

UC San Diego

Recent Work

Title

Household Need for Liquidity and the Credit Card Debt Puzzle

Permalink

<https://escholarship.org/uc/item/4c67r71r>

Author

Telyukova, Irina

Publication Date

2008-01-03

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: May 5, 2009

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 as the “credit card debt puzzle”, as it appears to violate the standard no-arbitrage condition. In this paper, I quantitatively evaluate demand for liquidity as an 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 in situations where 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 appears to have a sizeable unpredictable component. This would warrant holding positive balances in liquid accounts both for transactions and precautionary purposes. 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 between 85% and 104% of the households in the data who hold consumer debt and liquidity simultaneously, and for between 56 and 62 cents of every dollar held by a median household in the puzzle group. Thus the transactions and precautionary demand for liquidity appear to be a significant factor in accounting for the credit card debt puzzle.

*I am indebted to Victor Rios-Rull, Randall Wright, Jesús Fernández-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, USC, Notre Dame, Arizona State, Wash. U. St. Louis, Maryland, Queen’s, Western Ontario, University of British Columbia, the Board of Governors of the Federal Reserve, Federal Reserve Banks of New York, Cleveland, Philadelphia and Atlanta, European Central Bank, the Canadian Economic Association, the Midwest Macroeconomic Meetings, and the SED. I am grateful to the Jacob K. Javits Graduate Research Fellowship Fund and the Federal Reserve Board of Governors Dissertation Internship program for research support.

1 Introduction

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

Gross and Souleles (2002) were among the first to document this fact. They suggested several possible explanations for this behavior, two of which have been pursued in the literature 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, as they hold significant and positive financial and nonfinancial wealth. 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 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 may be many households whose behavior they are not likely to capture, for reasons mentioned above. In this paper, I offer a rigorous examination of an alternative 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 it can account for. I focus on the need for liquidity as the possible reason. The premise is that there are large parts of household monthly expenditures that cannot be paid for by credit card, so they must be paid by liquid instruments.¹ Such payments often are substantial in size, and include predicted expenses (such as mortgage and rent payments, utilities, babysitting and daycare services), as well as significant unpredictable ones (such as major household repairs, auto repairs and other types of emergencies).² Some of these are universally cash-only goods, while others may or may not be. For example, large 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) mention the idea of transactions demand for liquidity as a possible contributor to why people hold debt and money simultaneously, but dismiss it as insufficient for

¹I 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.

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

the purposes of explaining the puzzle. A careful quantitative analysis of the hypothesis presented here, however, is an involved exercise, from both theoretical and empirical perspectives, and it is crucial, because it allows us to evaluate the possibility that the puzzle may largely fit within the standard rational expectations framework, and hence appear much less puzzling in that context. Crucially, I am interested in understanding the nature and magnitude of not only transactions demand, but also *precautionary* demand for liquidity, and in linking it to the puzzle in question.

The main goal of this paper is to measure how much of the puzzle the liquidity need hypothesis can account for. Specifically, I answer the following two questions: (1) Can the need for liquidity 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. I explain in detail how the data analysis is performed, and show its robustness to assumptions that I make along the way.

Next, I develop a dynamic stochastic partial-equilibrium model of household portfolio choice, in order 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. 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

³In 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, in an analytically tractable way, 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 that is not analytically tractable - so I abstract from the reasons for why credit may

will show, the model has all the analytical implications important for addressing the credit card debt puzzle.

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 between 85 and 104% of the households who choose to revolve debt while holding money in the bank, and for a median such household, for 56-62 cents of every dollar it holds in liquid accounts. The ranges are given for two alternative calibrations, depending on the choice of the risk aversion parameter.

The main contributions of this paper are four. First, I carefully evaluate for the first time an intuitive answer to a standing puzzle, and find that it can account for the puzzle to a considerable degree. Debt puzzles of this nature have led the literature to challenge the standard consumption and saving models as incapable of explaining them. This paper can be seen partly as a response to this challenge. This does not mean that the alternative explanations involving opportunistic bankruptcy behavior or self-control considerations play no role, but I do find that a model that abstracts from such considerations goes a long way toward matching the facts. 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, this paper is the first, to my knowledge, to measure a new type of idiosyncratic risk - namely, expense risk that leads to precautionary liquidity demand. In light of the interest of the incomplete-market literature on idiosyncratic risk, this is an important step toward understanding the types of risk that households face. Fourth, in the process, I obtain new estimates of some parameters

be accepted in some markets, but not others, and simply take this fact as given.

of general interest - especially the elasticity of substitution between cash and credit goods - that have only been estimated in deterministic representative-agent models until now. My estimates suggest that idiosyncratic uncertainty affects these parameter estimates considerably. My estimate of the elasticity of substitution is significantly lower than what was previously obtained; my conclusion is that cash and credit goods tend to be complements rather than substitutes. This finding, explained in detail, is based on direct measurement in micro data of cash-good consumption, which has not been done previously, and the intuition for the result is clear-cut in the model.

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

The paper is structured as follows. In section 2, I characterize the credit card debt puzzle in the data, by studying the Survey of Consumer Finances and Consumer Expenditure Survey. Section 3 lays out the model and analyzes its properties. Section 4 discusses computation, and section 5 presents detailed information on the calibration strategy. Section 6 presents the results from the calibrated model, and section 7 discusses them. Section 8 concludes. Some details of the data are relegated to the appendix.

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. I am able to observe when balances are rolled over (distinguishing them from new purchases on the credit card), how large they are, and how often the household repays its balance in full. To define the puzzle group, as described below, I take

only the households that revolve debt *habitually*, that is, report repaying their balance off in full only sometimes or never. I am able to measure credit card debt precisely as well: I include only the amount of the balance due on the credit card left over after the last statement was paid - thus excluding, for example, recent purchases, or balances that were paid off. See appendix A.2 for more details on the subgroup selection.

The definition of liquid assets used in this paper includes checking accounts, savings accounts, and brokerage accounts (idle money in a brokerage house that is not being invested in stocks). This definition is motivated by what can be observed in the surveys. In particular, no data are collected on household cash holdings. I do not anticipate this to be a big problem, however: for the purpose of monthly transactions and larger purchases that I consider here, people are not likely to hold their liquidity in cash form, using their bank account actively instead to hold and transact with liquid assets via the use of checks, for example. The only households that are likely to be significantly affected by this data restriction are those in the borrower category, some of whom may not have bank accounts and are thus likely to be forced to hold cash. This, however, is not likely to impact my calculations regarding the puzzle itself, and may only bias the results slightly against me, in that some of the households I observe as borrowers may, in fact, have both debt and liquid assets that I cannot observe.

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 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 their debts are often repaid by their parents, as is well documented). Additional details of the surveys, the sample selection process, and the puzzle measurement methods are described in the data appendices A.1 and A.2.

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, 7 and 8 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.

2.1 Demographic and Asset Data

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 apparent that these groups are comparable in the two surveys, so that analyzing their consumption in the CEX and assets in the SCF is a

⁴To 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 procedure will make observed volatility of consumption of such goods appear smaller than it may be in reality, thus making the measurements I use err on the conservative side.

