, for any* x_{1t}, $\forall t$, $\exists\, \hat{z}_t \leq \bar{z}$ *such that* $c_{2t} < m_{2t}$ *for all* $z_t < \hat{z}_t$, *and* $c_{2t} = m_{2t}$ *otherwise.*
2. *The degree of partial insurance depends on relative returns to assets,* ϕ_{t+1}/ϕ_t, r_{t+1}, *as well as the state* x_{1t}.

Discussion. 1. The intuition is easily seen in a stationary equilibrium, although it carries through in any equilibrium of this problem. In a stationary equilibrium, $r_t = r\ \forall t$ and $\phi_{1t} = \phi_1\ \forall t$. Notice from (11) that

$$\beta \mathbb{E}_{z_t|z_{t-1}} \mathbb{E}_{s_{t+1}|s_t} u_1'(c_{1,t+1}) = \frac{u_1'(c_{1t})}{1+r}. \tag{13}$$

From this and (10), we get the following equation for m_{2t}:

$$\phi_1 u_1'(c_{1t}) = \frac{\phi_1 u_1'(c_{1t})}{1+r} + \phi_2 \sum_{\{z_i : c_{2t}(z_i) = m_{2t}\}} \Gamma(z_{t-1}, z_i)\mu(\cdot),$$

or equivalently,

$$u_1'(c_{1t})(\phi_1 - \frac{\phi_1}{1+r}) = \phi_2 \sum_{\{z_i : c_{2t}(z_i) = m_{2t}\}} \Gamma(z_{t-1}, z_i)\mu(\cdot). \tag{14}$$

Denote the right-hand side of (14) as

$$\Psi \equiv \phi_2 \sum_{\{z_i : c_{2t}(z_i) = m_{2t}\}} \Gamma(z_{t-1}, z_i)\mu(\cdot).$$

Ψ can be thought of as expected shadow value of relaxing a binding money constraint in the second subperiod, where $\mu > 0$ whenever the constraint binds. By Inada conditions on the utility function, we have $\Psi > 0$ as long as $1 + r > 1$, which implies that the constraint on c_{2t} binds in at least one state z if there is a wedge in returns between money and bonds/debt.

Now suppose that the agent knows that his next realization of z_t will be $z_t = \underline{z}$, the lowest realization. In this deterministic case, the agent chooses c_{1t}^d, c_{2t}^d and corresponding m_{2t}^d such that

$$\phi_1 u_1'(c_{1t}^d) = \phi_2 \underline{z} u_2'(c_{2t}^d),$$

where the equality comes from combining deterministic versions of the Euler equations for m_{2t} and c_{2t} (10) and (12). If the realization of the next preference shock is unknown, as in the current economy, then the agent solves, from these Euler equations,

$$\phi_1 u_1'(c_{1t}^s) = \mathbb{E}_{z_t|z_{t-1}} \phi_2 z_t u_2'(c_{2t}^s) > \phi_2 \underline{z} u_2'(c_{2t}^d).$$

From the last inequality, it is clear that $c_{1t}^s < c_{1t}^d$, while $m_{2t}^s > m_{2t}^d$ for any agent that is not borrowing-constrained, to keep all the Euler equations holding. In states $z_t > \underline{z}$, $c_{2t}^s(z_t) > c_{2t}^d(\underline{z})$.

To summarize, for any x_{1t}, there exists a cutoff level $\hat{z}_t \leq \bar{z}$, such that $c_{2t} < m_{2t}$ for $z_t < \hat{z}_t$, and $c_{2t} = m_{2t}$ otherwise.

2. Denote the agent's assets as $a_{1t} = \phi_1 m_{1t} - b_{1t}(1+r)$. By (14) and strict concavity of $u_1(\cdot)$, $\partial\Psi/\partial c_{1t} < 0$, and so $\partial\Psi/\partial a_{1t} < 0$. Also, $\partial\Psi/\partial r > 0$. That is, an increase in first-subperiod consumption increases the amount of insurance taken against the preference shocks, as does an increase in assets. At the same time, an increase in the cost of insurance r reduces the optimal amount of insurance, as long as $r > 0$.

I showed above that agents are constrained against achieving first-best in every realization of z_t since it is simply too costly, but that there is precautionary demand for money even if carrying money is dominated by repaying debt (or saving in b), so that for most states except the most constrained, for some z_t, $m_{1,t+1} > 0$ - agents will have positive liquid assets at the end of the period. As an aside, note that if there is no wedge in returns between the two assets, agents become indifferent between them, so one can insure completely against any realization of z_t as long as one holds any nonliquid assets (that is, $\sum_{z_i:c_{2t}(z_i)=m_{2t}} \mu(z_i) = 0$), while if the cost of insurance is extremely high ($r \to \infty$), agents may choose not to hold precautionary balances at all, so the money constraint would bind everywhere. Note also that if we fix s for any agent, (14) gives that more asset-wealthy people prefer to insure against preference shocks more fully - in other words, preference shocks become more important relative to income shocks, the more assets a household has.

I next show that an interior solution to the problem admits a wedge in returns between liquid assets and consumer credit, with the latter being more expensive. Since my analysis will continue in partial equilibrium, an alternative way to view this is that if prices are set such that consumer credit is more expensive than liquidity, an interior solution exists.

Property 3. ***Difference in rates of returns.*** *An interior solution to the household problem admits* $1 + r_{t+1} > \frac{\phi_{1,t+1}}{\phi_{1t}}$. *In stationary equilibrium,* $1 + r > 1$.

