

Why is credit utilization stable? Precaution, payments, and credit cards over the business cycle and life cycle

Scott L. Fulford and Scott Schuh*

April 2017

Preliminary version

Abstract

We show that credit utilization—the fraction of the available credit card limit used—is remarkably stable over the business cycle, the life-cycle, and for individuals in the short term. To understand this stable utilization, we build and estimate a model of life-cycle credit use in which credit cards have both a precautionary and payments value, uniting their monetary and revolving debt uses. Both uses suggest that increases in credit accompany increases in debt. We estimate the preferences necessary to match the observed life-cycle debt and consumption. The model successfully predicts stable utilization over the life-cycle, business cycle, and for individuals.

*Scott Fulford: Consumer Financial Protection Bureau; email: scott.fulford@cfpb.gov. I conducted some of this work while I was on the faculty at Boston College and a visiting scholar at the Consumer Payments Research Center at the Federal Reserve Bank of Boston. I would like to thank the Bank and the Center for their knowledge and help. Scott Schuh: Consumer Payment Research Center, Federal Reserve Bank of Boston; email: Scott.Schuh@bos.frb.org. We both thank David Zhang for his excellent research assistance. This paper has benefited from the comments of participants at the 2015 Canadian Economics Association, the 2015 Boulder Summer Conference on Consumer Financial Decision Making, the FDIC Consumer Research Symposium, the 2016 NBER Summer Institute Aggregate Implications of Micro Behavior workshop, and seminars at the Federal Reserve Bank of Boston, the Bank of Canada, and the Consumer Financial Protection Bureau. To ensure appropriate use of the data, Equifax required that we pre-clear results that used Equifax data before making them public. The views expressed in this paper are the authors' and do not necessarily reflect the official position of the Federal Reserve Bank of Boston, the Federal Reserve System, the Consumer Financial Protection Bureau, or the United States.

1 Introduction

As banks attempted to repair their balance sheets during the financial crisis of 2008-2009, they reduced the credit card limits of millions of Americans, wiping out nearly a trillion dollars in available credit, and reducing the average limit by about 40%. At the same time, Americans reduced their credit card debt by a similar amount and so the average credit utilization—the fraction of available credit used—was nearly constant from 2000 to 2015 (see figure 1). In aggregate, the debt reductions were approximately double the tax rebates from the Economic Stimulus Act (Parker et al., 2013) and the average fall in debt was more than \$1,000 dollars. Why did so many Americans pay back so much debt during a recession?

While the fall in credit during the crisis is the most dramatic example, we show that large changes in credit accompanied by similarly large changes in debt are common over the life-cycle and for individuals in the short term. Using a large panel from the credit bureau Equifax and collected by the New York Federal Reserve, we show that average credit card limits increase by more than 400 percent between age 20 and 30, and continue to increase after age 30, although at a somewhat slower rate. Yet since credit card debt increases at nearly the same pace, credit utilization declines only very slowly over the life cycle. Individuals also face substantial credit limit volatility, several times larger than income volatility (Fulford, 2015), but individual credit utilization is extremely persistent as well, with shocks dying out almost completely after about two years.

Credit cards can be used both for payments and to accumulate revolving debt, and we suggest that the reason for stable utilization is that both uses push towards stability, although in different ways. The first use is primarily monetary: a credit card may be a convenient way to pay for things, especially if it offers some kind of rewards. In this way, credit cards are similar to debit cards or checks which may take several days to clear and require an intermediary who promises to pay the merchant first and collect from the consumer later.¹ Convenience use of credit cards therefore takes

¹This aspect of credit cards, which involves the inter-relationship between credit and liquidity, has been studied recently by Telyukova and Wright (2008) and Telyukova (2013).

on some of the aspects of consumption. An increase in income that increases both consumption and the credit card limit will also increase convenience use, suggesting that utilization should be stable. The second use is that, having decided to pay for something, a consumer has the option to roll-over her debt from month to month. Credit therefore plays a precautionary role to smooth over shocks and a lifetime liquidity role to help bring forward consumption. Since such precautionary or buffer-stock wealth (Carroll, 1997; Deaton, 1991) is only there for insurance purposes, an increase in credit effectively makes people more wealthy, allowing them to spend more (Fulford, 2013). Both the monetary and precautionary functions of credit cards point to stable utilization since debt rises at the same time as credit.

To make the case that the combined payment and precautionary aspects of credit are important and explain the broad facts of credit card debt, we adapt a model of life-cycle consumption and savings (Gourinchas and Parker, 2002; Cagetti, 2003) to allow both for the use of credit cards for payments and the large changes in credit limits over the life-cycle that we have documented. Within the model, we allow the consumer to decide how much of current consumption to pay for with a credit card. Some payments are very convenient to use a credit card for, others less so. A key feature of credit card contracts allows us to identify the convenience value of credit cards: card holders who had revolved in the previous billing cycle lose the float or grace period to pay off the new purchases without interest. Our simple payments model is identified off of the difference in behavior between those who are revolving debt and those who are not. Using new data from the Diary of Consumer Payment Choice, we estimate that non-revolvers would be willing to pay 0.319 percent of all consumption to continue using credit cards. Since the value takes place on only the 17% of consumption paid for with credit cards, it suggests that this form of payment is very convenient, given the current payments infrastructure.

We then embed the value of payment choice in the life-cycle model which allows the accumulation of credit card debt and savings and estimate the preferences necessary to match the life-cycle profile of consumption or credit card debt using the Method of Simulated Moments (McFadden, 1989). Allowing for substantial heterogeneity of preferences, the model matches both the life-

cycle moments of credit card debt and consumption well. The estimates suggest that about half the population must be very impatient and care little about risk to hold the amount of revolving debt we observe.

Around half of credit card users are primarily using their cards for payments, and so allowing for both convenience and revolving use is necessary for understanding overall credit card debt. Building on the work of Gross and Souleles (2002), recent work estimates the response of debt to changes in credit and has noted the substantial heterogeneity in responses (Agarwal et al., 2015; Aydin, 2015). The heterogeneity comes both from the differences between revolvers and convenience users as well as different responses across the life cycle and fraction of credit used. Allowing for both type of uses, the dynamics of credit utilization from the model closely match the reduced form dynamics we estimate in the credit bureau data, suggesting that this kind of heterogeneity is necessary to understand responses to credit and other liquidity shocks.

Adding heterogeneous uses for credit to the life-cycle model also suggests a subtly different explanation for the hump shape of life-cycle consumption (Attanasio et al., 1999; Attanasio and Browning, 1995) than the combination of precaution and life-cycle savings suggested by Gourinchas and Parker (2002), instead hearkening to the older heterogeneous approach of Campbell and Mankiw (1989) and Campbell and Mankiw (1990). While precaution explains the level of consumption, it plays only a small role in the life-cycle shape in our estimates. Instead, those who hold revolving debt consume close to their income over the entire life-cycle, relying entirely on credit for smoothing purposes, and so all of their “savings” come from credit limit increases. The patient convenience users looks much like a standard Life-Cycle/Permanent Income Hypothesis consumer who consumes nearly the same amount over the entire life-cycle. The average of these two populations consume less than income over the life-cycle but with the distinct hump shape of consumption coming entirely from the income profile of the impatient hand-to-mouth population. Indeed, our estimate that revolving hand-to-mouth consumers compose about half of the population is very close to the estimates in Campbell and Mankiw (1990).

Using our model estimated to match the life-cycle of credit and consumption, we finally ask

how well it explains the business cycle of credit and its individual dynamics. This approach follows the standard hypothetico-deductive method (Deaton, 2010), in that the best evidence in favor of a model is if it helps explain new phenomena out of sample. We simulate a large population that matches the age distribution of the population of credit users and give them a large, unexpected, reduction in credit of the same size as occurred during the financial crisis. The model predicts the stable utilization that obtained over the period both in aggregate and at the individual level, despite the large rise in credit from 2000 to 2008 and the abrupt fall over 2009.

2 Literature

How consumers respond to changes in the availability of credit is rarely studied, despite the centrality of credit and debt in the financial lives of consumers and households. It is particularly hard to study changes in credit because credit—unlike debt, assets, or income—is only occasionally reported in surveys. Indeed, Zinman (2015) argues that household debt is understudied even within the field of household finance, itself understudied compared with other areas of finance. Other than some work on mortgages (Iacoviello and Pavan, 2013), this paper appears to be the first to study the life cycle of credit limits. Consequently, we view one of our major contributions as documenting how U.S. consumers’ credit and debt vary over the life cycle.

A small literature has examined changes in credit card limits, but has considered changes in limits only for single credit card accounts and has not considered the difference between convenience users and revolvers. The pioneering work is Gross and Souleles (2002), who use a panel from a single credit provider to look at how households respond to changes in interest rates and changes in available credit. Increases in credit limits were followed by increases in debt, particularly for consumers already close to their credit limit. Gross and Souleles (2002) also noted the “credit card puzzle,” the observation that some households pay high interest on credit card debts while at the same time earning low interest on liquid savings. Agarwal et al. (2015) use credit limit increases that are discontinuous with credit score to examine how exogenous changes in credit lim-

its matter. For consumers with the lowest scores, a dollar increase in limits is followed by 59 cents more debt within a year, while those with higher credit scores incur almost no increase in debt. Using the Equifax/CCP, Fulford (2015) demonstrates that there is substantial credit limit variability. Short-term credit volatility is larger than most measures of income volatility, and long-term credit volatility is much greater than long-term measures of income volatility. This volatility can help to explain the credit card puzzle. Ludvigson (1999) examines the response of consumption to credit volatility at the aggregate level. Leth-Petersen (2010) studies a Danish reform that allowed homeowners to use housing equity as collateral for the first time. The reform increased available credit, but it produced relatively moderate consumption responses. The response was strongest for the youngest households.

Since larger changes in credit are more likely when there is less financial intermediation in general, much of the recent work on changes in credit have studied developing countries. Fulford (2013) examines the short- and long-term consumption responses of consumers in a buffer-stock model following changes in credit and finds evidence consistent with the model in India following a massive banking expansion. In an structural estimation approach similar to our work, Kaboski and Townsend (2011) estimate a model of consumption, savings, and credit using an expansion of credit in Thailand. Karlan and Zinman (2010) examine the effects of expanding credit access to the marginal borrower in South Africa.

While changes in credit over the life cycle have been largely unstudied, much work examines the decision to save and consume over the life cycle. In much of this work, credit is intentionally pushed to the background. In the standard versions of the life-cycle or permanent income hypothesis, for example, the assumption that young consumers can smooth relies directly on the assumption that credit is readily available (Deaton, 1992). More recent life-cycle models take seriously that credit constraints may bind, but do not allow credit to vary over the life cycle. The typical assumption is that either there is no credit available, or that there is a fixed limit, so that only net assets need to be studied (see, for example, Gourinchas and Parker (2002)).

Some recent work has attempted to endogenize borrowing constraints, and much of this work

has direct life-cycle implications. Cocco et al. (2005) build a model of consumption and portfolio choice over the life cycle. Their work adds portfolio choice to the approach of Gourinchas and Parker (2002). As an extension, they introduce endogenous borrowing constraints. These constraints are based on the minimum value that income can take, since with limited enforcement, borrowers have to choose to pay back rather than face the penalty of default. Borrowing greatly affects portfolio choice, as young consumers borrow if the endogenous constraint permits, and only start investing in equity later in life. Lopes (2008) introduces a similar life-cycle model with default and bankruptcy. Lawrence (1995) appears to have been the first to introduce default in a life-cycle model. Athreya (2008) develops a life-cycle model with credit constraints, default, and social insurance. Relaxing default policy creates severe credit constraints among the young. Eliminating default relaxes credit constraints for the young and reduces consumption inequality, but increases consumption inequality for the old, who now can no longer default after sufficiently severe negative shocks.