⁵I 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
<i>Puzzle size:</i>		<i>Percent distribution</i>		
	SCF	5%	27%	68%
	CEX	7%	29%	64%
<i>Interest rates:</i>				
	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.

valid exercise. 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 or no observed liquid assets, and the rest have no significant credit card debt. Notice that these numbers imply that of all habitual credit card debt revolvers, 80 to 84% have some liquid assets that they could in principle use to pay down their debt! The last three rows of the table give average interest rates that households report paying on their credit card debt versus national interest average rates on checking and savings accounts.⁶ 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

⁶Very few of these households report paying “teaser” zero interest rates on the accounts on which they are revolving balances.

Table 2: Demographics

	Borrow	Borrow & Save	Save	Share in Population
	<i>% of subgroup with characteristic</i>			
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.

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 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.⁷

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

⁷This is confirmed in formal probit analysis, not presented here.

Table 3: Income and Asset Holding

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

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; it is computed by the author in the SCF.

(and equal to it in the median).⁸ 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 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

⁸A concern may arise that these numbers could be collected at the beginning of the month, say, when the paycheck has just arrived into the account. As per the Federal Reserve Board of Governors, which collects the data, SCF interviews are conducted throughout the month, and these asset numbers, averaged across households, thus represent a monthly average on the account. The Federal Reserve Board declined to release interview dates. Thus, I will treat liquidity measurements as monthly averages, and will carefully treat liquidity in the model to match the same average concept.

Table 4: Home Ownership by Subgroup

	Borrow	Borrow & Save	Save	Share in Population
	<i>% of subgroup with characteristic</i>			
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 town-house/condo association) are excluded.

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 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

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

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 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.

2.2 Consumption Data

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. I will also use these data to study volatility of liquid consumption. The first step is to separate out the group of goods that can legitimately be viewed as payable by liquid assets. There are data limitations, so that this measurement cannot

be perfect; I limit the measurement procedure along several dimensions, in order to achieve as careful and as conservative a measure as I can. Although survey data on consumer payment method choice are scant, 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 the 2004 wave 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 upwards the perception of volatility. 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.⁹ 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. This is not true for food, alcohol and tobacco; 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. The reason to include these goods is that they appear to be predominantly cash-only goods in the data. Why that is is a question outside of the scope of this paper, and is the subject of a separate payments studies literature. This issue is discussed in more detail in the appendix. I will show the sensitivity of some calculations to the inclusion and exclusion of these goods below. In any event, food, alcohol and tobacco are a

⁹The 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
		<i>U.S. Dollars</i>		
Liquid assets:	<i>Mean</i>	227	7,338	17,435
	<i>Median</i>	200	3,000	3,200
Monthly cash-only good cons:	<i>Mean</i>	1,561	2,106	1,665
	<i>Median</i>	1,369	1,890	1,433
Liquid assets/cons:	<i>Mean</i>	<i>0.1</i>	<i>3.4</i>	<i>10.0</i>
	<i>Median</i>	<i>0.1</i>	<i>1.5</i>	<i>2.0</i>

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

minority of the cash-only expenditure category.

More generally, the cash good group selection may be seen as conservative. No durable or semi-durable goods, such as appliance purchases or even clothing purchases, are included. Thus, for example, the cash-good category excludes many situations that may be reflections of emergencies that require liquid payment - such as an emergency purchase of (or downpayment on) a durable to replace - rather than repair - a broken durable, such as a car or an appliance. Similarly, medical payments, which include co-pays or other out-of-pocket expenses, some of which are unpredictable and may require a liquid payment - are not included either; the decision here was driven by the fact that medical expenses may be payable by credit card. Instead, many of the categories that are included - such as food, insurance premia, etc., are paid on monthly basis and are predictable. Thus, in measuring the volatility of cash-good consumption, using a lot of the “smooth” good categories, while excluding many that may reflect other types of emergencies besides repairs, will tend to understate my measurements of the uncertainty that households face, against which they may hold liquid assets. However, to convince the reader of the fact that the approach is conservative, the analysis below, as pertaining to volatility of liquid consumption, describes robustness checks on the cash-good measurement.

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.

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

	Borrow	Borrow	Save	All
	& Save			
	<i>Avg. cond'l st. dev.(%)</i>			
<i>Liquid consumption (residual), ε_{it}</i>				
Benchmark	22.8	23.5	24.5	24.1
Excluding food	27.7	27.9	30.1	29.2
Excluding food and property taxes	29.8	30.7	32.8	32.0
<i>Unpredictable liquid consumption, η_{it}</i>				
Benchmark	21.7	23.0	23.8	23.5
Excluding food	26.8	27.5	29.4	28.7
Excluding food and property taxes	29.0	30.1	32.2	31.4

Source: CEX. Conditional standard deviation: population average of household-level 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, household observables, and household fixed effect, in a model with AR(1) disturbances. The residual is taken as the idiosyncratic component of consumption. Percent measure: st.dev. of the log-residual multiplied by 100. Benchmark: see appendix in paper.

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 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 a cause of this precautionary behavior. I study volatility of liquid consumption in order to gauge whether there is evidence in the data of uncertainty that households hedge against specifically with liquid assets. 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 estimating the following model:¹⁰

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

This is a fixed-effects model with AR(1) innovations. The vector \mathbf{X} includes, depending on specification, household observables, such as age (a cubic), education, marital status, race, earnings, family size, homeownership status, as well as seasonal effects (a set of month and year dummies). Several specifications including different sets of these observables all produced nearly identical results. u_i is the household fixed effect. The residual ε_{it} is the idiosyncratic component of liquid consumption, and it further consists of a persistent component and a transitory component (this becomes relevant later when I model the shock as having a possible persistent component). The results for volatility of consumption below present standard deviation of the total residual ε_{it} . I show, further, that even when the persistent component is taken out, the volatility of the residual η_{it} remains strong.

Table 7, rows 1 and 4, show volatility of consumption, by household subgroup, in the cash-only good category, measured as average monthly standard deviation of the residuals ε and η , multiplied by 100 to convert it to percent terms. First, volatility is fairly large, ranging between 22.8 and 24.5 % for the total residual measure, and 21.7-23.8% for the non-persistent component.

¹⁰One important distinction, of course, between measuring income versus consumption uncertainty is that the measures of income volatility are often translated directly into measures of income shocks, while consumption volatility reflects only the *endogenous response* of the household to its idiosyncratic shocks, which may, in fact, be larger than the response. E.g. after a car breakdown, one may choose to make fewer repairs than would be necessary to bring the car back to the previous condition, to conserve the expense. This mapping between volatility and uncertainty will be discussed further in the Calibration section.

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. Indeed, expenses that pertain to home maintenance tend to be the most volatile in my data, while expenses such as food are the least volatile. The volatility we observe in cash-only good consumption may be a reflection of unexpected, and possibly large, spending shocks; households try to insure against them by holding extra liquidity in the bank.

