

Fulford, Scott L.; Schuh, Scott

Working Paper

Consumer revolving credit and debt over the life cycle and business cycle

Working Papers, No. 15-17

Provided in Cooperation with:

Federal Reserve Bank of Boston

Suggested Citation: Fulford, Scott L.; Schuh, Scott (2015) : Consumer revolving credit and debt over the life cycle and business cycle, Working Papers, No. 15-17, Federal Reserve Bank of Boston, Boston, MA

This Version is available at:

<https://hdl.handle.net/10419/171753>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Consumer Revolving Credit and Debt over the Life Cycle and Business Cycle

Scott L. Fulford and Scott Schuh

Abstract:

Little work has examined how unsecured consumer credit, such as the limit on credit cards, varies over the life cycle, and how consumers respond to changes in their ability to borrow over the short and long term. Using a large panel of credit accounts in the United States, we document that large movements in credit from 1999 to 2015 were accompanied by similar movements in debt, so utilization was nearly constant. Life-cycle variation in consumer credit is similarly large. Credit limits increase rapidly early in life, growing by more than 400 percent between the ages of 20 and 30. Debt grows almost as fast, however, so credit utilization falls slowly throughout the life cycle, reaching 20 percent only by age 70. Individual credit utilization is extremely stable despite the large life-cycle, business cycle, and individual volatility of credit. Stable utilization means that consumer debt is very sensitive to changes in credit limits. Distinguishing between consumers who revolve debt and those who use credit cards only for payments, we find that for revolvers nearly 100 percent of an increase in credit limits eventually becomes an increase in debt.

JEL Classifications: D14, D91, E21

Scott L. Fulford is an assistant professor at Boston College and a visiting scholar in the Consumer Payments Research Center of the research department of the Federal Reserve Bank of Boston. Scott Schuh is the director of the Center and a senior economist and policy advisor. Their e-mail addresses are scott.fulford@bc.edu and scott.schuh@bos.frb.org, respectively.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at <http://www.bostonfed.org/economic/wp/index.htm>. Scott L. Fulford would like to thank the Bank and the Center for their knowledge and help. Both authors thank David Zhang for his excellent research assistance. This paper has benefitted 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, and seminars at the Federal Reserve Bank of Boston and the Bank of Canada.

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 those of the authors and do not necessarily reflect the official position of the Federal Reserve Bank of Boston or the Federal Reserve System.

This version: October 2015

1 Introduction

Credit and debt dominate the financial lives of most U.S. households, and it is impossible to understand the consumption and savings decisions of these households without understanding the credit available to them and the debt they have accumulated. Credit card borrowing is the primary source of short-term consumption smoothing for U.S. households since the average household has very little liquid savings (Fulford, 2015b). Credit cards have become an important method of payment, so lacking access to credit limits payment options (Foster et al., 2013; Schuh and Stavins, 2014). More than 30 percent of individuals between the ages of 20 and 60 have an auto loan at any given moment, and close to 50 percent of 20–25 year-olds have a student loan (see figure 10). Paying down mortgage debt is the primary source of savings for the average household (Nothaft and Chang, 2005). The price of credit is the relevant opportunity cost (Zinman, forthcoming) and the availability of credit the relevant constraint for the average household when it considers its short-term consumption smoothing, long-term wealth accumulation, and housing consumption, payment choice, transportation, and human capital acquisition.

Despite the centrality of credit and debt in the financial lives of Americans, little is known about how U.S. consumers' access to and utilization of credit changes in the short and long term, and how changes in credit are related to changes in debt. While there is a large literature that considers which groups have access to credit and its cost, the literature that examines how credit or debt changes over time or with age is sparse. Yet consumer credit is extremely variable for individuals in the short term (Fulford, 2015a), and this variability fundamentally alters their savings and consumption decisions. How do changing credit conditions and use over the business cycle, over the life cycle, and for individuals relate to U.S. consumers' patterns of incurring, carrying, and paying off debt?

To answer this question, we use the Federal Reserve Bank of New York Consumer Credit Panel (CCP), which contains a 5 percent sample of every credit account in the United States from 1999 to 2014 from the credit reporting agency Equifax.¹ The panel nature of this data set is crucial, since it

¹For most of the econometric work, we use a subsample of the full 5 percent sample, since the smaller sample is

allows us to examine how short-term changes in credit and debt of an individual accumulate. The long panel then lets us examine how debt and credit evolve over the life cycle from these short-term changes. We supplement the CCP with the Survey of Consumer Payment Choice (Schuh and Stavins, 2014) and the Survey of Consumer Finances to draw a more complete picture. We focus on credit cards, since these have observable limits and are widely held, and we also briefly examine other forms of debt over the life cycle.

Credit and debt show extreme life-cycle variation, much larger than the changes in income or consumption over the life cycle (Attanasio et al., 1999). Between the ages of 20 and 30, credit card limits increase by more than 400 percent, and these continue to increase after age 30, although at a slower rate. Credit card debt increases at nearly the same pace early in the life cycle and it is only after age 50 that average credit card debt starts to decline. Other forms of debt, such as mortgages and auto loans, show similar life-cycle variation, while student loan debt peaks early in life. However, these types of debt do not have readily measurable limits. Aggregate consumer credit and debt show large business cycle variation as well, again substantially larger than the comparable income or consumption movements. For example, the average credit card limit fell by approximately 40 percent over the course of 2009.²

Despite this massive variation in credit and debt over the life cycle and business cycle, credit utilization is remarkably stable. Credit utilization is the fraction of available credit an individual is using. Credit utilization has held steady at just over 30 percent for the entire period from 2000 to 2014, even as there have been large swings in both credit and debts. Similarly, despite the 400 percent gains in credit and debt early in life, credit utilization falls very slowly over the life cycle. The mean credit utilization is 50 percent in the 20s and is still nearly 40 percent at age 50. Credit utilization falls below 20 percent only after age 70.

much more straightforward computationally and our confidence intervals are extremely small. Using a subsample will also allow us to perform checks of specification search and model selection bias by performing the same analysis on a different sample.

²See figure 2 and Ludvigson (1999) for a discussion of aggregate credit limit changes. Over the same period from their peak in approximately 2008:Q2 to the trough around 2009:Q2 both aggregate personal consumption and personal income fell by around 3.2 percent (see Federal Reserve Bank of St. Louis FRED, Personal Income and Personal Consumption Expenditure Series, <https://research.stlouisfed.org/fred2/>).

Individual credit utilization is extremely persistent as well, despite the individual volatility in credit limits documented by Fulford (2015a). We show both non-parametrically and through regressions that deviations from individual credit utilization disappear quickly, as individuals rapidly return to their typical individual-specific utilization. Credit utilization is not zero for most people, suggesting a strong tendency to use some of their available credit, and to return rapidly to their steady-state utilization following a shock to credit or debt.

Central to understanding the relationship between credit and debt is the fact that consumers use credit for distinct purposes. Some consumers use credit cards as a payment mechanism and pay their bills in full every month; these consumers are often called “convenience” users. Some consumers hold debt from month to month, and so are “revolvers.” We use basic consumption theory to help distinguish between convenience users and revolvers. For convenience users, reported debts are some fraction of consumption every month and so debts should take on the properties of consumption. If convenience users are smoothing their consumption well, then shocks to consumption should be unpredictable (see Deaton (1992) for an introduction and Blundell et al. (2008) for a recent application). For revolvers, on the other hand, credit card debt carries over from month to month by definition. Shocks to debt will therefore persist for revolvers.

Using these distinct models of debt to distinguish between convenience users and revolvers in a Finite Mixture of Regressions framework, we show that the fraction of revolvers declines slowly over the life cycle. As one would expect, shocks to credit utilization for revolvers decline much more slowly than the average: on average 83 percent of an increase in utilization remains after a quarter. The utilization of revolvers is remarkably constant over the life cycle: revolvers in their 20s use just under 60 percent of available credit; those in their 60s use just under 50 percent on average. Moreover, the pass-through of credit into debt for revolvers is nearly complete: following a 10 percent increase in credit limits, the debt of revolvers eventually increases by 9.99 percent. While for revolvers the pass-through is nearly complete at all ages and current levels of credit utilization, it occurs particularly rapidly for the young and those using much of their current credit.

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 (forthcoming) 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, 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 limits 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 (2015a) 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. 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.