This paper also overlaps with a larger literature examining the ways that individuals use the financial products available to them. Stango and Zinman (2009) examine how much a sample of U.S. consumers pays in fees and interest for financial products. Credit card interest is the biggest financial expense, although some households also pay significant overdraft fees. More than half of households could reduce fees substantially by moving among products. Agarwal et al. (2007) note that middle-aged adults pay the least in these kinds of avoidable fees, with the minimum paid at age 53. Zinman (2013) examines the question of whether markets over- or under-supply credit and concludes that, while there are models that suggest over-supply and models that suggest under-supply, there is not strong evidence for either.

3 Credit card use over the business cycle, life-cycle, and for individuals

Both credit and debt change substantially over the business cycle, the life cycle and for individuals in the short term. This section briefly discusses the context of consumer credit in the United States, introduces our main data sources, and presents some non-parametric and reduced form results. An earlier working paper version of this paper (Fulford and Schuh, 2015) provides additional descriptive statistics including additional evidence on the distribution of credit and on credit card holding by age using the Survey of Consumer Payment Choice (Schuh and Stavins, 2014). We then turn to a model in the next section to help make sense of these observations.

3.1 The data

The Equifax/NY Fed Consumer Credit Panel (CCP) contains a 5 percent sample of all accounts reported to the credit reporting agency Equifax quarterly from 1999 to 2014. For much of the analysis we use only a 0.1 percent sample for analytical tractability. Once an account is selected, its entire history is available. The data set contains a complete picture of the debt of any individual that is reported to the credit agency: credit card, auto, mortgage, and student loan debts, as well as some other, smaller, categories. Lee and van der Klaauw (2010) provide additional details on the sampling methodology and how closely the overall sample corresponds to the demographic characteristics of the overall United States population, and conclude that the demographics match the overall population very closely: the vast majority of the U.S. population over the age of 18 has a credit bureau account. The CCP also records whether an account is a joint or co-signed account.

Rather than capturing too few people, the main issue is that a sizable fraction of sampled accounts represent either incorrect information reported to Equifax or individuals who are only loosely attached to the credit system, either by choice or because they lack the documentation. For example, the accounts are based on Social Security numbers, so reporting an incorrect Social Security number can create accounts incorrectly attributed to an individual. For this reason, we

limit the sample throughout to accounts that include a listed age and show an open credit card account at some point from 1999 to 2014. This population has some access to credit and thus excludes those who could not obtain a credit card or chose not to apply for one. Depending on the analysis, we also limit the sample to those with current open accounts, debt, or limits.

Much of this paper focuses on credit cards, since they have explicit limits not readily observable in most other markets. However, the credit card limits reported to credit reporting agencies are at times incomplete. The Equifax/CCP reports only the aggregate limit for cards that are updated in a given quarter. Cards with current debt are updated, but accounts with no debt and no new charges may not be. To deal with this problem, we follow Fulford (2015) and create an implied aggregate limit by taking the average limit of reported cards times the total number of open cards. This method is exact if cards that have not been updated have the same limit as updated cards. Estimating the difference based on changes as new cards are reported and the limit changes, Fulford (2015) finds that non-updated cards typically have larger limits, and so the overall limit is an underestimate for some consumers. This issue is a concern mostly for consumers who are using only a small fraction of their available credit, since it is only cards not used at all that are not updated. For the consumers who use more of their credit and so may actually be bound by the limit, the limit is accurate, because all their cards are updated.

3.2 Credit and debt over the business cycle

Since 2000 overall credit and debt have varied tremendously. Figure 1 shows how the average U.S. consumer's credit card limit and debt have varied from 2000 through 2014. Although the Equifax data set starts in 1999, we exclude the first three quarters, since the limits initially are not comparable (see Avery et al. (2004) for a discussion of the initial reporting problems). The scale on the left is in logarithms, so proportional changes in debt have approximately the same importance as proportional changes in credit. Between 2000 and 2008 the average credit card limit increased by approximately 40 percent from around \$10,000 on average to a peak of \$14,000. Over 2009, overall limits collapsed rapidly before recovering slightly in 2012. Credit card debt shows a similar

variation over time. From 2000 to 2008, the average U.S. consumer's credit card debt increased from just over \$4,000 to just under \$5,000, before returning to around \$4,000 over 2009 and 2010.²

Utilization is much less volatile than credit or debt. The thick line in the middle of the figure shows credit utilization, the average fraction of available credit used. Credit utilization fell slightly from over 35 percent in 2000 to around 30 percent in 2006 before generally increasing again to 2010 and declining slightly since then. Although it is not immediately evident, the scale of figure 1 for credit and debt on the left is very similar to the scale for utilization on the right. A 1 percentage point change in utilization has the same vertical distance as a 1 percent change in credit or debt. Larger changes are less exactly comparable, but the distance between \$4,000 and \$5,000, for example, is still almost exactly 20 percentage points on the right axis. The similar scales mean that we can directly compare the relative changes over time in limits, debt, and credit utilization. Credit and debt vary together in ways that produce extremely stable utilization that has no obvious relationship with the overall business cycle. Such a relationship is not just mechanical: when credit was cut in 2009, families had the choice whether to maintain or increase their debt, and the cut in credit could have translated directly into an increase in utilization rather than a decrease in debt.

3.3 Credit and debt over the life cycle

We next examine how credit, debt, and utilization evolve over the life-cycle. Figure 2 shows the credit card limit and debt in the top panel and credit utilization in the bottom panel. Each line is for an age cohort that we follow over the entire time possible. The figure therefore makes no assumptions about cohort, age, or time effects. Credit limits increase very rapidly early in life, increasing by around 400 percent between the ages of 20 and 30. They continue to increase, although less rapidly, after age 40. Life-cycle variation dominates everything else in figure 2; while there is clearly some common variation over the business cycle, cohorts move nearly in line with

²The fall in debt is not because of charge offs in which the bank writes off the debt from its books as unrecoverable. The consumer still owes the charged off debt. Banks may eventually sell charge-off debt to a collection agency, in which case it may no longer appear as credit card debt within credit bureau accounts. Charge offs are not large enough to explain the fall in debt, although they did increase over 2009. See <https://www.federalreserve.gov/releases/chargeoff/delallsa.htm> for charge-off rates for credit cards.

age. Estimating a simple model that divides between common time and age effects makes, figure A-5 in the appendix makes this point explicit.³ Despite the very large variation over the business cycle evident in figure 1, the big changes are happening over the life-cycle.

Credit limits and debt combine to give the fraction of credit used, shown in the bottom panel of figure 2. Consumers with zero debt have zero credit utilization, and so are included in utilization but are excluded from mean credit, which includes only positive values.⁴ Credit utilization falls slowly from age 20 to age 80. On average, 20-year-olds are using more than 50 percent of their available credit, and 50-year-olds are still using 40 percent of their credit. Credit utilization falls to below 20 percent only around age 70.

The slow fall in credit utilization comes from two different sources over the life cycle. Early in life, credit utilization is high, as a substantial portion of the population uses much or all of its available credit. Credit increases somewhat more rapidly than debt, however, so credit utilization falls slowly. In mid-life, debt stabilizes, but credit limits continue to increase slowly. Finally, starting around age 60, average debt, conditional on having any, starts to decline, so credit utilization declines.

3.4 The individual evolution of utilization

The previous two sections show that despite very large changes in credit and debt over the life cycle and business cycle, credit utilization is remarkably stable. The aggregate data could be hiding substantial individual volatility in utilization, however. In this section we show that utilization

³Estimating a simple model that separates the variation between age and year allows us to make the importance of life-cycle variation even clearer. Figure A-5 shows the age and year effects from estimating a simple regression of the form:

$$\ln D_{it} = \theta + \theta_t + \theta_a + \epsilon_{it}, \quad (1)$$

where $\ln D_{it}$ is either log debt, log credit limits, or utilization, and allows these to vary between age effects θ_a and year effects θ_t but imposes common cohort effects. The excluded group is age 20 and year 2000, so each panel in figure A-5 starts at zero at age 20 and year 2000. The estimated effect is in log units, and so the scale of the figure suggests that variation over the life-cycle in credit is around 9 ($e^{2.5}/e^{0.3}$) times larger than over time, even with a massive credit contraction.

⁴The calculation in figure 2 are the average of log limits and log debts to match later analysis and so exclude zeros except for utilization. Figure A-6 in the appendix shows the fraction in each cohort who have positive credit and debt. Including the zeros would lower the average credit limit and debt, but actually makes the life-cycle variation larger.

for an individual is extremely stable; while different individuals have different credit utilization ratios that represent their individual steady state and these ratios may change slowly as they age, individuals return rapidly to their own typical ratio. Credit utilization is best characterized by fixed heterogeneity across individuals, and relatively small deviations for an individual over time. We present non-parametric results in appendix A and appendix figure A-7 and reach almost identical conclusions to the parametric estimates. The non-parametric results suggest that the simple linear dynamic reduced form model we employ is surprisingly accurate.

Table 1 shows how utilization this period is related to utilization in the previous period. We estimate regressions of the form:

$$v_{it} = \theta_t + \theta_a + \alpha_i + \beta v_{it-1} + \epsilon_{it}, \quad (2)$$

where $v_{it} = D_{it}/B_{it}$ is the credit utilization given the credit limit B_{it} and the current debt D_{it} , conditional on the credit limit $B_{it} > 0$, and age (θ_a) and quarter (θ_t) effects that allow utilization to vary systematically by age and year. Column 1 does not include fixed effects and so assumes a common intercept, column 2 includes quarter and age effects, while the other columns include individual fixed effects, quarter, and age effects.⁵

Without fixed effects, credit utilization is very persistent and returns to a non-zero steady state of approximately 40 percent utilization ($\alpha/(1 - \beta) = 0.38$). Note that this utilization is close to the average in figure 1, as it should be since they are estimated from the same data and the non-

⁵The combined age, year, and individual fixed effects in equation (2) are not fully identified. As in the age-cohort-period problem, it is impossible to fully identify all effects since there can be a trend in any one of age, time, or cohort, or split among all three, and any division is observationally equivalent, since birth cohort equals the year minus age. Put differently, if we estimate all individual effects, it is not possible to fully separate between getting older and a time shock. The size of the data set means that rather than estimating individual coefficients—sometimes referred to as nuisance parameters—we instead perform the standard within transformation by removing the mean from all variables in equation (2). The within transformation means that any additional restriction for identification must be on either the time or age effects. Rather than imposing the questionable assumption that two of the age or year effects are exactly equal—the implication of dropping more than one of the age or year dummies—we instead impose the restriction that there is no trend in the age effects. This restriction is innocuous in the sense that there can still be a trend with age, as individual effects that are older when we observe them can have larger θ_i , but that trend will appear in the individual effects rather than in the age effects. Moreover, the age effects can still allow life-cycle variation, but that variation must average out to zero. We implement this restriction following Deaton (1997, pp. 123–126) by recasting the age dummies such that $\hat{I}_a = I_a - [(a - 1)I_{21} - (a - 2)I_{20}]$, where I_a is 1 if the age of person i is a and zero otherwise.

parametric conditional expectation function shown in appendix figure A-7 is nearly linear. In the other columns, since the age, year, and fixed effects change the steady state, we do not report it, but it is important to realize that the average steady-state credit utilization is not zero.