Concerns may arise that even controlling for household observables, the nature of expenses may be such that it adds “lumpiness”, which looks like volatile consumption, to the extent that they are on goods that are durable or semi-durable, or goods that are neither, but perfectly predictable and consumed at a lower frequency than monthly. With respect to the first, I emphasize that I remove all durable and semi-durable goods from the cash good category - appliances, furniture, cars, houses, and even apparel, are not included here. With respect to the latter, one example that commonly comes up is property taxes or auto insurance, which - unlike other predictable liquid expenses such as rent or mortgage - are often paid on less frequent basis. The evidence on this is mixed - some households choose to pay these on monthly basis too, as options for “financing” auto insurance, for example, are available. However, to check the robustness of my measures to various assumptions regarding specific goods, I have looked at many different permutations of cash-good group measurements, taking out from the benchmark measure described above food/alcohol/tobacco, insurance payments, property tax payments, and other predictable expenses. I present results for three such permutations: in addition to the benchmark group in the top row, I also present (a) the benchmark minus food/alcohol/tobacco, and (b) group (a) minus, in addition, property taxes. As is evident, the more I exclude such predictable expenses, the more volatility of the remaining group increases. This suggests that the goods that may be paid infrequently do not seem to take away from volatility - at least, the evidence is consistently suggesting that the benchmark cash-good category gives by far the most dampened measure of consumption volatility. This is the group that I will use in further analysis, in an attempt to stay conservative in the measurements.

Finally, to answer the concern that the CEX may be more suitable for analysis of quarterly

Table 8: Unpredictable Volatility of Average Household Cash-Good Consumption, **Quarterly** Data

	Borrow	Borrow & Save	Save	All
	<i>Avg. cond'l st. dev.(%)</i>			
<i>Liquid consumption (residual), ε_{it}</i>				
Benchmark	20.5	18.4	20.1	19.6
Excluding food	28.0	24.7	29.0	27.7
Excluding food and property taxes	29.0	26.7	30.8	29.6
<i>Unpredictable liquid consumption, η_{it}</i>				
Benchmark	20.5	18.4	20.0	19.6
Excluding food	27.6	24.5	28.8	27.5
Excluding food and property taxes	28.6	26.5	30.7	29.4

data, I also present the same table (table 8) for quarterly consumption. What becomes difficult here is that now I only have 3 observations per household, so there may be some noise in the measurements - this is evident, in particular, in that we lose the monotonic relationship between consumption volatility and subgroup (now borrowers have higher volatility than borrower-savers, which may simply reflect the relatively small sample size of borrowers with only 3 data points per household). However, this exercise reassures that the measures of volatility of consumption based on monthly data are robust. They decrease a little on quarterly basis, as would be expected given that we are now aggregating monthly measures into smoother quarterly ones, but the decrease is small.

To sum up, data suggest that the credit card debt puzzle is significant in magnitude, but it appears that tying it to liquidity demand - both for transactions purposes and for precautionary reasons - is reasonable and may account for some of the puzzle. 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.¹¹ 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.

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

¹¹The question of why credit cannot be used is beyond the scope of this paper, as it is a question of understanding the supply side of consumer credit. There are several approaches to it in the macro literature in similar contexts: one is to assume spatial separation between the earner and the shopper, as in Stokey-Lucas-style cash-credit good models; another is to assume that agents are anonymous, as in money search models following Kiyotaki and Wright (1989). See Telyukova and Wright (2008) for a related model of money and credit that addresses the issue in more detail in a similar context.

portfolio. Let us assume that $s_t \in S$ is a discrete Markov process, with $S = \{\underline{s}, s_2, \dots, \bar{s}\}$, $\underline{s} > 0$. 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 $z \in Z = \{\underline{z}, z_2, \dots, \bar{z}\}$, 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. Note that the sequential timing structure in this model is not crucial for the results. The model could have the two markets co-existing in time, for example; the important feature is only that a household makes its portfolio decisions for the entire period in the beginning of it - which is realistic, given that liquid spending opportunities can arrive continually and randomly throughout the month and until its very end, while additional income does not.

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, $\lim_{c_{jt} \rightarrow 0} u'_j(c_{jt}) = \infty$ and $\lim_{c_{jt} \rightarrow \infty} u'_j(c_{jt}) = 0$. 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 trans-

parent. 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.¹²

I formulate the household problem recursively.¹³ 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:

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

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} .¹⁴

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

$$\begin{aligned}
V_2(s_t, z_t, m_{2t}, b_{2t}) &= \max_{c_{2t}} z_t u_2(c_{2t}) + \beta \mathbb{E}_{s_{t+1}|s_t} V_1(s_{t+1}, z_t, m_{1,t+1}, b_{1,t+1}) \\
s.t. \quad c_{2t} &\leq \phi_{2t} m_{2t} \\
m_{1,t+1} &= m_{2t} - \frac{c_{2t}}{\phi_{2t}} \\
b_{1,t+1} &= b_{2t}
\end{aligned} \tag{3}$$

¹²The analytical results I emphasize here do not hinge in any way on the separability assumption - all the results would go through even in the non-separable utility case. On the other hand, empirically, the interaction of the two consumption goods may play a part in the magnitude of the results, and it seems natural to expect that it is non-trivial in reality; I will take up this issue in the computational part of the paper.

¹³The 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.

¹⁴In computation, I will allow for an interest spread: interest rate on borrowing, $b_{2t} > 0$, will be higher than 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.

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. Note also that in this subperiod, no portfolio rebalancing can take place if a household experiences a low shock and has money left over at the end of the period. This restriction is meant to capture the continual nature of the unpredictable expenses in the data: since in reality, expense shocks could hit continually throughout the month, experiencing a low expense shock at any point during the month would not cause the household to spend the remainder of its precautionary liquid balances to pay off debt before the month is over.

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, $\forall x_{1t}$:

$$-u'_1(c_{1t}(x_{1t}))\phi_{1t} + \mathbb{E}_{z_t|z_{t-1}} V_{2m}(s_t, z_t, m_{2t}(x_{1t}), b_{2t}(x_{1t})) = 0 \quad (4)$$

$$u'_1(c_{1t}(x_{1t})) + \mathbb{E}_{z_t|z_{t-1}} V_{2b}(s_t, z_t, m_{2t}(x_{1t}), b_{2t}(x_{1t})) - \lambda(x_{1t}) = 0 \quad (5)$$

The envelope conditions of the first subperiod are:

$$V_{1m}(x_{1t}) = \phi_{1t} u'_1(c_{1t}(x_{1t})) \quad (6)$$

$$V_{1b}(x_{1t}) = -(1+r_t)u'_1(c_{1t}(x_{1t})) \quad (7)$$

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:

$$z_t u'_2(c_{2t}(x_{2t})) - \mu(x_{2t}) - \frac{\beta}{\phi_{2t}} \mathbb{E}_{s_{t+1}|s_t} V_{1m}(s_{t+1}, z_t, m_{2t} - \frac{c_{2t}(x_{2t})}{\phi_{2t}}, b_{2t}) = 0. \quad (8)$$

The envelope conditions are, after substituting in (6) and (7),

$$V_{2m}(x_{2t}) = \beta \mathbb{E}_{s_{t+1}|s_t} \phi_{1,t+1} u'_1(c_{1,t+1}(x_{1,t+1})) + \phi_{2t} \mu(x_{2t}) \quad (9)$$

$$V_{2b}(x_{2t}) = -\beta \mathbb{E}_{s_{t+1}|s_t} (1 + r_{t+1}) u'_1(c_{1,t+1}(x_{1,t+1})). \quad (10)$$