Discussion. Consider household Euler equations (10) and (11). For the majority of the households, the credit limit constraint does not bind, so that $\lambda(x_{1t}) = 0$, and for these households, the Euler equations give

$$\begin{aligned} u_1'(c_{1t}) &= \mathbb{E}_{z_t|z_{t-1}}\{\beta \mathbb{E}_{s_{t+1}|s_t} \frac{\phi_{1,t+1}}{\phi_{1t}} u_1'(c_{1,t+1}) + \frac{\phi_{2t}\mu_t}{\phi_{1t}}\} \\ u_1'(c_{1t}) &= \mathbb{E}_{z_t|z_{t-1}}\{\beta \mathbb{E}_{s_{t+1}|s_t}(1 + r_{t+1}) u_1'(c_{1,t+1})\} \end{aligned}$$

By property 2, $\mu_t(x_{2t}) > 0$ for some x_{2t}. Thus we have $\mathbb{E}_{z_t|z_{t-1}} \frac{\phi_{2t}\mu_t}{\phi_{1t}} > 0$, and so it is clear from comparing the right-hand sides of equations above that

$$\mathbb{E}_{z_t|z_{t-1}}\{\beta \mathbb{E}_{s_{t+1}|s_t} \frac{\phi_{1,t+1}}{\phi_{1t}} u_1'(c_{1,t+1})\} < \mathbb{E}_{z_t|z_{t-1}}\{\beta \mathbb{E}_{s_{t+1}|s_t}(1 + r_{t+1}) u_1'(c_{1,t+1})\},$$

and therefore,

$$\frac{\phi_{1,t+1}}{\phi_{1t}} < 1 + r_{t+1}.$$

In stationary equilibrium, this turns into

$$1 < 1 + r.$$

Property 3 and equation (14) give a complete characterization of agents' self-insurance behavior. Even for very good states x_{1t}, it is at most possible that agents carry *exactly* enough money to pay for consumption c_{2t} when the shock has its maximal realization, that is, they will never opt to carry more money that they would spend if $z_t = \bar{z}$.[13] By Inada conditions on $u(c_2)$, it is always optimal to have at least some consumption in the second subperiod, even in the lowest state realizations.

The above discussion leads us to consider the agents' behavior in regard to money and debt holdings. I show that the model generates the three subgroups in the population: borrowers, savers, and those who do both. The model thus replicates the credit card debt puzzle, at least qualitatively.

[13] Note also that in the model as it is written now, incomplete insurance against preference shocks implies that agents do not use cash holdings to insure against income shocks - these will instead, as is well-known, increase saving/ decrease borrowing in b_{2t}, relative to an economy with no uncertainty in income. This is because the model abstracts from cash advances, which can prompt income uncertainty to be a second channel to affect precautionary demand in money. See the Discussion section for further details.

Property 4. ***Optimality of different borrowing and saving behavior.***
In every period, there exist three subgroups of the population:

- *Borrowers have* $m_{2t} > 0$, $b_{2t} > 0$ *but* $m_{1,t+1} = 0$;
- *Borrowers and savers have* $m_{2t} > 0$, $b_{2t} > 0$ *and* $m_{1,t+1} > 0$;
- *Savers have* $b_{2t} \leq 0$, *while* $m_{2t} > 0$ *and* $m_{1,t+1} \geq 0$.

Of those who borrow in any given period, a positive measure of agents will borrow again in the next, that is, $b_{1t} > 0$ *and* $b_{2t} > 0$ *(debt revolving).*

> ***Discussion***. By property 2, $m_{2t} > 0$ for all agents in all states. Moreover, since partial insurance is optimal, for any asset level and some realizations of shock z_t, the money constraint binds, while for others it does not, so we have $m_{1,t+1} = m_{2t} - c_{2t} = 0$ for some (x_{1t}, z_t), while $m_{1,t+1} > 0$ for other (x_{1t}, z_t). Thus we have the money holding combinations for the three subgroups.
>
> It remains to show that $b_{2t} > 0$ for some states x_{1t}. Suppose household's assets a_t are at some very low level such that only minimal insurance is optimal, as given by (14), and we get $\mu(\cdot) > 0 \,\forall\, z$, so from Euler equations (10) and (12),
>
> $$\begin{aligned} \phi_{1t} u_1'(c_{1t}) &> \mathbb{E}_{z_t|z_{t-1}}\{\beta \mathbb{E}_{s_{t+1}|s_t} \phi_{1,t+1} u_1'(c_{1,t+1})\} \\ z_t u_2'(c_{2t}) &> \frac{1}{\phi_{2t}} \beta \mathbb{E}_{s_{t+1}|s_t} \phi_{1,t+1} u_1'(c_{1,t+1}). \end{aligned}$$
>
> That is, these agents value present consumption more than future consumption, and are willing to shift assets from tomorrow to today in order to reduce the inequalities. They are able to do so by borrowing, so we have $b_{2t} > 0$. In the next period, those who still have low assets will have to "repay" current debt by borrowing more, so they are revolving the present debt, and we have $b_{1,t+1} > 0$ and $b_{2,t+1} > 0$.

This last property shows that at different asset positions, it is optimal for the households in the model to engage in differing borrowing and saving behavior, thus potentially delivering the three subgroups that are observed in the data. It is important to note, however, that analytically it is impossible to say whether in the stationary distribution, households will actually find themselves at all of these asset positions. For example, we know that for some low level of assets a household will borrow. But we do not know whether any model household will actually have that low level of assets. This question can only be answered quantitatively, and in the next section, I show that such low levels of assets do in fact occur in the calibrated model.

To summarize, the model delivers all of the empirically desirable features of the credit card debt puzzle in the data: precautionary demand for money, existence of an equilibrium when credit is costly, and a subdivision of the population into three groups with liquid saving and borrowing behavior akin to those in the data. Note that the aggregate distribution of the population is plausible in this respect: people with very low assets and low shocks are borrowers, people in the middle of the asset and shock-history distribution are the puzzle group, while those at the top are savers only. Finally, it is important to note that households in the model will move in and out of the "puzzle" subgroup depending on their shock histories, so that no households would be in this situation permanently. I now calibrate and compute the model, in order to evaluate the power of the liquidity need hypothesis to account for the credit card debt puzzle.

