How important is variability in consumer credit limits?

Scott L. Fulford[*]

First draft: June 2010
This draft: May 2014

Abstract

Credit limit variability is a crucial aspect of the consumption, savings, and debt decisions of households in the United States. Using a large panel this paper first demonstrates that individuals gain and lose access to credit frequently and often have their credit limits reduced unexpectedly. Credit limit volatility is larger than most estimates of income volatility and varies over the business cycle. While typical models of intertemporal consumption fix the credit limit, I introduce a model with variable credit limits. Variable credit limits create a reason for households to hold both high interest debts and low interest savings at the same time since the savings act as insurance. Simulating the model using the estimates of credit limit volatility, I show that it explains all of the credit card puzzle: why around a third of households in the United States hold both debt and liquid savings at the same time. The approach also offers an important new channel through which financial system uncertainty can affect household decisions.

JEL classification: E21, D91, D14
Keywords: Credit card puzzle, intertemporal consumption, precaution, credit limits, household finance

1 Introduction

Three quarters of households in the United States hold a credit card. For the majority of these households the available credit is larger than their checking and savings accounts combined. In the short term, credit is the primary determinant of how much households in the United States can spend, not income or savings.

Yet this credit is extremely volatile. Using a large panel from the credit reporting agency Equifax, I show that Americans frequently lose access to credit and that credit limits on credit cards often increase and decrease. Credit limit volatility is substantially larger than standard estimates of income volatility. Moreover, in any given quarter 2.7 percent of individuals lose access to credit card borrowing entirely. Aggregate movements in credit limits and debt are large as well, as noted by Ludvigson (1999). For example, in the last quarter of 2008 one fifth of credit card holders had at least one credit card account closed, and overall credit limits fell by a quarter from 2008 to 2009.

Does this volatility affect household decisions? To answer this question I build a model of consumer decisions in which the credit limit can vary unexpectedly. While the typical model of consumption over time assumes that credit limits are fixed, the relatively minor change of allowing limits to vary affects consumer decisions substantially.

The key modeling insight in this paper is to consider the ability to borrow as part of a portfolio of assets that households use to smooth consumption. The value of this asset depends on whether it will be available when needed. Savings in a bank or under a mattress can be used in an emergency. The ability to borrow is much less certain. A lost wallet, identity theft, or the uncertainties of dealing with a large company that can alter the line of credit at any time mean that it is possible to lose access to credit unexpectedly. If consumers face uninsurable income or consumption shocks— times when their income goes down significantly or their need to spend goes up—the risk of being unable to borrow when times are bad creates a reason to hold low-interest savings while carrying expensive debt as a precaution against being unable to borrow. The savings act as insurance so that even in the worst case when the consumer cannot borrow, she can still consume.[1]

---

[1]See Carroll (2001) for a view of the development of and evidence behind precautionary models. A few of the

This modeling approach has two important implications. First, it explains the so-called "credit card puzzle." Around a third of households in the United States both carry large revolving credit card debt at high rates of interest and hold liquid assets which pay low or no interest. As suggested by Gross and Souleles (2002), such behavior seems very odd: Why not pay off the debt, which costs 14 percent per year, with the savings, which yield close to zero? The costs of not doing so are meaningfully large; for those that carry both credit card debt and liquid savings, the interest on the debt that could be paid off devours around 0.6 percent of monthly income.[2] Survey evidence suggests that many households are intentionally borrowing and saving at the same time (see section 5). Moreover, holding both debt and savings at the same time seems to be common among the poor all over the world (Collins et al. 2009, p. 49).

Using the estimates of credit limit volatility from the panel, I calibrate the model so that households hold as much credit card debt as they do in surveys. Doing so requires households to be fairly impatient, since credit card debt is expensive and American households hold a lot of it. I then examine how the households divide up their portfolio of credit and savings. The resulting distribution closely matches the joint distribution of savings and debt in the data. In particular, about the same fraction of households hold debt and savings at the same time in the model as in the data. A critical component of the calibration is heterogeneity in risk aversion. A large fraction of U.S. households hold approximately zero savings while holding substantial debt—a decision the model suggests is quite risky—and so to obtain the full range of behavior it is necessary to have some people willing to take those risks.

The model offers some simple predictions of who will borrow and save at the same time. Those who are more worried about the future tend to save more and so are less likely to need to borrow. If they have enough bad shocks that they want to borrow, however, they are much more likely to keep

key papers in a large buffer stock and precautionary savings literature are: Schechtman and Escudero (1977), Deaton (1991), and Carroll (1997). Fulford (2013a) examines the short and long-term consequences of a permanent change in the liquidity constraint. Ludvigson (1999) examines stochastic credit limits, although explicitly restricts the analysis to exclude borrowing and saving at the same time, suggesting that to do so is a "challenging direction" for research (p. 436).

[2]These calculations are made using the Federal Reserve 2007 Survey of Consumer Finances. See the section on the SCF for more detail. Zinman (2007) examines the distribution of the costs of the credit card puzzle or those who "borrow high and lend low," and finds that few households pay more than $10 per month.

some savings for precautionary purposes. Comparing across households using multiple rounds of the Survey of Consumer Finances, I find support for these predictions. I also examine some of the other explanations for the puzzle (Bertaut, Haliassos, and Reiter 2009; Laibson, Repetto, and Tobacman 2000; Lehnert and Maki 2002; Telyukova and Wright 2008) in light of the cross-sectional and time-series evidence.

The second implication of allowing credit limits to vary is what happens during times of financial crisis. Since credit makes up approximately two thirds of the resources available for immediate consumption of the average household, it is the single most important factor in determining the short-term budget constraint of households. Within a precautionary model such as the one introduced in this paper, credit is wealth in the short term since it can help smooth consumption. An increase in the chance of losing credit makes this wealth less valuable while a decrease in the credit limit reduces this wealth directly.

Financial uncertainty and credit reductions can thus directly affect consumer decisions even without affecting income or employment. Following the financial crisis of 2008, commercial banks reduced credit card limits overall by more than a quarter. During the worst two quarters, close to 20 percent of individuals lost a credit card account, including those with the best credit scores (see also the estimates in Jambulapati and Stavins (2013)). My estimates suggest that, after adjusting for age and risk, credit card limits declined nearly continuously from the end of 2008 through 2013. Using the calibrated model, I examine the effect of an increase in the probability of losing access to credit such as occurred in 2008–2009. With credit less valuable for consumption smoothing, consumers almost immediately "rebalance" their portfolios by increasing both savings and debt. Consumption declines slightly to accommodate the higher interest payments, but a pure increase in volatility has a relatively small aggregate effect. Coupled with a decline for all consumers in available credit, the effects are more dramatic: an immediate fall in consumption of around 2 percent following the reverse of the path shown by Fulford (2013a) for an increase in the credit limit. The fall in available credit reduces debt and so mostly cancels out the increase in debt from rebalancing the portfolio. While this paper focuses on consumers, doing so likely understates the impact of credit

volatility, since firms use short-term credit as well, and, faced with adjustment costs, may behave in similar ways.[3]

A central contribution of this paper is to estimate the credit volatility faced by individuals. There has been a great deal of work estimating and considering how income volatility affects consumer decisions, but this paper is the first to estimate credit volatility at the individual level and show how it varies over time. The key findings are that the short-term variance in credit limits is larger than most estimates of income variance, the long-term variance is much larger than the comparable long-term income variance, and individuals are much more likely to lose access to credit than to have no income. The second contribution is to link this volatility to household decision-making and show that it actually matters. Credit is vitally important to American households; it is also extremely variable. This variability affects both the amount of debt consumers hold and their savings, and explains one of the "puzzles" of household finance.

2 How variable are consumer credit limits?

2.1 Descriptive evidence

To examine credit limit volatility, I rely on two primary sources. First, I use a panel from 1999–2013 that contains a 0.1 percent sample of all individuals with a credit report at the credit-reporting agency Equifax, prepared by the Federal Reserve Bank of New York to form its Consumer Credit Panel.[4] Since the primary purpose for collecting these data is to help lenders price and monitor

---

[3]Firms use short-term credit as working capital to smooth inventory, meet unexpected demand, and make payroll, and the possibility of losing credit may cause them to adjust in similar ways. The buffer-stock theory of inventory is closely related to the consumer's problem (Deaton 1991; Deaton and Laroque 1992). Credit constraints can cause fluctuations in inventories (Kashyap, Lamont, and Stein 1994). Moreover, small businesses often use the personal credit of their owners as well (Zinman 2009) as is true across the world (Collins et al. 2009).

[4]The Equifax dataset is a random sample of consumers and gives the entire available history once an individual is selected. About 40 percent of the sample are in the data for all quarters from 1999–2013, but new accounts are continuously added so that the sample stays reasonably balanced between young and old. I limit the analysis to individual accounts that have an age listed and an open credit card account at any time over the entire sample. This includes many accounts that have an open credit card only briefly and otherwise have no changes in credit limits. 73 percent of accounts that have an age listed have an open credit card account at some point. Of the remaining 26 percent of accounts without an age, only 14 percent of them have an open credit card account at any time, so most would be dropped in any case. The accounts without ages appear on average for only about six quarters compared with 41 quarters for those with an age. This suggests that these accounts are not for actual people. They may represent

credit to individuals, it contains extensive information on credit limits and payment history. It has limited information on individual characteristics and beliefs, however, so I supplement the analysis with data from the Consumer Finance Monthly (CFM), a nationally representative survey conducted by the Center for Human Resources at The Ohio State University. The CFM asks a much richer set of questions about credit history and individual characteristics but is small and is not a panel. Jambulapati and Stavins (2013) examine both of these datasets to study whether the passage of legislation regulating the credit card industry affected the availability of consumer credit.

Table 1 shows the mean and median open credit card accounts as well as the percentage of gains and losses in the number of accounts for all account holders, by credit rating, and by age. Between 1999 and 2013 the mean number of open accounts was 2.2 and the median was 2. This includes many people who do not currently have an open credit card account although they may still have debt from an account that they can no longer use. Not surprisingly, the number of open accounts varies substantially with credit risk score.[5] The median individual with poor credit does not currently have an open credit card account although she had one at some point during the period. Losing an account is a big change in available credit for the large majority of individuals and often means a complete loss of the ability to borrow with a credit card. As individuals age and establish credit histories, the number of open accounts tends to increase.

In any given quarter 7.75 percent of individuals who have ever had an open credit card account between 1999 and 2013 lose an account, and 7.8 percent gain an account. The nearly identical gains and losses of accounts reflect the stable saturation of the credit card market. As figure 1 shows, the fraction of households with a credit card has been largely stable since 1992 even as the use of all forms of electronic payments has increased rapidly.

Account closures are often initiated by the individual, but a large portion come from the credit

business accounts of some sort, or are entry mistakes.

[5]The score is the proprietary Equifax score which is similar to but not necessarily the same as the FICO score. While the score is continuous, I use what appear to be standard industry breakdowns for credit scores: Excellent (52 percent) is 720 and above, Good (15 percent) is 680–720, Fair (9 percent) is 640–680, and Poor (24 percent) is below 640. The percentages are of individuals who have ever had an open credit card account between 1999 and 2013 and so these exclude both those who never used credit and those who cannot obtain it.

issuer. In the Consumer Finance Monthly,14.9 percent of consumers say they closed a credit card account during the previous year, and 5.4 percent report their bank closed an account.[6] Overall these yearly totals in account closure conform closely to the quarterly changes in the Equifax data. The survey evidence suggests that about three quarters of account closures are driven by the individual although not all of these closures are entirely voluntary. 13.3 percent report that they closed the account because the bank changed the terms of the card.[7] Including the proportion who closed an account because the terms changed suggests that the involuntary account closure rate is somewhere between 5.4 percent and 7.4 percent per year.

Banks can and do change credit limits without closing accounts. Changing the credit limit does not affect the repayment schedule or whether the account holder still has to pay back any accumulated debt (see appendix A). Credit limit changes in the Equifax data are less exact than the changes in open accounts for two reasons. The first is that the credit limits for accounts without an update over the previous quarter are not included in the data provided by Equifax. So the reported credit limits can decline if an individual is not using the card or the bank does not report any activity even though the line is still open and the individual could borrow on it. Since the account is still listed as open, it is possible to make adjustments for this reporting problem. I deal with this reporting problem systematically in the estimation section, but for the descriptive statistics I calculate the overall limit by multiplying the number of open accounts by the average limit from reported accounts. The second problem with credit limits is that for some accounts only the maximum amount ever borrowed is reported. That is because some cards do not have an explicitly set limit and because some lenders may not report the limit. There is no way to separate out such accounts in the data. This problem does not affect the estimates of the probability of gaining or losing credit and has an ambiguous affect on credit limit volatility, since once an upper bound is established, it does not change.

---

[6]The questions on account changes were asked over 2010 and 2011. I include all months together for a total sample of 3,098 who answered the questions, weighted to be nationally representative.

[7]Of individuals who closed an account: 27.8 percent report they did so because they did not use that credit card, 9.9 percent because they charged too much on the card or were worried they would, 13.3 percent because the bank or card issuer changed the terms of that card, and 49 percent closed it for some other reason.

Table 1 reports the percentage of accounts with negative, positive, and zero changes in credit limits. In any given quarter, two thirds of accounts had no change in credit limit. Reflecting the overall increase in available credit, 20 percent of accounts in any given quarter had an increase in limit, while 12 percent had a decrease. The mean decrease is a fall of 36 percent in the limit. The mean increase, conditional on the credit limit increasing by 200 percent or less, is 29 percent, but this figure is highly sensitive to how much of the the very long right tail is included. Those with poor credit scores were more likely to have had no change in credit because many were at zero already and they are less likely to have increases in limits. When they did have a decline in credit it tended to be large. The variance of the changes in credit limits was larger for young people than for older people. The young were more likely to have credit increases, but when they had a reduction in credit it was much larger.

Credit limit volatility varies over time. Figure 2 shows the mean number of open accounts, the fraction of individuals with an increase in accounts, and the fraction with a decrease over time. Figure 3 shows the fraction of all accounts with a positive limit, the fraction of accounts that gain access to credit, and the fraction that lose access each quarter. In each figure, the bold line shows the average over all accounts and the other lines show the averages for accounts that have different credit scores in each quarter. The large decrease in open credit card accounts starting in the last quarter of 2008 is the most important feature over the 14 years. A large number of banks reduced credit limits for many accounts, since these credit limits represented a potential liability.[8] While those with excellent credit are usually less likely to have an account closed, at the end of 2008 more than 20 percent of those with excellent credit had a decrease in open accounts. The continuing low mean number of open accounts after 2009 suggests that credit is still much less available than before the 2008 crisis.

---

[8]The Federal Reserve quarterly survey of senior loan officers in October 2008 found that 60 percent of banks had reduced credit card limits for nonprime borrowers, and 25 percent had reduced limits for prime borrowers (Federal Reserve Board 2008). While fall 2008 represented a severe contraction in credit, even the January 2010 survey showed 40 percent of banks reducing credit limits (Federal Reserve Board 2010).

2.2 Estimation

This section provides estimates of the extent of credit limit volatility that account for being cut off entirely. The estimates allow for systematic differences by age, time, geography, and the changes from accounts that are not reported.

Each individual has an evolving credit limit. The credit limit process has a Markov component that determines whether she can borrow at all, and a separate component that determines the actual limit when she can borrow. Much like typical estimates of income processes (Carroll and Samwick 1997), the credit limit process has a "permanent" (random walk) component and a transitory component. The parallel between income and credit limit is intentional; although card issuing banks are likely unable to adjust to the short-term income fluctuations of individuals, credit limits are typically set in proportion to income and so one might expect a similar evolution of credit limits.

The credit limit for person i at time t is:

$$B_{it} = D_{it}e^{X_{it}\beta_{it}}M_{it}W_{it}, \tag{1}$$

where $D_{it} \in \{0, 1\}$ determines whether person i can borrow at all; X_{it} contains age, date, geographical location, and credit risk; and the M_{it} and W_{it} are innovations specific to i. Defining $m_{it} = \ln M_{it}$ and $w_{it} = \ln W_{it}$, then the permanent component m_{it} evolves as $m_{it} = m_{it-1} + v_{it}$ and the individual innovations v_{it} and w_{it} are independent and identically distributed across individuals and over time. Both (log) innovations have mean zero and have variances σ_v^2 and σ_w^2, respectively.

Access to credit (D_{it}) evolves according to a simple Markov process: with probability p_{BN} someone who can currently borrow ($D_{it} = 1$) loses access next period ($D_{it+1} = 0$) and so transitions from the Borrowing state to the Not-borrowing state. With probability p_{NB}, someone who currently is in the Not-borrowing state transitions to Borrowing.

Since the Equifax data only report the limits of cards that are updated, the observed limit $\hat{B}_{it}$ is some fraction of the actual limit. If the ratio of the reported accounts to open accounts is f_{it},

then the observed limit is related to the actual limit by $\hat{B}_{it} = f_{it}^{\rho} B_{it}$, where ρ allows the limits of reported cards to have systematically higher or lower limits. Such a difference could arise if people systematically use their high limit cards but keep lower limit cards, which they do not often use, around for precautionary purposes.

Following Carroll and Samwick (1997) and Gourinchas and Parker (2002), I take a simple approach to estimating the various components of the credit limit process B_{it} by breaking up the estimation. So first I estimate the transition probabilities p_{BN} and p_{NB}. Then, conditional on being able to borrow, I estimate:

$$\ln \hat{B}_{it} = \rho f_{it} + X_{it}\beta + b_{it},$$

which gives the expected evolution of the credit limit by age, year, credit risk, geographical location, and reported cards, leaving individual innovations as the residuals. Finally, I estimate the transitory and permanent variances faced by individuals from the residual b_{it} process.[9]

Table 3 shows the estimates from this three-stage procedure and figure 4 plots the effects of age and quarter, whose coefficients I do not report separately. Credit limits increase quite rapidly with age in the twenties, slow in the thirties, and plateau after age 45 or so. This profile parallels income profiles, although it is also consistent with a process of learning who is a good credit risk. Credit limits overall are generally increasing, although with some aggregate volatility, until the end of 2008, and then decline rapidly. Adjusted for age and credit risk, they are still declining as of the end of 2013. The credit limits of cards that are not reported appear to have limits that are about one-third lower on average than the cards that are updated.