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 of either.

We organize our analysis to move from the aggregate to the individual. First, we describe the data. Then in section 3, we provide an overview of debt in the United States since 2000. In section 4, we describe how credit and debt evolve over the life cycle. We start non-parametrically, imposing no assumptions on the relative importance of age, year, cohort, or credit limit. Section 5 examines the evolution of individual credit utilization. In section 6, we introduce a model of the evolution of debt for convenience users and revolvers, and in section 7 we estimate the relationship between debt and credit, taking into account the difference between convenience users and revolvers.

2 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 (2015a) 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 (2015a) 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 they are not using one of their cards at all. 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.

While the Equifax/CCP gives a complete picture of consumers' debts, it does not reveal how the consumer is carrying or paying off that debt. For example, we cannot distinguish revolvers from convenience users directly. We use the Survey of Consumer Payment Choice (Schuh and Stavins, 2014) and the Survey of Consumer Finances to gain more insight into how people are using the debt we see.

The adoption and use of credit cards for payments has been relatively stable since the begin-

ning of the Equifax/CCP data in 1999. The share of consumers with a credit card has been approximately constant since 1989 at around 70 percent (Schuh and Stavins, 2015, p. 20). Information on how frequently credit cards are used is more recent. Figure 1 shows the share of consumers with a credit card by age and year from the Survey of Consumer Payment Choice. While the likelihood of having a credit card increases with age, within a given age the penetration has been relatively stable since 2008. As panel B illustrates, the share of payments made with a credit card has also been fairly stable for age groups over time. The oldest consumers, who are also the most likely to have a card, use it for approximately 30 percent of their transactions, while the youngest, who are the least likely to have one, use it for only 10 percent.

3 Credit and debt over the business cycle

Since 2000 overall credit and debt have varied tremendously. Figure 2 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. Later, in a slightly more sophisticated analysis that accounts for age and region, we show that the recovery in 2012 was driven largely by large dollar amounts at the top end of the credit distribution (see figure 8). 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.

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 2 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. The fact that utilization did not change much suggests the importance of credit constraints and a strong behavioral response to changes in credit limits.

While other credit markets do not have limits that are as clear, the aggregate debt in other markets is directly observable. Using the same Equifax/CCP data set, the Federal Reserve Bank of New York (2015) examines aggregate trends in households' debt. Housing debt predominates, especially after nearly doubling between 2000 and 2008, then falling somewhat more modestly since. Student loan debt increased rapidly after 2008. The other major debt categories are home equity loans, which rose and fell with mortgage debt, and auto loans, which have been fairly stable over the period.

4 Credit and debt over the life cycle

Both credit and debt change substantially over the life cycle. This section provides a non-parametric description of these changes, focusing on the both the level and the distribution of credit and debt.

4.1 The credit score over the life cycle

A useful summary of the availability of credit is the credit score. One of the services that a credit reporting agency provides, the score gives less-sophisticated parties, such as landlords, the ability to quickly understand the likely creditworthiness of a potential borrower. Equifax scores are based on a proprietary scoring model that produces a score ranging from 250 to 850, and are distinct from scores created by the Fair Isaac Corporation (FICO scores). In the Equifax/CCP data set we observe only the proprietary Equifax score.

Figures 3 and 4 show the mean and distribution of credit score changes by the consumer's age. Figure 3 is an entirely non-parametric approach to understanding the effects of age, birth cohort, and year on the mean credit score. Each line represents a single birth cohort over the entire time that the cohort members are in the CCP between the ages of 20 and 80. Very young or old cohorts might be observed for only several years. For example, the members of the cohort born in 1990 turn 20 in 2010 and so the cohort appears as a line on the chart from the ages of 20 through 25. We observe most cohorts for the entire 15 years, although which 15 years depends on the cohort's age. The information in the figure gives the cell mean for every cohort-age-year, and so could be presented with year or cohort on the horizontal axis instead of age. For example, if there are large changes that happen at the same place in every cohort line, then this suggests that there are large shared year effects. Similarly, if the cohort lines shift, then there are likely to be large cohort effects. As always with age-cohort-year analysis, it is impossible to distinguish fully between age, cohort, and year effects, since there can be a trend in any one of them. By showing age on the horizontal axis, the figures place visual emphasis on life-cycle changes, but since all of the information is present, they do not make the assumption that age matters while cohort or years do not. Especially for cohorts in their early 20s, figure 3 mixes changes with age for those consumers who already have accounts with the new entrants getting a credit account for the first time.³

³Our sample selection includes only accounts with an age listed and an open credit card account at some point over the entire sample. While this helps to avoid a separate selection problem of mixing those who never use a credit card with users, it creates an age selection problem. Young people who are only in our sample for a relatively short period of time, but may get a credit card later, may be less likely to be included. Similarly, those who are at the end of the age range, but may have had a credit card before the sample started, may be excluded.

The remarkable feature in figure 3 is how uniformly and nearly linearly the average score of a cohort increases with age. While there is some variation by year or cohort, where cohorts overlap they show remarkable uniformity. Each decade, scores increase approximately linearly by 25 points. The 20-year-olds have an average score of around 625, while 40-year-olds score 675, 60-year-olds 725, and 80-year-olds 775. The change is not entirely linear—scores seem to increase slightly more rapidly between the ages of 50 and 60—but the trend increase is extremely consistent across years, cohorts, and ages.

The nearly constant increase of the mean credit score hides substantial heterogeneity in scores at any age. Figure 4 shows the score at the 1st, 10th, 25th, 50th, 75th, 90th, and 99th percentile of the score distribution at each age. Unlike figure 3, figure 4 averages over all cohorts at a given age, and thus averages out any year or cohort effects. The median person at age 20 has a score of 650, while the 1st percentile at the bottom has a score below 400, and the 99th percentile at the top has a score around 750.

The steady increase in the average credit score in figure 3 comes from different parts of the distribution over the life cycle. Early on, the average growth of scores occurs as those individuals with already good credit increase their scores. The median score actually falls slightly between the ages of 20 and 25. This effect may come from changing composition: new individuals without any credit history may first establish credit relationships after age 20. These new accounts will typically have lower scores, since they are new, so they will bring down the median. Even after the decline in scores in the early 20s, the median or lower scores increase relatively slowly until age 40. After age 40, the 75th, 90th, and 99th percentiles grow more slowly, having hit the upper bound, while the median and lower scores start to increase more rapidly, particularly after age 40.

The combination of increases first at the top of the distribution and only later at the bottom of the distribution leads to a continuing increase in the dispersion of scores until age 40, and then a slow decrease. Figure 4, for example, shows the standard deviation of the credit score for each cohort by age. The standard deviation pulls in changes in the extremes, so there is more variation across cohorts and time, but nonetheless the overall life-cycle pattern of the standard deviation is

evident.

While there is a life-cycle component to the heterogeneity of credit scores, one surprising characteristic of the credit score distribution is how large it already is by age 20. One explanation for the large spread of scores is that the credit score partly reflects total life-cycle income, or at least the best market prediction of it, and thus captures the spread of life-cycle income, which is much greater than the spread of incomes at age 20. Put differently, college students have low income, but high potential life-cycle income, and so receive much more credit and higher credit scores than their current incomes would suggest. Another explanation for the large spread at age 20 is that at least some people inherit their credit rating from their parents by having a parent co-sign for a credit card, for example. Since credit is potentially an important source of consumption inequality, passing on credit ratings may be an important source of the intergenerational transmission of inequality.

4.2 Credit card debt and credit over the life cycle

The increase in credit scores with age allows for a potential increase in available credit, which increases opportunities to take on debt. This section examines the interplay between changing credit and debt over the life cycle in the credit card market, where both limits and debt are observable. Figure 5 shows the fraction of each cohort at each age that has a positive limit on a credit card and positive credit card debt. Having a positive limit in any given quarter is not a requirement for having debt, since the debt may have been accumulated previously when the consumer was able to borrow. Having a positive limit is increasing with age at least until age 60, and after that appears to fall on average. Interestingly, there appears to be a downward trend within cohorts after approximately age 30.