4 Computation and Calibration

4.1 Transformed Model

For the purposes of computation, I make some adjustments to the model. First, I make the utility function nonseparable, combining $u_1(c_{1t})$ and $z_t u_2(c_{2t})$ above into a new utility function, $u(c_{1t}, z_t c_{2t})$, and assuming that the utility function is thrice differentiable, strictly increasing and strictly concave in both arguments, and its third derivative is strictly positive. Inada conditions are also assumed to hold. The reason to make the utility function non-separable is that in reality there is likely to be an interaction between household spending on cash-only goods and spending on cash-or-credit goods, and this interaction is likely to have an important effect on results. Second, I introduce an interest spread for saving and borrowing, to match it in the data: borrowing on credit cards carries a much higher interest rate than saving in other financial assets does, on average. As this is a partial equilibrium model, these prices are set exogenously. Also, I normalize $\phi_{jt} = 1 \, \forall \, j, t$, which is innocuous given that I am not studying monetary policy-related issues, and in addition, I will focus on stationary equilibrium, so that all aggregate variables will be constant.

Finally, in order to reduce computation time, I reduce the state space in the first subperiod (no such possibility exists in the second). In particular, define assets ("cash-at-hand") to be, given assumptions on prices listed above:

$$\begin{aligned} a_t &= m_{1t} - b_{1t}(1 + r_t), \text{ where} \\ & r_t = r^b \text{ if } b_{1t} > 0 \\ & r_t = r^s < r^b \text{ if } b_{1t} < 0 \end{aligned}$$

The first-subperiod problem can be then rewritten as:

$$\begin{aligned} V_1(s_t, z_{t-1}, a_t) &= \max_{c_{1t}, b_{2t}, m_{2t}} \mathbb{E}_{z_t|z_{t-1}} V_2(s_t, z_t, m_{2t}, b_{2t}, c_{1t}) \\ \text{s.t. } c_{1t} + m_{2t} - b_{2t} &= s_t + a_t. \end{aligned} \tag{15}$$

Given all the adjustments, the second-subperiod problem becomes:

$$\begin{aligned} V_2(s_t, z_t, m_{2t}, b_{2t}, c_{1t}) &= \max_{c_{2t}} u(c_{1t}, z_t c_{2t}) + \beta \mathbb{E}_{s_{t+1}|s_t} V_1(s_{t+1}, z_t, a_{t+1}) \\ \text{s.t. } c_{2t} &\leq m_{2t} \\ a_{t+1} &= m_{2t} - c_{2t} - b_{2t}(1 + r_{t+1}), \end{aligned} \tag{16}$$

where the interest rate r is determined by whether or not the agent borrows or saves. As before, this problem is well-behaved and the solution exists, given the utility function specification and appropriate boundary conditions, which in practice amount to setting bounds on the constraint set that do not restrict the decision rules. 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 complete robust to the choice of method. Details are in the appendix.

4.2 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 consumption aggregator between the two consumption goods:

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

This choice satisfies all the necessary assumptions on the utility function listed above. The utility function gives three parameters to calibrate: α, ν and δ. 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. I calibrate the parameters of the income process outside the model, set δ to follow a standard choice in the literature, set the prices 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 in calibrating the model, 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 money holdings that households choose to keep. Thus, these key quantities of interest are left free to speak for the performance of the liquidity need hypothesis in accounting for the puzzle at hand.

The calibration of the income process is non-trivial in the context of this study. The standard calibration procedure of the income process parameters involves imposing an AR(1) process with normally distributed errors on income data from household surveys such as the PSID (e.g., Storesletten, Telmer and Yaron, 2004a, 2004b). However, micro data sources that have good measurements of income provide income data with an annual frequency only. Imposing an AR(1) process on annual data and using time disaggregation to get the monthly frequency leads to an extremely persistent monthly process with little variance, which generates little information about income uncertainty on a monthly basis. Thus, my approach has to depart from this practice. Instead, I pose a 3-state discrete Markov process as follows. The income states are chosen to match the relative average earnings of white-collar workers (s_3), blue-collar or service sector workers (s_2), and the value of unemployment benefits (s_1). This is one of several possible choices: for example, one could choose relative earnings of college-educated versus non-college-educated workers instead. I take the data on relative earnings above from the 2004 Bureau of Labor Statistics reports on earnings of full-time workers by occupation. Note that while one might like to have a greater number of income states, the key limitation in the number of states I can pose is that I have to calibrate the transition matrix between the income states, which prevents me from using, say, income quantiles - there are not enough relevant data at monthly frequency to compute a richer set of transition dynamics in question.[14] This is not a severe limitation in this context, as I discuss below.

In order to calibrate transition probabilities between income states, I use the following data. The average duration of unemployment in 2001, according to the BLS, was 13.5 weeks, which determines the monthly probability of exiting unemployment. I let the probabilities of exiting from unemployment into blue-collar and white-collar occupations to be determined by the shares of blue- and white-collar workers among the unemployed in the BLS data, which were 56% and 44% respectively in 2001.