Each quarter, 2.67 percent of people with credit in the previous quarter lose access, while 6.03 percent of those who did not have access gain it back. These estimates are the overall averages for 1999–2013 that are shown in figure 3. The higher probability of regaining credit leads to the large

---

[9]To estimate the variances of the individual innovations I use the structure of the debt process to provide an efficient and simple estimation approach as suggested by Carroll and Samwick (1997). Changes in the residuals $b_{it} - b_{it-1} = v_{it} + w_{it} - w_{it-1}$ have variances $\text{Var}[b_{it} - b_{it-1}] = \sigma_v^2 + 2\sigma_w^2$. Larger differences have variances given by: $\text{Var}[b_{it} - b_{it-d}] = d\sigma_v^2 + 2\sigma_w^2$. So a regression of the form $(b_{it} - b_{it-d})^2 = \sigma_v^2 d + 2\sigma_w^2 + \mu_{idt}$, where for each individual at each time every lag is an additional observation, provides a full GMM estimate of σ_v^2 and σ_w^2 using all available moments. In practice, since the assumption of a fully random walk is strong, I use only d between 1 and 12.

fraction who have access in any given quarter.

Beyond uncertainty of access, individuals face substantial volatility in limits even when they can borrow. That volatility is larger than typical estimates of the variance in income when calculated at a quarterly frequency. For example, the Carroll and Samwick (1997) estimate of the yearly transitory variance from the PSID is 0.044. This estimate implies a quarterly variance of 0.166 that is *smaller* than the estimated transitory credit limit volatility of 0.248.[10] Other papers (Gourinchas and Parker 2002; Storesletten, Telmer, and Yaron 2004) use or make similar estimates, while Sabelhaus and Song (2010) finds a somewhat higher temporary variance (0.119) using Social Security Administration records. So short-term credit limit volatility is at least as large as income volatility even without taking into account the large probability of losing access to credit entirely. Moreover, the long-term variance of credit limits is much larger than estimates of income variance (yearly 0.027 in Sabelhaus and Song (2010); 0.022 in Carroll and Samwick (1997)). While short-term fluctuations tend to cancel each other out over time, long-term fluctuations accumulate, so the implied yearly variance of permanent credit limit shocks is approximately four times the quarterly variance. Households face a great deal of credit volatility in both the short and long term.

3 A model of consumer spending and debt

This section presents a basic model of intertemporal consumption and demonstrates that an optimizing household will never hold both debt and cash at the same time with a certain borrowing limit, but may do so with a stochastic limit. For simplicity, it does not include several extensions used in the simulations.

The model departs from the standard intertemporal consumption model in two ways: First, it allows debt and cash consumption to be separate and decided separately. This means that there are two decision variables: consumption paid for using cash and consumption paid for using credit, and two states: debt and savings. In the standard model without a stochastic credit limit, debt

---

[10] Since the income process is transitory and independent, if the yearly variance is σ_{uy}^2, the quarterly variance is approximately four times greater, since independent quarterly fluctuations tend to cancel out. More precisely $\sigma_{uq}^2 = \log(4*(exp(\sigma_{uy}^2) - 1) + 1)$.

and cash wealth can be collapsed into a single state variable "assets" or "wealth" and so do not have to be treated separately. The second departure is to allow for a stochastic credit limit. With a stochastic limit, it is possible for a consumer to have debts greater than her credit limit, since she may have borrowed under a previously higher limit. It is therefore necessary to specify what happens in this situation, which is typically ignored since it never occurs in the standard model. I base the structure on the credit card agreements from major banks (see appendix A). Credit card issuers have the right to restrict credit limits at any time but cannot demand full repayment unless the account is in default. Those above their credit limit must pay at least the interest on the debt every period, and may pay more than that, but do not have to. So debt, for those who are above their credit limit, is nonincreasing.

3.1 The basic model

A household dynasty or infinitely lived individual seeks to maximize:

$$\max_{\{c_t\}_{t=0}^{\infty}} E_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t^w + c_t^b) \right], \tag{2}$$

where consumption in each time period is composed of two parts, consumption paid for with cash c_t^w and consumption paid for with debt c_t^b. Utility does not depend on how the consumption is financed, only on how much there is, so consumption this period is $c_t = c_t^w + c_t^b$.[11] At the start of each period a consumer receives stochastic labor income and the returns on cash wealth saved from the previous period, both of which are immediately deposited. Cash wealth w_t then evolves according to:

$$w_{t+1} = (1 + r_s)(w_t - c_t^w) + y_{t+1}, \tag{3}$$

where disposable income in period $t + 1$ is y_{t+1} with distribution $G(\cdot)$, whose support has a lower bound $y_l > 0$, and the safe rate of return on savings is $r_s > -1$. The consumer can never spend

---

[11]See Hartley (1996) for a model in which cash goods and luxury goods (which can be bought using either cash or credit) enter utility separately. Consumers in that model, whose only uncertainty is whether luxury goods will be available, do not hold both credit card debt and excess cash, an assumption that allows the value function to be reduced to a single state variable.

more cash than she has, so $c_t^w \leq w_t$, and so starting with $w_0 \geq y_l$, $w_t \geq y_l$ for all t.

Cash balances are necessarily positive because the consumer gets some sort of consumable resources each period, whether or not there is debt. This observation is the accounting explanation for the credit card puzzle: we should expect to see positive cash balances unless the timing of payments is such that wages can be paid directly into debt (if wages are garnished, for example). The real puzzle then is not whether the individual has positive cash holdings during the period, but whether she has cash savings from period to period that are positive $w_t - c_t^w > 0$, while at the same time carrying debts.

The evolution of debt is simpler, but the constraints on the choices are more complex. Debt b_t evolves according to

$$b_{t+1} = (1 + r_b)(b_t + c_t^b),\tag{4}$$

where $r_b > 0$ is the borrowing interest rate. If $c_t^b > 0$, then the consumer charges on the credit card and so debt increases. If $c_t^b < 0$, the consumer pays off some debt. The transfer takes place at the time of consumption—negative debt consumption requires higher positive cash consumption to reach the same utility.

Note that if a consumer buys something with a credit card, then pays it off with cash in the same time period, this transaction is a cash purchase within the model. Such convenience users are not carrying debts from month to month but instead are using the credit card as a payment mechanism.

I focus on the case where $r_b > r_s$, so borrowing is more costly than savings. Under these circumstances, the consumer might want to lend rather than borrow on the credit card. To keep this from happening, the lower bound on debt consumption is that it can only pay off current debt: $c_t^b \geq -b_t$. With $b_0 \geq 0$, this implies $b_t \geq 0$ for all t.

The upper bound on debt consumption has two parts. Debt-financed consumption can be up to the available credit in the current period, as long as current debt is not above the credit limit. If current debt is greater than the credit limit, then the consumer must pay at least the interest on the

current debt so the debt cannot grow. The full constraint on debt consumption is:

$$-b_t \leq c_t^b \leq \begin{cases} \frac{-r_b b_t}{1+r_b} & \text{if } b_t > B_t/(1+r_b) \\ \frac{B_t}{1+r_b} - b_t & \text{if } b_t \leq B_t/(1+r_b), \end{cases} \qquad (5)$$

where the credit limit $B_t \geq 0$. Note that the consumer can borrow only up to $B_t/(1+r_b)$, which ensures that $b_{t+1} \leq B_t$ next period. B_t is a limit on next period debt but is subscripted t to make it clear that it is a constraint on behavior in t and in the information set at t. How B_t changes or not over time, and how this affects the household's decisions, are considered in the next two sections.

Define debt holdings at the end of the period $D_t = (b_t + c_t^b)$ and cash savings at the end of the period as $S_t = (w_t - c_t^w)$. Note that both are non-negative for all t. The credit card puzzle is then the observation that for a large number of consumers $S_t > 0$ and $D_t > 0$ in the same period.

Since the constraints are crucial and complicated, the Lagrange approach (Chow 1997) is helpful to state the consumer's problem. Given $w_0 > 0$, $B_t \geq 0$, and $b_0 \geq 0$, the consumer's problem is:

$$\max_{\{c_t\}_{t=0}^{\infty}} E_0 \left[\sum_{t=0}^{\infty} \beta^t u(D_t - S_t + w_t - b_t) \right] \qquad \text{subject to:}$$

$$\lambda_t^w: w_{t+1} = (1+r_s)S_t + y_{t+1} \qquad\qquad \lambda_t^b: b_{t+1} = (1+r_b)D_t$$

$$\lambda_t^{D_0}: D_t \geq 0 \quad \lambda_t^S: S_t \geq 0 \qquad\qquad \lambda_t^D: \max\{B_t/(1+r_b), b_t/(1+r_b)\} - D_t \geq 0,$$

where the λ's have an attached complementary slackness constraint and are the multipliers from the associated Lagrange equation:

$$\mathcal{L} = E_0 \left[\sum_{t=0}^{\infty} \beta^t \Big(u(D_t - S_t + w_t - b_t) + \lambda_t^w((1+r_s)S_t + y_{t+1} - w_{t+1}) \right.$$

$$\left. + \lambda_t^b(b_{t+1} - (1+r_b)D_t) + \lambda_t^{D_0} D_t + \lambda_t^D(\max\{B_t/(1+r_b), b_t/(1+r_b)\} - D_t) + \lambda_t^S S_t \Big) \right].$$

3.2 Fixed credit limit

This section shows that no optimizing consumer who can borrow this period ($b_t < B$) has both $S_t > 0$ and $D_t > 0$ at the same time if the credit limit is not stochastic, so $B_t = B > 0$. The reason to show this mostly obvious result is that setting up the first-order conditions makes it obvious how the problem changes when B_t is stochastic in the next section and so it makes clear why credit limit variability is so important.

If $B_t = B$ for all t and $b_0 \leq B$, then the constraint on D_t can be simplified to $B/(1+r_b) - D_t \geq 0$. This simplification is crucial for what follows, and I show the implications of what happens when it does not hold in the next section. Consider the decision at t following a sequence of feasible, although not necessarily optimal, decisions $\{D_\tau, S_\tau\}_{\tau=0}^{t-1}$, so that $w_t \geq y_l$ and $0 \leq b_t \leq B$. The first-order necessary conditions are:

$$w_{t+1}: \beta E_t[u'(D_{t+1}^* - S_{t+1}^* + w_{t+1} - b_{t+1})] - \lambda_t^w = 0$$

$$b_{t+1}: \beta E_t[-u'(D_{t+1}^* - S_{t+1}^* + w_{t+1} - b_{t+1})] + \lambda_t^b = 0$$

$$D_t: u'(D_t^* - S_t^* + w_t - b_t) - \lambda_t^b(1 + r_b) + \lambda_t^{D_0} - \lambda_t^D = 0$$

$$S_t: -u'(D_t^* - S_t^* + w_t - b_t) + \lambda_t^w(1 + r_s) + \lambda_t^S = 0$$

$$\lambda_t^{D_0}: \lambda_t^{D_0} \geq 0, \ D_t \geq 0, \ \lambda_t^{D_0} D_t = 0$$

$$\lambda_t^D: \lambda_t^D \geq 0, \ B/(1 + r_b) - D_t \geq 0, \ \lambda_t^D(B/(1 + r_b) - D_t) = 0$$

$$\lambda_t^S: \lambda_t^S \geq 0, \ S_t \geq 0, \ \lambda_t^S S_t = 0,$$

which include the inequality constraints, nonnegativity constraints for the λ's, and complementary slackness conditions. Also necessary are the accumulation equations, which are the FOC's for λ_t^w and λ_t^b.

Suppose $D_t^* > 0$ and $S_t^* > 0$. Then, by the complementary slackness condition, both $\lambda_t^S = 0$ and $\lambda_t^{D_0} = 0$, and so using the FOC's for D_t and S_t gives:

$$\lambda_t^b(1 + r_b) + \lambda_t^D = u'(D_t^* - S_t^* + w_t - b_t) = \lambda_t^S(1 + r_s),$$

and substituting from the conditions for w_{t+1} and b_{t+1}:

$$(1+r_b)\beta E_t[-u'(D^*_{t+1}-S^*_{t+1}+w_{t+1}-b_{t+1})]+\lambda^D_t = (1+r_s)\beta E_t[u'(D^*_{t+1}-S^*_{t+1}+w_{t+1}-b_{t+1})]. \quad (6)$$

Since marginal utility is positive, $\lambda^D_t \geq 0$, and $(1+r_b) > (1+r_s)$, the equation above cannot hold, and so both $D^*_t > 0$ and $S^*_t > 0$ cannot be optimal. So the optimal path does not include instances where the consumer both leaves cash to accumulate at the low interest rate, and has debt that accumulates at the high interest rate in the same period.

3.3 Stochastic credit limit

When the credit limit is stochastic so that B_t is a t-measurable random variable whose support has a lower bound of 0, the optimal path may include both positive savings and debt. Intuitively, with a stochastic credit limit, when times are bad it may make sense to use debt to consume now rather than cash, since times may also be bad next period, but then the cash will be gone and borrowing may not be possible. Cash wealth acts as self-insurance, just as all wealth does in standard precautionary models. When B_t is stochastic, the decision to leave more debt for next period affects not just how much debt the consumer has next period, but also the likelihood and marginal utility cost (given by the shadow price λ^D_{t+1}) of the credit limit binding in the next period. To simplify the analysis, I restrict $B_t \in \{0, B\}$, so that the consumer can either borrow up to B or not at all.[12]

A stochastic credit limit means that the FOC for b_{t+1} must include that the constraint on D_{t+1} may bind, and so the decisions today affect the costs of the constraint binding in the next period. The FOC for b_{t+1} now reads:

$$b_{t+1}: \beta E_t[-u'(D^*_{t+1} - S^*_{t+1} + w_{t+1} - b_{t+1}) + 1(b^*_{t+1} > B_{t+1})\lambda^D_{t+1}/(1 + r_b)] + \lambda^b_t = 0,$$

---

[12]The assumption that $B_t \in \{0, B\}$ means that the decision about D_t and so b_{t+1} does not affect the distribution of $1(b^*_{t+1} > B_{t+1})$, since the restrictions on D_t mean that $b_{t+1} \leq B$ for all t. So the simplification allows me to ignore that having more debt might make it more likely that the credit limit will bind by increasing the probability that $b^*_{t+1} > B_{t+1}$.

where $1(b_{t+1}^* > B_{t+1})$ is a random variable that is 1 if $b_{t+1}^* > B_{t+1}$ and zero otherwise. The left-hand side of equation 6 gains an additional term: $-\lambda_{t+1}^D/(1 + r_b)1(b_{t+1}^* > B_{t+1})]$. This term means it is possible for the equality to hold and so the optimal path can include both D_t^* and S_t^* positive at the same time.

To gain some insight into the behavior that leads to D_t^* and S_t^* both being positive, rearrange equation 6 with the new term to equate the marginal costs to the marginal benefits of increasing both D_t and S_t and so entering the next period with (in expectation) both more debt and more savings:

$$(r_b - r_s)\beta E_t[u'(D_{t+1}^* - S_{t+1}^* + w_{t+1} - b_{t+1})] + \lambda_t^D = \beta E_t[1(b_{t+1}^* > B_{t+1})\lambda_{t+1}^D].$$

The price of higher debt and higher savings is the difference in interest rates $(r_b - r_s)$, which means there will be lower cash next period with an expected marginal utility cost in the next period given by $(r_b - r_s)\beta E_t[u'(D_{t+1}^* - S_{t+1}^* + w_{t+1}^* - b_{t+1}^*)]$. Since holding less cash and more debt decreases cash wealth next period, marginal utility increases. Increasing both cash and debt leaves consumption unchanged, and so has no cost this period (as long as the debt constraint λ_t^D is not binding).

The right-hand side gives the expected benefit in the next period of increasing both the cash and debt left to next period. It is the expected cost of the credit limit binding in the next period. When the credit limit binds, the household wants to consume more through debt, so the costs of the credit limit binding are high when marginal utility is high. The credit limit binds if past shocks have left savings low, current income is low, and the consumer has either accumulated substantial debts or the credit limit has decreased. If the period utility function displays a precautionary motive, $u''(\cdot) < 0$, then marginal utility is decreasing, although a borrowing constraint imposes a precautionary motive all by itself (Carroll and Kimball 2001). When the credit limit binds, then consumption is determined by the amount of cash savings, so increasing both savings and debt reduces the cost of the credit limit binding since marginal utility is decreasing. So the cost of the

credit limit binding tends to be high when marginal utility is high: with a bad shock (low y_{t+1}) and low cash wealth. In this case the costs of extra debt are small, while the benefits of increased cash-at-hand are large, and it is worth paying the carrying cost of a little extra debt to keep some cash as insurance.

From a household financial management point of view, when faced with variable credit limits, the first goal is to have some cash savings, and only then to pay down debt. It is only optimal to pay off the expensive debt once the household has sufficient cash savings to use in emergencies. When very cash poor, it may make sense for the household to use credit to fund consumption and so increase debt in order to build up cash savings.

While individual incomes and credit limits occasionally decrease, overall consumer credit and incomes have generally been increasing. The same analysis applies, however, with underlying growth that increases incomes and credit limits at the same rate, together with the restriction of the period utility to show constant relative risk aversion. Then the analysis proceeds in ratios (Carroll 2004; Deaton 1991): a low credit-limit-to-income ratio next period means that it may be worth having both the ratios of cash-to-income and debt-to-income positive this period. With growth, the consumer keeps cash as a precaution against the debt-limit-to-income ratio being low. Since growth tends to make consumers more impatient—the future will be better, so it makes sense to consume more now—it will tend to make borrowing more attractive, and so increase the proportion of the population that both borrows and saves at the same time. Ludvigson (1999) examines such a situation with stochastic credit limits, although she explicitly limits the analysis to exclude borrowing and saving at the same time.

While rising incomes may not be especially important for monthly consumption and debt decisions, the same approach of examining ratios is useful for comparing a population with very different incomes. Under the same conditions, those with different permanent incomes will make the same choices in ratios: a person who has a high wealth-to-permanent-income ratio will choose to consume the same fraction of income whether her permanent income is large or small. This property of the model allows me to compare the debt and savings decisions of those with both high

and low incomes in a consistent manner by considering their decisions relative to their incomes.