The next two columns report how credit utilization varies around an individual-specific α_i , by estimating using the within transformation. Nearly half of the overall variance in utilization comes from these fixed effects. In other words, we can think of the distribution of utilization as coming about half from factors that are fixed for an individual, allowing for common age and year trends, and half from relatively short-term deviations from the mean. A deviation from the mean utilization diminishes at a rate of about $0.353=1-0.647$ per quarter. And so, after a 10 percentage point increase in utilization, 6.47 percentage points remain in one quarter, 1.7 percentage points in a year, and less than 0.3 of a percentage point after two years. The last column suggests that the speed of return to individual utilization depends on age. For a 20-year-old, only a fraction 0.58 of the shock is left after one quarter, for a 60-year-old 0.70 of a shock remains. It is not obvious why the speed of return should increase with age. One possibility is that credit represents a much more important factor in the overall portfolio of young people, who typically have acquired few assets. A 60-year-old may have some other assets and is more likely to be a convenience user, and so may not target credit utilization as closely.

The estimates in Table 1 emphasize that credit utilization for an individual is very stable. While there are deviations from the long-term mean, these dissipate quickly and are largely gone within two years. Both the parametric and non-parametric evidence suggests that individuals have a strong tendency to return to their own credit utilization following shocks. Since credit utilization is not zero for most people, the results suggest a strong tendency to hold credit card debt for the long term. Since the parametric estimates include age effects, these deviations are around a life-cycle average, which the previous sections have shown is declining slowly, but persistently, with age. Both the slow decline of utilization with age and the quick return to individual credit utilization suggest that the pass-through from the credit card limit to credit card debt is large and occurs relatively rapidly. In the working paper (Fulford and Schuh, 2015), we show results on how debt and credit

co-evolve, rather than fixing their relationship by combining them into utilization. Relatively little is lost by simplifying only to utilization. Moreover, in a Granger Causality sense, the direction of causality moves primarily from changes in credit to change in debt.

4 A model of life-cycle consumption and credit card debt

We have demonstrated that there is a strong tendency for individuals to have stable utilization which falls slowly over the life-cycle. To understand these observations, this section describes a life-cycle consumption model similar to Gourinchas and Parker (2002) or Cagetti (2003) to which we add a payment choice and changing credit over the life-cycle. The next section estimates the parameters of the model based on the observed life-cycle profiles of debt and consumption and asks how well it explains the stable utilization.

4.1 The decision problem

From any age t a consumer seeks to maximize his or her utility for remaining life given her current resources and expected future income. With additively separable preferences, the consumer of age t with cash-at-hand W_t and current credit limit B_t maximizes the discounted value of expected future utility:

$$\max_{\{C_s, \pi_s\}_{s=t}^T} \left\{ E \left[\sum_{s=t}^T \beta^{s-t} u(\nu_s C_s) + \beta^{T+1} S(A_T) \right] \right\} \text{ subject to}$$

$$C_s \leq W_s$$

$$W_{s+1} = R(A_s)A_s + Y_{s+1} + B_{s+1}$$

$$A_s = W_s - B_s - C_s$$

$$\nu_s = \nu(\pi_s; I_s^R)$$

$$I_s^R = \{1 \text{ if } A_{s-1} < 0; 0 \text{ else}\}$$

where she gets period utility $u(\cdot)$ from consumption C_t which is made more valuable by ν_s depending on how she pays for that consumption. We will describe how she decides what fraction $\pi_s \in [0, 1]$ of her consumption to pay for with a credit card in more detail below, but consider ν_s to be the value in consumption units of her payment choice which depends on whether she has revolved the period before (I_s^R).

She discounts the future with a fixed discounted factor β . She lives for T periods where T is a random number that we will match to the actual life tables, but assume that she dies with certainty at age \tilde{T} . On death, she receives a final utility $S(\cdot)$ from left-over positive resources. In our base estimations, we set the bequest motive to allow for an annuity to heirs. Appendix B discusses the specific function. Recent work has disagreed over the importance of a bequest motive as opposed to other possible motives for keeping assets late in life such as long-term care and medical needs (De Nardi et al., 2010). Since we focus primarily on debts, our model and estimates are not well situated to distinguish between motives. Removing the bequest motive does not affect the results much, but requires the patient consumer to expect non-asset income after expenses to fall even more substantially late in life.

End of period assets A_s determine the rate of return for the next period:

$$R(A_s) = \begin{cases} R & \text{if } A_s \geq 0 \\ R_B & \text{if } A_s < 0, \end{cases}$$

with a borrower facing a higher rate than a saver $R_B \geq R$.

The payment decision Accompanying the consumption decision is a payment decision. In every period the consumer must decide the fraction of her consumption to pay for with a credit card π_s . The optimal fraction then makes consumption more valuable by an amount $\nu_s = \nu(\pi_s^*; I_{s-1}^R)$ where π_s^* is the optimal fraction and I_{s-1}^R is an indicator for whether she is revolving as of the previous period and so does not get the float that convenience users get. Since ν_s appears in the consumer's problem only in the utility function, not in the accumulation equation, it is not directly affecting

the budget constraint. Instead, we model the payment decision as one of convenience, in keeping with our dual approach to credit cards. For example, it is often more convenient to buy something online with a credit card. It might be possible to set up a direct bank withdrawal, or mail a check, or visit a bank to turn cash into a cashier's check, but these methods require additional steps or take longer. The primary value of credit cards as payments is that they are a widely accepted payment means. In addition, some credit cards offer rewards of various types that make payments using them more valuable in many ways. While some of these rewards, such as 1% cash back, could affect the budget constraint directly, others, such as airline miles or preferential access to event tickets, are primarily in the form of increased value of consumption.

Given this complexity, our approach is to make all payment choices purely hedonic, which simplifies the problem tremendously by allowing us to make the consumer's problem sequential. The consumer first decides how to pay for what she will consume. Then given ν_s , which makes consumption this period more or less valuable, she decides how much to consume, taking into account how her decisions this period will affect her payment and consumption decisions next period. In particular, the optimal amount to pay with a credit card this period directly depends on whether she left revolving debt from last period (I_{s-1}^R). There is no grace period on credit card payments for revolvers. Purchases start earning interest as soon as they are made if a consumer has unpaid credit card debt from the previous month. The consumer may choose to pay any new purchases off very soon, and so not owe much interest on them, of course, but for revolvers there is no free credit. The difference between the payments habits of convenience users and revolvers is crucial here. Unlike convenience users who get an interest free loan for all payments, and so may find credit card use very convenient, revolvers are paying interest on their purchases and so find them less convenient. Since the only state variable that matters for the payment decision is whether or not she is revolving, this approach implies that the payment decision takes on only two values $\nu^R = \nu(\pi^R; I^R = 1)$ for the optimal share by a revolver π^R and $\nu^C = \nu(\pi^C; I^R = 0)$ for convenience users with optimal share π^C . We will use this distinction and the different payment choices of revolvers and convenience users to identify a simple functional form for $\nu(\pi_s; I_{s-1}^R)$.

To model the payment decision, we assume that all of the goods and services that a consumer buys can be ranked in order of the value she gets from purchasing with a credit card. Some, such as goods online, are very convenient to buy with a credit card. Others, such as the rent or mortgage payments, are inconvenient or may come with some additional costs, and so a consumer will try to avoid paying for them with a credit card. Ordered this way, the marginal value of spending a larger fraction on a credit card is declining, and will eventually reach zero when it is inconvenient to spend more on a credit card. We assume that the decrease in the marginal value of convenience is linear at rate ν_1 and denote marginal convenience of the first consumption on a credit card as ν_0 . Figure 3 shows this marginal value. The marginal convenience value of the first small share of consumption on a credit card is very high, but as a consumer charges more and more, the value decreases. The optimum is where the marginal value of additional consumption on a credit card is zero.

Revolvers must pay interest, in addition to getting the convenience of the card, and so the starting marginal convenience for them is lower. If consumption is spread evenly over the month, then a revolver will pay additional interest of $(r_B/12)/2$ on her consumption that month.⁶ Revolvers start out at a lower initial value, and so their optimum payment share is lower in Figure 3, a prediction we see in the data and will discuss more for when we estimate this model in section 5.

Finally, we find the consumption value of this payment choice by taking the area under the curves and normalizing the value of other means of payment to 1 so that $\nu(0; \cdot) = 1$. Except for the value of bequests, the normalization does not affect the consumption decision since it affects the marginal value of consumption this period and in the future, leaving the tradeoff between them unchanged.

Although it is not required for the solution to the consumer's problem, it is useful to consider what the model says for convenience use and revolving credit. A consumer with negative assets is

⁶This formula comes from the way that annual credit card rates are reported and interest charged. The Annual Percentage Rate or APR is not a compound rate, and so it is appropriate to divide it by 12 to find the rate of interest. The financing charge on a credit card is calculated based on the average daily balance within a month, and so the financing charge on consumption spread evenly throughout a month is half the interest rate. Note that while the APR is not a compound rate, interest charges not paid off each month will compound both in reality and in our model.

borrowing and has debts that revolve from period to period. Since the payment decision is purely hedonic, it depends only on whether or not someone was revolving last period, and so there is an optimal revolver fraction of consumption on a credit card π^R convenience user fraction on a credit card π^C . The observed credit card debt at age t , which includes both new charges and previous debt is then:

$$D_{it}^R = \pi^R C_{it} - A_{it} \text{ if } A_{it} < 0$$

for a revolver (for whom $A_{i,t-1} < 0$). For convenience users on the other hand, assuming they are not transitioning into being revolvers, the observed credit card debt is just a fraction of their consumption:

$$D_{it}^C = \pi^C C_{it} \text{ if } A_{it} > 0.$$

The credit limit The limiting factor for consumption in any given period is cash-at-hand W_s which is determined by income Y_s , assets or debts from the previous period A_s , and the current credit limit B_s . The credit limit B_s evolves with permanent income:

$$B_s = b_s P_s,$$

where $b_s \geq 0$ is the possibly stochastic fraction of this amount that can be borrowed. This approach means that across consumers B_s will be in proportion to income P_s , but may follow a different average path over the life-cycle than income.⁷

⁷The consumer's problem as written, with W_s as a sufficient period budget constraint, implies that a consumer must immediately repay all debt over her limit if her credit limit falls. To see this, consider what happens if $B_{s-1} > 0$ and the consumer borrows leaving negative assets at the end of period $A_{s-1} < 0$. If $B_s = 0$, then assets at the end of period s must be weakly positive ($A_s \geq 0$), and so all debt has been repaid within a single period. A cut in credit limits implies an immediate repayment of debt in excess of the limit. This debt repayment when credit is cut below debt does not match credit card contracts, which do not require immediate and complete payment following a fall in credit (Fulford, 2015)). Instead, credit card borrowers can pay off their debt under the same terms, they just cannot add to it. However, allowing for such behavior means that there must be an additional continuous state variable since W_s and B_s no longer fully describe the consumer's problem, which adds substantially to the numerical complexity of the solution through the curse of dimensionality.

The income process Income or disposable income follows a random walk plus drift:

$$Y_{s+1} = P_{s+1}U_{s+1}$$

$$P_{s+1} = G_{s+1}P_sN_{s+1}$$

where G_{s+1} is the known growth rate from period to period, the “permanent” or random walk shocks N_{s+1} are independently and identically distributed as lognormal with mean one: $\ln N_{s+1} \sim N(-\sigma_N^2/2, \sigma_N^2)$; and the transitory shocks are distributed as:

$$U_{s+1} \sim \begin{cases} U_L & \text{with probability } p_L \\ \tilde{U}_s(1 - U_L p_L)/(1 - p_L) & \text{with probability } 1 - p_L, \end{cases}$$

where \tilde{U} is an i.i.d. lognormal with mean one: $\ln \tilde{U}_{s+1} \sim N(-\sigma_U^2/2, \sigma_U^2)$ and U_L is unemployment income as a fraction of permanent income. The structure of the shocks ensures that the expected income next period is always $G_{s+1}P_s$ since the mean of both transitory and permanent shocks is one.