Associated with the transition matrix Γ_s is the invariant distribution of agents across the three income states. Denoting this distribution as $\{\gamma_1^*, \gamma_2^*, \gamma_3^*\}$, I get two additional conditions: γ_1^* should equal the average monthly unemployment rate, which was 4.75% in 2001, and γ_2^* - the share of blue-collar workers among the employed, which was 43.5%. γ_3^* is the complement of the other two. Finally, I need to set one more parameter: I calibrate the probability of transitioning from a blue-collar job to a white-collar job to upward mobility rates for blue-collar workers, as computed by the BLS and reported by Gabriel (2003). The reported average monthly probability of an upward occupational move by a blue-collar worker was around 0.7% in 1998-1999. It is plausible that in 2001, this number might have declined slightly, due to a shift in economic conditions, but as I do not have specific information to that effect, I use this statistic here. The parameters of the resulting earnings process are reported in table 8.

This calibration has a clear limitation: it has no hope of capturing the top tail of the income distribution, nor indeed does it mimic the overall income inequality in the U.S. The top income level is only five times the lowest income level in this calibration, and over one-half of the population experiences the highest (also most persistent) income state. That is, the bottom tail is also clearly understated in the model. Whereas the Gini coefficient for income in the U.S. was 0.55 in 2001 in my sample, in the model it is only 0.15.

However, for the current exercise, it is not a significant problem. First, the credit card debt puzzle, as I demonstrated in my data analysis, is a phenomenon concentrated in the middle of the distribution, with mean and median incomes far below the top tail. Thus, understating the top tail of the distribution will restrict me from matching the top tail in the model, but

[14] In ongoing separate work, Telyukova and Klein (2008) are looking at ways to solve this problem, based somewhat on the methodology of Gervais and Klein (2006).

Table 8: Earnings Process

	Parameters	Value
Earnings states	$\{s_1, s_2, s_3\}$	$\{\ 0.2,\ \ 0.59,\ \ 1.0\ \}$
Transition matrix	$\Gamma(s_t, s_{t+1})$	$\begin{bmatrix} 0.706 & 0.163 & 0.131 \\ 0.021 & 0.971 & 0.008 \\ 0.010 & 0.007 & 0.983 \end{bmatrix}$
Invariant distribution in earnings	Γ_s^*	$\{\ 0.048,\ \ 0.414,\ \ 0.538\ \}$

that population group is not of most concern for the question at hand. Second, insofar as I understate the bottom tail of the data income distribution, I am biasing my results on the size of the borrower-and-saver group downwards - so the result of the computation can then be seen as a lower bound on what the model can account for. I will discuss this matter more in the results section.[15]

Aside from the income process, 8 parameters remain, including the prices. I choose the risk aversion parameter, $\delta = 2$, on the conservative side of the standard range in the literature. 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 in the SCF.

This leaves me with the discount rate β, the parameters of the consumption aggregator α and ν, and the preference process parameters, which have to be calibrated from the model. I estimate these parameters from within the model using a minimum-distance estimator based on the simulated method of moments. That is, the parameter vector is the minimizer of the sum of weighted squared distances between a relevant set of data moments and their simulated model counterparts.

[15] One final note on the income process: with the probabilities of exiting unemployment into white- and blue-collar jobs proportional to shares of white- and blue-collar workers among the unemployed, it is clear that a blue-collar worker in the model who becomes unemployed has a significant positive probability of becoming a white-collar worker. Thus, the uncertainty that blue-collar workers face is crucially understated in this broad calibration. This means that the model will have a harder time matching the properties of debt in the data - in particular, this is another source that will put a dampener both on the size of the puzzle group in the model, and possibly also on the liquidity holdings.

In the deterministic cash-in-advance literature, the parameters α and ν are typically calibrated from a money demand equation which is a direct implication of the first-order conditions of the problem when the cash-in-advance constraint always binds. Due to idiosyncratic uncertainty, however, these implications do not hold in my model (the money constraint does not always bind), and no closed-form counterparts exist. For the preference shock parameters, then, I assume that $\log(z_t)$ follows an AR(1) process, so the parameters to calibrate will be a persistence parameter ρ_z and standard deviation σ_z of this process, which I will then discretize 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.

The preference shock process is clearly not observed in the data, but the way households respond to these shocks is through their cash-only good consumption. Thus, the preference shock process has to match properties of consumption of cash-only goods in the data, namely its persistence (autocorrelation) and volatility (conditional standard deviation). In choosing the other targets for the estimation, I discipline the model by not setting as targets quantities that predetermine the size of the puzzle or liquidity holdings directly, since these are the quantities I seek to explain. I choose nine moments in total with this in mind; they are volatility of liquid consumption for each subgroup of households, autocorrelation of liquid consumption for the sample as a whole, mean cash-only good consumption relative to income for each of the subgroups, the mean debt-to-income and mean nonhousing wealth-to-income ratio in the population.

A word on the last two targets is in order. Aside from matching the micro-data properties of liquid consumption, which is the main focus of this exercise, it is also important to gauge how well the model does in reproducing some aggregate distributional statistics. Given that the focus of the paper is on the co-existence of assets and debt, the distributional statistics in question should focus both on the asset-to-income and debt-to-income ratio. There are many ways to measure these in the data, as there are many possible definitions of wealth. For the purposes of this estimation, I choose aggregate non-housing wealth and aggregate revolving unsecured debt (the majority of which is credit card debt), computed in the SCF.

The choice of non-housing wealth as a target, rather than total wealth, is tailored to the model - which, by design, has no role for a durable asset that provides consumption services.