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), $\forall x_{1t}, x_{2t}$:

$$\phi_{1t} u'_1(c_{1t}(x_{1t})) = \mathbb{E}_{z_t|z_{t-1}} \{ \beta \mathbb{E}_{s_{t+1}|s_t} \phi_{1,t+1} u'_1(c_{1,t+1}(x_{1,t+1})) + \phi_{2t} \mu(x_{2t}) \} \quad (11)$$

$$u'_1(c_{1t}(x_{1t})) - \lambda(x_{1t}) = \mathbb{E}_{z_t|z_{t-1}} \{ \beta \mathbb{E}_{s_{t+1}|s_t} (1 + r_{t+1}) u'_1(c_{1,t+1}(x_{1,t+1})) \} \quad (12)$$

$$z_t u'_2(c_{2t}(x_{2t})) = \beta \mathbb{E}_{s_{t+1}|s_t} \frac{\phi_{1,t+1}}{\phi_{2t}} u'_1(c_{1,t+1}(x_{1,t+1})) + \mu(x_{2t}) \quad (13)$$

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 (11) and (12), 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 (12) that

$$\beta \mathbb{E}_{z_t|z_{t-1}} \mathbb{E}_{s_{t+1}|s_t} u'_1(c_{1,t+1}) = \frac{u'_1(c_{1t})}{1+r}. \quad (14)$$

From this and (11), 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,

$$u'_1(c_{1t}) \left(\phi_1 - \frac{\phi_1}{1+r} \right) = \phi_2 \sum_{\{z_i: c_{2t}(z_i) = m_{2t}\}} \Gamma(z_{t-1}, z_i) \mu(\cdot). \quad (15)$$

Denote the right-hand side of (15) 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} (11) and (13). 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 (15) 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 ($r \rightarrow \infty$), 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, (15) 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 (11) and (12). 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

$$\begin{aligned} u'_1(c_{1t}) &= \mathbb{E}_{z_t|z_{t-1}} \left\{ \beta \mathbb{E}_{s_{t+1}|s_t} \frac{\phi_{1,t+1}}{\phi_{1t}} u'_1(c_{1,t+1}) + \frac{\phi_{2t}\mu_t}{\phi_{1t}} \right\} \\ u'_1(c_{1t}) &= \mathbb{E}_{z_t|z_{t-1}} \left\{ \beta \mathbb{E}_{s_{t+1}|s_t} (1 + r_{t+1}) u'_1(c_{1,t+1}) \right\} \end{aligned}$$

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

$$\mathbb{E}_{z_t|z_{t-1}} \left\{ \beta \mathbb{E}_{s_{t+1}|s_t} \frac{\phi_{1,t+1}}{\phi_{1t}} u'_1(c_{1,t+1}) \right\} < \mathbb{E}_{z_t|z_{t-1}} \left\{ \beta \mathbb{E}_{s_{t+1}|s_t} (1 + r_{t+1}) u'_1(c_{1,t+1}) \right\},$$

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 (15) 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 than they would spend if $z_t = \bar{z}$.¹⁵ 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.

Property 4. *Optimality of different borrowing and saving behavior.*

In every period, there exist three subgroups of the population:

¹⁵Note 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.

- 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 (15), and we get $\mu(\cdot) > 0 \forall z$, so from Euler equations (11) and (13),

$$\begin{aligned} \phi_{1t} u'_1(c_{1t}) &> \mathbb{E}_{z_t|z_{t-1}} \{ \beta \mathbb{E}_{s_{t+1}|s_t} \phi_{1,t+1} u'_1(c_{1,t+1}) \} \\ z_t u'_2(c_{2t}) &> \frac{1}{\phi_{2t}} \beta \mathbb{E}_{s_{t+1}|s_t} \phi_{1,t+1} u'_1(c_{1,t+1}). \end{aligned}$$

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

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 $u(c_{1t}, z_t c_{2t})$, 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 should not be ignored in calibration. 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 the 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:

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

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

$$\begin{aligned}
V_2(s_t, z_t, m_{2t}, b_{2t}, c_{1t}) &= \max_{c_{2t}} u(c_{1t}, z_t c_{2t}) + \beta \mathbb{E}_{s_{t+1}|s_t} V_1(s_{t+1}, z_t, a_{t+1}) & (17) \\
\text{s.t. } c_{2t} &\leq m_{2t} \\
a_{t+1} &= m_{2t} - c_{2t} - b_{2t}(1 + r_{t+1}),
\end{aligned}$$

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 completely robust to the choice of method.

5 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:

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

This choice satisfies all the necessary assumptions on the utility function listed above. The utility function gives three parameters to calibrate: α , ν and γ . β , the discount factor, is the fourth. The other parameters have to do with the shock processes on income and preferences, as well as prices. I calibrate the parameters of the income process outside the model, set γ to be in the standard range 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.

Some of the exogenously calibrated parameters are set as follows. I will compute the model twice, once with the risk aversion parameter $\gamma = 2$, and once with $\gamma = 3$, chosen in the lower and middle part of the standard range in the literature, in order to demonstrate some of the possible range of outcomes. The monthly interest rate on saving in nonliquid financial assets is set to match the annual rate of 4%, so that $r_s = 0.0033$. I set $r_b = 0.011$, which corresponds to the annual rate of 14%, the average interest rate paid on revolving credit card debt as reported by the debtors I observe in the SCF.

5.1 Income Process

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 4-state discrete Markov process as follows (see table 9). The income states are chosen to match the relative average earnings of white-collar workers (s_4), blue-collar and service sector workers (s_3), and the value of unemployment benefits (set at about 20% of the respective earnings) for white-collar and blue-collar workers (s_2 and 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. The reason I pose two unemployment states will be clear momentarily. 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. 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. This duration tends to be a bit longer for white-collar workers than for blue-collar workers, around 14.1 and 13 weeks respectively. Moreover, the shares of blue-collar and white-collar workers among the unemployed averaged 56% and 44% in 2001, respectively. I do not observe to which jobs the workers exit, so I assume that the vast majority of exits from unemployment is into same-sector jobs, allowing only small probabilities of transitioning to the other collar, probably understating the extent of such mobility a bit. This information together gives the probabilities of exiting white-collar and blue-collar unemployment and staying in it from month to month. There is no transition between the two unemployment states.

Associated with the transition matrix Γ_s is the invariant distribution of agents across the three income states. Denoting this distribution as $\{\gamma_1^*, \gamma_2^*, \gamma_3^*, \gamma_4^*\}$, I get two additional conditions: $\gamma_1^* + \gamma_2^*$ should equal the average monthly unemployment rate, which was 4.75% in 2001, and this is distributed between the two states according to the relative shares of blue- and white-collar workers among the unemployed, as detailed above. γ_3^* is the share of blue-collar workers among the *employed*, which was 43.5% in 2001. γ_4^* is the complement of the other three. These invariant shares give conditions to complete the first two columns of the transition matrix. Note that a white-collar worker cannot transition into blue-collar unemployment and vice versa. This clarifies why I have two unemployment states: this way, I avoid artificially high mobility of blue-collar workers into white-collar jobs via the state of unemployment, and vice versa.