4 The joint savings and debt distribution from the data and the model

Allowing the credit limit to be stochastic means that having positive debt and positive savings at the same time is possible for an optimizing consumer. Whether anyone decides to hold both debt and savings at the same time, however, depends on that person's preferences, relative returns, and the distribution of shocks. This section uses the estimates of credit limit volatility in section 2 to calibrate the model and compares it to the joint distribution of liquid savings and debt from the Survey of Consumer Finances (SCF). Different assumptions about the timing of what people report in the survey change the extent of the puzzle and so the exact extent of the credit card puzzle is unclear. While I calibrate the model based on the most reasonable definition, I put less emphasis on the parameters that give the best fit to the SCF, since these depend heavily on the data assumptions, and much more emphasis on whether the model can explain a wide range of behaviors

4.1 The debt and savings distribution from the Survey of Consumer Finances

The 2007 Survey of Consumer Finances asks questions about assets, debts, and income of a representative sample of a little over 4,000 households (see Bucks et al. (2009) for a description of the 2007 SCF and changes between 2004 and 2007). Weighted to be nationally representative, 73 percent of households had credit cards. Of the remaining 27 percent, around half did not apply for credit in the previous five years, one quarter applied and were rejected, and one quarter applied and were accepted but did not have a credit card at the time of the survey. The analysis concentrates on those currently with credit cards, since that is the approach taken by the literature in examining the credit card puzzle (see, for example, Gross and Souleles (2002)).

A simple glance at the data suggests that the "credit card puzzle" is nearly universal. In the SCF, 97 percent of households with a revolving credit card balance also held some liquid assets at some time during the month. Yet the near universal holding of liquid assets is partly an accounting artifact. The questions from the survey ask for the credit card balance after the household made

19

the last payment and for the household's holdings of liquid assets at the time of the survey—which may frequently be interpreted as the time of the last bank statement. This timing issue makes the true extent of the credit card puzzle imprecise, since when income is received, credit paid, and consumption occurs changes whether many households or only a few are making the portfolio decision not to shift assets from savings into paying down debt. Different assumptions are reasonable: Gross and Souleles (2002) calculate the puzzle by subtracting monthly income from liquid savings, an approach Telyukova and Wright (2008) suggest is conservative, since it depends on when consumption occurs.

Figure 5 helps illustrate the complexity of finding the correct accounting timing and how different reasonable approaches could give different answers. In both panels (a) is the amount left over after any payment on the credit card bill but before income or any consumption. The model timing has income coming at the beginning of the period, then the consumption, savings, and debt decisions at the end. In reality, income may come at any time during the month, and may even come at multiple times, and consumption need not take place at the same time that the credit card bill is paid, or all at once. We may, for example, observe (b), which includes credit card puzzle savings as well as income, instead of (a). So Gross and Souleles (2002) calculate (a) by subtracting monthly income from liquid savings. Telyukova and Wright (2008) suggest we might be observing credit card puzzle savings and income minus some consumption (c), and so underestimate the extent of the credit card puzzle. Yet if there is expected consumption after the credit card bill is due, but before the consumer expects to get income, then even liquid savings minus income may overstate the amount by which consumers hold positive savings.

To deal with this timing issue, I examine savings accounts and checking accounts separately and show several different definitions of the puzzle. Both checking and savings are very liquid in the sense that for most savings accounts the balance can be shifted to a checking account easily, or sometimes payments can come directly from the savings account.[13] Savings accounts offer

---

[13]In order to determine appropriate reserve requirements, Federal Reserve regulation D limits the number of ACH transactions of a savings account to no more than six per month (or statement cycle, see http://www.federalreserve.gov/bankinforeg/reglisting.htm) but does permit transfers from one account to another by the same depositor at the same institution, or withdrawals by mail, in person, by telephone, or at the ATM.

slightly higher interest with slightly less liquidity than checking accounts. Because the interest differential is typically small, there is little reason to transfer money from checking to savings if that money will be used to pay off a credit card balance in the same month. I therefore include all of reported balances in a savings account as part of credit card puzzle savings. In addition, some portion of a checking account may be held even as the consumer rolls over credit card debt. If income is directly deposited, then the checking account includes monthly income at some point in the month.[14]

Several approaches to measuring credit card puzzle savings and dividing the sample are shown in table 4. The first four rows divide the population into four groups: those with positive savings and positive debt, those with positive debt and no savings, those with positive savings and no debt, and those with no debt and no savings. Below the fraction in each group, I show the mean and median debt and savings-to-income ratios for each group and overall. To match the model and credit volatility data, I report multiples of quarterly income.

Different assumptions about accounting change the fraction with positive debt and positive savings dramatically, but the credit card puzzle never goes away. Between 24 percent and 59 percent of households fall into the credit card puzzle. Subtracting monthly income has a large effect on the fraction of households who hold no savings at all. Depending on the approach, between 1.7 percent and 36 percent of households had debt but no liquid savings, while between 0.5 percent and 14.6 percent held neither debt nor savings. Which end of the range depends on how much of the checking account is left after paying the credit card bill: none or all. My preferred approach is shown in column one: including checking balances only if they are greater than income, which gives a relatively conservative 41.4 percent of households with positive debt and liquid savings, and around 18.8 percent with debt and no savings. The joint histogram of savings and debt corresponding to

---

[14]I construct these amounts as follows from the 2007 SCF: The value in all checking accounts is the sum of all accounts (if positive amount) in SCF variables x3506, x3510, x3514, x3518, x3522, x3526, and the catchall x3529. Savings accounts are from SCF variables x3730, x3736, x3742, x3748, x3754, x3760, and the catchall x3765. Credit card debt comes from SCF variables x413, x421, x424, x427, and x430. The overall credit limit is asked only for all cards and is SCF variable x414. For income I use the household income SCF variable x5729 but replace it with "normal" income (x7362) if the respondent says income was not normal. To compare high and low-income households, all savings, checking, and debts, and credit limits are normalized by this income.

the first column is shown in panel A of figure 6.[15]

4.2 Computing the consumer's decisions

While the model is a relatively straightforward variation of a standard infinite-horizon consumption model, simulating the consumer's decisions for a given set of preferences and parameters is very difficult. There are two reasons for the difficulty. The first is the standard curse of dimensionality: there are two continuous state variables, and there are two continuous decisions, both of which are functions of both state variables. The added dimensionality makes the calculations more computationally intensive.

The dimension of the problem is compounded by a problem of scaling: in the short term the decision of how much to consume in total is much more important for utility than the portfolio decision of how that consumption should be financed. The problem is similar to trying to find the highest point on a knife-edged ridge that falls off steeply on either side and slopes only gently along the top of the ridge. An intelligent hiker knows that the best way to get to the top is to walk along the ridge and find the highest point. The dimensionality of the problem means that it is only possible to approximate at set grid points with functional-form approximations in between. This makes it very easy to fall off the ridge on either side. Since a wrong step means disaster, it is very easy to not find the optimum. In addition, along the threshold between regime changes (moving from zero to positive for any decision) the second derivative of the value function is not necessarily continuous. The scaling and discontinuity problems mean that optimization methods that rely on the second derivative—as most practically useful methods do—are often badly behaved, with wide swings between decisions, and frequent non-convergence. I use a hybrid approach that relies on a Newton solver initially, but then switches to a much slower, derivative-free method for the states that resist convergence. Appendix B discusses the complexities of finding the consumption functions for a given set of preferences, credit limit, and income processes. With the consumption

---

[15]The histogram is for the unweighted distribution, while the table is calculated using survey weights. The SCF overweights high-income earners in its sampling. While the analysis divides by income, and some high-income earners do both borrow and save at the same time, they are somewhat less likely to do so than those with lower incomes.

functions for a given set of parameters, I then draw multiple shocks for a large community of simulated individuals to find the ergodic distributions of wealth, debt, and consumption

4.3 Parameter choices and calibration

I calibrate the model to match the total amount of debt from the SCF. Preferences and the income and credit limit processes interact in complex ways in this model. I simulate the model using a range of parameters both to understand how different preferences and risks affect the decisions and also to demonstrate that the model's ability to generate savings and debt decisions similar to the data is not sensitive to small changes in parameters.

Many of the prices and limits are observable or well estimated elsewhere. I set the interest rate earned on savings r_s to 0.4 percent and adjust all interest rates for inflation of 2.5 percent so that the real return on savings is negative.[16] I set the interest paid on debt r_b to 13.60 percent.[17]

There are two sources of uncertainty within the model: the income process and the credit limit process. For both processes I exclude variation from a "permanent" or random walk component. That uncertainty is not a central component of short-term decisions (one cannot effectively smooth over a permanent shock) and allowing for a random walk requires an additional continuous state-variable, exponentially increasing the computational complexity of the model.

I base the credit limit process on the estimates of losing and gaining access to credit in section 2 and table 3. The credit limit, conditional on being able to borrow, is one-quarter of yearly income and does not vary stochastically. Focusing only on access leaves out the extensive volatility of

---

[16]The average interest paid on a no-frills Wells Fargo savings account was 0.15 percent in 2006 (the year the SCF was in the field), while it was 0.615 percent for a Bank of America account with a minimum deposit of $20,0000. The average interest earned is somewhere between. The CPI was 2.5 percent in 2006 according to the Bureau of Labor Statistics.

[17]The average credit card interest rate reported in the SCF is 14.22 percent while it was 14.73 percent based on Federal Reserve series G19 (Commercial Bank Interest Rate on Credit Card Plans NSA). Since the model does not have default, the actual rate that consumers pay on their debt is somewhat lower. Edelberg (2006) examines the risk-based pricing of interest rates on different types of consumer loans, including credit cards. The basic idea is to use the probability of default from the PSID and map that probability by demographic type into different interest rates reported in the SCF. In 1998 when the mean in the SCF for credit card rates was 14.52 percent, she finds the zero bankruptcy risk rate to be approximately 13.9 percent. Since so few people actually default, the actual rate is close to the default rate. I adjust the 2007 SCF interest rate of 14.22 percent downward by 0.62 percentage points. This is a smaller adjustment than made in Angeletos et al. (2001), who adjust for default by 2 percentage points.

credit limits that occurs when the limit is positive and so somewhat understates the importance of credit limit volatility. Getting cut off from credit entirely is far more important, however. Being unable to borrow is a potential constraint for everyone, while having a lower credit limit is a constraint only for those with significant debts.[18]

The income process is the ultimate source of marginal utility shocks in the model. Income follows an independent and identically distributed lognormal process with variance 0.166 that implies a year transitory income variance of 0.044 to match Carroll and Samwick (1997) (see footnote 10). At a quarterly frequency this leads to a very volatile income process, but the income process also needs to incorporate all of the other marginal utility shocks that may occur: cars breaking down, sickness, or the need to move. While these types of shocks are difficult to estimate well, they may be more important than income at a quarterly or monthly level. [19]

Preferences display constant relative risk aversion (CRRA). I show the results for multiple preferences but combine them using the estimates of heterogeneity of risk aversion in Kimball, Sahm, and Shapiro (2009). Their estimates suggest somewhat higher risk aversion than the single-parameter estimates used elsewhere.[20]

---

[18]Allowing for credit limit volatility beyond whether or not one can borrow turns out to make the already-difficult problem of simulating the consumer's decision much harder. When the credit limit uncertainty is only whether the individual is able to borrow, all of the importance of changing limits occurs at the edge of the debt state space. With internally evolving limits, the consumer with some debt has to make a decision about whether to increase debt, taking into account that with more debt there is a possibility that she will face a lower limit next period. With a discretized credit limit process, this means that there is a jump in the probability of being constrained as debt increases, which means a kink in the value function (at least for intermediate value functions: it is possible that the fixed-point value function may smooth out the kink). The simulations I attempted that allowed for limit volatility were not well behaved: they either would not convergence or on convergence had decision functions that were not smooth. While for the ones that did converge I did not find that allowing for the extra volatility made much difference, that may have been because convergence only occurred when that volatility was not important.

[19]According to the National Highway Safety Administration there were just over 6 million car accidents in 2007 (National Highway Trafic Safety Administration 2009), while there were 112 million households in 2007 (Kreider and Elliott 2009), implying that there was on average slightly more than one car accident for every 20 households in 2007. There were 2.4 million deaths from all causes, which suggests that on the order of 2 percent of households had to pay for a funeral (see Center for Disease Control http://www.cdc.gov/nchs/fastats/deaths.htm, accessed 2 May 2010). Other large consumption events that are unexpected sufficiently long in advance include weddings, births, and moving from one residence to another. While these may be known about in advance, they also may represent shocks when they are first discovered.

[20]The heterogeneity estimated by Kimball, Sahm, and Shapiro (2009) is for six constant relative risk aversion parameters absorbing the full population: 1.4 (6.6 percent), 2.2 (13.7 percent), 2.8 (15.0 percent), 3.5 (15.6 percent), 4.2 (18.2 percent), and 6.7 (30.9 percent). That implies an arithmetic mean of 4.19 and a median of 4.2 (although the average risk aversion will not produce the average behavior of heterogeneous preferences). A coefficient of relative risk aversion around 2 is common, as in for example, Storesletten, Telmer, and Yaron (2004), so these preferences

Finally, I allow the yearly discount factor to vary and calibrate the model so that the heterogeneous consumers use the same fraction of their available credit as in the 2007 Survey of Consumer Finances: 23.3 percent as shown in table 4. I target the fraction used rather than the very close debt/income ratio (0.22), since the fraction used does not have the additional variance generated by income differences and so is more precisely estimated and less susceptible to outliers and under-reporting.[21] Within the model they are exactly the same.

Calibrating the model to get consumers to hold enough debt requires a yearly discount factor of 0.794. Figure 7 shows the effects of varying the discount factor. The horizontal dotted line shows the savings and debt decisions estimated from the 2007 SCF (in column 1 of table 4). The vertical dotted line shows the calibrated discount factor. The thick line shows the weighted distribution, while the other lines show the distribution for specific risk preferences. Reducing the discount factor tends to increase debt and so the total debt line (second row, right column) crosses the debt line from the SCF at a discount factor of 0.794.

The discount rate necessary to get consumers to carry enough debt is higher than normally assumed for models of long-term savings but is reasonable given the much higher frequency of the model and the high degree of risk aversion. Getting consumers to carry enough debt while also saving is a challenge. For example, Laibson, Repetto, and Tobacman (2007) calibrate a model that has both credit card debt and retirement savings with quasi-hyperbolic discounting and find that a short-term discount rate of 40 percent is necessary to get consumers to carry enough credit card debt. Since the model in this paper focusses on short-term smoothing rather than long-term accumulation, it is reasonable for the discount factor to be high. Second, given the relatively risk-averse preferences from Kimball, Sahm, and Shapiro (2009), consumers naturally want to hold savings and keep credit available since they face a risky income environment. Getting them to hold enough debt requires them to discount the future heavily. Since with CRRA preferences the

---

imply a high degree of risk aversion and precaution.

[21]The respondents from the SCF tend to under-report their credit debt (Zinman 2009). My comparisons with the Equifax data suggest they also understate their limits relative to what banks report. That leaves their credit utilization not as far off as the total credit or debt. Since the joint distribution of savings and debt is crucial, I do not try to adjust the credit amounts, or add in cash (Federal Reserve Bank Notes) that are also left out of the SCF. Cash holdings are approximately $138 (Foster, Schuh, and Zhang 2013).

coefficient of relative risk aversion (γ in the figures) is the inverse of the intertemporal elasticity of substitution, the coefficient and discount factor are intimately related. It is often difficult to distinguish between caring more about future bad draws and caring more about the future. A less risk-averse population would not need to discount the future so highly in order to hold the debt we see in the data.

Surprisingly, since the model was *not* calibrated to match this moment, the calibrated model explains all of the fraction of households that have positive debt and positive savings. The calibrated model tends to predict too few people making the very risky decision to not have any savings. It also predicts too little savings. That is not a surprise: precautionary models typically have trouble explaining large-scale accumulation for life-cycle or investment purposes (see, for example, Carroll (2001)). Nonetheless, the model predicts around three-quarters of the liquid savings actually accumulated in the SCF and so actually explains liquid savings reasonably well.

4.4 The full distribution

One substantial advantage of this model is that it can generate a full distribution of debt and savings. Figure 6 shows how well the calibrated model does in generating joint savings and debt distribution. Panel (A) shows the joint histogram from the 2007 SCF, while panel (B) shows the simulated distribution. The model does a remarkably good job of explaining the broad facets of the data. Substantial fractions of the population have both debt and savings at the same time, just have savings, or just have debt. While the calibrated distribution is not as spread out as the survey, it still captures much of the diversity of savings and debt, including the approximately log-normal shape of the joint distribution.

The simulated distribution also matches some of the comparative features of the data well. In the SCF, those who have both debt and savings have, on average, slightly less debt than those who have debt alone, but have much lower savings. Within the model, savings are a way of self-insuring against bad shocks and against being unable to borrow. But having lots of savings protects against both, and so those who are wealthy now can use accumulated savings to smooth

26

consumption, while still having savings left over—they do not need to use debt to maintain savings as insurance. So using both savings and debt occurs only among the less wealthy, although not among the poorest, who from preferences or necessity, do not maintain any liquid wealth. This feature of the data is matched by the simulations. Such a distribution cannot be explained by alternative explanations based on inattention or mistakes. Mistakes or inattention are more costly for poorer people, since marginal utility is higher, and so one would expect to see those with lower savings making mistakes less often.

Since the model has a kink in the interest rate between borrowing and saving, some people do end up neither borrowing nor saving, but many fewer than in the data. The large fraction with neither debt nor savings in the data is mostly a statistical artifact, however. First, the SCF does not track holding of cash bills. So it seems likely that most of the people with no "savings" actually do have some cash. The average cash holdings is not large—about $138 according to (Foster, Schuh, and Zhang 2013)—but those with no bank savings or checking may be more hold more than the average. Second, as discussed in section 4.1, nearly everyone with credit card debt had a positive bank account balance at some point during the month. Subtracting income, as in column 1 of table 4, surely over-adjusts for some people, artificially leaving them with zero savings, as is evident by comparing column 1 with the other columns.

4.5 Parameter variations

The discount rate. Figure 7 shows how the credit and debt decisions vary as the yearly discount factor varies. The fraction who hold both debt and savings declines, largely because people hold more savings and less debt and so a smaller fraction of the population is actively borrowing.