A large fraction of account holders have positive debt at any given point in time. This large fraction is slightly misleading, however, since the Equifax/CCP does not allow us to distinguish new charges from debt acquired previously. Some credit card users pay off their entire balance every month. Others may roll debt over from one month to the next and so are using the

revolving credit aspect of credit cards. Both have unpaid debt at any given point in time, so they are indistinguishable in the data, but other work makes it clear that the use patterns of convenience users and revolvers are distinct. Using the 2008–2014 Survey of Consumer Payment Choice, we examine the fraction of credit card users who do not revolve debt from month to month out of all of those who either used a credit card in the past month or rolled over debt from a previous month. Figure 6 shows the fraction in different age groups who reported that they did not revolve credit card debts from month to month, but do use their cards. Between the ages of 25 and 50, only around 35 percent of the credit card-using population—those who would show up in figure 5 with positive debt in the Equifax/CCP data—were not revolving at least some of their debt from month to month. Starting at age 55, the fraction of non-revolvers slowly increases to 55 percent by age 70. Put differently, even at age 70, 45 percent of consumers who used a credit card in the last month did not pay off their debt in full that month, and so were actively borrowing.

Credit card limits are not quite the same as a credit line, since the card-issuing bank can reduce the credit limit without notice. With enough lead time, a consumer may be able to increase her credit limit by applying for additional credit or requesting a limit increase, so the limits existing at any given time may not represent the full amount of credit card credit a consumer could obtain. Nonetheless, the limit is the maximum amount a consumer can borrow at a given instant, and so represents at least a short-term constraint. The cost of reaching or exceeding the limit is potentially large, since alternative sources of credit, such as payday loans, are likely to be even more expensive and may not be available.

Figure 7 shows the credit card limit and debt in the top panel and credit utilization in the bottom panel. Since both credit and debt are shown on a log scale, they are conditional on being positive. Figure 5 shows the fraction of consumers who have a positive limit or positive debt. 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. There appear to be common trends within each cohort, suggesting that there are common factors that affect all cohorts equally.

Early in the life cycle, credit card debt increases with credit limits and continues to increase

until the age of 50. After age 60, credit card debt starts falling. On average, however, 70- and 80-year-olds have more credit card debt than 20-year-olds. Some of that debt is convenience use, but it illustrates the extent to which credit is an integral part of the financial life of people across all ages, as well as the importance of credit limits.

Credit limits and debt combine to give the fraction of credit used, shown in the bottom panel of figure 7. 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. Credit utilization falls continuously 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. These increases could come from increases in income, from timely payment of debt, and from other factors that may be reflected in improving credit scores. Finally, starting around age 60, average debt, conditional on having any, starts to decline, so credit utilization declines. While the credit and debt cohorts show only those with positive credit or debt, figure 5 shows the fraction in each cohort who have positive credit and debt. Averaging in the zeros would lower the average credit limit and debt, but actually makes the life-cycle variation larger.

Estimating a simple model that separates the variation between age and year allows us to make the importance of life-cycle variation even clearer. Figure 8 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 . While the figures discussed earlier made no assumption about the relative importance of birth cohort, year, and age, the estimates assume that birth cohort does not matter. The excluded group is age 20 and year 2000, so each panel in figure 8 starts at zero at age 20 and year 2000.

Life-cycle variation dominates everything else in figure 8. The scale of the figures varies, so it is easy to miss that the variation of credit card limits and debt over the life cycle is around 10 times larger than over the time. If all of the figures were on the same scale, there would appear to be no variation over time, and the fall in credit utilization with age would not be as obvious, so we use different scales. Allowing for the log scale, the increase from age 20 to age 30 in credit limits is about 1.5 in log units or approximately 450 percent ($e^{1.5}$), and the increase is over 300 percent for debt. The scale of the right-hand columns showing limits, debts, and utilization over time are relatively similar. Credit utilization varies only slightly over time, while aggregate credit card limits and debts are quite volatile. A similar conclusion holds for credit utilization with age. Although credit utilization falls steadily with age, that stability masks much larger changes in credit and debt.

While figure 8 shows the average credit and debt for each cohort, the distribution of credit and debt is changing across ages as well. Figure 9 shows the percentiles and standard deviation of the credit card limit (panel A), credit card debt (panel B), and credit card utilization (panel C), including all cohorts at a given age from 1999 to 2014. The median line in panel A for credit card limits, for example, shows that 50 percent of the population with a positive credit limit at that age has a credit limit below the line. The credit and debt distributions exclude those with zero debt or zero credit limits, but the credit utilization distribution includes those with zero debt and a positive limit.

Over the life cycle, inequality in credit limits starts out very high, increases somewhat to approximately age 40, and then decreases slightly, as shown in figure 9 panel A. The credit limit and debt percentiles are largely parallel with age from the 50th to the 99th percentile. The distance in logs between the 99th percentile and the 50th is almost the same at age 20 as it is at age 50. The

lower percentiles grow much more slowly, so inequality increases up to age 50. After age 50, there is some compression of the distribution of credit limits, as the median draws closer to the 99th percentile, and the 10th and 25th percentiles continue to grow.

Inequality in credit card debt is large initially, and then increases slightly over the life cycle, as shown in panel B of figure 9. Conditional on having any, debt increases rapidly in the 20s for all percentiles, reaching a peak around age 50 before declining. The same evolution occurs for every percentile of the distribution, so the spread of debt is relatively constant.

The distribution of credit utilization has a somewhat different evolution, as shown in the bottom panel of figure 9. Up through age 60, at least 10 percent of the population with a current positive limit has debt that exceeds that limit. The median utilization looks much like the mean utilization in figure 7. Median credit utilization is close to 50 percent at age 20 and falls very slowly. At least 10 percent of the population at any age is not using any of their available credit. The standard deviation of credit utilization is largely stable over the life cycle, declining only slightly from age 20 to age 60, but somewhat faster after.

4.3 Other debts over the life cycle

Since credit scores and credit cards interact with other forms of credit and debt, it is useful to understand the life-cycle variation in other debt. Figure 10 shows the cohort-age profiles for mortgage debt, student loan debt, and auto loans. The first column shows the fraction of each birth cohort who have that particular debt, the second the mean log debt, conditional on having any. Unsurprisingly, there are strong life-cycle components to all three. Few people at age 20 have a mortgage, and the mortgages they do have are typically much smaller than the mortgages of older people. By age 40, approximately 40 percent of individuals have mortgage debt. The fraction of households with mortgage debt may be higher, if not everyone over age 20 in a household is on the mortgage, or lower, if it is more common to hold mortgages jointly. The average size of mortgages appears to decline with age, although there is a strong trend toward larger mortgages over time as well, even with the decline after 2007, so the cohort age graphs make it difficult to disentangle

life-cycle, age, and cohort effects.

Student loans are mostly taken out by the youngest ages, and therefore display a distinct life-cycle pattern. Between the ages of 40 and 60 over this period around 20 percent of the population had positive student loan debt, of a size that seems to have diminished only very slowly with age. One reason may be that parents are taking on debt for their children, which may explain the surprising upward trend in the size of the debt of some cohorts in middle age. Student loans also display clear cohort differences. Of the youngest cohorts in our data set—those aged approximately 20 to 25 between 2010 and 2015—nearly half had student loan debt.

Following a rapid increase in the 20s, about 35 percent of individuals had an auto loan. The size of auto loans increases slightly until age 40, then decreases slightly, but auto debt is relatively constant in size over the life cycle compared with other markets.

The only other major debt held by U.S. households is revolving loans, usually based on home equity. Home equity loans are large in total dollar value (see Federal Reserve Bank of New York (2015)), but are not as widely held as credit card debt. The amount of revolving debt increases rapidly until age 50, and then declines, following the same pattern as credit card debt in figure 7, so we do not show separately the cohort figures for revolving debt.

5 Individual credit, debt, and credit 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. We begin by describing the distribution of utilization and how the distribution varies over the life cycle. Next, we examine non-parametrically how utilization changes at the individual level, and finally we estimate a series of regressions on how individual credit utilization evolves. The basic conclusion is that utilization 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.

Since credit utilization is a ratio, there are two things that can change it: credit (the denominator) and debt (the numerator). In the final section we examine how credit and past debt influence future debt. Credit utilization changes both as individuals change their consumption and savings behavior, thereby increasing or decreasing their debts, and as card-issuing banks increase and decrease credit limits depending on age and credit management. Debts seem to be highly related to credit, with around 90 percent of credit passing through into higher debt eventually, consistent with findings for utilization.