Recursive formulation and iso-elastic preferences. We assume that sub-utility displays Constant Relative Risk Aversion (CRRA):

$$u(C) = \frac{C^{1-\gamma}}{1-\gamma}.$$

With CRRA preferences, it is possible to normalize the problem in terms of permanent income P_t at any given age (see Carroll (2012) for a more complete discussion). Using lower case to represent the normalized value: $c_t = C_t/P_t$, $w_t = W_t/P_t$, and $a_t = A_t/P_t$, and rewriting in recursive form,

the problem is equivalent to:

$$\begin{aligned}
v_t(w_t, b_t) &= \max_{c_t} \left\{ u(\nu_t c_t) + E_t[\beta_{t+1}(G_{t+1}N_{t+1})^{1-\gamma} v_{t+1}(w_{t+1}, b_{t+1})] \right\} \text{ subject to} \\
c_t &\leq w_t \\
w_{t+1} &= R_{t+1}(I_{t+1}^R) a_t + U_{t+1} + b_{t+1} \\
a_t &= w_t - b_t - c_t \\
\nu_t &= \nu(I_t^R) = \max_{\pi_t} \nu(\pi_t; I_t^R) \\
I_t^R &= \{1 \text{ if } a_{t-1} < 0; 0 \text{ else}\}
\end{aligned}$$

where $R_{t+1}(I_{t+1}^R) = R/(G_{t+1}N_{t+1})$ if $a_t \geq 0$ and $R_{t+1}(I_{t+1}^R) = R_B/(G_{t+1}N_{t+1})$ if $a_t < 0$. The expectation at t includes the possibility of death before T and the certainty of death at \tilde{T} leaving a bequest worth $\beta_{t+1}s(a_t)$ where $s(\cdot)$ is the bequest function normalized by P_t .

The initial distribution of income and wealth, and late life income Finally, there are several important decision parameters that affect initial distributions and decisions late in life. We assume that the initial distribution of the wealth/permanent income ratio is log-normal with a location parameter λ_0 and variance parameter that matches the variance of permanent income shocks. We allow for a fall in disposable income in retirement or old age by assuming that after R periods permanent income falls to λ_1 of its value at age $R - 1$. Such a fall in disposable income may also come from increased medical or care needs, and so is a flexible way of allowing for the possibility that wealth needs may be higher late in life because income after necessary expenses may be lower.

Model frequency Because of data and computational constraints, much of the structural consumption literature has been limited to examining decisions made at a yearly frequency. Yet consumption decisions must be made more frequently than yearly. If smoothing within the year is perfect, then the frequency should not matter. However, the logic of the model and the data suggest that people do occasionally hit their budget constraint, which directly implies that ignoring

decisions made within the year may miss important facets of consumer behavior. In addition, if we want to understand whether the model can match the quarterly dynamics of individual and aggregate credit utilization, it must be have at least a quarterly frequency.

4.2 Numerical solution

With the problem written recursively, we proceed through backward recursion to find a numerical approximation of the consumer's problem. For a given sent of parameters, once $v_{T+1}(A_T, 0)$ is given, it is possible to find an approximation of $v_T(w, b)$, and use the approximation of $v_T(w, b)$ to find $v_{T-1}(w, b)$. The solution to each period's value function is a consumption function $c_t(w, b)$. We follow several standard steps (see Carroll (2012) for a more in depth discussion of many of these approaches). First, we discretize the log-normal shocks using a Gauss-Hermite quadrature which turns the integration in the expectation function into a summation over discrete states. Since the income process is surely not exactly log-normal there is no gain or loss in accuracy from doing so; we are simply replacing one approximation of shocks with another.

Second, we follow the method of endogenous gridpoints (Carroll, 2006) to find the optimal consumption that leads to end of period assets a_t at a number of gridpoints for a_t and b_t . Very elegantly, it is then possible to find optimal consumption which leaves this amount of assets $c_t(w, b)$ at the endogenous gridpoints for w simply using the accounting identity $a_t = w_t - b_t - c_t$, and so avoid a computationally costly numerical root-finding approximation entirely. More precisely, if the consumer has not consumed all available cash-at-hand for the next period, and so is not strictly constrained by the credit limit, then the standard first order conditions and the Euler equation imply that:

$$u'(\nu(I_t^R)c_t) = E_t[\beta_t R_{t+1}(a_t)(G_{t+1}N_{t+1})^{1-\gamma}u'(\nu_{t+1}c_{t+1}(w_{t+1}, b_{t+1}))]$$

where despite its subscript $\nu_{t+t} = \nu(I_{t+1}^R)$ is determined entirely by the choice of whether to leave positive or negative assets for the next period. Given the next period consumption function, it is

straightforward to find the optimal consumption that leaves end of period assets a_t as:

$$c_t^a(a, b, I_{t+1}^R) = (1/\nu_t) \left(E_t \left[\beta \hat{\beta}_{t+1} R_{t+1}(a) (G_{t+1} N_{t+1})^{1-\gamma} (\nu_{t+1} c(R_{t+1}(a)a + U_{t+1} + b_{t+1}, b_{t+1}))^{-\gamma} \right] \right)^{-1/\gamma}. \quad (3)$$

For a vector of end of period assets \vec{a} , it is nearly costless to find the optimal consumption at a vector of endogenous points for cash-at-hand $\vec{w} = \vec{a} + b + c_t^a(\vec{a}, b, I_{t+1}^R)$ is the amount at which consuming $c_t^a(a, b, I_{t+1}^R)$ and leaving a for next period is optimal. We linearly interpolate between these points to find an approximation of the consumption function. Note that in addition to the current state of assets and the credit limit, the consumption function is also a function of whether the consumer is revolving by having negative assets last period. While it is not a continuous state, the addition of another state variable complicates the solution since we must find the optimal consumption for both revolvers, who find consuming less valuable since they cannot pay for it in quite as convenient way, and convenience users. For the most part, someone who is not revolving this period will not be revolving next period and so $\nu_t = \nu_{t+1}$ and the payment choice does not affect the consumption decision directly. It does, however, make revolving somewhat more costly.

Since the consumer's problem includes both an externally imposed credit limit as well as interest rates that differ depending on whether assets are positive or negative, there are several additional complications. The first is that the standard Euler equation does not hold when the consumer is against her credit limit and so spends all available resources, since she would like to spend more today but cannot (Deaton, 1991). This problem is relatively easy to deal with, however, by including the inflection point, or the last point at which the Euler equation holds and so assets left for the next period are $-b_t$. For any cash-at-hand less than $w^* = c_t^a(-b, b)$, the consumer consumes all cash-at-hand, so $c_t(w, b) = w$ if $w \leq w^*$. The second problem is that the interest rate differential introduces a step in the consumption function since there are two solutions to equation (3) for $a = 0$. One, the limit with assets approaching zero from below, uses the borrowing rate R^B and the other uses the saving rate R . The economic intuition is that leaving zero assets for the next period is optimal at a high borrowing rate well before it is optimal at low savings rate. For cash-at-hand be-

tween these two points, the consumer has a marginal propensity to consume of one since the return on savings is not high enough to induce her to save, but the cost of borrowing is sufficient to keep her from borrowing and so additional resources go straight to consumption. To deal with this issue, the endogenous gridpoints include two points where $a = 0$: the first $c_t^B = c_t^a(0, b; R^B)$ the solution to equation (3) when $a = 0$ using R^B and $w_t^B = 0 + b + c_t^B$, and the second $c_t^F = c_t^a(0, b; R)$ and w_t^F . Between the points (w_t^B, c_t^B) and (w_t^F, c_t^F) the consumer has a marginal propensity to consume of one.

Figure A-2 shows the consumption function at different ages with a constant ability to borrow over the life cycle of one fifth of income.⁸ Several points are worth discussing. First, the consumption function generally falls with age. This occurs as the consumer plans for retirement when having accumulated a large amount of savings is valuable. Since we set $b_{T+1} = 0$, and the remaining wealth is divided among the many years of retirement, consumption out of available resources $c_{T+1}(w_{T+1})$ is quite low. Planning for this high demand for savings means after age 45, consumers seek to accumulate by spending much less from available resources. Early in the life-cycle, retirement is far away, and short-term shocks dominate, and so the consumption function at 26 is very similar to the one at 35 and 45. Second, for low cash-at-hand below w^* the marginal propensity to consume is one. Between w^* and w^B the consumer is leaving debt for next period, and so is paying a high interest rate R^B . Between w^B and w^F the consumer does not want to borrow, but the return on savings is not high enough, and so she leaves zero assets and has a marginal propensity to consume of one. This kink in the consumption function implies that there can be a positive fraction of consumers who hold exactly zero assets. The distance between w^B and w^F depends on the interest rate differential, with a wider differential implying a larger distance.

⁸The other relevant parameters are a savings rate $R = 1.03$ and borrowing rate $R^B = 1.14$, a coefficient of relative risk aversion of $\gamma = 2$, a discount parameter $\beta = 0.9$ and other parameters and retirement functions matching Gourinchas and Parker (2002).

5 Estimation

This section describes the data used to estimate the model, the estimation procedure, and the results. Along with information on the income process described below and the payments parameters, the life-cycle portion of the model is described by only four parameters. We allow for preference heterogeneity by allowing for two sub-populations with different preferences. Of course, additional preference heterogeneity is possible, but our results show that this is the minimum heterogeneity necessary and we prefer this simple heterogeneity since it makes obvious the contribution of different populations. There are thus three forms of heterogeneity in the estimated model: age as people make different decisions at different ages; idiosyncratic heterogeneity as people are hit with different shocks and so have different assets and incomes; and preference heterogeneity as different sub-groups react differently to shocks.

For each sub-group we estimate the discount rate β , coefficient of relative risk aversion γ , and the fall in income at the end of life λ_1 . To combine groups we estimate the share of group A, f^A in the population. Income and initial wealth must be jointly determined with f^A to keep the simulated mean income and initial wealth equal to their observed counterparts. To allow one group to have a higher income than the other on average, we estimate the multiple of the average permanent income earned by group A, ζ^A , so that if group A earns half of the population average income $\zeta^A = 0.5$. Together f^A and ζ^A directly determine ζ^B . Similarly, we estimate the initial liquid funds to income ratio of group A λ_0^A , and so f^A determines λ_0^B .⁹ Finally, we estimate the convenience value of credit cards for payments for revolvers ν^R and convenience users ν^C . We estimate all of the other parameters of the decision in a separate first stage. We therefore estimate 11 parameters jointly $\theta = \{\gamma^A, \beta^A, \lambda_0^A, \lambda_1^A, \gamma^B, \beta^B, \lambda_1^B, f^A, \zeta_0^A, \nu^R, \nu^C\}$.

We estimate the parameters of the highly non-linear model that best match the life-cycle evolution of debt and consumption using the Method of Simulated Moments (MSM) of McFadden (1989). For a given set of parameters $\theta \in \Theta$, and first stage parameters χ such as the interest

⁹In order for the average income of the combined populations to equal the average observed income $f^A \zeta^A + f^B \zeta^B = 1$ which, since $f^B = 1 - f^A$, implies that $\zeta^B = (1 - f^A \zeta^A)/(1 - f^A)$. Similarly, since $f^A \lambda_0^A + f^B \lambda_0^B =$ Initial liquid funds to income then $\lambda_0^B = ((\text{Initial liquid funds to income}) - f^A \lambda_0^A)/(1 - f^A)$.

rate and income process estimated separately, we numerically find consumption functions at each age. Then with an initial distribution of initial assets, income, and credit limits, we draw a starting value from this distribution for each $i = 1$ to N consumers. Each consumer makes a consumption decision, leaving assets for next period. For each consumer, we draw the next period income and credit limits from their distributions and find his or her consumption and then assets for the next period. This process proceeds until the last period, generating for every age the consumption and assets for all N consumers.