What remains is the lower-right 2x2 matrix, which gives persistence of white-collar and blue-collar jobs and the transition between them. 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

Table 9: Earnings Process

	Parameters	Value
Earnings states	$\{s_1, s_2, s_3, s_4\}$	$\{ 0.14, 0.2, 0.589, 1.0 \}$
Transition matrix	$\Gamma(s_t, s_{t+1})$	$\begin{bmatrix} 0.667 & 0 & 0.332 & 0.001 \\ 0 & 0.694 & 0.002 & 0.304 \\ 0.021 & 0 & 0.972 & 0.007 \\ 0 & 0.012 & 0.006 & 0.982 \end{bmatrix}$
Invariant distribution in earnings	Γ_s^*	$\{ 0.026, 0.021, 0.414, 0.538 \}$

as I do not have specific information to that effect, I use this statistic here. This completes the third row of the matrix, and together with the invariant distribution, allows for the completion of the fourth row as well.

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 seven times the lowest income level in this calibration. On the other hand, 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.

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 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. Third, insofar as the top outliers in the data affect the difference between the mean and median puzzle households (as presented in the data section), I will particularly focus on the analysis of the model's results pertaining to the median household, rather than the mean, thus removing the results' dependence on the calibration of the distributions' top tail.

5.2 Idiosyncratic Preference Risk

The remaining parameters – the discount rate β , the parameters of the consumption aggregator α and ν , and the preference process parameters – are calibrated together, from within the model, using a minimum-distance estimator based on the simulated method of moments. In this subsection, I describe in some detail the calibration of preference shocks. In the subsection that follows, the remaining parameters are described.

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

The preference shock process is clearly not observed in the data, but the way households respond to these shocks is, through their liquid consumption. Thus, the preference shock process has to match properties of consumption of cash-only goods in the data, namely its persistence (measured as autocorrelation) and volatility (conditional standard deviation). In the calibration targets, conditional standard deviation is computed by subgroup, so in total, I get four calibration targets for the shock process.

Finally, as I described in the data section, the properties of liquid consumption – crucially the volatility of consumption – are sensitive to how the cash-good group is computed. Specifically, of all the measures that I have examined, the benchmark measure (the most inclusive) produces the smallest volatility of consumption; in addition, the exclusion of many expenses that may embody extreme events (like medical shocks and so on) is likely to constrain this measure further. In the estimation, I use this most conservative measure - the benchmark - so as to discipline the model in the best way possible. But one should keep in mind that this may be one reason for why the results of the model are a lower bound.

In order to convince the reader that normal disturbances are a reasonable assumption for the shocks, and also that the calibration does not overstate the tail shocks through such a




Figure 1: Liquid Log-Consumption Residuals, Benchmark Cash Good Measure, All Households

representation, figure 1 plots the consumption residual ε_{it} , together with a nonparametric kernel estimator of its density (thick red line) and the corresponding normal approximation (thin green line). It is very important to recall here that the right tail of the calibration (the highest preference shock(s)) is key for determining money holdings - it is the highest shock(s), which cause the money constraint of the household to bind, that drive the amount of precautionary demand for liquidity. Thus, the calibration of the right tail of the shock distribution is important for the results on liquidity demand. The figure is presented here only for the benchmark measure of cash goods, although the graph looks similar - just with wider tails and lower middle density - for other measures (excluding food, property taxes, etc.) mentioned in the data section. The graphs do not change if the residuals are considered by subgroups of households. Finally, for the transitory component η_{it} and for quarterly CEX, the graphs look very similar, and are not shown here for space considerations.

While this graph represents households' response to the expense shocks rather than the shocks themselves, the figure is instructive. As is apparent, normal distribution approximates the actual residuals fairly well. It understates the density of consumption at the mean (which will be corrected by the fact that I will match the autocorrelation of consumption as a targeted moment), but overstates consumption very near the mean (within one standard deviation), and understates the tails of the density. The tails of this consumption distribution are very wide,

especially the right tail, which extends to about 12 standard deviations away from the mean; we may speculate that the distribution of the actual shocks thus has even wider tails. I treat this right tail very conservatively in the calibration, however. First, in my discretization of the shocks by the Tauchen method, I restrict the five discrete states to fall only within *two* standard deviations of the mean. This approximates the AR(1) well, but is clearly an understatement of the actual tails of the distribution, particularly the right tail. Also, I only target up to the second moment of the residual distribution in the data, without attempting to match the tails. If I were to match the tails, the discrete shocks would have to include a much higher top realization than the one that I allow for currently. Once again, I impose this restriction to keep the calibration very conservative, with the goal of getting a disciplined, if understated, answer to the quantitative question; but it is important to keep in mind that this discipline will affect directly the implications of the model for the magnitude of liquidity demand.

5.3 Remaining Parameters and Mapping to Targets

Together with ρ_z and σ_z , I also estimate β , α and ν . For this purpose, in addition to the properties of liquid consumption mentioned above, I choose four more targets – mean cash-only good consumption relative to income for each of the subgroups, and the mean revolving debt-to-income ratio in the population. Although all eight calibration targets jointly determine all five parameters, there is also a fairly direct mapping between the targets and the parameters. The consumption-to-income ratios by subgroup help pin down α . β is pinned down by the debt-to-income ratio. The time series properties of liquid consumption help determine the shock process, as discussed above, and also help determine ν .

The estimation of the parameter ν – which determines the degree of substitutability between liquid and nonliquid consumption – deserves some attention. Previous estimates of the parameter of substitutability between liquid and nonliquid consumption come from deterministic cash-credit good models, in which the cash-in-advance constraint always binds, so that (aggregate) cash-good consumption equals (aggregate) money demand. A direct implication of this class of models is that the elasticity of substitution between cash and credit goods can be measured as the regression coefficient that gauges sensitivity of aggregate money demand to the gross nominal interest rate. This parameter turns out to be quite high in the data (0.79-0.84, see

e.g. Chari, Christiano, Kehoe 1991), which means that the estimated elasticity of substitution between cash and credit goods is also high (4.76-6) - that is, cash and credit goods appear to be substitutes. (This practice in existing literature could be misleading, since it maps liquidity demand in a model with no precautionary demand to money demand in the data, where there may be significant precautionary motive present, so that aggregate liquid consumption may not in fact equal aggregate money demand.) In my model, instead, cash-good consumption and money demand are distinct, due to the presence of idiosyncratic risk, where the cash-in-advance constraint is often not binding. As a result, the model no longer has a closed-form implication for the parameter ν . In the individual shock realizations where the constraint does not bind, liquid consumption is not sensitive to the nominal interest rate (see Telyukova and Visschers, 2009, for rigorous analysis of this fact). Thus, once idiosyncratic preference uncertainty is introduced into the model, aggregate liquid consumption becomes much less sensitive to nominal interest rates than money demand. Insofar as the elasticity of substitution can still be linked to this sensitivity, this reasoning leads to the conclusion that the elasticity will be much lower in this model than in the deterministic models used previously - this sensitivity is tied, in particular, to the probability that the money shock binds, which is around 6.7% in this calibration. I am able to construct a direct measure of household cash-good consumption from the micro data, and use the properties of this series, together with the interest spread, to pin down ν . Note that cross-sectional variation between the subgroups along the ratios of consumption to earnings and standard deviation of liquid consumption is linked to the preference shock process, but that process is restricted to be the same for all groups. What contributes to the cross-sectional variation is the interest spread r_b/r_s , where each subgroup faces a different tradeoff with respect to this ratio (borrower-savers, for instance, are faced with the spread, while borrowers and savers are not, in any given period). But the interest ratio is fixed exogenously and is time-invariant. Given this, the cross-sectional variation in targets that are computed by subgroup pins down ν , by distinguishing between the three subgroups of households in terms of the marginal opportunity cost of holding money that they face.