This section does not distinguish between revolvers and convenience users, so, while it estimates population average effects, which may be appropriate for changes that affect everyone, the estimates may miss important heterogeneity, which provides insight into different types of behavior. Section 6 introduces an approach for dividing the population into revolvers and convenience users, and 7 provides estimates for these groups separately.

5.1 The distribution of credit utilization

Figure 11 shows histograms of credit utilization for broad age groups. Credit use predominates until very late in life. The median person is always borrowing, although at the end of life she is not borrowing much. Among users in their 20s, a large portion are close to or exactly at their credit limit. More than 10 percent are actually over their limit, since a previously higher limit has since been reduced. As people age, the distribution slowly shifts leftward, putting more and more of the population at low or zero credit use. Over age 60, around three-quarters of the population is using less than 25 percent of their available credit. Conditional on using more than a small amount of available credit, the spread is fairly even, with about the same fraction using 30 percent of their credit as use 60 percent or 90 percent. It is important to note again that we cannot distinguish directly between revolving debt and convenience debt, and so an important part of the shift in credit utilization may be between these two separate types of use.

A useful way to read the histograms in figure 11 is that there are, broadly speaking, two populations mixed together: a population that uses almost none of its credit, and a population that can use anywhere from 20 to 100 percent of its available credit about evenly. As people age, more and more fall into the first group, using little of their credit. But, conditional on using more than 20 percent of credit, the distributions from ages 20 to 30 to ages 60 to 80 are very similar: mostly flat with a peak around 80 percent. While we cannot distinguish cleanly between them, it seems likely that most of the group using less than 20 percent of its credit are convenience users, while those using more are mostly revolving debt. Of course, the histograms do not show dynamics. It might be the case that an individual uses little of her credit in one quarter, then 50 percent, then little credit, and that the frequency of increases in credit use decreases with age. Instead, we show in the next two subsections that individuals have extremely persistent credit utilization.

5.2 Changes in credit utilization: non-parametric evidence

Figure 12 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 whose overall distribution is shown in figure 11. Figure 12 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 0.39.⁴ The same conclusion is evident from the conditional changes comparing utilization this quarter to a year from now and to 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.

It is important to note that 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 12 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 12 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

⁴Since the conditional expectation of utilization next quarter given this quarter is $u_{t+1} = 0.041 + 0.896u_t$ the steady-state utilization is $0.39=0.041/(1-0.896)$.

percent of the increase in credit in two years.

5.3 Changes in utilization: parametric estimates

In this section we examine how credit utilization changes from quarter to quarter parametrically. In Figure 12, the conditional expectation functions are surprisingly linear, although at longer horizons there appears to be an inflection at the steady state. The effect of being 20 percentage points above the mean is nearly the same as being 20 percentage points below it. The linearity suggests that we can summarize the changing utilization relationship extremely parsimoniously. 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.⁵ Column 1 does not include fixed effects and so assumes a common intercept, as in the top row of figure 12; column 2 includes year and age effects, while the other columns include fixed effects and so are the equivalent of the bottom row of the figure.

Without fixed effects, credit utilization is very persistent and returns to a non-zero steady state, $\alpha/(1 - \beta) = 0.38$. Note that this utilization is close to the average in figure 2, as it should be since they are estimated from the same data and the conditional expectation function in figure 2 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 steady-state credit utilization is not zero.

Given the heterogeneity of utilization shown in the distributions in figure 11, we next allow individuals to return to their own utilization rather than to a single common steady state. 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

⁵The combined age, year, and individual fixed effects are not identified. We drop one of each and use the normalization on the age effects discussed in section 5.4.

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. The distribution of these fixed effects matches the distribution of credit utilization in figure 11 closely, since deviations from an individual's steady state disappear rapidly. A deviation from the mean 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. We explore these implications in the next two sections.

5.4 Individual changes in credit and debt

This section examines how debt changes for an individual and how these changes are related to changes in credit and debt in the past. We begin with a series of regressions that do not distinguish between convenience users and revolvers to understand the average effects from the data. Later, we will use the insights from section 6 to distinguish between convenience users and revolvers. However, all of the estimates are very precisely estimated conditional expectations, and so do not require the model to illuminate how debt and credit are related.

The basic specification we employ is:

$$\ln D_{it} = \theta_i + \theta_t + \theta_a + \alpha \ln D_{it-1} + \beta \ln B_{i,t-1} + \epsilon_{it}, \quad (3)$$

where D_{it} is credit card debt, $B_{i,t-1}$ is the credit card borrowing limit observed at the end of the quarter, and we allow for individual-specific levels of log credit card debt, common time shocks (θ_t), common age shocks (θ_a), and for past debt. Since we observe the credit limit and debt only once a quarter, it is necessary to make a decision about timing. The relevant constraint for the debt accumulated at time t is the credit limit that existed between $t - 1$ and t . Since consumption and credit limits can change continuously, the relevant binding constraint on additional debt is the last one to hold. An increase in limits just before the end of a quarter still allows an increase in debt. We therefore measure $D_{i,t}$ and $B_{i,t-1}$ in the same quarter, even though their subscripts are different.

The coefficient β then determines how quickly a shock to credit card debt dissipates back to the individual long-term effect, which is given by $\theta_i + \theta_t + \theta_a + \beta \ln B_{it}$. The effect of a change in credit limits is β within one quarter, and $\beta/(1 - \alpha)$ in the long term. In more advanced specifications we will allow α and β to change with age and with credit utilization so that, for example, older people or those close to using all of their available credit may react differently to a change in the limit. Since we are still mixing convenience users with revolvers, we do not attach a behavioral interpretation to equation (3), which we will start to do in section (6), but are instead interested in

characterizing the overall evolution of debts and limits.

Equation (3) is not identified since it contains individual, time, and age effects. 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 (3). 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.⁶

The functional form in equation (3) with logs necessarily excludes both those who have zero credit limits and those with zero debt this period or the previous period, since the log of zero is undefined. Equation (3) therefore estimates the response of those with debt to changes of limits, conditional on having debt and a positive credit limit. The fraction of accounts with a positive credit limit and positive debt changes with age in figure 5. A standard approach to understanding the implications of excluding zeros is to use a slightly different transformation by giving everyone a small amount. In some specifications, we explore the implications of this conditionality by giving everyone \$10 in both credit and debt so that, rather than being undefined, these individuals are included as having nearly zero debt and credit. We estimate similar specification in levels as well.

Table 2 shows the results of estimating several variations of equation (3). Column 1 shows

⁶We 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.

the base specification, column 2 excludes individual fixed effects but still includes overall cohort effects, column 3 gives everyone \$10 in credit and debt and so includes those with no credit or no debts. At the bottom of the table we calculate the long-term effect of a permanent increase in credit $\beta/(1 - \alpha)$.

The pass-through of credit into debt, adjusting for age, occurs rapidly and is nearly 90 percent in the long term. Using column 1, a 1 percent increase in credit is associated with a 0.4 percent increase in debt within one quarter, and 0.86 percent in the long term. Not including fixed effects in column 2 changes the persistence of debt, but not the ultimate pass-through of credit, shown at the bottom of the table. Including those with zero debt and limits in column 3 seems to reduce the immediate effect slightly—perhaps because credit limit changes for those using none of their credit do not matter as much—but the long-term change in debt is still close to 90 percent of the change in credit.

Finally, we begin to examine whether there is an important feedback mechanism from debt to credit. Column 5 in table 2 shows the impact of past credit and debt on current credit. Allowing for fixed effects, deviations from the long term are fairly persistent, with 85 percent of a deviation still existing within a quarter. Debt has a small positive impact on credit. A permanent 1 percent increase in debt results in a 0.007 percent increase in credit in one quarter, using the estimates in column 5, and a 0.45 percent increase in the long term. It seems likely that this small positive effect is hiding substantial heterogeneity: card issuers may be happy to raise the limit for those they think are a good risk when these consumers want to acquire more debt, but may actually reduce the limit for riskier types acquiring more debt. Either way, the average effect is very small, so focusing primarily on the impact of credit on debt does not miss much.

6 Credit, debt, and consumption

The stability of credit utilization over the short and long term indicates that debt and credit evolve together. This section derives the implications of changing debt and credit for future debt for

convenience users and revolvers. The debts of revolvers and convenience users evolve in distinct ways, and in the next section we use the difference to help separate them in the data.