A given set of model parameters can thus generate a life-cycle distribution of consumption, debt, and savings. With these distributions, it is possible to calculate moments that describe the distribution and attempt to find the parameters that minimize the difference between these moments and their observable counterparts. Gourinchas and Parker (2002), for example, calculate the risk aversion and discount parameters that make the calculated mean consumption over the life-cycle match the hump shape observed in the Consumer Expenditure Survey. Cagetti (2003) instead matches the accumulation of wealth over the life-cycle as observed in the Survey of Consumer Finances.

More formally, for a given $\theta \in \Theta$, and first stage parameters χ estimated below, we can define the difference between an empirical moment and a simulated moment as $g_j(\theta; \chi)$ for each of J total moments. For example, one of the empirical moments we use is the mean credit card debt at age 30 estimated from the Equifax/NY FED, and so one moment is:

$$g_j(\theta; \chi) = (1/I) \sum_{i=1}^I D_{30}^i - (1/N) \sum_{k=1}^N \hat{D}(\theta; \chi)_{30}^k$$

where $\hat{D}(\theta; \chi)_{30}^i$ is the simulated debt of person i at age 30 given parameters θ , and D_{30}^k is the debt of person k in Equifax/NY Fed CCP at age 30. These are averaged over the I people in the CCP sample, and the N people we simulate. The MSM then seeks to minimize the weighted square of these differences:

$$\min_{\theta \in \Theta} g(\theta; \chi)' W g(\theta; \chi) \quad (4)$$

where $g(\theta; \chi) = (g_1(\theta; \chi), \dots, g_J(\theta; \chi))$ and W is a $J \times J$ weighting matrix. We generally use the optimal weighting matrix proportional to the inverse variance of the empirical moments, which gives more weight to better estimated moments.

5.1 The empirical moments and first stage estimates

We estimate the model to provide the best fit to two life-cycle moments: (1) the observed credit card debt over the life-cycle from the NY Fed/Equifax Consumer Credit Panel described in section 3 and the (2) household consumption over the life-cycle from the Consumer Expenditure Survey from 2000-2014. Since our observed credit data is for the individual account holder, rather than the household, we adjust household consumption by dividing by the number of adults in the household. We allow for some unobserved taste changes over the life-cycle by adjusting consumption for the number of children in the household.¹⁰

To be identified, the model requires a number of other ancillary parameters estimated elsewhere. We adjust the standard errors of the MSM for these first stage estimates using their variance-covariance matrices following Laibson et al. (2007) who also update Gourinchas and Parker (2002) by allowing for the empirical moments to have different numbers of observations.

We use data from the Diary of Consumer Payment Choice to find the fraction of consumption on a credit card by revolvers and convenience users, and discuss the identification of the payments problem below. We estimate the average life-cycle growth (G_t) of income based a fifth order polynomial best fit of after tax income per adult household member from the Consumer Expenditure Survey. The raw data and the fitted lines are in appendix figure A-8. Similarly, we take a fifth degree polynomial estimate of the average credit limit per account from Equifax/NY Fed CCP to form B_s . Not smoothing these two budget constraints makes little difference to the overall esti-

¹⁰Formally, we estimate

$$\ln(C_{i,t}/\text{Adults}_{i,t}) = \theta_a + \theta_t + \beta \text{Children}_{i,t} + \epsilon_{i,t}$$

and then calculate average household consumption per adult at each after removing the effect of children at the individual level. The removing the implied consumption effect of children has a surprisingly small effect. Figure A-8 shows the unadjusted and adjusted consumption. Children raise expenditures per adult household member slightly from ages 35-45, but the adjustment is small.

mates, but introduces distracting jumps in life-cycle consumption and debt as consumers respond to short-term budget changes.

We estimate the initial liquid funds to income ratio of the population from the SCF for households age 20-27. Liquid funds are liquid financial assets (savings, checking, money market, stocks, bonds) plus available credit (credit limit minus debt). Divided by one quarter normal income, we calculate that these households have an average of 1.28 quarters of normal income in available liquid funds.

While the population average income follows the observed life-cycle path, individual incomes vary based on their idiosyncratic shocks. We use the estimates of the annual income process from Gourinchas and Parker (2002), which are updates of Carroll and Samwick (1997), calculated from the Panel Study of Income Dynamics. We adjust these volatilities for quarterly dynamics so that four quarterly shocks combine to produce the same variance as one yearly shock. A close approximation is that the quarterly transitory variance must be four times the annual variance since quarterly shocks average out, while the quarterly permanent variance must be one fourth the yearly variance since permanent shocks stack.

We observe two prices directly, although there is likely greater heterogeneity in the prices actually paid than we incorporate in the model. We set the interest rate on debt $R_b = 14.11$ based on the average revolving interest rate over the period, adjusted slightly for the lack of explicit default risk in the model.¹¹ Since there is only one riskless liquid asset other than the ability to borrow, the appropriate rate on savings is not obvious since we would like to capture the returns that people expect to receive on their savings. We therefore set the return on savings as the average return on an all bond portfolio from 1926 to 2015 as calculated by the mutual fund company Vanguard of 5.4 percent. We adjust both borrowing and saving prices for the geometric average inflation rate from 2000-2015 of 2.15 percent. In the expected growth over the life-cycle we also

¹¹We follow Fulford (2015) in using the Edelberg (2006) adjustment for default. The Federal Reserve series G19 (Commercial Bank Interest Rate on Credit Card Plans NSA) average over the period is 14.73 percent (the average credit card interest rate reported in the SCF is 14.22 percent). Based on calculating the risk of default from the PSID, Edelberg (2006) calculates the zero bankruptcy risk rate would be 0.62 percentage points lower, a smaller adjustment than in Angeletos et al. (2001), who adjust for default by 2 percentage points.

include expected real aggregate growth of 1.5%, the average compounded rate from 1947 to 2015 from the BEA (2009 chained dollars GDP per capita).¹²

5.2 Identification and estimation of the payments problem

The structure of the consumer's problem means that the payment decision is only influenced by the consumption decision through whether or not the consumer was revolving as of the period before. We can thus find the solution to the payments problem first, and then allow the solution to the payment problem to influence the consumption problem. The two parameters of the payment problem are exactly identified by two moments we observe in the Diary of Consumer Payment Choice: the fraction of consumption that revolvers and convenience users pay for with a credit card, along with the observable market rate on borrowing on a credit card.

Table 2 shows the fraction of all expenditures over a three day period that the nationally representative sample of consumers from the Dairy of Consumer Payment Choice put on a credit card. The average consumer pays for 17.2 percent of consumption with a credit card. Revolvers pay for slightly less at 15.6 percent and convenience users pay for slightly more at 18.2 percent. Dividing the population by age suggests that there may be more significant differences in the usage patterns for the young, but those in middle age have a fairly stable pattern of credit card use, with convenience users spending slightly more than revolvers at most ages. Interestingly, both revolvers and convenience users over 65 tend to spend more on a credit card. Perhaps this change comes from a shift in the type of expenditure after retirement.

The difference between revolvers and convenience users then exactly identifies the payment

¹²While each of these parameters is volatile and different agents may experience different prices, there is no estimation uncertainty about them and so we do not adjust the MSM variance-covariance matrix for them. Put a different way: the "population" average interest rate is not an estimate and so does not have a standard error (it may be a draw from a "super-population" process, but our economic agents live in this population, and so the only relevant price is the one they face). Instead, these parameters may be miss-specified, which is a modeling problem, not an estimation one. Economic agents may face uncertainty about their rate of return, inflation, and economic growth which can be modeled, although within our model all agents take these as given. Similarly, there are multiple ways of saving with different rates of return, and a single rate of return only imperfectly captures this heterogeneity.

We thank Chris Carroll for pointing out that even if we remove trends from life-cycle profiles, the economic decision of the agent includes expected aggregate growth.

model as figure 3 illustrates. Since revolvers lose the float on purchases, their payments use of a credit card is shifted down by an amount equal to the loss of float. The difference in the optimal fraction on a credit card π^R and π^C identifies the slope and intercept of the two lines. We use a simple Delta method expansion to calculate the standard errors for the payment parameters ν_0 and ν_1 . Appendix C shows these calculations and Table 2 shows the estimated coefficients with an interest rate on borrowing of 14.11 percent adjusted for inflation of 2.15 percent (see discussion in section 5.1).

The model then directly gives the convenience value of credit credit cards. At close to 12 percent real borrowing rate, the convenience value of using a credit card for payments over other methods is worth 0.319 percent of consumption to convenience users and 0.235 percent to revolvers. This calculation is the consumer surplus of credit cards as a payment mechanism and so does not directly calculate welfare. The average convenience depends on the fraction of revolvers and convenience users. In addition, the calculation does not take into account the costs of operating the payments system or the producer surplus from additional sales made because some purchases are more convenient, or the gains to the processors, network operators, and banks. Nonetheless, a 0.319 percent multiple of all consumption is a substantial number, and since it is produced from only the 18.2 percent of consumption paid for with a credit card, suggests that the value that consumers get from using credit cards for the purchases they do use them for is very large. The value of the intercept ν_0 , for example, suggests that for the most valuable purchases, using a credit card has a convenience value of 4.1 percent of all consumption for these purchases. For comparison, if all convenience consumers received the equivalent of one percent cash back on their purchases with credit cards, the implied consumer surplus is 0.182 percent of consumption. This calculation likely overstates the direct value of rewards since not all cards offer rewards, but suggests that in the neighborhood of half of the convenience value from credit cards comes from direct rewards and the other half from their value as a convenient payment mechanism.

5.3 Identification and estimation of the life-cycle model

Using the first stage estimates of the payments problem and the other parameters, we next estimate the full life-cycle problem and then discuss the variation that helps identify the different parameters.¹³ Table 3 shows the model estimates while figure 4 shows the how debt and consumption vary over the life-cycle. Since the scale of the two panels of figure 4 is in logs, the MSM is essentially finding the parameters so that the sum of the squared differences between the predicted consumption and debt lines is as small as possible. It is clear that given the constraints of the optimization model, the model estimates can successfully capture the life-cycle profiles of debt and consumption. To do so, the model suggests that about half the population (f^A) must be impatient and almost risk neutral. This half of the population, which the figure and tables call population A, starts life with a substantial amount of debt (λ_0^A) and has substantial revolving debt throughout the life-cycle. In order to match the amount of debt and consumption, the estimates suggest that this population has a slightly higher income than the average (ζ^A) and so it is consuming almost all of its income in any given quarter. This population thus lives essentially hand-to-mouth over the entire life-cycle. It relies on credit for all of its smoothing, although its average utilization is high through much of the life-cycle (the third panel in figure 4), and so it is not keeping much credit for precautionary reasons either.

The estimates suggest that the other half of the population is relatively patient and risk averse. Its risk aversion is lower than many estimates which are often around 2, but are higher than the estimates in Gourinchas and Parker (2002). Population B is too patient to ever want to hold much debt and quickly repays the initial debt which the estimates suggest it holds (λ_0^A). After about age 30, this population only uses credit cards for payments.¹⁴ Since the estimates suggest that this population expects to receive little income after expenses (λ_1^B) in late life and it is relatively

¹³In robustness checks, it appears that the quadratic function in equation (4) has multiple local minima. We discuss what appears to be the best local minimum here, but expect that in future versions we will provide greater evidence that it is indeed the best minimum based on starting the estimation at a grid of starting points. We also plan to characterize other minima and how they help understand identification.