A word on the last target, the debt-to-income ratio. It is meant to gauge how well the model does in reproducing a dimension of the aggregate economy. Given that the focus of the paper is on the part of the population who have revolving debt, it seems like the natural statistic to

target. There are many ways to measure this target in the data, as there are many possible definitions of debt. For the purposes of this estimation, I choose aggregate revolving unsecured debt (the majority of which is credit card debt), computed in the SCF. In principle, debt could be measured as unsecured and secured revolving debt - allowing to include home equity lines of credit, for example - but as the SCF in 2001 does not show significant uptake of such lines of credit, and since my model has no role for housing or collateral more generally, I use only unsecured revolving debt as a measure.¹⁶

There is one other important detail in my calculation of the aggregate target ratio of debt 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 the aggregate target to the income calibration in the following way: I compute and target in estimation the debt ratio in the bottom 75% of the 2001 U.S. *income* distribution, rather than in the whole population. Note that doing so does not change the debt dispersion in the data all that significantly, so that the average debt-to-income ratio is only slightly above what it would be in the population as a whole. For example, revolving debt-to-income ratio in the whole population is about 5% - compare that to the measure for the bottom-75 % of the income distribution used here of 5.8 %.¹⁷

In sum, I estimate the five parameters within the model based on eight moments. For each set of parameters in the minimization process, the procedure solves the model, simulates a 502-month panel of 100,000 households, computes the moments from it, and compares them with the moments in the data. The complexity of the problem prevents me from using gradient-based minimization methods. Thus, for the minimization I use the simplex method of Nelder

¹⁶The 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.

¹⁷The fact that the aggregate target concerns 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, with each subgroup 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 would likely produce some increase in the optimal household liquidity holdings as well.

Table 10: Calibration Targets - Data and Model

	Target	Data	Model ($\gamma = 2$)	Model ($\gamma = 3$)
<i>Liquid consumption/earnings ratio</i> ^a : Borrowers		0.719	0.881	0.928
Borrowers & savers		0.618	0.621	0.602
Savers		0.629	0.630	0.641
<i>Autocorrelation (annual) of log liquid consumption</i> :		0.226	0.225	0.228
<i>Percent cond'l. st. dev. of log liquid consumption</i> ^a : Borrowers		0.228	0.138	0.111
Borrowers & savers		0.235	0.230	0.236
Savers		0.245	0.265	0.274
<i>Mean debt/income ratio</i> ^{b,c}		0.058	0.058	0.058

Notes: (a) The cash-only good series is the benchmark measure. (b) Debt is measured as revolving unsecured debt. (c) The moment is computed for the bottom 75% of the U.S. *income* distribution.

and Mead (1965), parallelized at parameter level as suggested by Lee and Wiswall (2007). The weighting matrix is the identity matrix in the first step, subsequently adjusted to correct for moments computed with highest variance (those moments that concern the borrower group, which is smallest in the data). Data covariances of the moments in question are not possible to compute in this exercise, since the moments come from two different data sets.

6 Results

6.1 Model Fit and Resulting Parameters

From this point forward, I will discuss the results of two calibrations, with the risk aversion parameter set alternately to 2 and 3. In each case, all the parameters are re-estimated to the same targets. In order to assess the fit of the calibrated model, table 10 presents the target moments in the data and the model. As discussed above, I have eight targets and five 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 to judge the fit of the model. The calibrated model fits most targets closely. The crucial moments concern the borrower-and-saver group and the debt-to-income ratio. The saver group is also important. 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 $\gamma = 3$ calibration creates more dispersion between the subgroups, so the estimated parameters match the borrower-saver group best, and diverge a bit more for the saver group. This concerns particularly the standard deviation of liquid consumption, which is higher in the model than it is in the data for the saver group. That is, in the model the savers are a bit more responsive to the expense shocks than they appear to be in the data. For the borrower-saver group, instead, the model with $\gamma = 2$ understates that volatility slightly, while it gets it exactly right with $\gamma = 3$.¹⁸

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 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 (with understated tails), will have trouble generating enough of the former. The latter group, those with the worst preference shock realizations, is 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 in the model, and their expenses are nearly

¹⁸In case there is a concern that the difference in results between the two calibrations is created by this difference in standard deviations, in the process of estimation, many parameter combinations were evaluated where the standard deviation was allowed to vary quite a bit; the results remain very robust to this variation, subject to a particular choice of γ .

Table 11: Calibration

	Parameter	Value 1	Value 2
Interest rates	r_s	0.0033 (annual $r_s = 0.04$)	0.0033
	r_b	0.0107 (annual $r_b = 0.14$)	0.0107
Risk aversion/IES	γ	2.0	3.0
Discount rate	β	0.9924 (annual $\beta_a = 0.9130$)	0.9913 ($\beta_a = 0.9009$)
Consumption aggregator parameters	α	0.6270	0.6238
	ν	-1.5176	-1.5204
Preference shock process: AR(1) with discretization	ρ_z	0.4727	0.4724
	σ_z	0.5670	0.5175

constant, so volatility is understated relative to the data; their liquid expense-to-income ratio is likely to be overpredicted as well.

The final target, the average debt-to-annual-income ratio, is matched perfectly in both calibrations.

Table 11 presents the resulting parameterization. The discount factor in the two calibrations is equivalent to 0.9-0.91 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 0.62-0.63, which once again confirms that they are an important part of a household's expenditures - not surprising given that payments to mortgages and household repairs fall 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 complements rather than substitutes, with the elasticity of substitution of 0.4. 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 the same as my estimate and this is encouraging. Their ν , as I mentioned above, tends to be on the order of 0.79-0.84, producing the elasticity of substitution on the order of 4.5 to 6. I already discussed in the calibration section the reasons for why it can be reasonably expected that a model where the cash-in-advance

constraint does not bind for the majority of households will have this elasticity of substitution much lower: once aggregate liquid consumption is disconnected from aggregate money demand, the sensitivity of liquid consumption to the nominal interest rate is apt to drop significantly relative to previous estimates, leading to a drop in the elasticity parameter. My estimation results, using micro data, confirm this.

Finally, the estimates of the preference process are of importance, since this study presents a new (and to my knowledge, first) effort to quantify unobservable idiosyncratic uncertainty to preferences from microdata specific to liquidity needs. The estimated monthly AR(1) parameter on $\log(z)$ is around 0.47. 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 0.47 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 0.51-0.57, with the lower parameter needed in the model with the higher risk aversion coefficient. 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. With higher curvature, the precautionary motive is enhanced, hence lower variation of the underlying shock is needed to accomplish the observed variation of consumption in the data. 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 relevant data dimensions; the moments that are not as well matched in 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.