One could argue that the fact that households can borrow against their home equity would be a reason to consider housing as a financial asset too - but in the 2001 data, there is not very significant evidence of liquidity borrowed against housing, so I claim that treating housing as primarily a durable nonfinancial asset is the correct approach in the current model context. Similarly, debt could be measured as unsecured and secured revolving debt - allowing to include home equity lines of credit - but as the SCF in 2001 does not show significant uptake of such lines of credit, I use only unsecured debt as a measure.[16]

There is one other important detail in my calculation of the aggregate target ratios of debt and wealth to income. Because the income calibration does not represent the top tail of the income distribution well, the calibrated model will by design have a hard time matching the average asset statistics for the whole population, and attempting to do so may bias estimates. To account for this limitation, I map my targets to the income calibration in the following way: I compute and target in estimation the debt and wealth ratios in the bottom 75% of the 2001 U.S. *income* distribution, rather than in the whole population. Note that doing so does not reduce the wealth dispersion in the data all that significantly, so that the average wealth-to-income ratios are only slightly below what they would be in the population as a whole. For example, nonhousing wealth-to-income ratio in the whole population is 2.14 - compare that to the measure for the bottom-75 % of the income distribution used here of 1.69. The difference is not staggering.[17]

Finally, as I mentioned in the data section and appendix, the properties of liquid consumption are sensitive to whether or not food, alcohol and tobacco are included in the calculations. Specifically, including them increases the average liquid consumption-to-income ratio for the households, but also decreases the measured volatility of liquid consumption, since food is one

[16] The uptake of home equity lines of credit (HELOC's) surged significantly at the end of 2001, as a result of record-low interest rates and many households refinancing their mortgages, at which time they were offered HELOC's for free. This increase, at 30% a year, lasted until 2005 or so, according to the Federal Reserve Board. The 2001 survey data were thus collected too early to reflect this upsurge.

[17] The fact that the aggregate targets concern only the bottom 75% of the population may raise the question of why the other targets are not calculated for the same subsection of the data sample. The reason is that much of the other analysis is done by subgroup, which are each already located somewhere on a specific subsection of the income distribution. I have computed all the targets for the bottom-75% however, to find that consumption-to-income ratios will increase for all the subgroups in question to around 0.70-0.72, while autocorrelation of log liquid consumption and its standard deviation remain unchanged. Thus, changing the targets to bottom-75% would actually favor my model, because an increase in average consumption-to-income ratios will produce some increase in the optimal household liquidity holdings as well.

of the more invariant expenditure components. The reason I single out this category is that one is generally able to pay for food, alcohol and tobacco by credit in most places - so it is not strictly a cash-only good. Yet, survey data overwhelmingly suggest that people choose to pay for these goods using liquid assets. The question of why is a matter for other research, and there is a payments literature that addresses it. For my purposes, I let data dictate that food is a cash good, and thus food is included in the estimation.

In sum, I estimate the five parameters within the model based on nine moments. For each set of parameters in the minimization process, the procedure solves the model, simulates a 252-month panel of 100,000 households, computes the moments from it, and compares them with the moments in the data. The complexity of the problem prevents me from using gradient-based minimization methods. Thus, for the minimization 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, and the wealth-to-income ratio). Data covariances of the moments in question are not possible to compute in this exercise, since the moments come from two different data sets.

4.3 Model Fit and Resulting Parameters

In order to assess the fit of the calibrated model, table 9 presents the target moments in the data and the model. As discussed above, I have 9 targets and 5 parameters: this overidentification means that I do not have enough instruments to match all of the moments perfectly, but the closeness of the match allows me to judge the fit of the model. The calibrated model fits most targets closely. The crucial moments concern the saver and borrower-and-saver groups. For these, both the liquid consumption-to-income ratio and the volatility of cash-only good consumption are matched quite well, and further, autocorrelation of liquid consumption, measured across subgroups, is matched nearly perfectly.

The borrower subgroup presents a challenge to the model, as is evident in the targets: the model does not match the borrowers' characteristics particularly closely, overpredicting their liquid consumption-to-income ratio, and significantly underpredicting the volatility of their liquid

Table 9: Calibration Targets - Data and Model

Target	Data	Model
Liquid consumption/earnings ratio^a*:* Borrowers	0.719	0.872
Borrowers & savers	0.618	0.597
Savers	0.629	0.613
Autocorrelation (annual) of log liquid consumption:	0.226	0.227
Percent cond'l. st. dev. of log liquid consumption^a*:* Borrowers	0.201	0.115
Borrowers & savers	0.210	0.207
Savers	0.226	0.227
Mean debt/income ratio^{b,c}	0.058	0.059
Mean wealth/income ratio^{b,c}	1.685	0.607

Notes: *(a)* The cash-only good series includes food. *(b)* Wealth is measured as non-housing wealth; debt is revolving unsecured debt. *(c)* The moment is computed for the bottom 75% of the U.S. *income* distribution.

consumption. The reason is clear. It is very difficult to match the group I call "borrowers" in the data with that group in the model. In the data, these are households that report having no, or very little, liquidity. In practice, what the survey data do not measure is holdings of cash, and there is likely a number of households whose liquidity holdings are much higher than what we observe, due to it being held as cash rather than in a bank account. In addition, for the 12 months of expenditures in the data, I only have one annual observation of the household's asset position - so I cannot observe the household's subgroup status changes from month to month. This may overstate the duration of borrower status for some households, and may thus make their time series characteristics appear closer to other groups than they are in reality. In the model, the only households that appear as borrowers are those who get hit by a binding expense shock, so that they spend all of their money by the end of any given month. This includes two types of households: those who perpetually hold very little liquidity, so that their liquidity constraint binds in (nearly) all preference shock realizations, and those who hold sizeable liquidity but encounter the worst shock realization. The model, due to properties of the income calibration, will have trouble generating enough of the former - as discussed above, there are not enough people in the lowest tail of the distribution, relative to the data. The latter group, those with the worst preference shock realizations, are unlikely to stay in the borrower category for long. This makes it difficult to measure the time-series characteristics of the borrowers' expenses over time:

the former group dominates, and their expenses are nearly constant, so volatility is understated relative to the data; their liquid expense-to-income ratio is likely to be overpredicted as well.