6.1 Debt for convenience users

Many of those with credit card debts in our data set are actually convenience users who are using credit as a convenient payment mechanism but plan to pay off their entire debt before being charged interest. They show up as having debt but are not well described by equation (6), since the debt is not representative of their asset positions. Using the SCPC, figure 6 shows that such convenience users are around 40 percent of the credit card-using population early in the life cycle, and that the proportion rises with age. Instead, we assume that convenience users charge some stochastic fraction $\omega_{i,t}$ of their consumption to their credit card each month:

$$D_{i,t} = \omega_{i,t} C_{i,t}.$$

Standard intertemporal theory suggests that changes in consumption should follow a martingale with age- and time-specific drift terms that capture changes in tastes or the decision environment.⁷

Then:

$$\ln C_{i,t} = \eta_i + \eta(\text{age}_{it}) + \eta_t + \epsilon_{it}^C,$$

where, by intertemporal optimization, the consumption shock $E_{t-1}[\epsilon_{i,t}^C] = 0$. Further, if the choice or ability to consume with a credit card is stochastic around an individual-specific mean so that

$$\omega_{i,t} = \omega(\text{age}_{i,t}) + \omega_i + \omega_t + \epsilon_{i,t}^\omega,$$

then log credit card debt for a convenience user can be written as

$$d_{i,t} = \tilde{\eta}_i + \tilde{\eta}(\text{age}_{i,t}) + \tilde{\eta}_t + \epsilon_{i,t}, \tag{4}$$

⁷See Hall (1978) for the original formulation, Deaton (1992) for an extended discussion, and Blundell et al. (2008) for a more recent version that incorporates the life cycle and uncertainty.

where $\epsilon_{i,t} = \epsilon_{i,t}^C + \epsilon_{i,t}^\omega$. Neither past debt nor past limits appear in equation (4). Past debt does not matter because debt is not retained from the previous period, and for someone who is smoothing marginal utility from period to period, shocks to consumption in the past should not predict shocks to consumption in the future. By omitting limits, we are implicitly assuming that convenience users are not credit constrained. Evidence from the Survey of Consumer Finances (see figure 13) suggests that convenience users are typically using only a small fraction of their available credit.

Alternatively, dividing by the credit limit $B_{i,t}$,

$$\text{Credit utilization}_{i,t} = \nu_{i,t} = D_{it}/B_{it} = \omega_{i,t}C_{i,t}/B_{i,t}, \quad (5)$$

should not predict future utilization beyond individual and age drift terms.

6.2 Debt for credit revolvers

We next model a revolver who is using credit cards for debt rather than as a payment mechanism. By definition, debts change from period to period according to the standard accounting accumulation equation:

$$D_{t+1} = (1 + r)(D_t - Y_t + C_t), \quad (6)$$

where Y_t is income, C_t is consumption, and t is either age or time, two concepts that are indistinguishable for an individual. Equation (6) does not make any behavioral assumptions about how consumption is decided. A revolver pays off debt if her income is greater than her consumption, and she accumulates debt if her consumption exceeds her income.

Even without putting additional structure on the evolution of income and consumption, the accumulation equation directly implies that past debts must impact future debts for revolvers. Dividing by the credit limit, the accumulation equation implies a relationship for credit utilization like the relationship in equation (2), in which the past credit limit predicts current credit limits. Moreover, since debt accumulates for revolvers, the current limit may directly affect current consumption, suggesting that a relationship described by the basic reduced-form equation (3), in which

both past debt and current limits appear, is appropriate for revolvers. The key difference between revolvers and convenience users is that for revolvers past debts matter.

The rest of this section then goes one step further and gives a behavioral interpretation to the coefficients of equation (3) for revolvers, although the estimates do not require the additional interpretation to be meaningful estimates of the dynamics of debt and credit. A central question in economics is how consumption changes with previously accumulated debt or assets D_t , the current credit limit B_t , and current income Y_t . One approach taken by, for example, Aiyagari (1994), is that consumption is a function of available resources or cash-at-hand, $W_t = Y_t + B_t - D_t$. The consumption function may vary with age as expectations of future income change (Carroll, 2001), and so:

$$C_t = C_t(W_t) = C_t(Y_t + B_t - D_t).$$

Notice that using cash-at-hand treats credit limits as equivalent to liquid savings within the consumption function. This assumption is inherent in using cash-at-hand in the consumption function. It does not imply that credit limits are necessarily binding this period for revolvers, but instead that revolvers take into account the fact that by consuming more and increasing their debt, they are reducing their available cash-at-hand for the future. A useful feature of the cash-at-hand assumption is that it implies a specific and testable relationship between the effects of credit and debt. a relationship that we develop below.

Appendix A.1 shows that taking a log linear approximation of the accumulation equation

$$d_{i,t+1} = \mu_i + \mu_t + g(\text{age}_{i,t}) + \alpha d_{i,t} + \beta \frac{b_{i,t}}{\nu(\text{age}_{i,t})} + \epsilon_{i,t}, \quad (7)$$

where the lower case indicates logs, and where $\alpha = (1 + r)(1 - m)$, $\beta = (1 + r)m$ and $m = C'(W^*)/C(W^*)$ is the elasticity of consumption with respect to changes in cash-at-hand, measured at the steady state of cash-at-hand. $\nu(\text{age}_{i,t})$ is the age-specific average credit utilization. The

appendix also shows that a similar relationship holds for the dollar value of debts:

$$D_{i,t+1} = \mu_i + \mu_t + g(\text{age}_{i,t}) + \alpha D_{i,t} + \beta B_{i,t} + \epsilon_{i,t},$$

where $\alpha = (1 + r)(1 - M)$, $\beta = (1 + r)M$, and M is the marginal propensity to consume out of cash-at-hand at its steady state.

For revolvers, a regression of debt on previous debt and the active credit limit provides a test that credit limits are important and recovers the marginal elasticity of consumption m or the marginal propensity to consume M out of liquid wealth. While the model assumes that credit limits matter, it is possible that the credit limit is not an important constraint on consumer choices. The consumer may not find credit limits salient, particularly if they are not binding today. Alternatively, the consumer may be able to raise the limit easily, so it may not represent a true constraint. Similarly, the model does not allow for alternative assets, since if the consumer is borrowing, she must not be saving. In reality, consumers do keep a small amount of liquid assets (Gross and Souleles, 2002), and some have substantial illiquid assets. If these assets are easily substitutable for consumer credit, then the credit limit will not matter as much. Then, a simple test for whether the credit limit matters for consumption and debt is $\beta > 0$.

Moreover, the accumulation equation approximations predict that the effect of changing credit is closely related to the impact of past debt, so $\alpha + \beta = (1 + r)$. This prediction is useful for understanding both the economic content of the model and its empirical implications. While it is possible that credit limits do not matter, debt certainly does, since it directly affects the intertemporal budget constraint whether or not credit limits ever bind. Put differently, while the consumer has to decide whether to adjust behavior when the credit limit changes and may decide to ignore changes that are not binding today, the creditors will insist and can enforce a change in behavior following an increase in debt. While $\beta > 0$ implies that credit constraints matter, if $\beta = (1 + r - \alpha)$, then changes in debt have the same impact as changes in assets or income.

Why might credit be less salient than assets? Perhaps it is not as well reported or remembered.

Perhaps its volatility means that it is less valuable than a savings or checking account (Fulford, 2015a). Then a change in credit will have a smaller impact than a change in debt or an income shock. We can back out a value for salience by assuming that in the accumulation equation (6) only a fraction σB_t of credit matters for consumption decisions. Then, given estimates of α and β and an appropriate interest rate r , the salience of credit compared to assets is $\sigma = \beta/(1 + r - \alpha)$.

Finally, we can allow some important forms of nonlinearity to matter. Individuals whose cash-at-hand is low may have a much higher marginal propensity to consume. We can capture such nonlinear effects by allowing the marginal propensity to consume to vary with the individual credit utilization $v_{it} = D_{it}/B_{it}$, so that consumers whose budget constraints bind tightly this period may have a different response than those whose constraints are less binding. Similarly, the marginal propensity to consume may change with age. The effects of both credit utilization and age can be flexibly captured by allowing for functions of age and utilization to alter the MPC through interactions:

$$d_{i,t+1} = \mu_i + \mu_t + g(\text{age}_{i,t}) + \alpha(\text{age}_{i,t}, v_{i,t})d_{i,t} + \beta(\text{age}_{i,t}, v_{i,t})\frac{b_{i,t}}{\nu(\text{age}_{i,t})} + \epsilon_{i,t}, \quad (8)$$

where the prediction for any given age and utilization is still $\alpha(\text{age}_{i,t}, v_{i,t}) + \beta(\text{age}_{i,t}, v_{i,t}) = (1 + r)$.