Probability of losing access. Increasing the frequency with which people lose access has a somewhat different effect on decision-making as is shown in figure 8. As the probability of losing access increases, the value of having savings increases, since credit is less likely to be available when you need it. So the fraction with positive debt and savings increases while the fraction without any savings decreases. Even for a very wide range of probabilities, however, total savings and

27

debt vary relatively little. That is because even as credit becomes less valuable as a precautionary wealth, it becomes more valuable as the marginal source for smoothing if it is available. Hit with a bad shock, more consumers borrow if they can in order to preserve their valuable savings. So even as the fraction who cannot borrow in any given quarter is increasing, debt increases slightly as well.

Income volatility. Figure 9 shows how the credit and debt decisions vary as the variance of income increases. Higher income variance increases the value of both remaining credit and savings, so people tend to hold more savings and lower debt.

The importance of preference heterogeneity. Preference heterogeneity is a necessary component for understanding the debt and savings distribution. The data show that a substantial portion of the population has debts but little or no savings. That is a very risky decision given the debt process and income process: there is a fairly high probability that someone with no savings will be cut off from credit and have a bad shock. Accepting those risks requires relatively low prudence and risk aversion. There are also many households that hold debt and savings at the same time. This behavior implies they are unwilling to take such risks. Figures 7 and 8 show that almost all of the population with debt but no savings is composed of those with low risk aversion, while most of the population holding debt and savings is highly risk averse. So getting both a substantial portion of the population to hold debt and savings at the same time and a substantial portion of the population to have only debt requires preference heterogeneity. For just the right combination of credit risks and incomes, a single coefficient of relative risk aversion could deliver some people saving and borrowing and others just saving. Since the credit card puzzle has been relatively stable for at least 20 years, it is not credible to believe that it rests on a set of very specific parameters. Allowing for preference heterogeneity means that there are no rapid changes as risks or preferences change, since as one population is shifting entirely to borrowing and saving, another is willing to borrow without saving.

5 Who is borrowing and saving?

This section examines who is both borrowing and saving and how that has changed over time. Given the evidence, it then briefly reviews alternative explanations of the credit card puzzle and suggests that the other explanations, while a possible motivation for some households, are unlikely to explain much of the puzzle.

5.1 Who is borrowing and saving?

As figure 1 shows, the credit card puzzle has been relatively stable for at least 20 years even as the acceptance and use of credit cards has grown tremendously. There was a significant dip from 2007 to 2010 as the fraction borrowing and saving declined from 42.4 percent to 35.9 percent. The overall stability of the credit card puzzle across time is not a problem for the explanation used in this paper: the puzzle is always defined as why those with credit cards choose to borrow and save at the same time. Since most consumption is "cash consumption" as people spend their incomes, there do not have to be many opportunities to spend using credit in order for credit to offer a useful form of smoothing.

Who is borrowing and saving at the same time? Table 5 shows how the probability of borrowing and saving changes with household characteristics, using multiple waves of the SCF. The reported values are the marginal probabilities (not the logistic coefficients) from a multinomial logistic regression estimating the marginal probability (at the mean) of having (1) positive debt and positive savings, (2) positive debt, no savings, (3) positive savings, no debt, or (4) no savings, no debt. The first column estimates the marginal probability of being in (1) when all four possibilities are available; the second estimates the marginal probability of borrowing and so being in (1) or (2) and not (3) or (4); the third asks if the household is borrowing, what is the marginal probability it is in (1) as opposed to (2). The difference between the three columns is crucial since the model predicts that in general more patient and more prudent people are more likely to be saving. That makes them less likely to be borrowing and saving at the same time since they are less likely to borrow, as the difference between the high and low risk aversion lines in figures 7 and 8 illustrates.

Conditional on having enough bad shocks that they want to borrow, the very prudent always keep some savings available and so are more likely to borrow and save at the same time if they are borrowing.

The model distinction appears to provide a good breakdown of who is borrowing and saving at the same time as shown in table 5. While older households are less likely to borrow and save at the same time since they are less likely to borrow, if they do borrow they are more likely to have some savings. Households that say they are willing to take more risks to get more returns are less likely to borrow and save—a safer but lower-return option. Those who say they want a larger amount of emergency funds compared to their income—a ratio that Fulford (2013b) shows encapsulates all of the precautionary preferences—are less likely to borrow but more likely to save if they do borrow. The more willing a household is to use credit, the more likely it is to borrow, but conditional on borrowing, the less likely it is to save. Having a spouse or a larger household makes the household less likely to borrow and save, but it does not seem to have much of an effect conditional on borrowing. Higher-income people borrow less, but when they do, they tend to keep some money in savings. Homeownership—equity in a large illiquid asset—does not have much of an effect.

5.2 Alternative explanations of the credit card puzzle

Accounting. The most obvious, although least interesting, explanation is that the puzzle is a mere accounting issue: wages or salary may be deposited directly in the bank before being used to pay down debt. Then, at its extreme, if wages are paid the day after the due date on a credit card statement, there could be large balances for the entire month even though all available cash is used to pay down the credit balance every month. Even allowing for a month of gross total household income to be kept in liquid accounts, however, Gross and Souleles (2002) find that more than one-third of credit card borrowers keep more than this amount. Section 4.1 and table 4 examine the effects of different assumptions about timing in the data and show that at least one quarter of consumers must be borrowing and saving at the same time even by the strictest accounting

30

definition.

Mistakes. While mistakes may explain the puzzle for some households, the evidence suggests that a substantial portion of households are intentionally borrowing and saving rather than just making mistakes. When asked directly in the nationally representative Consumer Finance Monthly to evaluate the statement: "Using extra money in a bank savings account to pay off high interest credit card debt is a good idea," 21 percent say it is a good idea, 12 percent are unsure, and 67 percent say it is not a good idea. Even within the model, keeping debt and savings at the same time is the right decision for only some consumers, since those with a great deal of savings should pay off their expensive debts. If respondents interpret the question as asking whether it is good to pay off debts right now, then the answers correspond reasonably closely to the overall proportion who are saving and borrowing at the same time. Panel B of table 2 reports that the percentage who have both positive debt and savings at the same time is nearly identical between the groups who say it is a good idea and those who say it is not.

Of course, American households are not always well informed about their finances, but where respondents are likely to have direct experience with a decision, it seems they can find the correct answer. The CFM asks other questions about financial literacy. I show the full breakdown of two questions with clear, but not necessarily obvious, answers in panel B of table 2. Many households will have direct experience buying cars and only 10 percent answer incorrectly when given several financing options and asked which was the least costly. On the other hand, more than two-thirds could not pick the best insurance policy for a young family from the arcane and non-self-explanatory insurance types. When they did not know the answer, respondents were generally willing to say so. Since households have experience with the costs of borrowing on a credit card and have to make the decision regularly, a reasonable interpretation is that a substantial portion of those who are borrowing and saving at the same time are doing so intentionally. The costs of doing so are large enough to be noticeable: my estimates suggest that for those who are both borrowing and saving the cost is around 0.6 percent of monthly household income; Zinman (2007) finds comparable costs.

Merchant acceptance. The most compelling explanation is based on transaction uncertainty and so relies on similar precautionary preferences to those in this paper. Telyukova and Wright (2008) suggest that the credit card puzzle is just a new version of a much older question: why do people hold money, which pays no interest and may have a negative return in the presence of inflation, when they could earn a positive return in the bank? They propose a model in which some transactions cannot be paid for with credit since they take place anonymously and so must be settled on the spot. While some consumption items that are generally paid in cash, such as rent, are predictable, others such as emergency visits from the plumber, are not. The unpredictability of cash needs encourages consumers to keep cash, even while they maintain credit card debt. Telyukova (2013) calibrates a similar model that allows for a division into goods paid for with cash and goods that can be paid for with credit. A similar approach is used by Masters and Rodríguez-Reyes (2005) to explain why countries with similar levels of technology can have very different credit card use and acceptance rates.

Yet the merchant acceptance explanation actually explains a different puzzle: why anyone carries cash in the form of Federal Reserve notes. Cash pays no interest, can be stolen, and must be replenished, and so is dominated as a payment mechanism by credit cards, which have limited liability if stolen, warranty protections, and often offer rewards of some sort. It makes sense that consumers still carry cash in their pockets, along with credit or debit cards, on the chance that they will want to consume from a merchant who does not take a credit card.

Another explanation is needed to explain the portfolio decision, however. That decision takes place when the household decides not to pay as much of the credit card bill as it can. Figure 1 shows how payment methods and the credit card puzzle have evolved over time. The fraction of households in the credit card puzzle has been approximately constant since 1992. During the same time the acceptance of credit cards and their use in transactions has increased substantially. While in 1992 few grocery stores accepted credit cards, now almost all do (Evans and Schmalensee 2005). Electronic payments (whether credit or debit) grew from approximately 22 percent of non-cash transactions per person in 1995 to 67 percent of non-cash transactions in 2006. Moreover,

the ability to get a cash advance on a credit card removes any need to keep money in a savings account to pay merchants who will not accept credit. While cash advances are not often used compared to regular credit (they are more expensive, and represent about 0.8 percent of credit card debt by value), Gerdes (2008) points out that they are likely used for emergencies since the average cash advance is much higher than the average ATM withdrawal. The checks sent by credit card companies to use for cash advances even suggest their use for emergencies. That the credit card puzzle has held constant at the same time as the precautionary need based on transactions suggested by Telyukova and Wright (2008) and Telyukova (2013) has largely disappeared strongly suggests that another explanation is needed. The Telyukova and Wright (2008) explanation is consistent with other trends in payments use, such as the precipitous decline in traveler's checks and what may be a decline in the amount of cash people are carrying as merchant acceptance for electronic payments grows.[22]

Bankruptcy. Trying to use the bankruptcy system, which protects some assets, may help explain the behavior of some borrowers (Lehnert and Maki 2002). It seems likely that at least some households that expect to go bankrupt are preparing their assets to do so. Bankruptcy is still relatively uncommon (Lawless 2013) and so it seems that this explanation cannot explain more than a small fraction of households. The vast majority of households that are borrowing and saving will not go bankrupt soon enough to make paying the extra interest worthwhile. Moreover, there was a large change in bankruptcy laws in 2005 that made it substantially harder to discharge personal debts while keeping other assets (DeLaurell and Rouse 2006). Bankruptcy filings were much higher than normal in the year before the change, as forward-looking consumers rushed to file under the old regime, and much lower afterwards (Lawless 2013). Fewer assets are protected under the new regime so there is less reason to borrow and save. Yet the fraction of households borrowing and saving at the same time barely changed from 2001 to 2004 when it should have been increasing, and it did not fall from 2004 to 2007 when it should have been decreasing if bankruptcy were an important reason to borrow and save (see figure 1).

---

[22]The Survey of Consumer Payments Choice tracks the use of cash for the last several years (Foster, Schuh, and Zhang 2013). The incomplete data that exist for the long term suggest that cash use is falling (Gerdes 2008).

Self-control. Self-control issues are a different possible explanation. Hyperbolic discounting can explain a separate puzzle: why credit card borrowers also hold illiquid assets as a self-control mechanism (Laibson, Repetto, and Tobacman 2000), but it has difficulty explaining the short-term portfolio decision between assets that are comparably liquid. Haliassos and Reiter (2007) and Bertaut, Haliassos, and Reiter (2009) propose a model in which one portion of the household, the accountant, attempts to control another part, the shopper, by limiting the credit line available. While such preferences could coexist within one person, they are more likely in a household that has to make joint decisions and so must reconcile potentially different preferences. Table 5 compares the marginal probability of borrowing and saving when there is a spouse in the household. Households with a financially integrated partner are about 2 percent more likely to borrow and save, but, conditional on borrowing, such households are not much more likely to borrow and save than just to borrow. Such dual-preference self-control issues may play a role, but it seems unlikely that they explain the puzzle for most households.

6 What happens when credit limit uncertainty increases?

This section examines the portfolio and consumption dynamics following an increase in the risk of losing access to credit. Consumers who are both borrowing and saving at the same time are sensitive to the risk of losing credit, since that is why they keep liquid savings in the first place. Figures 1 and 2 show that at the end of 2008 the likelihood of losing credit increased substantially. Moreover, overall limits declined by at least a quarter over 2009, as is clear in figure 4. Using the calibrated model, this section examines the impact of these changes.

Figure 10 shows several different paths of adjustment. First, it examines the path of debt and savings following a permanent increase in the probability of losing access from the baseline of 2.67 percent to 4.2 percent as occurred at the end of 2008. Over the long term that means a move from the vertical dashed line in figure 8 to the right edge of the graph. The solid line in figure 10 shows the path in between the two long-term distributions. The second path in figure 10 shows the debt and savings if the increase lasts for only four quarters but consumers treat the increase as

permanent when it occurs. The third and fourth paths allow for a decrease in the credit limit of all consumers by 25 percent as well as an increase in the probability of losing credit.

Given the calibration in section 4, an increase in the probability of losing access has only a very small effect on consumption (note the smaller scale of the consumption panel that is necessary to show the effect at all). Instead almost all of the effect is on the distribution of credit and savings. Since relatively few people are constrained, consumers can immediately rebalance their portfolios. Savings and debt immediately increase, while consumption shifts only slightly to reflect the additional carrying cost of debt. The rebalancing is reversed immediately if the probability falls again as it does in the second scenario.

Including a decrease in credit limits has a much more pronounced effect. Lower available credit effectively lowers the precautionary wealth of all consumers, while increasing the probability of losing access makes the wealth less valuable. Lower wealth cannot be solved by a simple one-period rebalancing. Instead, consumers want to acquire additional wealth, which requires lower consumption to increase savings and decrease debt. Over time, as they acquire additional savings and hold less debt, consumption increases again. Over the long term, since consumers hold less debt, consumption may actually increase. This path is the reverse of what happens following an increase in credit limits (Fulford 2013a). The effects of a decrease in the credit limit and an increase in the probability of losing credit limits somewhat offset each other.

The last scenario fits the evolution of credit over the period the most closely: there is a brief increase in the probability of losing access accompanied by a permanent reduction in all credit limits. The decline in credit limits prompts consumption to fall by about 2 percent, and over time for debt to fall and savings to slightly increase. Although initially more people borrow and save at the same time, that effect goes away once the increase in the probability of losing access does, leaving the fraction falling slowly as savings increase. Qualitatively, these changes seem to describe the evolution of credit and wealth based on the 2007–2009 SCF panel.

6.1 Understanding changes between 2007 and 2009

The paths in figure 10 help to explain the changes in household finances following the 2008 crisis. The Survey of Consumer Finances reinterviewed the same households in 2009 as in 2007. While 42 percent of households had credit limits that were lower by $1,000 or more in 2009 than in 2007, 40 percent reported credit limits that were higher by more than $1000. Yet the responses of these two groups are startlingly different when their credit limits change. Table 6 shows summary statistics for each group in 2007 and 2009. In 2007, other than their credit limits and debt, these groups are quite similar; they had a similar likelihood to have been turned down for credit, similar incomes, and similar amounts of savings. Those who lost credit were slightly older. For those who lost credit the decrease was from an average credit limit of $43,000 to less than $20,000, while those who gained went from $22,000 to $41,000. The groups were largely similar beforehand and average incomes barely changed. Afterward, the two groups have switched places in terms of savings if they do not hold debt, debt if they hold it, and credit limit ratios.

Losing or gaining credit is closely related to increases and decreases in credit card debt. Those whose credit limits increased also increased their borrowing substantially, while those whose credit limits decreased also decreased their borrowing. What is not obvious from the table is that nearly 14 percent of households in 2007 had credit limits in 2009 below their credit card debt in 2007. About 60 percent of these households had credit limits cut to below $1000 and so were largely if not entirely cut off from credit. The households whose credit limit was cut below their debt are responsible for the vast majority of deleveraging: their average debt drops from more than $14,000 to about $2,700, while it barely changes for other households who also lost credit. These households are also responsible for the large drop in the fraction of households that are borrowing and saving at the same time among those who lost credit; many of them have completely paid back any debts and can no longer borrow substantial amounts.

These results suggest that credit variability is key to understanding the debt and savings decisions of households. The changes from 2007 to 2009 are more consistent with the permanent reduction in available credit for a large portion of the population as in the last path in figure 10,

36

although clearly there is a great deal of heterogeneity. The group that lost credit was not poorer on average as table 6 shows, but they had large declines in credit limits, while others had substantial gains. Moreover, a substantial portion of the population was cut off entirely, and so benefited directly from having liquid savings. That suggests that just as cyclical declines in income are not distributed evenly across the population, decreases in credit limits are concentrated as well.

7   Conclusion

Consumer credit variability is very important for household decisions. Credit volatility is large in both the short and long term. Credit volatility matters since American households depend on their credit—it is the largest immediately available resource available for most households—and borrow a lot. Using a model that can allow for credit limit changes, this paper shows that uncertain credit limits affect the ways that households save, borrow, and consume. For example, the so-called credit card "puzzle" is only a puzzle because of the modeling assumption that consumers face a simple world where everything except their incomes is certain. The model also offers a way that financial uncertainty can affect household decisions. Financial crises matter even for households that are protected by deposit insurance: one of their largest assets for consumption smoothing is the ability to borrow and there are no guarantees that it will be there when they need it.

As credit has become a more important tool for households and debt has increased, there has been too little examination of the volatility of credit and the consequences of that volatility for household decisions. There is increasing regulatory attention over consumer uses of credit in the United States. Yet it is impossible to design good regulations and informational campaigns without understanding why and how consumers use credit.[23] Moreover, the common focus on the interest

---

[23]For example, the fourth of five tips for "Getting the most from your credit card" in the consumer information section of the Federal Reserve Board of Governors website suggests "If you can't pay your balance in full each month, try to pay as much of the total as you can. Over time, you'll pay less in interest charges–money that you will be able to spend on other things, and you'll pay off your balance sooner" (http://www.federalreserve.gov/consumerinfo/fivetips_creditcard.htm, accessed 7 May 2010). If consumers keep liquid savings as a precautionary measure while borrowing to fund current consumption, then telling them to pay down expensive credit is bad advice, even if well intentioned. Interestingly, the financial gurus appear to understand the risk. Dave Ramsey's seven steps to financial peace suggests building a $1000 emergency fund first, then paying off debt as step two (Ramsey 2010). That advice follows the clear hierarchy of payments shown in the model.

rate may be mistaken. The most important welfare issue around credit may not be how much it costs, but whether it is available when households need it. That many households turn to extremely expensive services, such as payday lending, that are readily available when needed may indicate not poor financial management or self-control issues, but that less expensive lines of credit are not always available.