6.2 The Credit Card Debt Puzzle

As mentioned before, I left the magnitudes of interest for answering the central question of this paper untargeted in calibration. 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

Table 12: Results - Subgroup Size (Percent)

	Data	Model ($\gamma = 2$)	Model ($\gamma = 3$)
Borrowers	5.2	2.5	1.8
Borrowers & savers	27.1	28.3	23.0
Savers	67.7	69.3	75.3

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 12 gives the size of the three subgroups in the data and the model. In the model, the size of the borrower-and-saver group is between 23 and 28% of the population, while in the data, it is 27%. Thus, the model accounts for between 85% and 104% of the puzzle group, overpredicting its size very slightly for the calibration with the lower risk aversion parameter. Further, the model slightly overstates the size of the saver group, at the cost of understating the size of the borrower group. The reasons for why the borrower group is underpredicted was already discussed above. Primarily, the model's borrower group consists only of those who are constrained at the end of the month, while in the data, there may be some households who have very few liquid assets throughout the month and year; these households are not captured by the model, since it is never optimal to hold zero liquidity in this model.

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 interviewed throughout the month. This cannot apply to the borrower group, however: it is not likely that *all* households in this group truly never hold liquid assets during the month, given their average liquid spending documented above, so many of these may 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 to map to the observed amount in the data; it should be kept in mind, however, that the data money holdings for the borrower group might be more aptly compared to end-month liquidity

Table 13: Results - Liquid-Asset-to-Income Ratio, **Monthly Average**, Median Household

	Data	Model ($\gamma = 2$)	Model/Data	Model ($\gamma = 3$)	Model/Data
Borrowers	0.10	0.41	4.10	0.43	4.30
Borrowers & savers	0.79	0.44	0.56	0.49	0.62
Savers	0.88	0.55	0.62	0.59	0.67

holdings of this group, which is equal to 0.

In table 13, I present average monthly liquid asset holdings relative to income for a median household in each 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 strongly biased from the outset toward medians, rather than means. The “Model/Data” columns translate 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 between 56 and 62 cents of every dollar held by the median borrower-saver household in the data, depending on the calibration. This range is 62-67 cents for the median household in the saver group. For the borrowers, the model generates over 400% of the money holdings in the data, based on the average monthly liquidity holdings. As discussed above, it may be more appropriate for the borrower group to compare the data number to the end of the month liquid holdings, which is 0 in the model. ¹⁹

7 Discussion of the Results

The results presented above lead to the conclusion that demand for liquidity - for predetermined expenses and precautionary reasons - is a key factor in accounting for the credit card debt puzzle.

In light of this analysis, the puzzle appears much less puzzling.

One part of analysis presented here cannot be made entirely criticism-proof: namely, the

¹⁹For the sake of completeness, the model’s performance for the *average* household in the borrower-saver group in the data is as follows. I emphasize that the model is not calibrated to speak to the upper tail and hence averages; the liquidity/earnings ratio in the data for the borrower-saver group is very different for the median (0.8) versus mean (1.7) household, while in the model they are much closer to each other. For the borrower-saver group and monthly average liquidity holdings, the model generates between 32 and 34 cents of the dollar held by the average borrower-saver household in the data. Under a different, more disperse, calibration of income, for example, the model would do a much better job for the average data household as well.

measurement of liquid consumption, and particularly, of its volatility. The data are limited as to the information they give along this dimension, and assumptions have to be made along the way in order to complete the measurement. However, I have tried to be explicit as to the assumptions that I made, and to argue for these assumptions based on robustness of the measures. I showed that the group that I use as cash-only consumption goods are a conservative group, especially when it comes to measuring volatility of this consumption and mapping it to expense shocks. On the one hand, many of the goods included in this group are “smooth”, in the sense that households consume relatively constant amounts of these goods every month, in a pre-determined way. As I showed in the data section, as I begin to omit goods that are of the more predictable nature - even those that might not be paid on a monthly basis, such as property taxes and insurance payments - the volatility of the measured residual goes up, not down, and the increase is very significant. On the other hand, I also omit many situations from this group that may represent true emergencies that may be paid by liquid assets, such as durable purchases that happen as a way to replace, rather than repair, an unexpectedly broken unit (like car or appliance), medical expenses, etc. Finally, the discretization of shocks that I use to represent the underlying preference uncertainty process is limited to realizations of only two standard deviations away from the mean, and I target only up to the second moment of the liquid consumption distribution, rather than matching the wide tails, which I presented in the calibration section. My conclusion from these points is that the expense uncertainty that households may face is possibly much higher than the measurements that I used. My choices were based on giving the model strict discipline; to the extent that the shocks in the data are possibly a lot higher, optimal precautionary liquidity demand will also be much higher.

A related concern that sometimes arises is that even if an expense shock comes along of the nature described here, the expense itself can be postponed. I emphasize that I calibrate the expense shocks to micro data. If households choose to postpone some of the expenses that arise, this will be reflected in lower volatility of liquid consumption and possibly higher persistence. Since I measure both in the data, and use both as targets for my model, I do not believe this to be a concern for the exercise.

There are several other ways in which the current results on liquidity demand by the borrower-saver group (and others) may be seen as a lower bound. For example, money de-

mand 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.²⁰ 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.

Finally, and in my view 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 in the face of income risk. For example, if one should lose one's job and paycheck, one still needs to pay the mortgage each month. 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, as of 2001, was a cash advance from a credit card, which would incur a withdrawal fee of several percent of the amount withdrawn, and *in addition*, would incur an interest rate much higher than that

²⁰For 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 56-62% of the total median balance, which translates to \$1,680-1,860, with the minimum balance requirement, I would be capturing 84-93% of "operational" liquidity balance of such a household.

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 provide a precautionary motive for holding liquidity 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 the expansion of an already large state space. Thus, it is beyond the scope of this paper to implement such an extension in full. However, in computing two-period versions of the benchmark and the extended models, I found that income uncertainty adds to liquidity demand significantly: in the example, liquidity demand increased by 30-50% relative to the benchmark case, depending on the exact asset specification. Although these results should be taken as only indicative absent a full recalibration, the numbers would translate to matching between 72 and 85% of the liquidity holdings of the median “puzzle” household. These results are highly suggestive that the liquidity hypothesis, although already a strong candidate for the explanation of the puzzle, may do even better when all the sources of uncertainty and money demand (in the case of minimum balances) are taken into account. I reiterate that this does not mean that previously proposed explanations do not play a role in accounting for the puzzle - for some households they may even be the more accurate explanations. All I endeavored to show is that a model that abstracts from behavioral, self-control, or strategic bankruptcy considerations, when carefully calibrated, goes a long way toward accounting for the facts of the puzzle in the data.

While I do not want to overstate the case here, in that the discussion of these additional

channels of precautionary liquidity demand would have to be more rigorous and nuanced (many tradeoffs need to be evaluated), I think it is reasonable to treat the results in this exercise as a potentially lower bound on how large precautionary demand for liquidity may be, and on their link to the credit card debt puzzle.

8 Conclusion

This paper presents the first available rigorous examination of liquidity demand as an explanation for the credit card debt puzzle. I examine the hypothesis that there is a significant share of household expenditures each month that cannot be paid by credit card, so that households need to keep liquidity in the bank at all times to pay for these expenditures. The data suggest 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 measurements of magnitudes of unobservable idiosyncratic uncertainty and the elasticity of substitution between cash and credit goods in micro data. Further, I find that, depending on the calibration, the hypothesis successfully accounts for all or nearly all (85%-104%) of the households who revolve debt while having money in the bank, and for a median such household, it accounts for 56-62 cents of every dollar held in liquid assets. There is a set of compelling reasons to view these results as a lower bound. Thus, even though there are likely households for which alternative explanations along

the lines of time inconsistency or strategic bankruptcy behavior are valid, or even dominant, it should be apparent that liquidity demand - including precautionary demand for liquidity - is a factor that contributes significantly toward accounting for the puzzle.