7 Separating convenience users from revolvers

In this section, we take the modeling insights from the previous section and use them to help divide the estimates from section 5 into the effects for revolvers and convenience users. During the course of the life cycle, the average consumer slowly shifts from revolver to convenience user, consistent with the estimates shown in figure 6 from the SCPC. For revolvers, utilization and debts take longer to return to steady state, and the pass-through of credit into debt is nearly 100 percent at all ages and rates of credit utilization.

We begin by examining utilization for revolvers and convenience users. The basic idea is to use the data to separate the population statistically into those who at a given period are more likely to

be convenience users and those who are more likely to be revolvers. We employ a Finite Mixture of Regressions model (Faria and Soromenho, 2010), sometimes also called a latent class model, depending on the discipline and application. McLachlan and Peel (2000) provide a more complete treatment.

Since we cannot observe directly who is a convenience user, the observed data represent a combination of revolvers and convenience users, where each observation is one or the other, but we cannot observe this latent class. We take the model as:

$$\begin{aligned} \text{Convenience users:} \quad & \nu_{i,t} = \theta_a^C + \theta_t^C + \theta_i^C + \epsilon_{i,t}^C && \text{with density } f_C(\nu_{i,t}|X_{i,t}; \theta^C, \sigma^C) \\ \text{Revolvers:} \quad & \nu_{i,t} = \theta_a^R + \theta_t^R + \theta_i^R + \beta\nu_{i,t-1} + \epsilon_{i,t}^R && \text{with density } f_R(\nu_{i,t}|\nu_{i,t-1}, X_{i,t}; \theta^R, \sigma^R), \end{aligned}$$

which implies the joint density of the observed data:

$$H(\nu_{i,t}|\nu_{i,t-1}, X_{i,t}; \Theta) = p^C f_C(\nu_{i,t}|X_{i,t}; \theta^C, \sigma^C) + (1 - p^C) f_R(\nu_{i,t}|\nu_{i,t-1}, X_{i,t}; \theta^R, \sigma^R),$$

where p^C is the unconditional probability that any observation is from a convenience user, which is not directly observable. Since the mixing probabilities p^C are unobserved, maximizing the sum over all i and t of $\ln H$ requires also maximizing over the unobserved probabilities p^C . This problem is very difficult to maximize jointly, and instead the standard approach is to use the EM algorithm, which alternates between estimating the parameters for revolvers and convenience users conditional on p^C , and p^C conditional on the parameters. We start round j of the algorithm with an estimate of the probability that each observation is a convenience user $w_{i,t}^{j,C}$, where $p^{j,C}$ is the average over all i and t of $w_{i,t}^{j,C}$. The initial $w_{i,t}^{0,C}$ are assumed to be uniformly distributed between 0 and 1. Then, 1) with the $w_{i,t}^{j,C}$ as weights, we use Weighted Least Squares to estimate each of the models for convenience users and revolvers independently; and 2) then we update the weights and $p^{j+1,C}$ using Bayes' rule based on the new estimates from each model so that for iteration j :

$$w_{i,t}^{j+1,C} = \frac{p^{j,C} f_C(\nu_{i,t}|\theta^{j,C})}{p^{j,C} f_C(\nu_{i,t}|\theta^{j,C}) + (1 - p^{j,C}) f_R(\nu_{i,t}|\theta^{j,R})},$$

and $p^{j+1,C}$ is the average of the new posterior weights for each observation. The two steps alternate until the overall likelihood converges.

We assume that the densities $f_c(\cdot)$ and $f_R(\cdot)$ are normal and impose the additional structure that the conditional likelihood for a convenience user follows the same conditional likelihood of convenience as in the SCF for age and credit utilization. The density for a convenience user is then: $f_C(\nu_{i,t}|X_{i,t}; \theta^C, \sigma^C) = p^{SCF}(\text{age}_{it}, \nu_{it})\phi(\nu_{it}|X_{it}; \theta^C, \sigma^C)$, where $\phi(\cdot)$ denotes the density of the normal distribution with mean $\theta_a^C + \theta_t^C + \theta_i^C$ and variance $(\sigma^C)^2$. We estimate $p^{SCF}(\text{age}_{it}, \nu_{it})$ using a logit with a cubic for age and utilization based on the combined 1995–2010 waves of the SCF. Figure 13 shows the implied fraction of convenience users over age and utilization using the SCF estimates of $p^{SCF}(\text{age}_{it}, \nu_{it})$, but the distributions of age and utilization in the CCP. The likelihood of being a convenience user is falling with age and utilization. Finally, to estimate the full model we first remove the fixed effects θ_i^C and θ_i^R by removing the mean from all variables, which substantially speeds estimation, but implies that for a given individual $\theta_i^C = \theta_i^R$. The problem of the collinearity of age, year, and fixed effects remains. We allow a trend in the fixed effects, and unrestricted time effects, but estimate only a square and cube of age in the convenience or revolver log-likelihoods.

Figure 13 shows the fraction of convenience users at different ages that come from the EM estimates for utilization (using the posterior weights for each observation). The fraction of convenience users is increasing with age and decreasing with utilization, closely matching the SCF, which is not a surprise since we used the estimates from the SCF in forming the individual densities. The SCF estimates may understate the number of convenience users at early ages, however, as suggested by the SCPC in figure 6.

Figure 14 then shows the average utilization of revolvers and convenience users with age. It uses the converged weights to estimate the utilization as a local polynomial function of age. The slow decline of utilization observed in the raw data is coming mainly from the switch to convenience use with age. Conditional on still being a revolver, the average utilization declines very slowly from around 60 percent in the 20s to 50 percent in the 60s. The average utilization of con-

venience users also declines slowly. The main factor explaining the overall decline in utilization is the decline in the fraction of revolvers, as the population slowly shifts from the top line of revolvers to the bottom line of convenience users. The last column of table 1 shows the relationship of past utilization to current utilization for revolvers. Since convenience users return immediately, in expectation, to their long-term mean, shocks to credit utilization are more persistent for revolvers.

Using the same mixture model approach, we estimate the relationship between credit and debt using equation (7) for revolvers and equation (4) for convenience users. Column 4 of table 2 shows the estimated effects of debt and credit changes for revolvers. Debt returns to its steady state more slowly than when estimated over all credit users, as one would expect. The immediate impact of credit is lower, but because of the persistence of debt, the long-term impact of a change in credit for revolvers is nearly 100 percent. Using the posterior weights from the EM algorithm, columns 6 and 7 of table 2 estimate how past debt and limits affect current debts for convenience users and revolvers. While the effect of debt on limits for all users in column 5 is small and positive, for convenience users there is a negative effect, and for revolvers a slightly larger positive effect. One possible reason is that revolvers may be more willing to request higher limits the more debt they are using.

We can also allow the debt response to past debt and current limits to vary flexibly with credit utilization and age.⁸ Since the coefficients from such a regression are difficult to interpret on their own, we show the marginal effect evaluated over ages 20 to 70 and utilization from 0.1 to 1 in figure 15 for revolvers.