References

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. 2001. "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *The Journal of Economic Perspectives* 15 (3):pp. 47–68.

Bertaut, Carol C., Michael Haliassos, and Michael Reiter. 2009. "Credit Card Debt Puzzles and Debt Revolvers for Self Control." *Review of Finance* 13 (4):657–692.

Bucks, Brian K., Arthur B. Kennickell, Traci L. Mach, and Kevin B. Moore. 2009. "Changes in U.S. Family Finances from 2004 to 2007: Evidence from the Survey of Consumer Finances." *Federal Reserve Bulletin* 95:A1–A56. Available: http://www.federalreserve.gov/pubs/bulletin/2009/pdf/scf09.pdf, accessed 2 May 2010.

Carroll, Christopher D. 1997. "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112 (1):1–55.

———. 2001. "A Theory of the Consumption Function, with and without Liquidity Constraints." *Journal of Economic Perspectives* 15 (3):23–45.

———. 2004. "Theoretical Foundation of Buffer Stock Saving." NBER Working Paper No. 10867.

Carroll, Christopher D. and Miles S. Kimball. 2001. "Liquidity Constraints and Precautionary Saving." NBER Working Paper No. W8496.

Carroll, Christopher D. and Andrew A. Samwick. 1997. "The nature of precautionary wealth." *Journal of Monetary Economics* 40 (1):41–71.

Chow, Gregory C. 1997. *Dynamic Economic: Optimization by the Lagrange Method*. New York: Oxford University Press.

Collins, Darly, Jonathan Morduch, Stuart Rutherford, and Orlanda Ruthven. 2009. *Portfolios of the Poor*. Princeton: Princeton University Press.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica* 59 (5):1221–1248.

Deaton, Angus and Guy Laroque. 1992. "On the Behavior of Commodity Prices." *Review of Economic Studies* 59 (1):1–23.

DeLaurell, Roxane and Robert Rouse. 2006. "The Bankruptcy Reform Act of 2005: A New Landscape." *The CPA Journal* Available: http://www.nysscpa.org/cpajournal/2006/1106/essentials/p36.htm (accessed 28 April 2012).

Edelberg, Wendy. 2006. "Risk-based pricing of interest rates for consumer loans." *Journal of Monetary Economics* 53 (8):2283 – 2298.

Evans, David S. and Richard Schmalensee. 2005. *Paying with Plastic: The Digital Revolution in Buying and Borrowing*. Cambridge: MIT Press, 2nd edition ed.

Federal Reserve Board. 2008. "Senior Loan Officer Opinion Survey on Bank Lending Practices." Available http://www.federalreserve.gov/boarddocs/SnLoanSurvey/200811/fullreport.pdf, accessed 30 April 2010.

———. 2010. "Senior Loan Officer Opinion Survey on Bank Lending Practices." Available http://www.federalreserve.gov/boarddocs/SnLoanSurvey/201002/fullreport.pdf, accessed 30 April 2010.

Foster, Kevin, Scott Schuh, and Hanbing Zhang. 2013. "The 2010 Survey of Consumer Payment Choice." Research Data Reports 13-2, Federal Reserve Bank of Boston. URL http://www.bostonfed.org/economic/rdr/2013/rdr1302.pdf.

Fulford, Scott. 2013a. "The effects of financial development in the short and long run: Theory and evidence from India." *Journal of Development Economics* 104:56–72.

———. 2013b. "The Preaution of the Rich and Poor." Working Paper 814, Boston College. URL http://fmwww.bc.edu/EC-P/wp814.pdf.

Gerdes, Geoffrey R. 2008. "Recent payment trends in the United States." *Federal Reserve Bulletin* 94:A75–A106.

Gourinchas, Pierre-Olivier and Jonathan A. Parker. 2002. "Consumption Over the Life Cycle." *Econometrica* 70 (1):47–89.

Gross, David B. and Nicholas S. Souleles. 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics* 117 (1):149–185.

Haliassos, Michael and Michael Reiter. 2007. "Credit Card Debt Puzzles." CFS Working Paper No. 2005/26.

Hartley, Peter R. 1996. "Value function approximation in the presence of uncertainty and inequality constraints: An application to the demand for credit cards." *Journal of Economic Dynamics and Control* 20:63–92.

Jambulapati, Vikram and Joanna Stavins. 2013. "The Credit CARD Act of 2009: What Did Banks Do?" Public Policy Discussion Paper 13-7, Federal Reserve Bank of Boston.

Judd, Kenneth. 1998. *Numerical Methods in Economics*. Cambridge: The MIT Press.

Kashyap, Anil K, Owen A. Lamont, and Jeremy C. Stein. 1994. "Credit Conditions and the Cyclical Behavior of Inventories." *The Quarterly Journal of Economics* 109 (3):565–592.

Kimball, Miles S., Claudia R. Sahm, and Matthew D. Shapiro. 2009. "Risk Preferences in the PSID: Individual Imputations and Family Covariation." *American Economic Review: Papers and Proceedings* 99 (2):363–68.

Kreider, Rose M. and Diana B. Elliott. 2009. "Americas Families and Living Arrangements: 2007." US Census Bureau Current Population Survey Reports. Available: http://www.census.gov/population/www/socdemo/hh-fam.html, accessed 2 May 2010.

Laibson, David, Andrea Repetto, and Jeremy Tobacman. 2000. "A Debt Puzzle." NBER Working Paper W7879.

———. 2007. "Estimating Discount Functions with Consumption Choices over the Lifecycle." Working Paper 13314, National Bureau of Economic Research. URL http://www.nber.org/papers/w13314.

Lawless, Bob. 2013. "Bankruptcy Filing Rate Decline Is Historically Anomalous." Blog post 8 July, Credit Slips. http://www.creditslips.org/creditslips/2013/07/bankruptcy-filing-rate-decline-is-historically-anomalous.html,accessed 4 March 2014.

Lehnert, Andreas and Dean M. Maki. 2002. "Consumption, Debt and Portfolio Choice: Testing the Effect of Bankruptcy Law." Federal Reserve Board, Finance and Economics Discussion Working Paper 2002-14.

Ludvigson, Sydney. 1999. "Consumption and Credit: A Model of Time-Varying Liquidity Constraints." *Review of Economics and Statistics* 81 (3):434–447.

Masters, Adrian and Luis Raúl Rodríguez-Reyes. 2005. "Endogenous credit-card acceptance in a model of precautionary demand for money." *Oxford Economic Papers* 57 (1):157–168.

Miranda, Mario J. and Paul L. Fackler. 2002. *Applied Computational Economics and Finance*. Cambridge: MIT Press.

National Highway Trafic Safety Administration. 2009. "2008 Traffic Safety Annual Assessment Highlights." Available: http://www-nrd.nhtsa.dot.gov/Pubs/811172.PDF, accessed 2 May 2010.

Ramsey, Dave. 2010. "The Seven Baby Steps: Begin your journey to financial peace." Available http://www.daveramsey.com/new/baby-steps/, accessed 21 December 2010.

Sabelhaus, John and Jae Song. 2010. "The great moderation in micro labor earnings." *Journal of Monetary Economics* 57 (4):391 – 403.

Schechtman, Jack and Vera L. S. Escudero. 1977. "Some results on an 'Income Fluctuation Problem'." *Journal of Economic Theory* 16:151–166.

Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron. 2004. "Consumption and risk sharing over the life cycle." *Journal of Monetary Economics* 51 (3):609 – 633.

Telyukova, Irina A. 2013. "Household Need for Liquidity and the Credit Card Debt Puzzle." *Review of Economic Studies* 80 (3):1148–1177.

Telyukova, Irina A. and Randall Wright. 2008. "A Model of Money and Credit, with Application to the Credit Card Debt Puzzle." *Review of Economic Studies* 75:629–647.

Wiens, Travis. 2009. "Complex Method of Optimization." Available: http://www.mathworks.com/matlabcentral/fileexchange/25428-complex-method-of-optimization, accessed 7 May 2010.

Zinman, Jonathan. 2007. "Household Borrowing High and Lending Low Under No-Arbitrage." Available: http://www.dartmouth.edu/~jzinman/Papers/Zinman_BHLL_apr07.pdf (accessed 8 May 2012).

———. 2009. "Where is the missing credit card debt? Clues and implications." *Review of Income and Wealth* 55 (2):249–265.

A   The structure of credit card agreements

This section briefly reviews credit card account agreements. Card-issuing banks can change limits or revoke credit at any time, but the card holder is still responsible for any remaining debt. The card-issuing bank cannot demand full repayment unless the debtor is in default. The Consumer Financial Protection Board now requires credit card agreements to be posted with them. I examined numerous agreements for major card-issuing banks in June 2013 (contracts can change, but are subject to consumer credit regulation at both the state and federal level). For most of them the language is clear and unequivocal that they can cancel or reduce a credit limit at any time and that on doing so the borrower still has to repay based on the agreement. For example, the Capital One Customer Agreement states "We may also increase, decrease, restrict, or cancel your credit limit on any Segment at any time. This will not affect your obligation to pay us"

The repayment schedule is the subject of a number of regulations and in general cannot be changed unilaterally. A minimum payment is typically necessary each month, but the entire balance may become due immediately if the account falls into default: "If your Account is 180 days past due, part of a bankruptcy proceeding or otherwise charges off, the entire balance is immediately due and payable" (Capital One Customer Agreement). Similarly, the Bank of America agreement for all of their Visa, MasterCard, Preferred, Gold, and Platinum cards under the heading "When We May Require Immediate Repayment" states that they can demand payment of the outstanding balance if an account is in default. Default occurs if the borrower fails to make the "Total Minimum Payment," which is any fees, all interest, and 100 percent of the balance. The agreement states that Bank of America "may suspend or close your account or otherwise terminate your right to use your account. We may do this at any time and for any reason." Other banks use similar language. CitiBank asserts its right to reduce the credit limit at any time, and, while promising to notify you, says that the change may take place before you receive notice. The agreement says that only in default can they demand immediate payment. The Chase Bank Slate Visa-MasterCard agreement asserts they can "cancel, change or restrict your credit availability at any time," but can close the account and require full repayment only if the account is in default.

B   Simulations

The two continuous decision variable, two continuous state variable, multiple discrete state variable, infinite horizon, stochastic, optimization problem described has a number of complexities that make it difficult to simulate. The first is that the "curse of dimensionality" means that a discrete approximation of the continuous state variables that is fine enough to capture the changes in behavior in the model quickly grows in the memory required. I use cubic spines to approximate the value function, reducing the number of states necessary to achieve a good approximation.

A second difficulty is that with a quarterly frequency the discount factor is close to one (the conversion is $\beta_{\text{monthly}} = (1 + \beta_{\text{annual}})^{-1/4}$). Function iteration techniques therefore tend to converge slowly, since initial conditions matter for many iterations.

The form of the problem creates its own difficulties. The function iteration step which takes a function V_{t+1} to a function V_t:

$$V_t(w_t, b_t, B_t) = \max_{X \in \Omega(w_t, b_t, B_t)} u(X) + \beta E_t V_{t+1}(w_{t+1}(X), b_{t+1}(X), B_{t+1}),$$

where $X = (c_t^w, c_t^b)$ and $\Omega(w_t, b_t, B_t)$, is the constraint set. The consumer cares relatively little about debt versus cash consumption, but cares a great deal about their sum, total consumption. The maximization problem in the function iteration step for a given set of state variables tends to look like a ridge rather than a mountain: along a locus that keeps total consumption constant the function increases slowly to an optimum allocation of consumption to debt and cash. Stepping off from that ridge—increasing or decreasing total consumption—makes a big difference in utility, while moving along it the change in utility is relatively small. This scaling problem leads to a tendency to overshoot in one direction or another, and makes single derivative "method of steepest ascent" approaches tend to oscillate instead of converge. One solution would be to rewrite the problem so that the decision variables are total consumption C_t and a portfolio allocation c_t^w or c_t^w/C_t. The problem with this approach is that it is no longer possible to write the constraints for each variable in terms of the state alone: they depend on the other decision variable (for example, the portfolio allocation is bounded by total consumption so that it is impossible to consume more than total consumption as cash consumption). Such constraints require a different optimization approach than the standard linear complementarity problem.

A final difficulty with the optimization is that intermediate value functions are not necessarily everywhere second order differentiable. Using information on the second derivative of the objective function is standard in many optimization methods such as Newton and quasi-Newton methods (Judd 1998; Miranda and Fackler 2002), and is important in this problem because of the tendency to oscillate. Along regime changes for V_{t+1}, where the decision variables go from being constrained to unconstrained, the first derivative tends to exhibit a kink, and so the second derivative becomes discontinuous. Close to the states (w_t, b_t) for which (w_{t+1}, b_{t+1}) is at a regime change, the second derivative is very irregular since the smooth approximation of V_{t+1} cannot approximate discontinuities well. Increasing the number of nodal points for the splines does not solve the problem, but only more precisely finds the states (w_t, b_t) with the problem. The ridge nature of the problem makes method of steepest ascent (Judd 1998, p. 111) frequently fail to converge, and so, after allowing convergence as far as possible, I switch to a derivative-free method. The options available for derivative free optimization of a continuous problem with multiple dimensions and constraints tend to be slow, since they cannot rely on any information about curvature to determine the step size. I use the "Complex" method implemented by Wiens (2009) which is the linear complementarity version of the simplex method for multiple dimensions.

Table 1: Quarterly volatility in consumer credit limits

| | Open credit card accounts | | | | | Credit limits | | | | | |
| | Mean | Median | Percent of individuals with change: | | | Percent of individuals with change: | | | Change in limit (%) | | |
			Neg.	Zero	Pos.	Neg.	Zero	Pos.	if neg.	if pos.	S.D.
Overall	2.17	2	7.75	84.44	7.82	12.02	67.56	20.41	-36.43	28.68	30.20
Excellent Credit	2.55	2	7.65	85.37	6.98	12.73	66.72	20.55	-29.44	26.86	25.66
Good Credit	2.56	2	7.14	82.21	10.65	12.29	60.72	26.99	-31.04	30.01	30.62
Fair Credit	1.90	1	6.06	84.34	9.60	9.85	68.00	22.15	-35.89	31.55	32.58
Poor Credit	1.17	0	8.97	83.87	7.16	11.15	73.51	15.34	-57.35	31.68	41.34
Age < 20	0.70	1	3.23	85.53	11.24	5.10	73.74	21.15	-72.68	56.02	42.30
Age 20-30	1.44	1	6.15	85.13	8.71	9.11	69.38	21.51	-48.47	36.27	36.26
Age 30-40	2.05	1	8.02	83.16	8.82	12.03	65.66	22.31	-39.30	29.78	32.74
Age 40-50	2.42	2	8.62	82.67	8.71	13.37	64.02	22.62	-34.54	26.94	30.18
Age 50-60	2.66	2	8.67	83.10	8.22	13.82	64.20	21.98	-31.71	25.63	27.93
Age 60-70	2.63	2	8.24	84.48	7.27	13.39	66.96	19.65	-31.14	25.99	26.53
Age >70	1.78	1	6.47	88.92	4.61	9.98	77.06	12.96	-38.27	28.40	26.73

Notes: Sample is 0.1 percent of all individuals in the United States with an open credit card account at any time between 1999 and 2013. Credit limit increases greater than 200 percent are set to 200 percent . Source: Equifax and the Federal Reserve Bank of New York Consumer Credit Panel.

Table 2: Credit and financial decision making in the Consumer Finance Monthly

A: Credit volatility	(%)
Report a bank closed a credit card account in last 12 months	5.4
Report respondent closed a credit card account in last 12 months	14.9
If the account holder closed an account, why?	
You did not use that credit card	27.8
You charged too much on the card	7.7
You were worried that you would charge too much	2.2
The bank or card issuer changed the terms of that card	13.3
You closed your account for some other reason	49.0

B: Financial decision making

Question: Using extra money in a bank savings account to pay off high interest credit card debt is a good idea.

	Don't know	True	False
Percent of respondents	12.0	66.9	21.0
Percent both debt >0 & savings>0	22.3	39.0	38.8

Question: To reduce the total finance costs paid over the life of an auto loan, you should choose a loan with the . . .

	Don't Know	Lowest monthly payment	Longest repayment term	Shortest repayment term
Percent of respondents	14.3	6.3	3.8	75.7
Percent both debt >0 & savings>0	25.1	36.4	32.0	30.2

Question: Which policy provides the most coverage at the lowest cost for a young family . . .

	Don't Know	Renewable term life	Whole life	Universal life
Percent of respondents	35.4	35.3	17.0	12.2
Percent both debt >0 & savings>0	27.6	29.5	32.5	34.9

Notes: The proportions with savings and debt positive are based on whether the respondent revolves credit card debt and has a positive balance in the savings and checking account and so is most comparable to definition 4 in table 4 from the SCF. The overall percentages are 43 percent positive debt-positive savings, 7.9 percent positive debt-no savings, 45 percent positive savings-no debt, 4.2 percent no debt-no savings and so are very comparable to the SCF. Source: The Consumer Finance Monthly.

Table 3: Quarterly estimates of credit limit volatility

Coefficient	Estimate	S.E
Probability lose access p_{BN}	0.0267	0.00008
Probability gain access p_{NB}	0.0603	0.00018
Non-reported cards ρ	0.619	0.0039
Credit risk	4.31E-03	1.03E-04
(Credit risk)^2	1.84E-06	7.86E-08
Permanent variance σ_v^2	0.247	0.0005
Transitory variance σ_w^2	0.248	0.0016
Observations	5963594	
Individual accounts	126101	

Notes: Estimates of equation 1. Also included in the level of credit regressions are age and quarter effects shown in figure 4 as well as ZIP code indicators (not reported). Sample is 0.1 percent of all individuals in the United States with an open credit card account at any time between 1999 and 2013. Source: Author's calculations from Equifax and the Federal Reserve Bank of New York Consumer Credit Panel.