The two final targets concern the average debt-to-annual-income and wealth-to-annual-income ratios. The debt-to-income ratio is matched nearly perfectly, but the average wealth-to-income ratio is underpredicted significantly. The performance of the model along the average asset dimension is not surprising. There is only one parameter that largely determines how much wealth and how much debt there is in the economy: the discount factor β. This one parameter cannot match both moments - the higher the value of β, the higher the wealth in the economy, but necessarily, the lower the debt. Further, it is a common problem in this type of heterogeneous-agent models that the wealth distribution is difficult to match well with simple asset structures like in my model. Models that try to get the distribution of wealth to be disperse enough have to be explicit about the richness of features of the U.S. asset market, including the Social Security system, business and home ownership, high and stochastic capital gains, and others (see, for example, Quadrini and Ríos-Rull, 1997). These features are clearly beyond the scope of this model, nor are they the focus of the exercise. Therefore, in this context it is acceptable that by construction, the model matches the amount of debt in the economy well, and average wealth not nearly as well. The results of interest here, which concern liquid assets, will not be impacted by the total amount of wealth, largely nonliquid, accumulated by the model households.

Table 10 presents the resulting parameterization. The discount factor is equivalent to 0.9234 in annual terms. The parameters of the utility function are in themselves of interest and a contribution of this paper: to date, to my knowledge, in liquidity-based models, these parameters were estimated in deterministic cash-in-advance models only. α, the weight on cash-only goods in the CES utility function, is 0.58, which once again confirms that they are an important part of a household's expenditures - not surprising given that I consider payments to mortgages and household repairs under this heading. The parameter that measures elasticity of substitution between cash and credit goods is approximately $\nu = -1.5$, that is, cash and credit goods are compliments, rather than substitutes. Comparing these with other estimates from the literature, Chari et al (1991) and others after them find an estimate for α of around 0.62, which is very close to my estimate and this is encouraging; however, their ν tends to be on the order of 0.79-0.84.

Table 10: Calibration

	Parameter	Value
Interest rates	r_s	0.0033 (annual $r_s = 0.04$)
	r_b	0.0107 (annual $r_b = 0.14$)
Risk aversion/IES	δ	2.0
Discount rate	β	0.99338 (annual $\beta_a = 0.92339$)
Consumption aggregator	α	0.58377
parameters	ν	-1.49669
Preference shock process:	ρ_z	0.46287
AR(1) with discretization	σ_z	0.48449

It appears, then, that idiosyncratic uncertainty makes a big difference for the substitutability parameter and, in my view, brings it closer to what one might intuitively expect.

Finally, the estimates of the preference process are of importance, since this study presents a new attempt to quantify unobservable idiosyncratic uncertainty from microdata specific to liquidity needs. Notice that the role of these parameters is central in determining the properties of liquid consumption in the model: persistence of consumption is related to persistence of the preference shocks, and its variability - to the volatility of shock states. The estimated monthly AR(1) parameter on $\log(z)$ is 0.46. The AR(1) specification is flexible, encompassing anything from a very persistent shock process to an i.i.d. one. The high outlier preference shock states are likely to be extreme events, as consumption patterns in the data would suggest, so we would expect their persistence to be low. The estimate of 0.46 suggests that the extreme realizations of the shock are relatively rare and rarely persist for more than one period. The standard deviation of the shock process is estimated at 0.48. 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, and it is interesting to note that variability of the shock itself is more than twice the variability of observed liquid consumption.

Based on the analysis of the calibration targets, the parameterization described above produces a realistic economy in terms of its mapping to the data; the moments that are missed in

Table 11: Results - Subgroup Size (Percent)

	Data	Model
Borrowers	5.2	1.4
Borrowers & savers	27.1	19.7
Savers	67.7	78.9

estimation are missed for reasons external to the model and by design, and with consequences not central to the puzzle in question, which is discussed in the next section.

5 Results

As mentioned before, I left the magnitudes of interest for answering the central question of this paper untargeted in calibration. The model is mapped to the data based on quantities unrelated to the results of interest, and this freedom allows me to measure exactly how much of the puzzle is accounted for by the liquidity need hypothesis with preference uncertainty as the main driving force. To measure this, I focus on the size of the subgroups (borrowers, borrowers-and-savers, and savers), as well as liquidity holdings that each subgroup optimally chooses.

Table 11 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 19.7% of the population, while in the data, it is 27.1%. Thus, the model accounts for 73% of the puzzle group. It overstates the size of the saver group, at 79% relative to the 68% in the data, and understates the size of the borrower group, putting it at 1.4% instead of 5.2%. The intuition for these facts is similar to that discussed in the calibration section. First, income calibration understates the amount of uncertainty that households face in the data, and is likely to underpredict how many people choose to borrow as a result - which leads to the underprediction of the numbers of both borrowers and borrowers-and-savers in the model. Second, 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, not captured by the model by construction.

In order to measure liquid assets, I have to define what the money holdings observed in the data are. As discussed, a cross-sectional average of money holdings in the SCF reflects an average monthly amount of money in the bank accounts, since households are continually

Table 12: Results - Liquid-Asset-to-Income Ratio, Average During Month, **Median** Household

	Data	Model	Model/Data
Borrowers	0.10	0.39	3.90
Borrowers & savers	0.79	0.43	0.55
Savers	0.88	0.49	0.56

interviewed throughout the month. This cannot apply to the borrower group, however: it is not likely that households in this group truly never hold liquid assets during the month, given their average liquid spending documented above, so these must be households observed at the end of the month who have drawn down all of their liquid assets, most likely due to binding resource constraints. Since in the model I observe money holdings at two points during the month, rather than just one, I study average monthly money holdings for all households. However, I separate out the borrower group by looking at end-month liquid holdings, since in the model, as in the data, no household will have zero liquid assets at the beginning of the month.