References

- [1] Aiyagari, S. R. and M. Gertler. 1991. "Asset Returns with Transactions Costs and Uninsurable Individual Risk." *Journal of Monetary Economics*, 27(3), 311-331.
- [2] American Bankers Association / Dove Consulting. "2004 Study of Consumer Payment Preferences".
- [3] Athreya, K. 2004. "Fresh Start or Head Start? Uniform Bankruptcy Exemptions and Welfare." Federal Reserve Bank of Richmond.
- [4] Ausubel, L. 1991. "The Failure of Competition in the Credit Card Market." *The American Economic Review*, 81, 50-81.
- [5] Bertaut, C. and M. Haliassos, 2002. "Debt Revolvers for Self Control." University of Cyprus.
- [6] Baumol, W.J. 1952. "The Transactions Demand for Cash: an Inventory Theoretic Approach." *Quarterly Journal of Economics*, 66, 545-556.
- [7] Brito, D. and P. Hartley, 1995. "Consumer Rationality and Credit Cards". *Journal of Political Economy*, 103(2), 400-433.
- [8] Browning, M. and A. Lusardi, 1996. "Household Saving: Micro Theories and Micro Facts." *Journal of Economic Literature*, 34, 1797-1855.
- [9] Bruce, L. 2002. "A Graphic Look at Deposit Rates, 1998-2002." bankrate.com.
- [10] 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.
- [11] Cargill, T. and J. Wendel, 1996. "Bank Credit Cards: Consumer Irrationality versus Market Forces?" *The Journal of Consumer Affairs*, 30, 373-389.
- [12] 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.

- [13] Chari, V. V., L. Christiano and P. J. Kehoe, 1991. "Optimal Fiscal and Monetary Policy: Some Recent Results". *Journal of Money, Credit and Banking*, 23, 519-539.
- [14] Chatterjee, S., D. Corbae, M. Nakajima, J.V. Rios-Rull, 2007. "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default." Forthcoming, *Econometrica*.
- [15] Duca, J. V., and W. C. Whitesell, 1995. "Credit Cards and Money Demand: A Cross-Sectional Study." *Journal of Money, Credit and Banking*, 27, 604-623.
- [16] Durkin, T., 2000. "Credit Cards: Use and Consumer Attitudes, 1970-2000." *Federal Reserve Bulletin*, September.
- [17] Durkin, T., 2002. "Consumers and Credit Disclosures: Credit Cards and Credit Insurance." *Federal Reserve Bulletin*, April.
- [18] Eisefeldt, A., 2003. "Smoothing with Liquid and Illiquid Assets". Kellogg School of Management.
- [19] Erosa, A. and G. Ventura, 2002. "On Inflation as a Regressive Consumption Tax." *Journal of Monetary Economics*, May.
- [20] Evans, D. and R. Schmalensee, 1999. *Paying with Plastic: The Digital Revolution in Buying and Selling*. MIT Press.
- [21] Federal Reserve Board. Statistical Release G.19, Consumer Credit. Various issues.
- [22] Fernández-Villaverde, J. and D. Krueger, 2004. "Consumption over the Life Cycle: Facts from the Consumer Expenditure Survey." University of Pennsylvania.
- [23] Gabriel, P., 2003. "An examination of occupational mobility of full-time workers". *Monthly Labor Review*, September.
- [24] Gervais, M. and P. Klein. 2006. "Measuring Consumption Smoothing in CEX Data". Mimeo, University of Western Ontario.
- [25] Gourieroux, C. and A. Monfort, 1996. *Simulation-Based Econometric Methods*, Oxford University Press.

- [26] 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.
- [27] Gross, D. and N. Souleles, 2001. "An Empirical Analysis of Personal Bankruptcy and Delinquency." Wharton School, University of Pennsylvania.
- [28] Haliassos, M. and M. Reiter, 2003. "Credit Card Debt Puzzles." University of Cyprus.
- [29] Hayashi, F. and E. Klee, 2003. "Technology Adoption and Consumer Payments: Evidence from Survey Data." *Review of Network Economics*, 2, 175-190.
- [30] Hopenhayn, H. and E. Prescott. 1992. "Stochastic Monotonicity and Stationary Distributions for Dynamic Economies," *Econometrica*, 60(6), 1387-1406.
- [31] Jasper, M. 2000. *Legal Almanac: Credit Cards and the Law*. Oceana Publications, Inc.
- [32] Johnson, K.W., 2004. "Convenience or Necessity? Understanding the Recent Rise in Credit Card Debt." *Finance and Economics Discussion Series*, Federal Reserve Board.
- [33] 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.
- [34] Lacker, J. and S. Schreft. 1996. "Money and credit as means of payment". *Journal of Monetary Economics*, 38(1), 3-23.
- [35] Laibson, D., A. Repetto and J. Tobacman, 2002. "A Debt Puzzle." Harvard University.
- [36] Lee, D. and M. Wiswald, 2007. "A Parallel Implementation of the Simplex Function Minimization Routine". Forthcoming, *Computational Economics*.
- [37] 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.
- [38] Lucas, R., E. Prescott and N. Stokey, 1989. "Recursive Methods in Economic Dynamics". Harvard University Press.

- [39] Meisenheimer, J. and R. Ilg, 2000. Looking for a 'Better Job': The Job Search Activity of the Employed. *Monthly Labor Review*, September.
- [40] Nelder, J. A., and R. Mead. 1965. "A Simplex Method for Function Minimization". *The Computer Journal*.
- [41] 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.
- [42] 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.
- [43] Schreft, S. 1992. "Transaction Costs and the Use of Cash and Credit". *Economic Theory*, 2(2), 283-96.
- [44] Shimer, R. 2005. "The Cyclical Behavior of Equilibrium Employment and Vacancies." *American Economic Review*, 95(1), 25-49.
- [45] Stavins, J., 2000. "Credit Card Borrowing, Delinquency, and Personal Bankruptcy." *New England Economic Review*, July/August.
- [46] Storesletten, K., C. I. Telmer and A. Yaron, 2004a. "Consumption and Risk Sharing Over the Life Cycle". *Journal of Monetary Economics*, 51, 609-633.
- [47] Storesletten, K., C. I. Telmer and A. Yaron, 2004b. "Cyclical Dynamics in Idiosyncratic Labor-Market Risk". *Journal of Political Economy*, 112, 695-717.
- [48] Tauchen, G., 1986. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions". *Economics Letters*, 20, 177-181.
- [49] Telyukova, I.A. and L. Visschers, 2009. "Precautionary Demand for Money in a Monetary Business Cycle Model". Mimeo, UCSD.
- [50] Telyukova, I. A. and R. Wright, 2008. "A Model of Money and Credit, with Application to the Credit Card Debt Puzzle." *Review of Economic Studies*, 75(2), 629-647.
- [51] U.S. Census Bureau, 2003. *Statistical Abstract of the United States*.

- [52] 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.
- [53] 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, child care, 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.