Several important changes with age and credit utilization for revolvers are clear from figure 15. First, at all utilization levels, the effect of the credit limit is decreasing with age, while the

⁸We estimate the following functional form for α and an identical one for β :

$$\alpha(\text{age}_{it}, v_{it-1}) = \alpha + \alpha_0 v_{it-1}^{(0)} + \alpha_1 v_{it-1}^{(1)} + \alpha_2 v_{it} + \alpha_3 v_{it-1}^2 + \alpha_4 v_{it-1}^3 + \alpha_5 \text{age}_{it} + \alpha_6 \text{age}_{it}^2 + \alpha_7 \text{age}_{it}^3 + \alpha_8 v_{it-1} * \text{age}_{it} + \alpha_9 v_{it-1}^2 * \text{age}_{it} + \alpha_{10} v_{it-1} * \text{age}_{it}^2 + \alpha_{11} v_{it-1}^2 * \text{age}_{it}^2,$$

where $v_{it}^{(0)}$ is 1 if utilization is 0, and 0 otherwise, $v_{it}^{(1)}$ is 1 if utilization is greater than 1.1 and 0 otherwise, and $v_{it}^2 = v_{it} * v_{it}$. Note that, like the credit limit, the credit utilization rate is measured as the credit limit at the end of the period divided by the debt at the beginning.

effect of the past debt is increasing. Put differently, older people are less sensitive to credit limit changes, but are more sensitive to having debt. Second, at any age or utilization rate the sum of the estimated credit effect and estimated debt effect is nearly constant and close to one, implying a relatively constant and high salience of credit. Third, the long-term effect of credit on debt is close to one and nearly constant across all ages and utilization rates. Panel (C) calculates the long-term effect at each age and utilization (essentially dividing the effect at each age in panel (A) by one minus the marginal effect of debt in panel (B)). Young people reach that long-term state faster, but for all revolvers the long-term effect of credit on debt is nearly 100 percent. Fourth, at any age the effect of utilization is extremely nonlinear. The effects of credit and past debt are nearly identical for those using between 0.1 and 0.7 of their credit, but then change rapidly as individuals get closer to using all of their credit. Credit utilization matters a lot only when the consumer is close to her credit limit. Finally, age and credit utilization do not seem to interact. Excluding the age-credit utilization interactions would not change this picture appreciably, because the different lines connecting credit utilization are nearly parallel for different ages.

8 Conclusion

Available credit appears to be the driving factor of debt in both the short and long term. Separating convenience users from revolvers, we find that for revolvers a 10 percent increase in credit is followed by a 1.3 percent increase in debt within one quarter and a 9.99 percent increase in debt over the long term. For those revolving debt, long-term credit and debt are closely related. The fact that debt follows available credit so closely explains the extreme stability of overall credit utilization. Despite very large changes in credit and debt over the life cycle and in aggregate since 2000, overall credit card utilization varies little. Moreover, we show that individual credit utilization is also remarkably stable, as individuals rapidly return to their steady-state credit use.

The large pass-through of credit into debt and the large life-cycle variation of credit have important implications for savings and consumption over the life cycle. One of the startling facts

of consumer finance is how little households save when they are not forced to do so by Social Security or mortgage payments. Indeed, Gourinchas and Parker (2002) estimate that the average household starts acquiring assets for retirement only when its head is in her late 40s. We show that credit increases rapidly in the 20s and continues to increase into the 30s. These increases help to explain why consumers save so little early in life. The large response to credit changes in our data suggests that early in life many consumers may be constrained by their credit limits. Moreover, we show that, in expectation, the credit limit for the average person in her 20s will be much higher next year. Since for precautionary purposes credit is a form of wealth (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 consumers, 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. Not taking this variation into account leads to a misunderstanding of the financial status of U.S. households.

In future work we expect to introduce models of life-cycle saving that allow for precautionary accumulation in the short term, planning for retirement in the long term, the acquisition of illiquid assets such as housing, and the use of credit for payments. Such models have implications for both debt acquisition and the impact of credit. The estimates in figure 15 suggest that the effects of debt and credit change over the life cycle, and we expect that future work will examine what models are consistent with the observed changes.

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.
- Aiyagari, S. Rao. 1994. “Uninsured Idiosyncratic Risk and Aggregate Saving.” *Quarterly Journal of Economics*, 109(3): 659–84.
- 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.
- Avery, Robert B., Paul S. Calem, and Glenn B. Canner. 2004. “Credit Report Accuracy and Access to Credit.” *Federal Reserve Bulletin*, Summer: 297–322.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. “Consumption Inequality and Partial Insurance.” *American Economic Review*, 98(5): 1887–1921.
- Carroll, Christopher D. 2001. “A Theory of the Consumption Function, with and without Liquidity Constraints.” *Journal of Economic Perspectives*, 15(3): 23–45.
- 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.
- Deaton, Angus. 1992. *Understanding Consumption*. Oxford: Clarendon Press.
- Deaton, Angus. 1997. *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. Baltimore: Johns Hopkins University Press.
- Faria, Susana, and Gilda Soromenho. 2010. “Fitting mixtures of linear regressions.” *Journal of Statistical Computation and Simulation*, 80(2): 201–225.
- Federal Reserve Bank of New York. 2015. “Quarterly Report on Household Debt and Credit: February 2015.” microeconomic studies, Federal Reserve Bank of New York Research and Statistics Group.
- Foster, Kevin, Scott Schuh, and Hanbing Zhang. 2013. “The 2010 Survey of Consumer Payment Choice.” Research Data Reports 13-2, Federal Reserve Bank of Boston.
- 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. 2015a. "How important is variability in consumer credit limits?" *Journal of Monetary Economics*, 72: 42 – 63.
- Fulford, Scott L. 2015b. "The surprisingly low importance of income uncertainty for precaution." *European Economic Review*, 79: 151–171.
- 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.
- Hall, Robert E. 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*, 86(6): 971–87.
- 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.
- 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.
- McLachlan, Geoffrey, and David Peel. 2000. *Finite Mixture Models*. New York: John Wiley and Sons.
- Nothaft, Frank E., and Yan Chang. 2005. *Refinance and the Accumulation of Home Equity Wealth*. Chap. Refinance and the Accumulation of Home Equity Wealth: 71–102: The Brookings Institution.
- 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.
- Schuh, Scott, and Joanna Stavins. 2015. "The 2013 Survey of Consumer Payment Choice." Research Data Reports 15-4, 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.
- Zinman, Jonathan. 2013. "Consumer Credit: Too Much or Too Little (or Just Right)?" Working Paper 19682, National Bureau of Economic Research.
- Zinman, Jonathan. forthcoming. "Household Debt: Facts, Puzzles, Theories, and Policies." *Annual Review of Economics*.

A Appendix

A.1 Derivation of debt revolver accumulation equation

From equation (6) we have:

$$\frac{D_{i,t+1}}{1+r} - D_{i,t} + Y_{i,t} = C_{i,t}(Y_{i,t} + B_{i,t} - D_{i,t}).$$

Let $Y_{it} = P_{i,t}U_{i,t}$, where $P_{i,t} = E_{t-1}[Y_{i,t}]$ is the long-term or permanent component of income, given age t , and $B_{i,t}^*$ the expected credit limit at age t for a given individual. Then define $D_{i,t}^*$ as the debt at which, given a credit limit $B_{i,t}^*$ and income realization $Y_{i,t} = P_{i,t}$, consumption is equal to income minus interest payments and so debt is not increasing or decreasing:

$$\frac{-rD_{i,t}^*}{1+r} + P_{i,t} = C_{i,t}(P_{i,t} + B_{i,t}^* - D_{i,t}^*).$$

A first-order expansion around $W_{i,t} = D_{i,t}^* + B_{i,t}^* + P_{i,t}$ then gives:

$$D_{i,t+1} \approx (1+r)M_{i,t}B_{i,t} + (1+r)(1-M_{i,t})B_{i,t} + M_{i,t}Y_{i,t},$$

where $M_{i,t} = C'_{i,t}(W_{i,t}^*)$ is the marginal propensity to consume out of liquid cash-at-hand at its steady state. If $M_{i,t}Y_{i,t}$ can be well captured by individual fixed effects, age effects, and year effects, then a regression of the form:

$$D_{i,t+1} = \mu_i + \mu_t + g(\text{age}_{i,t}) + \alpha D_{i,t} + \beta B_{i,t} + \epsilon_{i,t},$$

where $\alpha = (1+r)(1-M)$ and $\beta = (1+r)M$ are the average of the $M_{i,t}$, and $\epsilon_{i,t}$ captures the approximation error, the unobserved income shocks not explained by age, individual, and time, and the differences from the average α and β . $M_{i,t}$ may vary with age or overall credit utilization as well.