Table 4: Credit Card Puzzle in the 2007 SCF

Savings definition		Full sample				Age≤50
		[1]	[2]	[3]	[4]	[1]
Both debt >0 & savings>0	(%)	41.4	36.6	24.2	58.5	33.3
Debt >0, no savings	(%)	18.8	23.7	36.0	1.7	16.0
Savings >0, no debt	(%)	32.3	25.2	28.0	39.3	40.2
No debt, no savings	(%)	7.5	14.6	11.8	0.5	10.5
Savings/quarterly income	(mean)	0.93	0.66	0.87	1.08	1.50
	(median)	0.07	0.01	0.00	0.26	0.17
Debt/quarterly income	(mean)	0.22	0.22	0.22	0.22	0.24
	(median)	0.00	0.00	0.00	0.00	0.00
Debt/inc if debt>0 & savings>0	(mean)	0.52	0.49	0.52	0.51	0.62
	(median)	0.20	0.18	0.17	0.20	0.22
Debt/inc if debt>0 & no savings	(mean)	0.48	0.54	0.50	0.33	0.60
	(median)	0.22	0.23	0.23	0.18	0.24
Savings/inc if debt>0 & savings>0	(mean)	0.86	0.74	1.32	0.77	1.41
	(median)	0.21	0.18	0.39	0.26	0.31
Savings/inc if savings>0 & no debt	(mean)	1.95	1.70	2.43	1.56	2.73
	(median)	0.58	0.46	0.89	0.42	0.98
Credit limit/inc	(mean)	1.37	1.37	1.37	1.37	1.90
	(median)	0.53	0.53	0.53	0.53	0.86
Credit used (%)	(mean)	23.3	23.3	23.3	23.3	16.0
	(median)	5.6	5.6	5.6	5.6	0.0
Carry cost of extra debt (% of inc)	(mean)	0.60	0.53	0.80	0.65	0.76
	(median)	0.24	0.22	0.33	0.33	0.30
Available resources credit (%)	(mean)	60.0	67.1	69.6	42.0	60.3
	(median)	73.9	87.7	92.3	45.2	73.0

Notes: The savings definitions are: [1] Savings + (Checking - Income) if (Checking-Income)>0; [2] Savings account only; [3] Savings + Checking - Income; [4] Savings + Checking. Income is monthly income in the savings definition (to match the credit card bill), but the Savings/Income ratios are multiples of *quarterly* income to match the simulations and panel data. The last column restricts the sample to households whose head is age 50 or under. Savings and debt are greater than zero if they are greater than 0.01 percent of income. Includes only households with credit cards. All calculations are survey weighted. Source: the 2007 Survey of Consumer Finances.

Table 5: Logit marginal effects of both simultaneous borrowing and saving

| | Marginal probability | | If Debt> 0 |
	Debt & save > 0	Debt> 0	Debt & save > 0
Age/10	-0.176***	-0.252***	-0.0166
	(0.0148)	(0.0135)	(0.0206)
$(Age/10)^2$	-0.272***	-0.588***	0.203***
	(0.0445)	(0.0396)	(0.0561)
Size of household	0.0145***	0.0283***	-0.00631
	(0.00339)	(0.00348)	(0.00425)
Spouse/partner in household	0.0213**	0.0109	0.0148
	(0.0104)	(0.0101)	(0.0140)
Log normal income	-0.0427***	-0.128***	0.0850***
	(0.00562)	(0.00545)	(0.00922)
Own current home?	-0.0361	-0.0581**	0.0109
	(0.0262)	(0.0231)	(0.0368)
When making investment decisions	-0.00743**	-0.0186***	0.00918**
shop around for best terms	(0.00289)	(0.00280)	(0.00387)
Attitude to financial risk/returns	-0.00733	0.0222***	-0.0347***
	(0.00518)	(0.00494)	(0.00700)
log target emergency savings/income	-0.0257***	-0.0502***	0.0138***
	(0.00323)	(0.00305)	(0.00431)
Good idea to buy with credit?	-0.0171***	-0.0210***	-0.00507
(1 good, 3 , 5 bad)	(0.00262)	(0.00250)	(0.00345)
Is it all right to borrow money to			
—pay for a vacation?	0.0651***	0.103***	-0.00769
	(0.0111)	(0.0115)	(0.0139)
—cover living expenses	0.0186**	0.0481***	-0.0233**
when income is cut?	(0.00804)	(0.00770)	(0.0107)
—purchase a fur coat or jewelry?	0.0207	0.0286*	0.00444
	(0.0161)	(0.0167)	(0.0204)
—purchase a car?	0.0778***	0.0634***	0.0552***
	(0.0124)	(0.0111)	(0.0160)
—finance education expenses?	0.0173	0.0101	0.0203
	(0.0124)	(0.0113)	(0.0164)
Year effects	YES	YES	YES
Observations	21,887	21,887	9,782

Notes: The marginal probabilities from a multinomial (non-ordered) logistic regression of those with credit cards using the SCF from 1995 to 2010 (except the 2009 panel). The first column shows the change in the marginal probability of being in (1) positive debt, positive savings, where the other options are (2) positive debt, no savings, (3) positive savings, no debt, (4) no savings, no debt. The second column shows the marginal probability of being in either (1) or (2); the third, the marginal probability of being in (1) conditional on being in (1) or (2). The shop around for financial terms varies from 1 (a little) to 5 (a great deal). The attitude to financial risk varies from 1 (low risk/low returns) to 4 (high risk/high returns). The regression uses only a single imputation.

Table 6: Changes between 2007 and 2009

	Lose Credit		No Change		Gain Credit	
	2007	2009	2007	2009	2007	2009
Frac. of all households	0.42		0.18		0.40	
Frac. both debt >0 & savings>0	0.44	0.30	0.36	0.31	0.44	0.41
Frac. debt >0, no savings	0.18	0.12	0.27	0.16	0.17	0.16
Frac. savings >0, no debt	0.30	0.45	0.30	0.41	0.33	0.38
Frac. no debt, no savings	0.07	0.14	0.07	0.13	0.05	0.05
Debt/inc if debt>0 & savings>0	2.85	1.57	0.90	0.92	1.35	2.04
Debt/inc if debt>0 & no savings	2.19	1.90	1.01	1.42	1.23	2.84
Savings/inc if debt>0 & savings>0	3.12	2.64	2.06	2.29	2.04	2.09
Savings/inc if savings>0 & no debt	7.00	4.87	5.32	4.92	5.56	7.47
Credit card debt	6,378	3,508	2,453	2,234	3,925	6,335
Credit card limit	43,194	19,748	16,425	16,293	21,550	40,495
Credit card limit/inc	9.48	3.99	3.52	3.70	4.19	8.62
Checking account	10,819	7,508	7,258	5,564	10,561	7,831
Savings account	28,044	29,729	17,221	26,594	27,714	29,792
Age of head	52.59	54.57	48.82	51.12	48.97	51.08
Log income	8.51	8.51	8.26	8.32	8.56	8.65
Been turned down for credit?	0.18	0.12	0.22	0.19	0.16	0.12
Fraction resources credit	0.75	0.58	0.65	0.57	0.65	0.76

Notes: Losing or gaining credit is reporting a credit limit that was more than $1000 larger or smaller in 2009 than in 2007. The savings definition for the savings and debt-to-income ratios is [1] Savings + (Checking - Income) if (Checking-Income)>0. Checking and savings accounts are as reported in the survey. Income is converted to monthly, so ratios are the number of months of income. Whether a household has been turned down for credit means in the last five years in 2007, while in the last two years in 2009, and is included to compare across populations, not time. The fraction of resources credit calculates how much of the resources immediately available to the household is credit: (credit limit - credit debt)/(savings (using definition [1]) + credit limit - credit debt).

Figure 1: Trends in use of electronic payment methods and household credit use
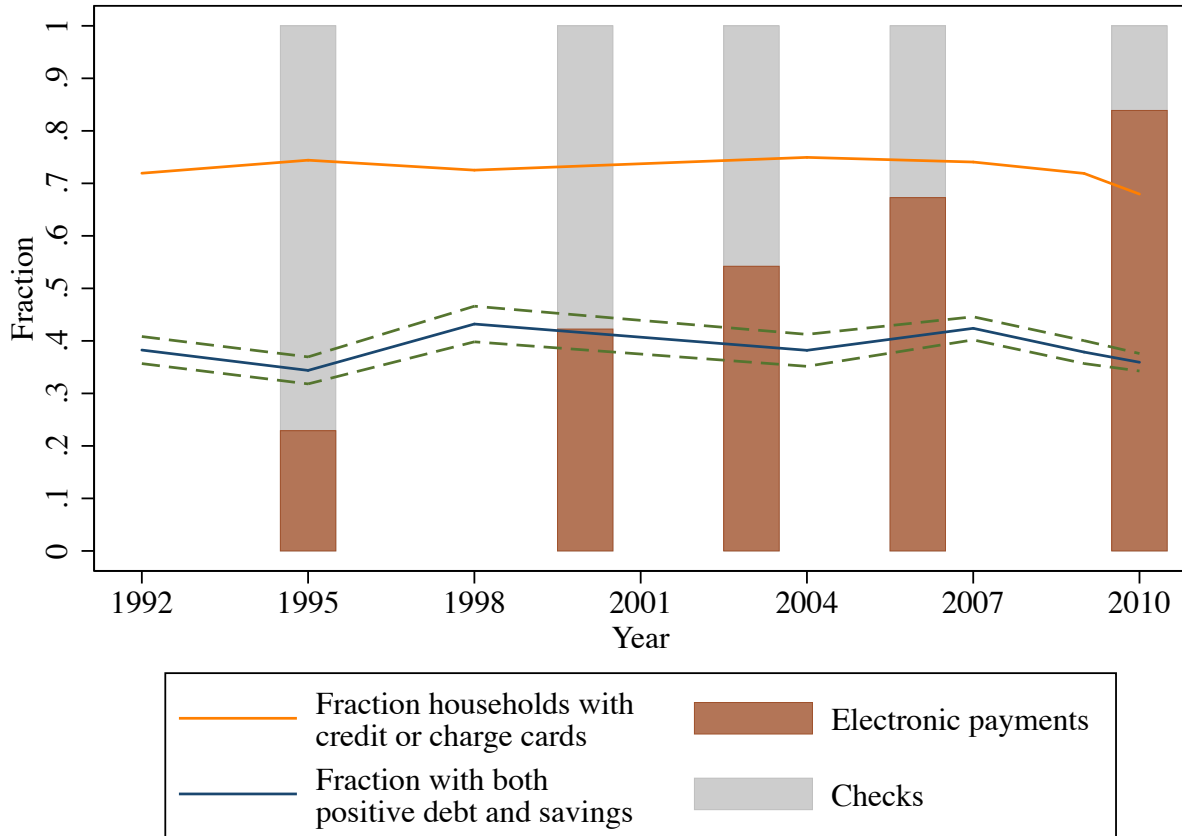
Notes: The left axis shows the fraction of households who hold a charge or credit card, and the fraction that maintain both positive debt and savings balances (using method 1 defined in column 1 of table 4 and section 4.1). 95 percent confidence intervals (accounting for multiple imputation by the SCF) are shown by dashed lines. The right axis shows the number of non-cash payments per person each year made using electronic payment methods (credit cards, debit cards) and checks. Source: SCF, Gerdes (2008) and Foster, Schuh, and Zhang (2013) for the 2010 payments. The 2010 value is from the Survey of Consumer Payment Choice and likely captures a much larger number of payments than the 2007 and earlier numbers.

Figure 2: Quarterly credit card account closures and openings over time
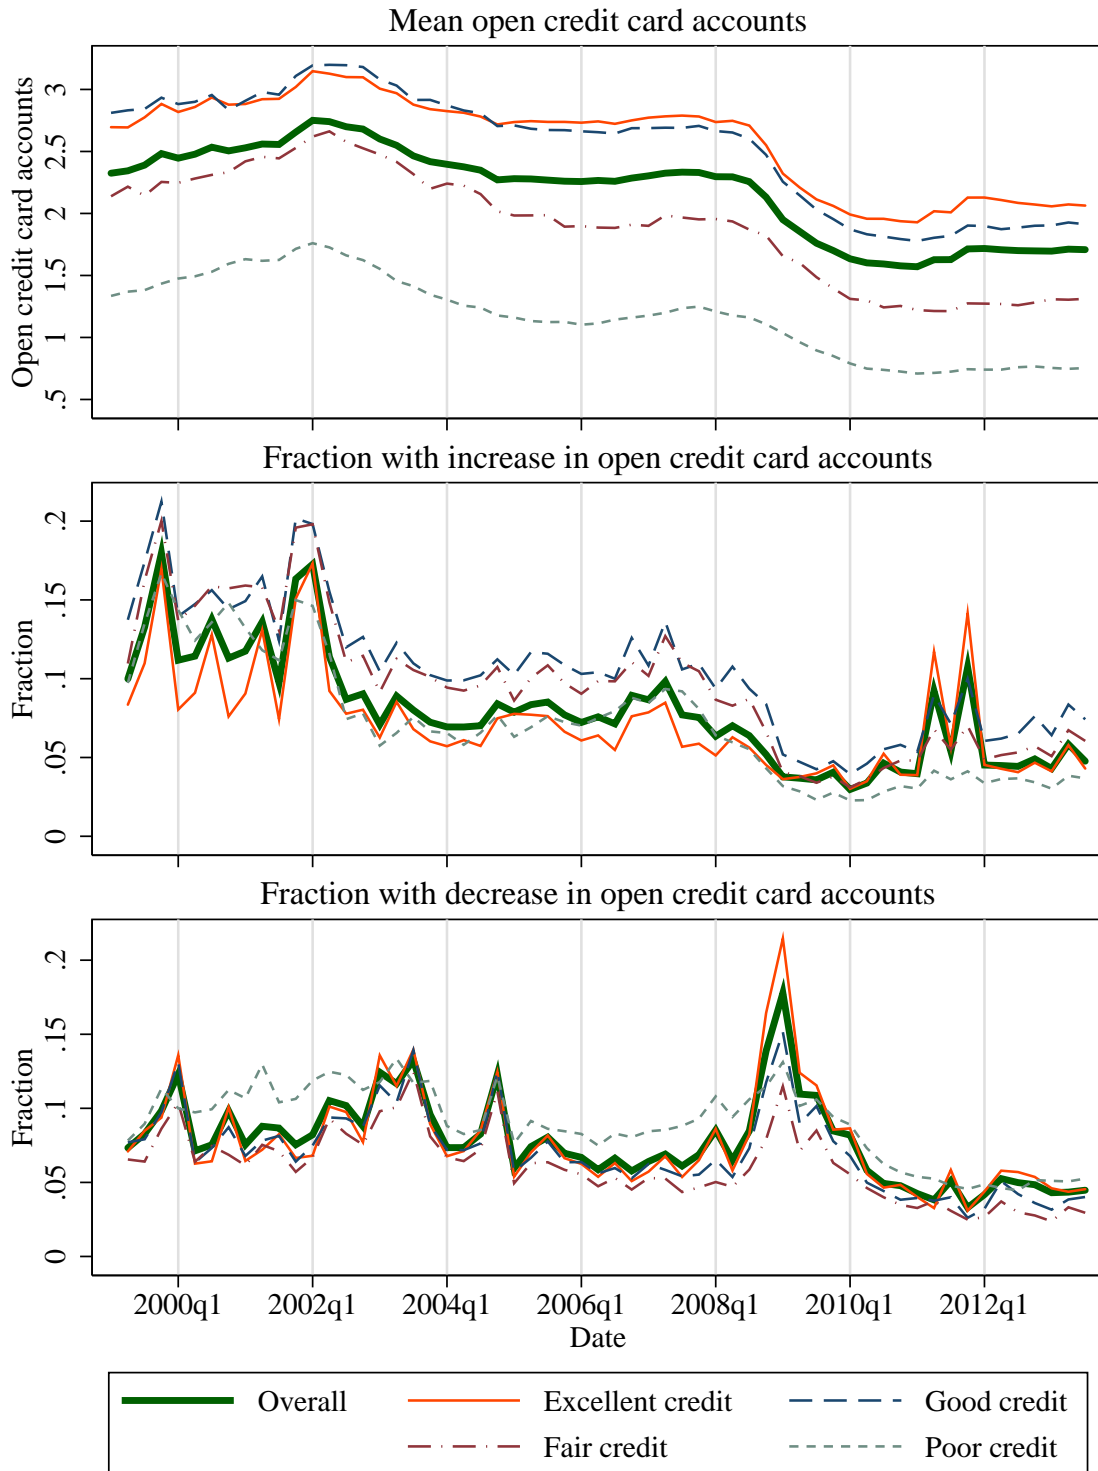
Notes: The open credit card accounts is for individuals who have ever had an open account from 1999 to 2013 and so includes many people who do not currently have a card. Source: Equifax and the Federal Reserve Bank of New York Consumer Credit Panel.