¹⁴The added debt from convenience use of credit cards is one month worth of consumption (1/3 of quarterly consumption) times the estimated rate of consumption on a credit card for a convenience user from table 2.

patient, almost the entire life-cycle is spent acquiring assets for old age. Consumption is therefore relatively flat over the entire life-cycle since it is attempting to smooth over this fall in income after expenses.

Since this is a highly non-linear model, all moments are typically used to identify all parameters, but it is useful to understand how different sources of variation identify the parameters. Both the consumption and debt that we observe over the life-cycle are averages over the entire population, and so the model is identified off of the average of the two populations. The estimates show that for the model to produce as much debt as in the data, approximately half of the population must be relatively impatient and not much concerned about debt. Since this population lives nearly hand-to-mouth, its consumption closely follows income over the life-cycle. To get an average consumption profile in which consumption is below income for much of the life-cycle therefore requires the other half of the population to be relatively patient and risk averse. This population is too patient to keep much debt at high interest rates, and pays off almost all of credit card debt it has early in life by its mid 20s. Most of the life-cycle dynamics of debt therefore come from the impatient population.

While the level of consumption and debt comes from the average, the life-cycle profiles are largely determined by only one of the populations. Since the patient population carries almost no debt, the profile of credit card debt largely identifies the preferences of the impatient population, their initial wealth, and their expected residual income late in life. The discount rate of this population is largely identified by the real interest on credit card debt; if the population is too patient, then it will not accumulate enough debt, too impatient and it will acquire too little debt. For this impatient population, the hump shape of debt comes from increases in credit limits early in life which allows it to increase its debts, and the fall in income late in life which makes carrying debt more painful. The risk aversion is identified by how much credit this population keeps as a buffer.

The more patient and risk averse population carries little debt past its 20s, so its preferences are largely identified by the consumption profile. The impatient population A has a strong hump in consumption as it follows income. For the average consumption profile to be below average

income, the patient population B must have a relatively flat consumption profile. Its preferences for risk (γ^B), discount rate (β^B), and expected late life income after expenses (ζ^B) are determined by this shape. The discount rate is also closely identified by the rate of return on savings.

The estimated model matches other life-cycle profiles such as the utilization and the wealth path well as shown in the remaining panels of figure 4. Credit utilization starts high and slowly falls over the life-cycle. The estimates suggest that utilization is on average slightly lower than we observe in the data. To examine the evolution of wealth, which may be negative, we take the log of wealth after giving everyone \$10,000, which allows us to consider the full distribution in a single graph. The model estimates predict slightly less wealth accumulation over the life-cycle than the SCF, but predicts the same trend increase and flattening after age 55. The model was not estimated to match these two profiles, and so being able to successfully predict their evolution suggests that the model is capturing important facets of life-cycle decision making. The heterogeneity in preferences is key here. Gourinchas and Parker (2002) estimate parameters to match the consumption profile and vastly under-predict wealth accumulation, while Cagetti (2003) estimates parameters to match the wealth profile, but needs a high degree of risk-averse that he cannot capture the consumption profile.

The model predicts that the fraction revolving will be approximately constant over the life-cycle, while the SCF suggests it should decline nearly continuously.¹⁵ The population average fraction revolving in the model is approximately correct, but since the impatient population A is always in debt, and the patient population B almost never revolves, the fraction revolving does not change much over the life-cycle. It is possible that additional heterogeneity across the population could replicate this behavior by allowing part of the population to be on the edge between revolving and not, with the balance shifting over the life-cycle. The problem with this approach is that the difference between the savings interest rate and the borrowing rate is so large that it is only for

¹⁵Since the SCF sometimes has trouble with debts (Zinman, 2009), and is at a household rather than individual level, there is reason to question whether the SCF fraction revolving is the best benchmark. Our estimates from the Dairy of Consumer Payments Choice in appendix figure A-1 suggest that the the fraction revolving is approximately constant until age 50, but then declines steadily. While this profile is closer to the model prediction, the model is still not predicting a movement out of revolving late in life.

very particular preferences that someone who is not willing to borrow late in life would be willing to borrow early in life. Such particular preferences are very sensitive to changes in interest rates, which suggests that credit use is sensitive to interest rates. The stability of utilization in 1 despite large changes in interest rates over the period suggests that the answer must lie elsewhere. Instead, it seems likely that there is some sort of preference heterogeneity over the life-cycle. Young people who think themselves invincible may have low risk aversion, since they have not seen the bad consequences of poor planning, but become increasingly patient and risk averse as they age and experience bad shocks.

The last two columns of table 3 show estimates that use a different weighting matrix and turn off the endogenous payment choice by making all consumers get the revolving value from credit use. Optimal weights give each moment a weight proportional to how well it is estimated. Since the debt moments are estimated from large sample of administrative data, the estimation tends to put more weight on minimizing the debt residuals. An alternative weighting scheme is to allow optimal weighting within each of the credit and debt moments, but renormalize so that the average moment for debt and consumption gets the same weight. This weight matrix produces nearly identical estimates. Similarly, the endogenous payments choice does not much affect the estimates. The reason is that since so few people switch from revolving to convenience use, the value an individual gets from credit card consumption this period is almost always the same as next period. In marginal utility the value of consumption on a credit card does not affect the tradeoff between today and the future. It makes not revolving slightly more valuable, but since hardly anyone is on the edge, this value does not matter. Including credit cards as a payment mechanism is important, however. The debts we observe in the credit bureau data include both revolving and convenience use and not accounting for the convenience use would make the average debt too high.

5.4 Model predictions for the business cycle and individuals

In this section, we take the model and estimated parameters and ask how well it predicts phenomena outside the life-cycle that it was not estimated to match. A good way to think of these results

is that they provide both an out-of-sample examination of how good the model estimates are and whether the model can successfully explain other phenomena. We take this examination as distinct from our discussion in the previous section of whether the model can replicate other life-cycle moments such as utilization and wealth which help validate the model. Showing that the model can help understand individual and business cycle dynamics, on the other hand, provides a test of the model predictions.

To perform this test we simulate a large population with the same age profile as in Equifax from age 24 to 74. We allow for the constant real growth of 1.5 percent the consumers in the model assume, and adjust the dollar values for the average inflation rate. Finally, to mimic the fall in credit limits that started in the last quarter in 2008 and went through 2009, we introduce a fall in credit of 40 percent to one sixth of the population for six quarters. This particular experiment is the simplest way to get the approximately 40 percent drop in credit limits spread over more than a year evident in figure 1. It is not a full replication of the changing environment. In particular, it does not include a fall in income or a decline in the confidence about future income growth.

The individual dynamics of credit utilization from the simulations closely match the dynamics from the credit bureau data. In table 1 we showed that, once we control for fixed unobserved heterogeneity with fixed effects, shocks to utilization disappear quickly, with 64.7 percent of a shock surviving each quarter (the third column). The last column performs exactly the same regression on the simulated data. The simulated consumers experience the large unexpected fall in credit in over 2009, and the expected increase over the life-cycle, but the only unexpected credit volatility that they face comes because credit is proportional to volatile permanent income. Since volatility in income is much less than volatility in credit (Fulford, 2015), the consumers in the model face less credit volatility than actual consumers over the time period. Nonetheless, their response to changes in credit limits are very close to actual consumers; the estimated model captures the dynamics of credit utilization nearly exactly with 68.3 percent of a shock persisting to the next quarter.

We finally examine the business cycle dynamics with which we started the paper. The bottom panel of figure 1 shows the aggregate response of the simulated consumers to the 40 percent

fall in credit introduced over six quarters. Credit increases at the same 1.5 percent real rate as income, plus 2.1 percent for average inflation. Model credit growth is slightly slower than the actual credit growth over the period suggesting that pegging credit to income does not fully capture the aggregate growth. Since we already capture the much larger life-cycle changes, the remaining unexplained aggregate growth during the period is unimportant for individual decisions. Since consumers expect credit growth, their debts grow at the same time, and credit utilization is stable despite the growth before and after the crisis just as in the data. The simulated population has a slightly lower credit utilization over the life-cycle (see figure 4), and so a lower credit utilization in aggregate.

During the crisis, debt quickly adjusts to the fall in credit, and so utilization is much smoother than either credit or debt. As the individual dynamics show, while shocks at the individual level disappear quickly in both the model and data, it still takes several quarters for consumers to fully adjust their debt and savings to a 40 percent fall in credit. The excessive smoothness of utilization in the credit bureau data suggests that there may be additional features of the period not captured by the simple simulated shock, since the simulated data closely matches the individual responses to a fall in credit. The simulation does not include any decline in income which, if expected to be permanent, would also decrease debt. The fall in credit was also more concentrated among some consumers. While initially the decline was concentrated among those with high credit scores, later declines were to those with lower credit scores (Fulford, 2015), whose responses to a fall in limits are even faster than average (see our working paper (Fulford and Schuh, 2015) for more evidence on heterogeneity of responses).

6 Conclusion

Credit is more volatile than income over both the business cycle and life-cycle. Credit increases very quickly early in life, yet since debt follows these increases, credit utilization decreases only slowly over the life-cycle. Credit utilization is also remarkably stable across the business cycle

despite a nearly 40 percent drop following the financial crisis in 2008. We show that this stability comes from individual decisions as shocks to credit utilization quickly disappear.

We introduce and estimate a model of life-cycle consumption, saving, and debt that explains this stability. Credit plays a precautionary role in the model, and so consumers manage it just as they would any buffer of savings. Credit cards are also used for payments and this form of debt is smoothed like consumption. Both aspects points to stable credit use. We estimate the payments aspect using the differing credit card use for payments of revolvers and convenience users and find that the payments value of credit cards is worth 0.319 percent of consumption for convenience users. Estimating the rest of our model off of the life-cycle of debt and consumption, we show that it can produce the life-cycle profiles of consumption and debt. It also predicts the stable utilization over the business cycle and for individuals. While the estimated model does capture the slow fall in utilization over the life cycle, it predicts a constant fraction of the population revolving, rather than a steady decline over the life cycle. With fully forward looking but very impatient consumers, life-cycle concerns are not sufficient to get those willing to borrow in early life to stop borrowing at all late in life. Instead, the decline suggests that some additional form of learning or preference changes is needed.

One of the startling facts of consumer finance is how little the median household saves, particularly in liquid form. The increases in credit over the life-cycle help to explain why consumers save so little early in life. In expectation, the credit limit for the average person in her 20s will be much higher next year. Since credit is a form of wealth for precautionary purposes (Fulford, 2013), the young are effectively becoming wealthier through credit increases, reducing their need to save. In middle age, on the other hand, many households have substantial debt. For these households, saving should be mostly about paying down previous debt. Paying off credit card debt has a riskless return that averages 14 percent, which no other asset class can match. The large life-cycle variation in credit and debt suggests why the average household has little in positive assets beyond a small emergency fund and illiquid housing until very late in life: it generally has substantial and increasing credit, and can save by paying down debt. Not taking this variation into account leads

to a misunderstanding of the financial status of U.S. households.