In table 13, I present liquid asset holdings relative to income by subgroup, in the data and in the model. I focus on the median household in both the model and the data, since, for reasons having to do with the difficulty of matching the upper tails of both the income and wealth distributions, the model is biased from the outset toward medians, rather than means. The last column translates the model's results into per-dollar amounts relative to the data. In particular, for the median household in the puzzle group, the model matches 55 cents of every dollar held by the median puzzle household in the data. This number is 56 cents for the saver group. Notice that for the borrowers, the model matches 390% of the money holdings in the data. The reason, again, is that borrowers in the data are households we observe with near-zero liquidity holdings. Yet, in the model, nobody chooses to have zero liquid holdings, because of Inada conditions and optimality of self-insurance. Thus, this number in the data is actually hard to compare to the model - I present the result for completeness only.

6 Discussion of the Results

There are several ways in which the current results on the size of the borrower-and-saver group and liquidity holdings may be seen as a lower bound. First, as discussed in the calibration section, the properties of the income calibration, necessarily restricted by the lack of micro-data of sufficiently low frequency, are likely to understate the size of the puzzle group in the model, as well as, possibly, liquidity holdings. Second, 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 holder 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 minimum balance 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 I can account for will rise by the share of the total account balance that such a minimum balance captures in the data.[18] 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.

Third, and most importantly, 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. For example, if one should lose one's job and paycheck, one still needs to pay the mortgage each month. And in reality, unlike in the model, getting liquid assets on the spot from any source other than the bank account is actually very costly. The most available method, at least in 2001, was a cash advance from a credit card, which would incurs

[18]For example, for a median household with a \$3,000 liquidity holding, if the minimum balance on its account were \$1,000, then that \$1,000 would be unusable for daily expenses, assuming the household wants to avoid fees, which can be sizeable. Thus, I would only have to account for \$2,000 in this account. Since my model matches 55% of the total median balance, which translates to roughly \$1,650, with the minimum balance requirement, I would be capturing 82.5% of "operational" liquidity balance of such a household.

a withdrawal fee of several percent of the amount withdrawn, and *in addition*, would incur an interest rate much higher than that on credit card purchases (20-25% versus 14%, on average). Further, this interest would begin accumulating immediately upon withdrawal, without a grace period. There is also an additional cost which is that if a household has a balance on a credit card and has a cash advance on it, any payment applied toward the card account goes toward the lower-interest balance first. Thus, a household without a paycheck would find itself in an extremely costly borrowing situation if it did not have extra money in the bank. Borrowing from sources other than credit cards, such as bank loans, is also costly: bank loans and real estate loans tend to be large lump sums, and involve significant opening/closing costs and time delays.

The idea, then, is that both preference (expense) and income uncertainty, both of which are present in the data, may easily provide a precautionary motive for liquidity holdings for most households. I have disentangled the influence of one. Extending the model by adding a direct cost of transfers from consumer credit to liquidity, and recomputing and recalibrating it to quantify how all the costs of borrowing affect demand for liquidity, is a worthwhile but difficult exercise, in that the model thus extended becomes much more difficult to solve due to additional non-convexities and an expansion of an already large state space. Thus, it is beyond the scope of this paper to implement such an extention. But I did compute two-period versions of the benchmark and the extended models, and this exercise suggests that income uncertainty adds to liquidity demand significantly (in the example, which is only indicative, it increased by 30-50% relative to the benchmark case, depending on the exact asset specification). Thus, it should be emphasized that the liquidity hypothesis may not just be a strong explanation for the puzzle, but indeed most of the explanation.

7 Conclusion

This paper has presented a new explanation for the credit card debt puzzle, the phenomenon that many U.S. households who revolve expensive credit card debt also keep significant low-return liquid assets in the bank, without using them to pay off the debt. 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. It is crucial that there is a significant unpredictable component to these expenses, so households not only hold the money for pre-committed expenses, but also have an additional stock of liquidity to insure against such unexpected spending needs. 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 careful measurement of how much of the puzzle can be accounted for by the liquidity need hypothesis. After documenting the puzzle carefully in the data, I pose a dynamic stochastic model of household portfolio choice with two types of idiosyncratic uncertainty timed in such a way 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 moments in the model, in such a way that none of the quantities I target in calibration are related to quantities of interest in accounting for the puzzle. The parameter estimates are in themselves of interest, providing insights into measuring unobservable idiosyncratic uncertainty and substitutability between cash and credit goods in micro data. Further, I find that the hypothesis successfully accounts for 73% of the households who revolve debt while having money in the bank, and for a median such household, it accounts for at least 55 cents of every dollar held in liquid assets. There is a set of compelling reasons to view these results as a lower bound, so that the hypothesis here presented can be assumed to account for the majority of the puzzle at hand. Even though there may be households for which alternative explanations along the lines of time inconsistency or strategic bankruptcy behavior are valid, or even dominant, for this particular puzzle, I am inclined to conclude that the need for liquidity in the data is strong enough to account for the majority of it.

References

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

[2] Athreya, K. 2004. "Fresh Start or Head Start? Uniform Bankruptcy Exemptions and Welfare." Federal Reserve Bank of Richmond.

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

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

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

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

[7] Browning, M. and A. Lusardi, 1996. "Household Saving: Micro Theories and Micro Facts." *Journal of Economic Literature*, 34, 1797-1855.

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

[9] Budría Rodríguez, S., J. Díaz-Giménez, V. Quadrini and J.V. Ríos-Rull, 2002. "Updated Facts on the U.S. Distributions of Earnings, Income and Wealth". *Federal Reserve Bank of Minneapolis Quarterly Review*, 26(3), 2-35.