Figure 3: Quarterly probability of gaining and losing credit
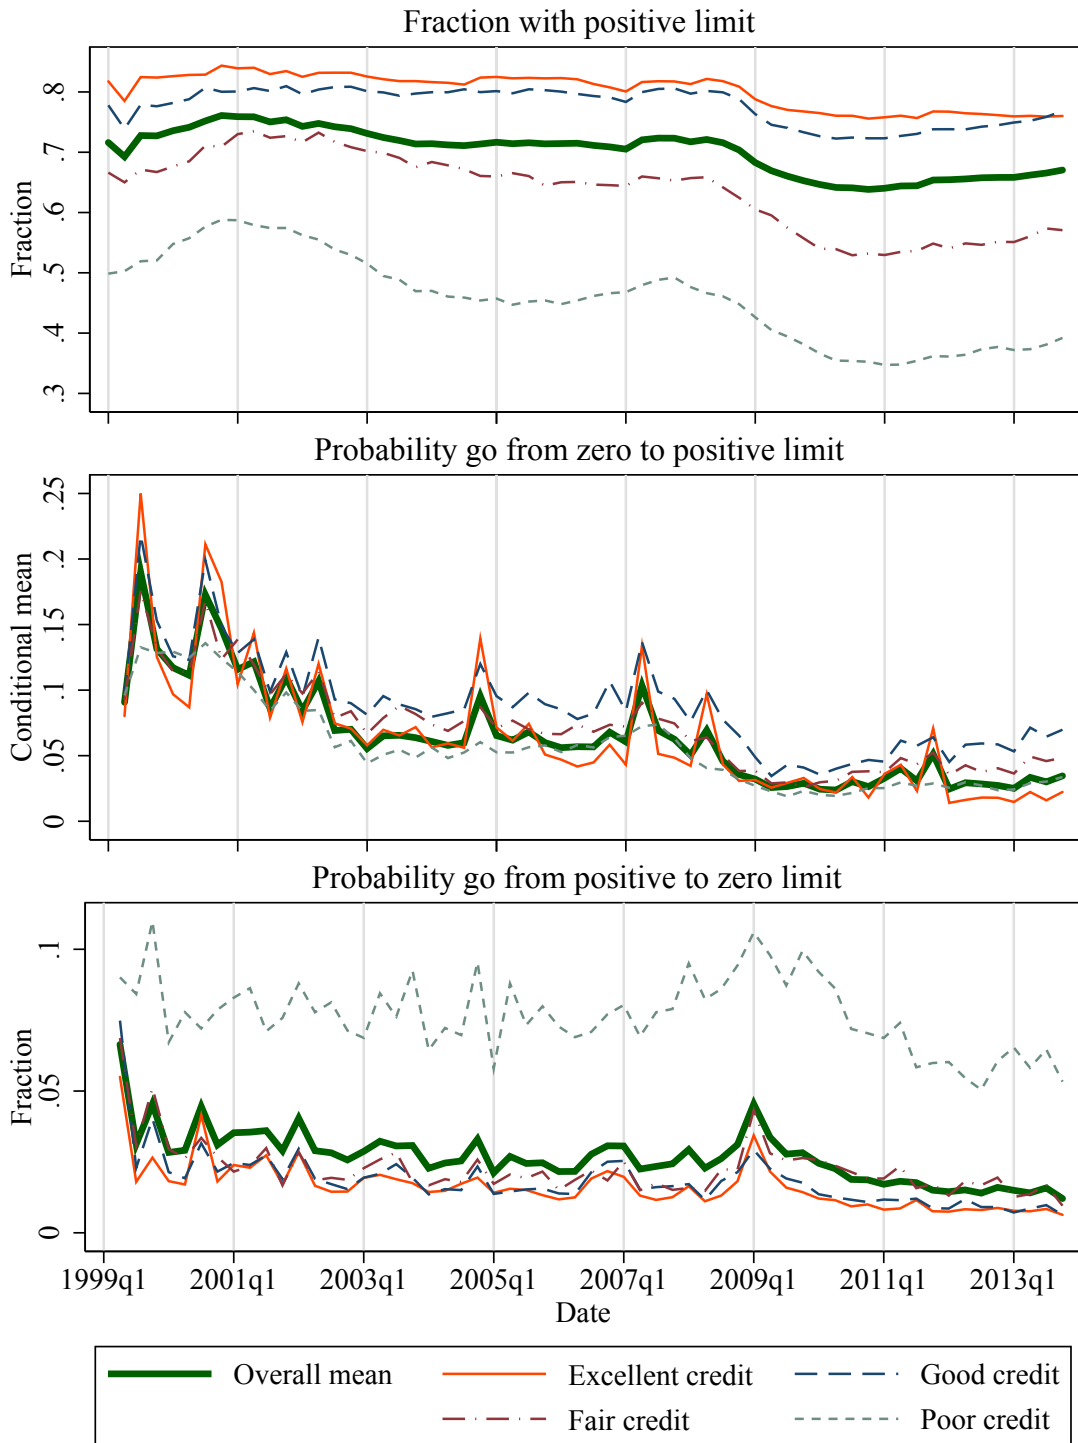
Notes: The fraction with positive limit includes all individuals who have ever had an open account from 1999 to 2013 and so includes many people who do not have a card in any given quarter. Source: Equifax and the Federal Reserve Bank of New York Consumer Credit Panel.

Figure 4: Credit limits by age and time from estimation
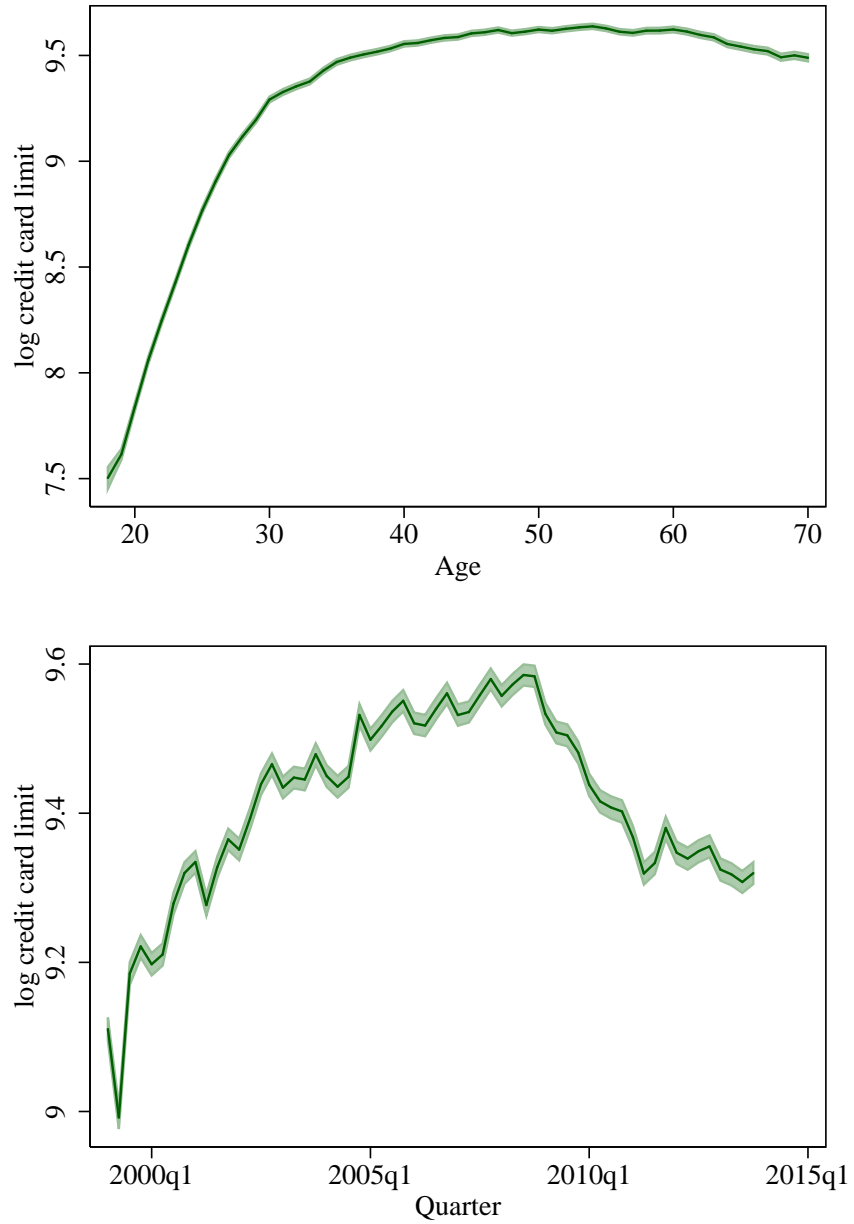
Figure 5: Possible timing of consumption, income, and credit card payments

(A) Model timing
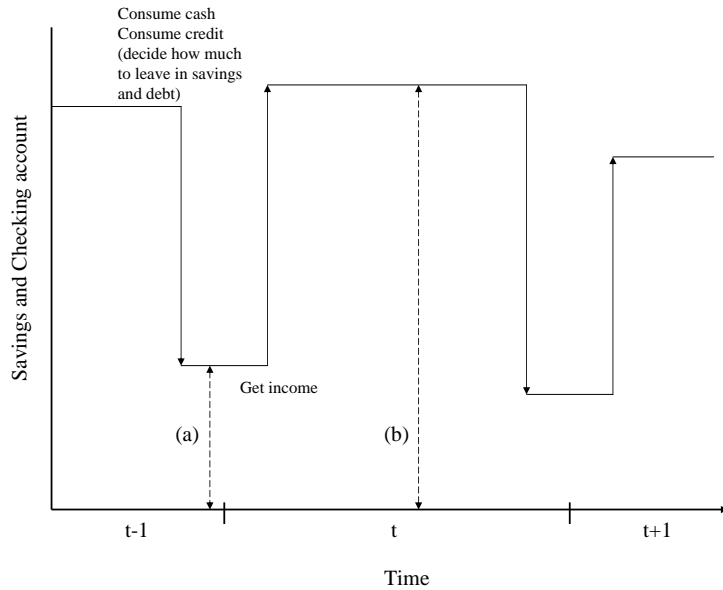
(B) An alternative timing
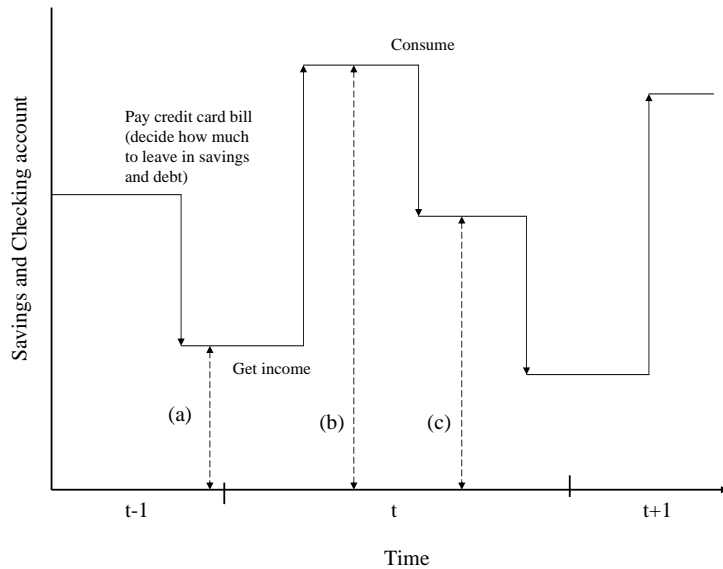
Notes: In both panels (a) is the amount left over after the credit card bill has been paid (possibly not completely leaving debt for the next period) but before income or any consumption. The model timing has income coming at the beginning of the period, then consumption and the savings/debt decision at the end. In reality, income may come at any time in during the month, and consumption need not take place at the same time as the credit card bill is paid. We may observe (b) or (c) instead, of (a).

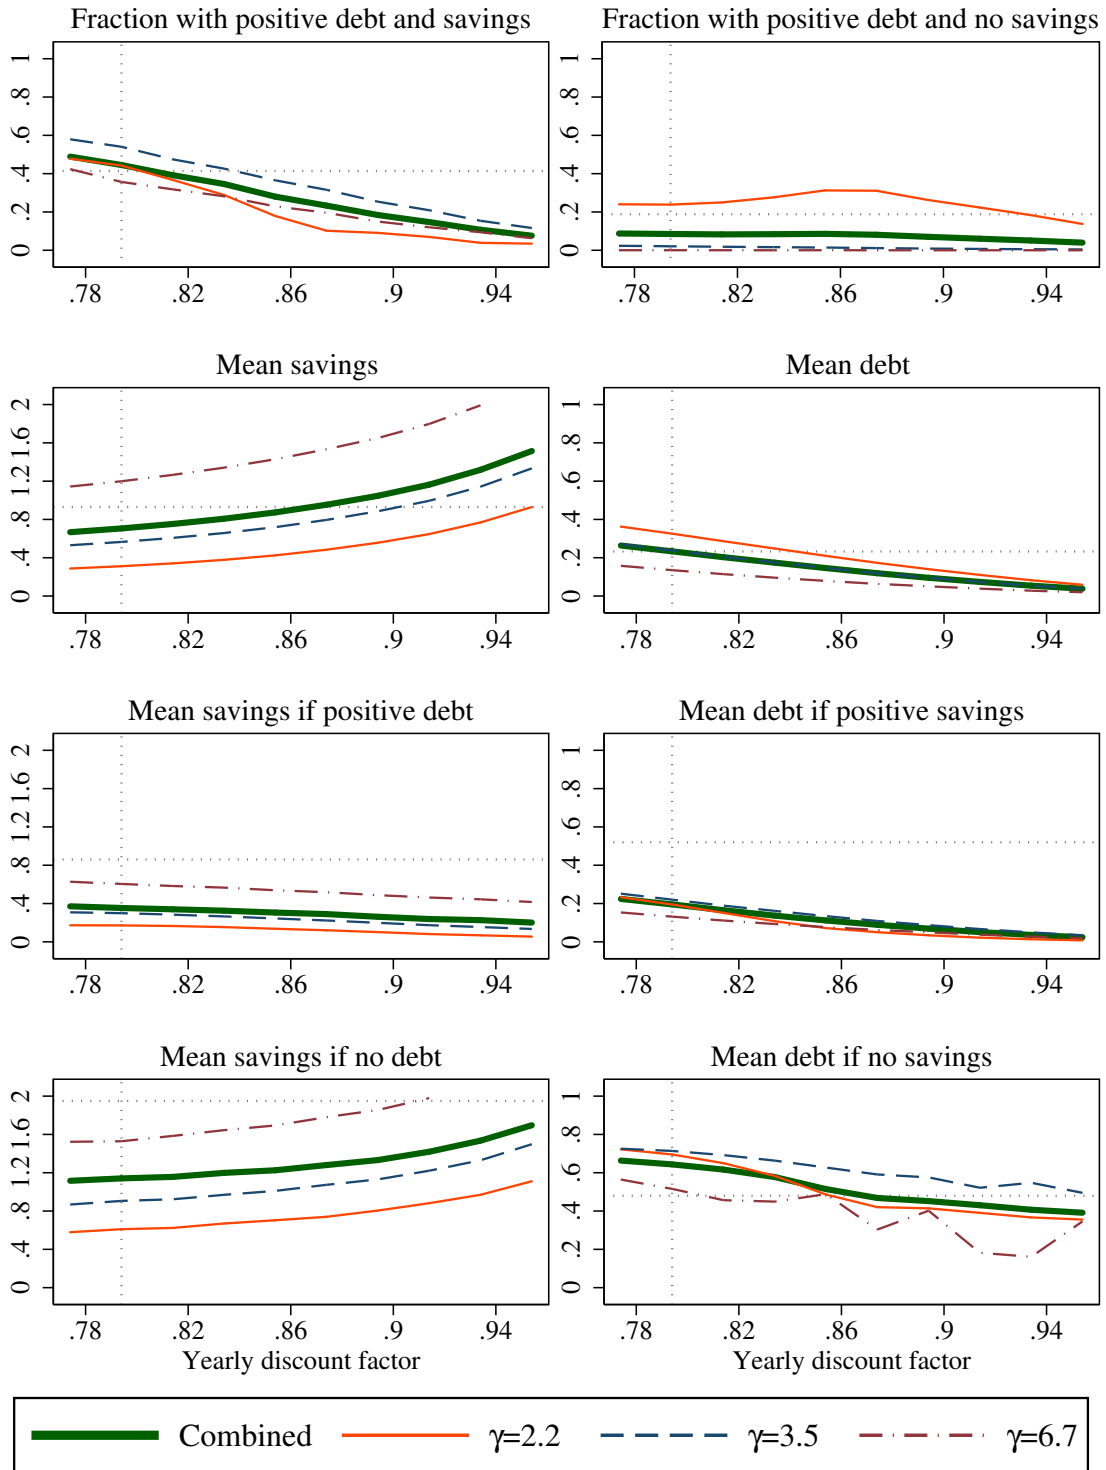Figure 6: Joint histograms of savings and debt
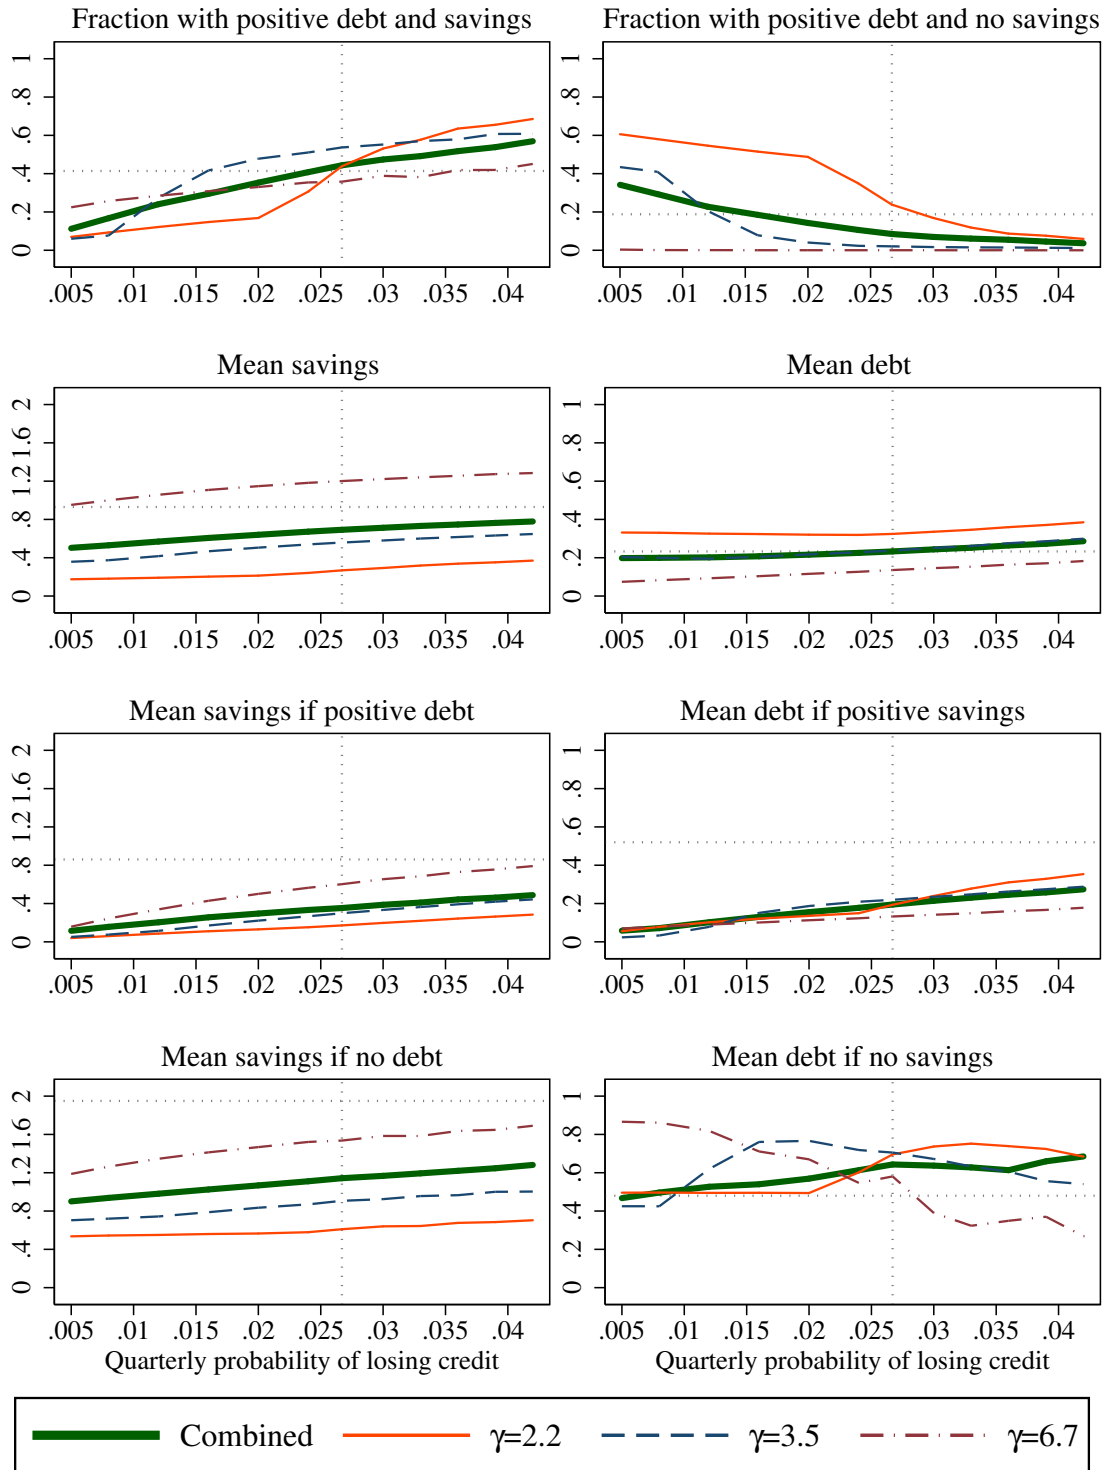(A) From the 2007 SCF
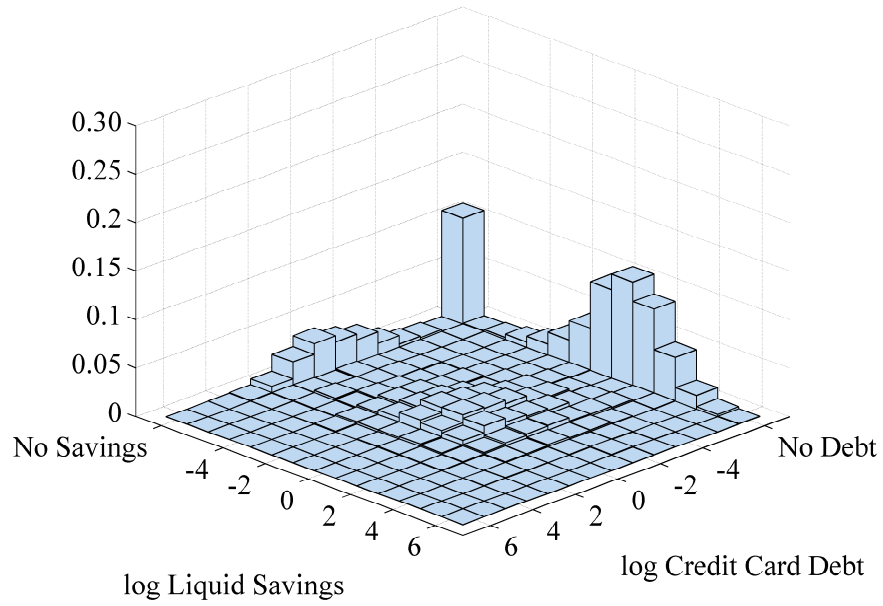
(B) Best simulated fit
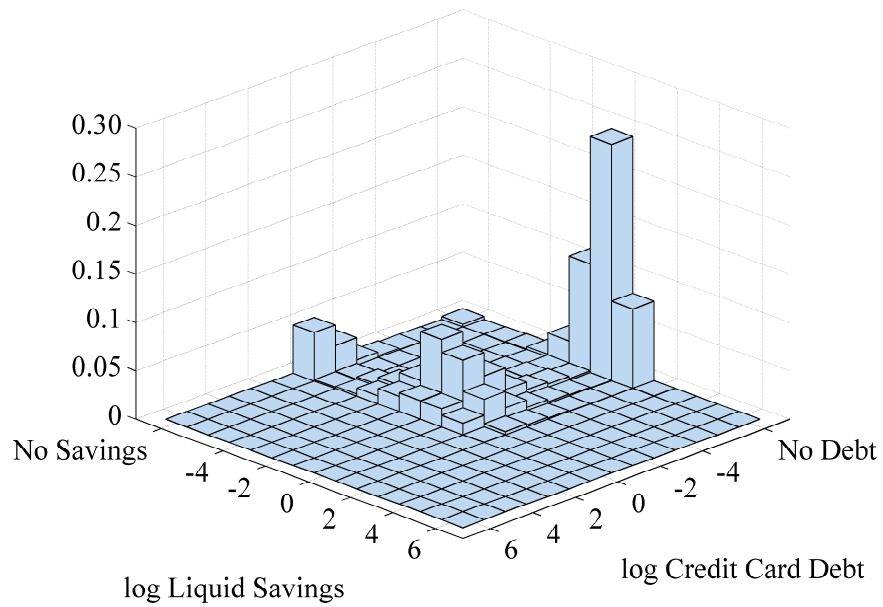
Notes: Panel (A) is from 2007 SCF with savings defined as Savings account + (Checking -Income) if (Checking -Income)>0. Panel (B) is from calibrated distribution in section 4.