References

- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. 2015. "Do Banks Pass Through Credit Expansions? The Marginal Profitability of Consumer Lending During the Great Recession." Technical report, SSRN.
- Agarwal, Sumit, John C. Driscoll, Xavier Gabaix, and David Laibson. 2007. "The Age of Reason: Financial Decisions Over the Lifecycle." Working Paper 13191, National Bureau of Economic Research.
- 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.
- Athreya, Kartik B. 2008. "Default, insurance, and debt over the life-cycle." *Journal of Monetary Economics*, 55(4): 752 – 774.
- Attanasio, Orazio P., James Banks, Costas Meghir, and Guglielmo Weber. 1999. "Humps and Bumps in Lifetime Consumption." *Journal of Business & Economic Statistics*, 17(1): pp. 22–35.
- Attanasio, Orazio P., and Martin Browning. 1995. "Consumption over the Life Cycle and over the Business Cycle." *The American Economic Review*, 85(5): pp. 1118–1137.
- Avery, Robert B., Paul S. Calem, and Glenn B. Canner. 2004. "Credit Report Accuracy and Access to Credit." *Federal Reserve Bulletin*, Summer: 297–322.
- Aydin, Deniz. 2015. "The marginal propensity to consume out of liquidity: Evidence from a randomized controlled trial." 2015 Meeting Papers 270, Society for Economic Dynamics.
- Cagetti, Marco. 2003. "Wealth Accumulation Over the Life Cycle and Precautionary Savings." *Journal of Business and Economic Statistics*, 21(3): 339–353.
- Campbell, John Y., and N. Gregory Mankiw. 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." *NBER Macroeconomics Annual*, 4: 185–216.
- Campbell, John Y., and N Gregory Mankiw. 1990. "Permanent Income, Current Income, and Consumption." *Journal of Business & Economic Statistics*, 8(3): 265–79.
- Carroll, Christopher D. 1997. "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics*, 112(1): 1–55.
- Carroll, Christopher D. 2006. "The method of endogenous gridpoints for solving dynamic stochastic optimization problems." *Economics Letters*, 91(3): 312 – 320.
- Carroll, Christopher D. 2012. "Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems." Available: <http://www.econ2.jhu.edu/people/ccarroll/SolvingMicroDSOPs/>, accessed 3 June 2015.


- Carroll, Christopher D., and Andrew A. Samwick. 1997. "The nature of precautionary wealth." *Journal of Monetary Economics*, 40(1): 41–71.
- Cocco, João F., Francisco J. Gomes, and Pascal J. Maenhout. 2005. "Consumption and Portfolio Choice over the Life Cycle." *Review of Financial Studies*, 18(2): 491–533.
- De Nardi, Mariacristina, Eric French, and John B. Jones. 2010. "Why Do the Elderly Save? The Role of Medical Expenses." *Journal of Political Economy*, 118(1): pp. 39–75.
- Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica*, 59(5): 1221–1248.
- Deaton, Angus. 1992. *Understanding Consumption*. Oxford: Clarendon Press.
- Deaton, Angus. 1997. *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. Baltimore: Johns Hopkins University Press.
- Deaton, Angus. 2010. "Understanding the Mechanisms of Economic Development." *The Journal of Economic Perspectives*, 24(3): pp. 3–16.
- Edelberg, Wendy. 2006. "Risk-based pricing of interest rates for consumer loans." *Journal of Monetary Economics*, 53(8): 2283 – 2298.
- Fulford, Scott. 2013. "The effects of financial development in the short and long run: Theory and evidence from India." *Journal of Development Economics*, 104: 56–72.
- Fulford, Scott L. 2015. "How important is variability in consumer credit limits?" *Journal of Monetary Economics*, 72: 42 – 63.
- Fulford, Scott L., and Scott Schuh. 2015. "Consumer revolving credit and debt over the life-cycle and business cycle." April, In preparation.
- 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.
- Iacoviello, Matteo, and Marina Pavan. 2013. "Housing and debt over the life cycle and over the business cycle." *Journal of Monetary Economics*, 60(2): 221 – 238.
- Kaboski, Joseph P., and Robert M. Townsend. 2011. "A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative." *Econometrica*, 79: 1357–1406.
- Karlan, Dean, and Jonathan Zinman. 2010. "Expanding Credit Access: Using Randomized Supply Decisions to Estimate the Impacts." *The Review of Financial Studies*, 23(1): 433.
- Laibson, David, Andrea Repetto, and Jeremy Tobacman. 2007. "Estimating Discount Functions with Consumption Choices over the Lifecycle." Working Paper 13314, National Bureau of Economic Research.

- Lawrence, Emily C. 1995. "Consumer Default and the Life Cycle Model." *Journal of Money, Credit and Banking*, 27(4): 939–954.
- Lee, Donghoon, and Wilbert van der Klaauw. 2010. "An Introduction to the FRBNY Consumer Credit Panel." Staff Report 479, Federal Reserve Bank of New York, Available: http://newyorkfed.org/research/staff_reports/sr479.pdf, accessed 6 April 2015.
- Leth-Petersen, Søren. 2010. "Intertemporal Consumption and Credit Constraints: Does Total Expenditure Respond to an Exogenous Shock to Credit?" *The American Economic Review*, 100(3): 1080–1103.
- Lopes, Paula. 2008. "Credit Card Debt and Default over the Life Cycle." *Journal of Money, Credit and Banking*, 40(4): 769–790.
- Ludvigson, Sydney. 1999. "Consumption and Credit: A Model of Time-Varying Liquidity Constraints." *Review of Economics and Statistics*, 81(3): 434–447.
- McFadden, Daniel. 1989. "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration." *Econometrica*, 57(5): 995–1026.
- Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review*, 103(6): 2530–53.
- Schuh, Scott, and Joanna Stavins. 2014. "The 2011 and 2012 Surveys of Consumer Payment Choice." Research Data Reports 14-1, Federal Reserve Bank of Boston.
- Stango, Victor, and Jonathan Zinman. 2009. "What Do Consumers Really Pay on Their Checking and Credit Card Accounts? Explicit, Implicit, and Avoidable Costs." *The American Economic Review Papers and Proceedings*, 99(2): 424–429.
- Stepner, Michael. 2013. "BINSCATTER: Stata module to generate binned scatterplots." Statistical Software Components, Boston College Department of Economics, October.
- 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.
- Zinman, Jonathan. 2009. "Where is the missing credit card debt? Clues and implications." *Review of Income and Wealth*, 55(2): 249–265.
- Zinman, Jonathan. 2013. "Consumer Credit: Too Much or Too Little (or Just Right)?" Working Paper 19682, National Bureau of Economic Research.
- Zinman, Jonathan. 2015. "Household Debt: Facts, Puzzles, Theories, and Policies." *Annual Review of Economics*, 7: 251–276.

Figure 1: Credit card limits, debt, and utilization: 2000–2014
Observed limits, debts, and utilization




Model prediction given fall in credit limits




Notes: The left axis shows the average credit card limits (top line) and debt (bottom line). Note the log scale. The right axis shows mean credit utilization (middle line) defined as the credit card debt/credit card limit if the limit is greater than zero. Source: Authors' calculations from Equifax/NY Fed CCP.

Figure 2: Credit card limits, debt, and credit utilization





Notes: Each line represents the average credit card limit, debt, and utilization of one birth cohort, 1999–2014. Source: Authors' calculations from Equifax/NY Fed CCP.

Figure 3: Marginal value of additional consumption on a credit card




Notes: This figure illustrates how the marginal value of additional consumption on a credit card changes with the share spent on a credit card. The top line is for convenience users who put a share π^c of consumption on a credit card. The bottom line for revolvers is shifted down by the amount $-r_b/24$ since revolvers lose the float on payments made using credit cards. They therefore put a smaller share on their credit cards π^R .


Figure 4: Consumption and debt over the life-cycle: model estimates
 Estimation moments: Debt Estimation moments: Consumption




Estimation predictions: Utilization




Estimation predictions: Fraction revolving



Estimation predictions: Wealth path



Estimation predictions: Variance of debt



Notes:

Table 1: Credit utilization

	Equifax/NY Fed CCP				Model	
	Credit utilization _t					
Credit utilization _{t-1}	0.874*** (0.000876)	0.868*** (0.000892)	0.647*** (0.00131)	0.647*** (0.00139)	0.514*** (0.00441)	0.683*** (0.000329)
Credit utilization _{t-2}				0.0156*** (0.000643)		
Credit util _{t-1} × Age					0.00314*** (9.93e-05)	
Constant	0.0479*** (0.000461)					
Observations	347,642	347,642	347,642	332,696	347,642	4,752,000
R-squared	0.741	0.743	0.429	0.444	0.431	0.489
Fixed effects	No	No	Yes	Yes	Yes	Yes
Age and year effects	No	Yes	Yes	Yes	Yes	Yes
Number of accounts			10,451	10,103	10,451	102,564
Frac. Variance from FE			0.477	0.467	0.498	0.292

Notes: The sample includes 0 credit utilization but excludes individual quarters where the utilization is undefined since the limit is zero and when utilization is greater than 5 (a very small fraction, see distributions of utilization in figure). All columns include age and year effects, with age effects normalized to have zero trend when fixed effects are included. Source: Authors' calculations from Equifax/NY Fed CCP.

Table 2: Fraction of consumption on a credit card and value for payments

	Fraction on Credit card	Std. error	Std. dev.
All consumers	0.172	0.0082	0.310
All revolvers	0.156	0.0130	0.283
All convenience users	0.182	0.0105	0.324
Revolvers			
Age < 25	0.234	0.0760	0.301
25 <= Age < 35	0.157	0.0287	0.299
35 <= Age < 45	0.134	0.0344	0.282
45 <= Age < 55	0.142	0.0213	0.142
55 <= Age < 65	0.144	0.0207	0.144
Age >= 65	0.202	0.0399	0.312
Convenience Users			
Age < 25	0.065	0.0233	0.209
25 <= Age < 35	0.169	0.0255	0.169
35 <= Age < 45	0.155	0.0228	0.291
45 <= Age < 55	0.173	0.0215	0.327
55 <= Age < 65	0.216	0.0242	0.345
Age >= 65	0.243	0.0271	0.350
Model Estimates			
Level ν_0	0.035	0.0216	
Slope ν_1	0.194	0.1259	
Implied value of Credit Card use (in percent of consumption)			
Revolvers	0.235%	0.1512	
Convenience users	0.319%	0.0962	

Notes: Author's calculations from the Diary of Consumer Payment Choice. The standard errors are calculated by bootstrapping.

Table 3: Model estimates

Population A			
CRRA γ^A	0.154	0.164	0.176
	(0.015)	(0.018)	(0.019)
Discount β^A	0.889	0.889	0.889
	(0.000)	(0.000)	(0.000)
Initial wealth λ_0^A	-5.320	-2.725	-2.734
	(1.031)	(0.091)	(0.094)
Late life inc. λ_1^A	0.515	0.522	0.545
	(0.009)	(0.009)	(0.008)
Population B			
CRRA γ^A	0.936	0.868	0.958
	(0.202)	(0.217)	(0.249)
Discount β^B	0.961	0.963	0.962
	(0.003)	(0.003)	(0.004)
Initial wealth λ_0^B	-0.072	-0.090	-0.120
	(0.142)	(0.138)	(0.155)
Late life inc. λ_1^B	0.070	0.079	0.080
	(0.074)	(0.068)	(0.072)
Share A f^A	0.557	0.557	0.557
	(0.001)	(0.002)	(0.002)
Inc. mult. A ζ^A	1.161	1.193	1.206
	(0.025)	(0.027)	(0.027)
Weights	Optimal	Equal	Optimal
Endogenous payments	Yes	Yes	No

Notes: Standard errors in parentheses. Equal weights are optimal within each of the life-cycle debt and consumption moments, but equal weight for each set of moments. Optimal weights are the inverse of the variance of each individual moment. For a brief description of each moment and the estimation method see the beginning of section 5.