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

[11] Castañeda, A., J. Díaz-Giménez, and J. V. Ríos-Rull. 1998. "Exploring the Income Distribution Business Cycle Dynamics". *Journal of Monetary Economics*, 42, 93-130.

[12] 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.

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

[14] 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.

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

[16] Durkin, T., 2002. “Consumers and Credit Disclosures: Credit Cards and Credit Insurance.” *Federal Reserve Bulletin,* April.

[17] Eisfeldt, A., 2003. “Smoothing with Liquid and Illiquid Assets”. Kellogg School of Management.

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

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

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

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

[22] Gabriel, P., 2003. “An examination of occupational mobility of full-time workers”. *Monthly Labor Review*, September.

[23] Gervais, M. and P. Klein. 2006. “Measuring Consumption Smoothing in CEX Data”. Mimeo, University of Western Ontario.

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

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

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

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

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

[32] 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.

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

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

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

[36] 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.

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

[38] Meisenheimer, J. and R. Ilg, 2000. Looking for a 'Better Job': The Job Search Activity of the Employed. *Monthly Labor Review*, September.

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

[40] 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.

[41] 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.

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

[43] Shimer, R. 2005. "The Cyclical Behavior of Equilibrium Employment and Vacancies." *American Economic Review*, 95(1), 25-49.

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

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

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

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

[48] Telyukova, I. A. and R. Wright, 2008. "A Model of Money and Credit, with Application to the Credit Card Debt Puzzle." Forthcoming, *Review of Economic Studies.*

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

[50] 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.

[51] 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, and who held credit card debt some time during that period. 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. As an aside, note that it is well-known that debt information tends to be underreported in the SCF (Evans and Schmalensee, 1999), but this serves to my advantage, since at worst it understates the size of the puzzle in the data, or the amount of debt that households hold. 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 fairly reliably 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 (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

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
Childcare, 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

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.

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?"

Expenditures on food, alcohol and tobacco deserve special attention. In separating out the cash-only category, it was important to make a decision regarding goods that consumers mostly *choose* to pay by liquid instruments, even though credit cards may be an option. For example, it is clear from the survey, as well as other general payment method studies by the Federal Reserve, that households tend to prefer to pay for essentials, such as food, by a check, debit card, or cash. However, in most supermarkets, credit cards became an option in the mid-1990's; a more questionable category is food in restaurants, since many smaller fine restaurants opt not to accept credit cards. A second issue is that in the CEX, these good groups are goods for

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

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, babysitting, public transportation, health insurance; 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, etc.
Durables	Households furnishings and major appliances, vehicle purchases
Education	Tuition and fee expenses, textbook purchases

which the question in the survey asks the household to remember a monthly average spent over the last three months, rather than an accurate expenditure in each month. This would tend to depress the measure of consumption volatility of whichever group food is included in. A further discussion of this group is in the text.

The resulting categories are presented in table A.3.3.

Appendix B Computational Algorithm

The problem is divided into an outer maximization, which corresponds to the first-subperiod household problem, and an inner maximization, which is the second-subperiod choice of consumption given preference shock realization. As the inner maximization has only one control variable, it is solved by approximating the first-order-condition. Value function iteration can be used instead, and produces the same results. The algorithm is as follows.

1. Discretize the state space: grids are made on m, b, c_1 and a. Shock spaces are discrete, as described in calibration methodology. I use linear interpolation for approximating value functions in between grid points.

2. Guess the value function, $V_1^0(s_t, z_{t-1}, a_t)$. For the first-order condition in c_{2t}, numerically compute the derivative of the value function with respect to liquid holdings:
$$V_{1m}^0(s_t, z_{t-1}, a_t) = \frac{V_1^0(s_t, z_{t-1}, a_t + \Delta) - V_1^0(s_t, z_{t-1}, a_t)}{\Delta}$$

3. *Inner Maximization.* Given the above guess, solve the second-subperiod problem for each state $(s_t, z_t, m_{2t}, b_{2t}, c_{1t})$ in two steps:

   - Assume the constraint does not bind so that $\mu(s_t, z_t, m_{2t}, b_{2t}, c_{1t}) = 0$. Then solve the first-order condition with $\mu_t(\cdot) = 0$:
   $$z_t u_2(c_{1t}, z_t c_{2t}(x_{2t})) = \beta \mathbb{E}_{s_{t+1}|s_t} V_{1m}^0(s_{t+1}, z_t, a_{t+1}) + \mu_t(\cdot) \tag{17}$$
   The right-hand side is computed using the numerical derivative of the previous guess of the value function, interpolated between points on the m grid.
   - Check that the constraint is satisfied, that is $c_{2t} < m_{2t}$. If it is, record the solution as $(c_{2t}^*, \mu^*) = (c_{2t}, 0)$.
   - If the constraint is not satisfied, set the constraint to bind: this gives $c_{2t}^* = m_{2t}$. Then use (17) to solve for $\mu(\cdot)$. Record the solution as $(c_{2t}^*, \mu^*) = (m_{2t}, \mu)$.

4. *Outer Maximization.* For each state (s_t, z_{t-1}, a_t), perform a maximization via discretization on (m_{2t}, b_{2t}), to maximize the value function at that point. Denote the maximized value $V^1(\cdot)$.

5. Convergence check. If $V^1(\cdot) \approx V^0(\cdot)$ for all states, then we have the solution. Else, update the value function's new guess as the last iteration's computation, $V^0(\cdot) = V^1(\cdot)$ along with its numerical derivative, and restart at step 3.

The algorithm further employs monotonicity and concavity properties of the value function, as well as, periodically, Howard's acceleration algorithm between the maximization steps. The algorithm converges in around 19-25 iterations, depending on the parameters.