Figure 7: Savings and debt as the discount factor increases

Notes: The horizontal dotted line is the equivalent mean or fraction from the 2007 Survey of Consumer Finances. The vertical dotted line shows the parameter for the baseline calibration. γ is the coefficient of relative risk aversion. See section 4.

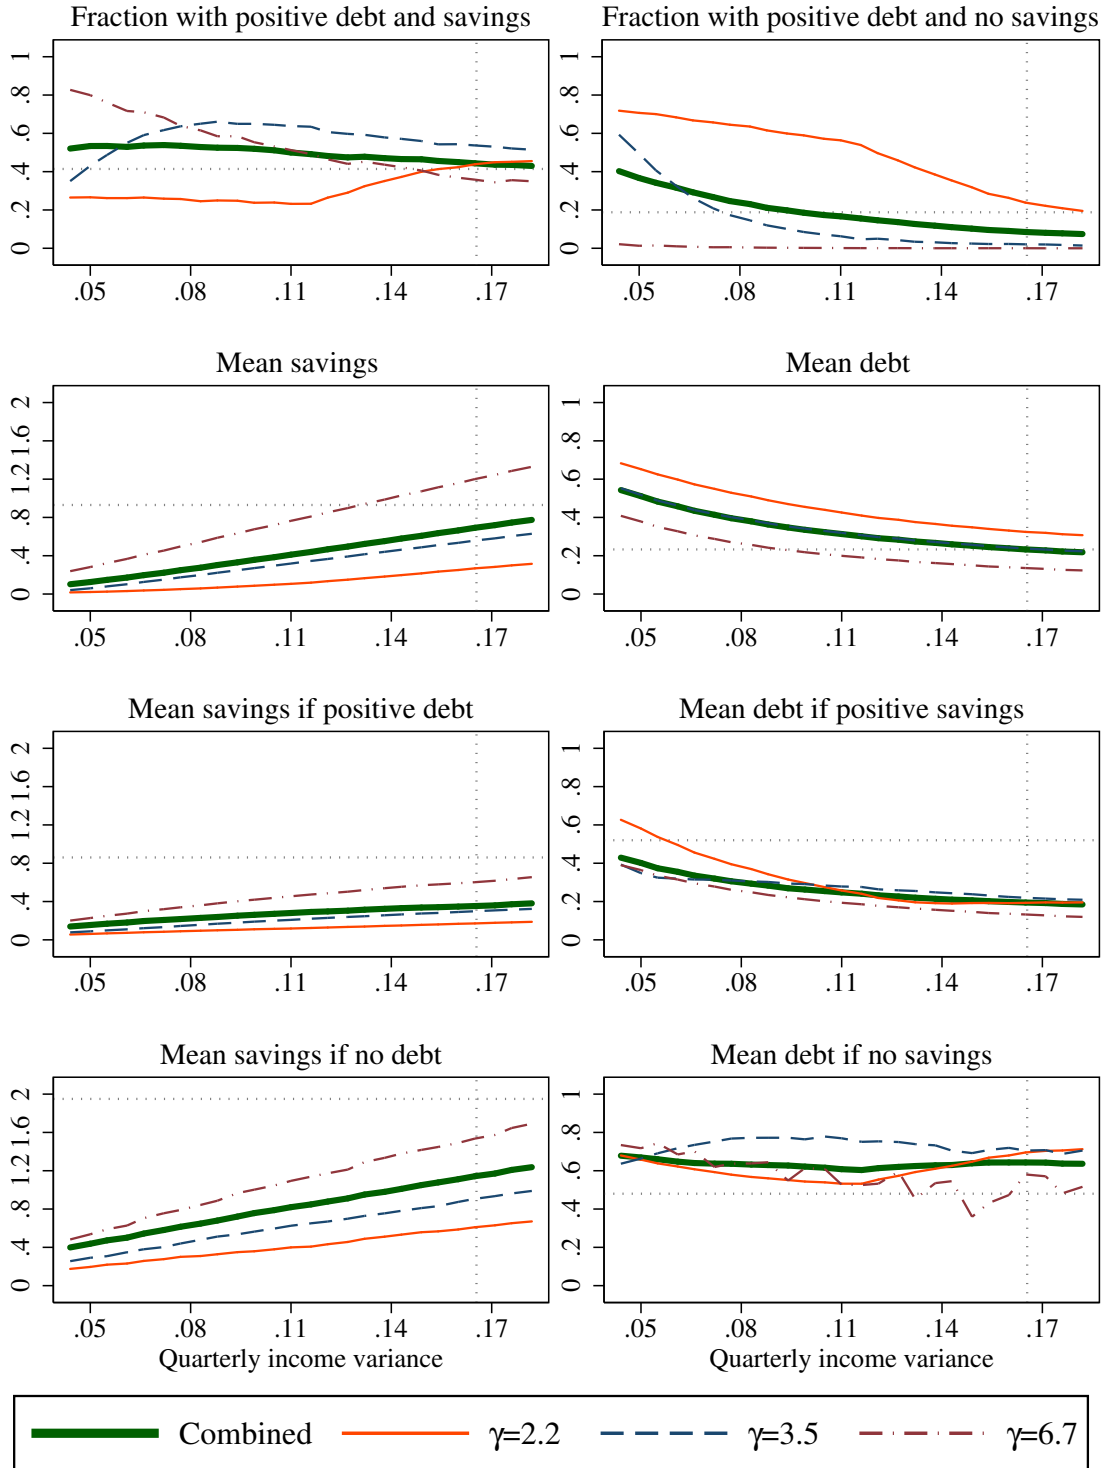Figure 8: Savings and debt as the probability of losing access to credit increases

Notes: The horizontal dotted line is the equivalent mean or fraction from the 2007 Survey of Consumer Finances. The vertical dotted line shows the parameter for the baseline calibration. γ is the coefficient of relative risk aversion. See section 4.

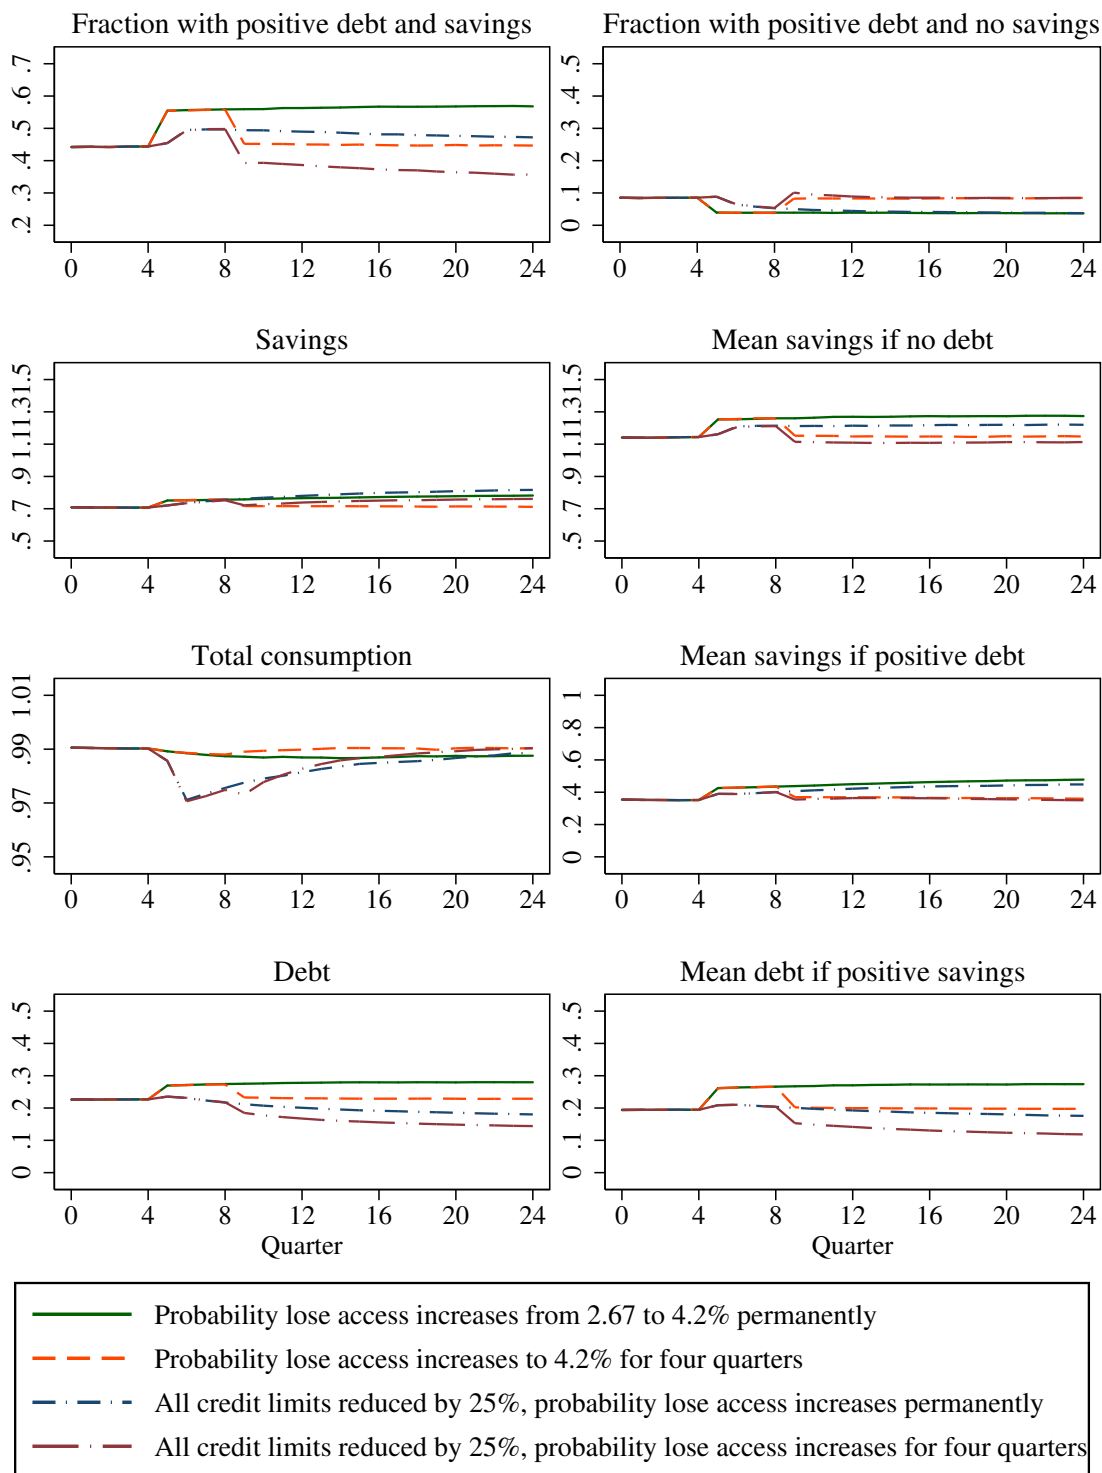Figure 9: Savings and debt as the income variance increases

Notes: The horizontal dotted line is the equivalent mean or fraction from the 2007 Survey of Consumer Finances. The vertical dotted line shows the parameter for the baseline calibration. γ is the coefficient of relative risk aversion. See section 4.

Figure 10: Comparative dynamics of increasing the probability of losing access to borrowing

Fraction with positive debt and savings

Fraction with positive debt and no savings

Savings

Mean savings if no debt

Total consumption

Mean savings if positive debt

Debt

Quarter

Mean debt if positive savings

Quarter

——— Probability lose access increases from 2.67 to 4.2% permanently

— — — Probability lose access increases to 4.2% for four quarters

—·—·— All credit limits reduced by 25%, probability lose access increases permanently

—··—··— All credit limits reduced by 25%, probability lose access increases for four quarters

Notes: Shows the evolution of debt, savings, and consumption following a change in the probability of losing access to credit in quarter 4. See section 6. The baseline calibration matches the amount of debt in the SCF described in section 4.