Appendix

For Online Publication

A Changes in credit utilization: non-parametric evidence

Figure A-7 shows conditional mean scatter plots of credit utilization in one quarter against credit utilization in the next quarter, in the next year, and in two years. The top row shows the mean in the future, conditional only on having the utilization shown on the x-axis in that quarter. The bottom row instead takes the within transformation and allows for age and year effects. It therefore shows how far from the individual's average credit utilization she is in the next quarter, conditional on differing from her average utilization by the amount on the x-axis this quarter. In other words, if an individual is 10 percentage points above her typical utilization in one quarter, how far will she be on average in the next quarter, next year, and in two years? Each dot contains an equal portion of the sample. Figure A-7 thus captures the relationship between utilization today and in the future without imposing any parametric assumptions. Each panel also shows the best fit line for the conditional means and the estimated coefficients.

The top panels show that credit utilization is highly persistent and does not tend to zero on average. Credit utilization this quarter is typically very close to credit utilization next quarter, since the conditional means are typically very close to the 45-degree line. For example, on average if a person is using 40 percent of her credit this quarter, she will be using about 40 percent of her credit next quarter. Looking closely, however, average credit utilization is higher next quarter for those using less than 20 percent of their credit, and lower for those using more than 80 percent of their credit. The best fit line through the conditional means suggests that credit utilization is not trending to zero. Instead, the long-term steady-state utilization is 0.39.¹⁶ The same conclusion is evident from the conditional changes comparing utilization this quarter to a year from now and to

¹⁶Since the conditional expectation of utilization next quarter given this quarter is $u_{t+1} = 0.041 + 0.896u_t$ the steady-state utilization is $0.39=0.041/(1-0.896)$.

two years from now. Those consumers using less than approximately 40 percent of their available credit this quarter are using more of their credit in one year and in two years, those using more than 40 percent of their credit are using less of their credit on average within one year and two years. The steady-state credit utilization is around 40 percent (evident by finding where the conditional expectation function crosses the 45 degree line), although the movement toward the steady state is fairly slow.

On average individuals do not trend to zero utilization, nor to using all of their credit. Conditional on using zero credit this quarter, credit utilization is nearly 5 percent within one quarter and nearly 8 percent in a year. On the other hand, the average person using all of her credit in one quarter is using less than 90 percent of it in a year.

The second row of figure A-7 allows individuals to return to their own mean and adds substantial nuance. The reason credit utilization is so persistent in the top row is that it appears that individuals have their own mean to which they actually return quite rapidly. The speed of the return is evident from the slopes of the lines. Only two-thirds of a shock to utilization remains after one quarter and 13 percent remains after two years.

Even if individuals return very rapidly to their own means, it is important to note that those means are not zero. Credit utilization is persistent in the top row of figure A-7 because individuals are typically quite close to their own mean credit utilization. Since credit utilization is the ratio of debt and credit, the stability of credit utilization implies that an individual with an increase in credit has increased her debts by 33 percent of the increase in credit within one quarter, and 87 percent of the increase in credit in two years.

B Bequests

In our base estimation, we give consumers who die with positive resources some extra utility from bequests. We have also examined the consequences of allowing consumers to place no value on bequests. Since our estimation uses debt and consumption moments, bequests are not an important

factor in the estimation since the median person is leaving very little in bequests. We have found that any reasonable approach to introduce bequests makes only a very minor difference in our estimates. This appendix outlines one flexible approach, and discusses how the parameter estimates change when not including bequests.

The difficulty with introducing bequests in a model with both debt and savings is that while positive assets left to heirs may provide utility, unsecured debts are not passed on. Since it is possible to die with negative assets, zero bequests are possible. Simply including bequests in the sub-utility function will therefore produce negative infinite value from leaving no bequest, which consumers will counterfactually act to avoid by never being in debt. Instead, we model the bequest motive as the consumer considering the marginal utility a bequest to her heirs will bring them on top of their own incomes. The consumer dies with certainty at age \bar{T} , and with some probability before that. To value bequests reasonably, we take positive assets as providing an annuity value on top of some base income that heirs already get that is a function of the consumer's permanent income, and the bequest value as the present value over a lifetime of these additional assets. The value of bequests is then:

$$S(A_t) = \left(\sum_{s=0}^{\bar{T}} \beta^s \frac{(\zeta P_t + r_B(1 - I_{t+1}^R)A_t)^{1-\gamma}}{1 - \gamma} \right),$$

where the consumer only leaves a bequest if she has positive assets at the end of the period (I_{t+1}^R is 0 if $A_t \geq 0$). The parameter ζ determines the marginal utility of bequests and can be thought of as how much more or less income children have than parents.

This structure is important since people can decide to die in debt, and so leave nothing. If children are much better off than parents, then the marginal utility of leaving a bequest is low, and parents will generally only leave a large bequest accidentally. In our baseline estimation with bequests, we set $\zeta = 2$ which implies that at the average age of death children are earning twice as much as their elderly parents, but we have also estimated ζ , and found that while values around 2 are reasonable, ζ is very poorly identified given the moments we use. Even very late in life,

the likelihood of death in the next year is fairly low, and so the exact form of the bequest motive, including removing it, does not affect decisions meaningfully for most of the distribution. The reason to use finite children's lifetime is to allow for possibly very patient parents with β close to or greater than 1.

C Identification of the payments model

This section shows how to identify the payment model parameters and standard errors from observable moments. It then calculates the consumer surplus and its standard errors. We observe:

$$\pi^R = \frac{1}{N} \sum_{i=1}^N \pi_{i,t}^* | I_{i,t-1}^R = 1$$

the average expenditures by revolvers on a credit card, and similarly π^C the average for convenience users. Denote our estimates of the standard errors of these means as σ_R and σ_C . Then the intercept for the average convenience user is just

$$\pi^C = \nu_0 / \nu_1$$

and for a revolver it is

$$\pi^R = (\nu_0 - r^B/24) / \nu_1$$

where r^B is the APR interest charged on payments, which have an average daily balance of half of the months consumption. Solving for ν_0 and ν_1 gives:

$$\nu_1 = \frac{r^B/24}{\pi^C - \pi^R}$$

and

$$\nu_0 = \pi^C \nu_1 = \frac{(r^B/24)\pi^C}{\pi^C - \pi^R}$$

The presence of a difference of two random variables whose supports may overlap in the denominator of the transformed variables makes calculating their variances potentially tricky. Since $\pi^C - \pi^R$ is may be close to 0, then ν_1 and ν_0 may be very large, which is a different way of saying that the model is not identified if there is not a difference in the average behavior of convenience users and revolvers. We calculate the standard errors of the transformed variables using the delta method, which avoids this issue by examining only small changes around the optimum, and so does not consider the highly non-linear increase around $\pi^C - \pi^R = 0$. For small changes ϵ^C and ϵ^R around π^C and π^R :

$$\nu_1 \approx (r^B/24) \left(\frac{1}{\pi^C - \pi^R} - \frac{\epsilon^C - \epsilon^R}{(\pi^C - \pi^R)^2} \right).$$

Since π^R and π^C are independent, the variance of ν_1 is approximately:

$$Var[\nu_1] \approx \left(\frac{r^B/24}{(\pi^C - \pi^R)^2} \right)^2 (\sigma_C^2 + \sigma_R^2).$$

Taking the same expansion for ν_0 including the covariance of the numerator and denominator:

$$Var[\nu_0] \approx \left(\frac{\pi^C r^B/24}{(\pi^C - \pi^R)^2} \right)^2 \left(\sigma_C^2 + \left(\frac{\pi^R}{\pi^C} \right)^2 \sigma_R^2 \right).$$

Finally, the total additional convenience value of using a credit card over the alternatives for a convenience user is just the area under the curve:

$$\nu^{C=} \nu(\pi^C; I_{i,t-1}^R = 0) = (\pi^C \nu_0)/2 = \frac{(r^B/48)(\pi^C)^2}{\pi^C - \pi^R}$$

and for a revolvers is:

$$\nu^R = \nu(\pi^R; I_{i,t-1}^R = 1) = \frac{\pi^R}{2} \left(\frac{(r^B/24)\pi^C}{\pi^C - \pi^R} - \frac{r^B}{24} \right).$$


Taking an expansion around π^C and π^R yields:

$$Var[\nu^C] \approx \left(\frac{r^B}{48}\right)^2 \left(\frac{2\pi^C}{\pi^C - \pi^R} - \left(\frac{\pi^C}{\pi^C - \pi^R}\right)^2\right)^2 \sigma_C^2 + \left(\frac{\pi^C}{\pi^C - \pi^R}\right)^2 \sigma_R^2$$

and


$$Var[\nu^R] \approx \left(\frac{r^B}{48}\right)^2 \left(\frac{\pi^C}{\pi^C - \pi^R} - \frac{\pi^C \pi^R}{(\pi^C - \pi^R)^2}\right)^2 \sigma_C^2 + \left(\frac{\pi^C}{\pi^C - \pi^R} - \frac{\pi^C \pi^R}{(\pi^C - \pi^R)^2} - 1\right)^2 \sigma_R^2.$$

Figure A-1: Fraction of convenience users of credit cards, conditional on use




Notes: Each dot shows the fraction in that five year age group (labeled by the youngest member) who did not revolve credit card debt from month to month, 'conditional on using their credit card. 95 percent confidence intervals in bars. Source: Authors' calculations from the 2008–2014 Surveys of Consumer Payment Choice.

Figure A-2: Consumption functions over the life-cycle with borrowing




Notes: Shows the consumption function for a consumer at different points of the life-cycle. At w^* the consumer no longer consumes all cash-at-hand (which includes all available credit). At w^B the consumer no longer borrows, but does not save anything between w^B and w^F since the interest rate for borrowing is higher than the rate for saving. After w^F , the consumer saves (at the lower rate of return). The parameters to produce these functions are a savings rate of return $R = 1.03$ and borrowing rate $R^B = 1.14$, a coefficient of relative risk aversion of $\gamma = 2$, a discount parameter $\beta = 0.9$ and other parameters and retirement functions matching Gourinchas and Parker (2002).

Figure A-3: Credit card adoption and use for payments by age and year
 (A) Credit card adoption rate




(B) Share of payments using credit cards




Source: Authors' calculations from the 2013 Survey of Consumer Payment Choice.


Figure A-4: Credit card limit, debt, and credit utilization distributions by age
 (A) Credit Card Limits



(B) Credit Card Debt




(C) Credit utilization




Notes: Each line is the percentile of credit limit at that age, conditional on having a positive credit limit on a log scale. For example, the 90th percentile line shows that 10 percent of the population (with a positive credit limit) have a limit larger than that line. Source: Authors' calculations from Equifax/NY Fed CCP.

Figure A-5: Credit card limits, debt, and utilization: age and year effects




Notes: Each line shows the estimated age or year effects from equation (1). Note the different scales. Source: Authors' calculations from Equifax/NY Fed CCP.

Figure A-6: Fraction with positive credit card limit and debt by cohort and age
Fraction with positive limit Fraction with positive debt




Notes: Each line represents the fraction with positive credit card limits or debt of one birth cohort, 1999–2014. Source: Authors' calculations from Equifax/NY Fed CCP.

Figure A-7: Changes in credit utilization in one quarter, one year, and two years



Notes: Each point in the top row shows the mean credit utilization in the future conditional on being in the bin with a mean credit utilization on x-axis today. The bottom row shows the conditional relationship between deviations from the individual mean utilization over the entire sample, adjusting for age and year. Source: Authors' calculations from Equifax/NY Fed CCP using the program binscatter (Stepner, 2013).

Figure A-8: Consumption and income over the life-cycle from the Consumer Expenditure Survey



Notes: Shows the average consumption and income at each age from the Consumer Expenditure Survey, pooling all surveys from 2000-2014. Consumption is the total household expenditures divided by the number of adults. Adjusted consumption removes the estimated effect of children. Income is after-tax income, and its smoothed version is based on a quintic from age 24 to 81. Since the CEX pools income and consumption after age 81 (or 83 in later years), ages 81 on the average over 81.