The assumptions necessary for the linear expansion in levels to provide a good approximation are strong, particularly comparing across many individuals with very different incomes and debt. A more flexible expansion involves taking logs and expanding around $D_{i,t}^*$, $B_{i,t}^*$, and $P_{i,t}$, and canceling constants using the steady-state equation gives a first-order approximation:

$$\frac{D_{i,t}^*}{(1+r)} d_{i,t+1} - D_{i,t}^* d_{i,t} + P_{i,t} \ln U_{i,t} \simeq m_{i,t}(P_{i,t} \ln U_{i,t} + B_{i,t}^* b_{i,t} - D_{i,t}^* d_{i,t}),$$

where $b_{i,t} = \ln B_{i,t}$, $d_{i,t} = \ln D_{i,t}$, and $m_{i,t} = C'_t(W_{i,t}^*)/C_{i,t}(W_{i,t}^*)$ is the elasticity of consumption with respect to cash-at-hand at the steady-state cash-at-hand $W_{i,t}^*$. Rearranging gives:

$$d_{i,t+1} \simeq (1+r)(1-m_{i,t})d_{i,t} + (1+r)m_{i,t} \frac{B_{i,t}^*}{D_{i,t}^*} b_{i,t} + (1+r)(m_{i,t}-1) \frac{P_{i,t}}{D_{i,t}^*} \ln U_{i,t}.$$

Define $m_{i,t} = m + \epsilon_{i,t}^m$. If the target ratio of credit limit to debt can be expressed as a function of

deviations from the average utilization at each age $(B_{i,t}^*/D_{i,t}^*) = 1/\nu(\text{age}_{i,t}) + \epsilon_{i,t}^\nu$, then:

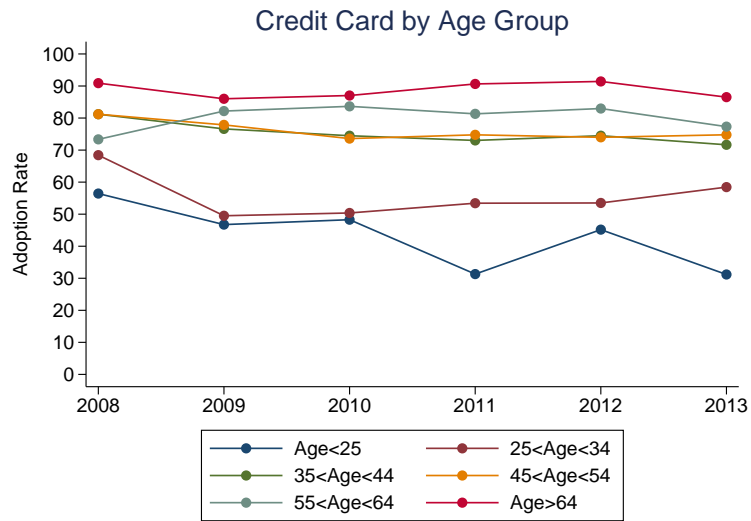
$$d_{i,t+1} = (1+r)(1-m)d_{i,t} + (1+r)m \frac{b_{i,t}}{\nu(\text{age}_{i,t})} + \epsilon_{i,t}^s,$$

where $\epsilon_{i,t}^s$ captures the random coefficients and the unpredictable income component. Following Blundell et al. (2008), suppose that idiosyncratic and age-specific drift factors are well captured by an individual effect and age effects (or functions) so that the structural and approximation error $\epsilon_{i,t}^s = \mu_i + \mu_t + g(\text{age}_{i,t}) + \epsilon_{i,t}$, then

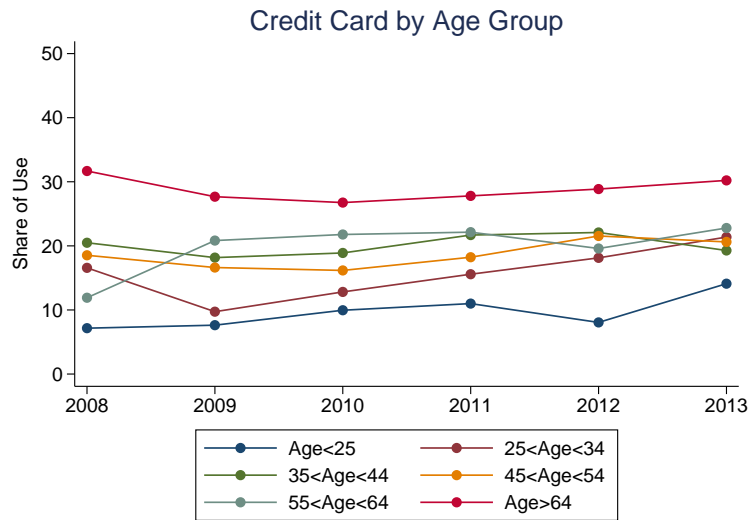
$$d_{i,t+1} = \mu_i + \mu_t + g(\text{age}_{i,t}) + \alpha d_{i,t} + \beta \frac{b_{i,t}}{\nu(\text{age}_{i,t})} + \epsilon_{i,t},$$

where $\alpha = (1+r)(1-m)$, $\beta = (1+r)m$ and $E[\epsilon_{i,t}] = 0$.

Figure 1: Credit card adoption and use for payments by age and year
 (A) Credit card adoption rate

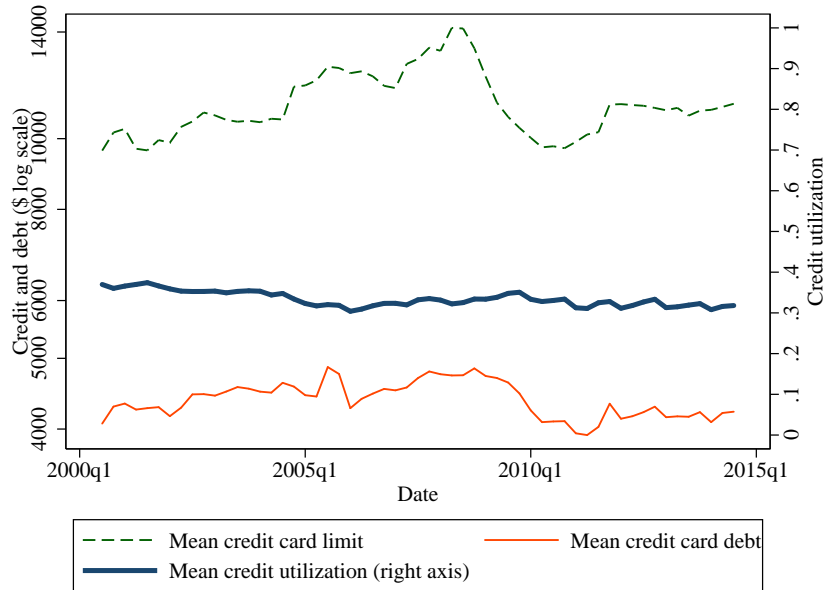


(B) Share of payments using credit cards



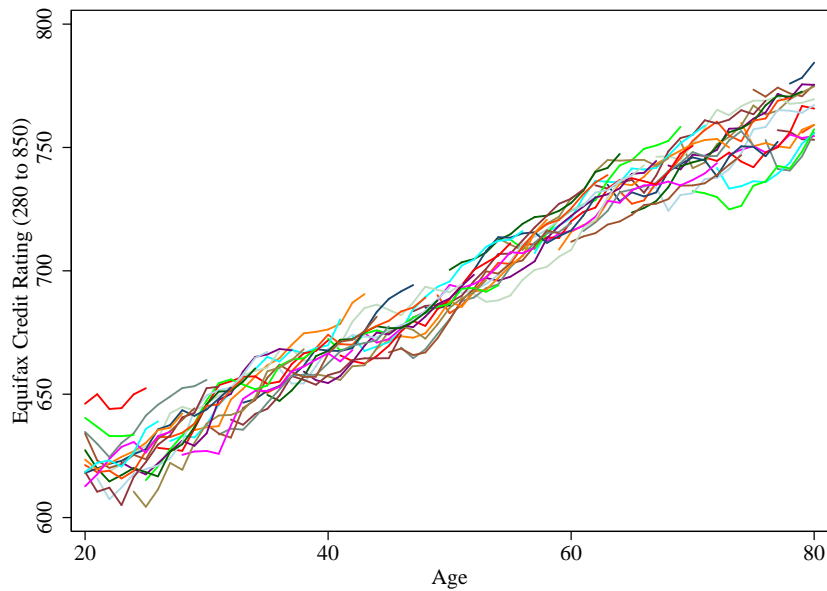
Source: Authors' calculations from the 2013 Survey of Consumer Payment Choice.

Figure 2: Credit card limits, debt, and utilization: 2000–2014



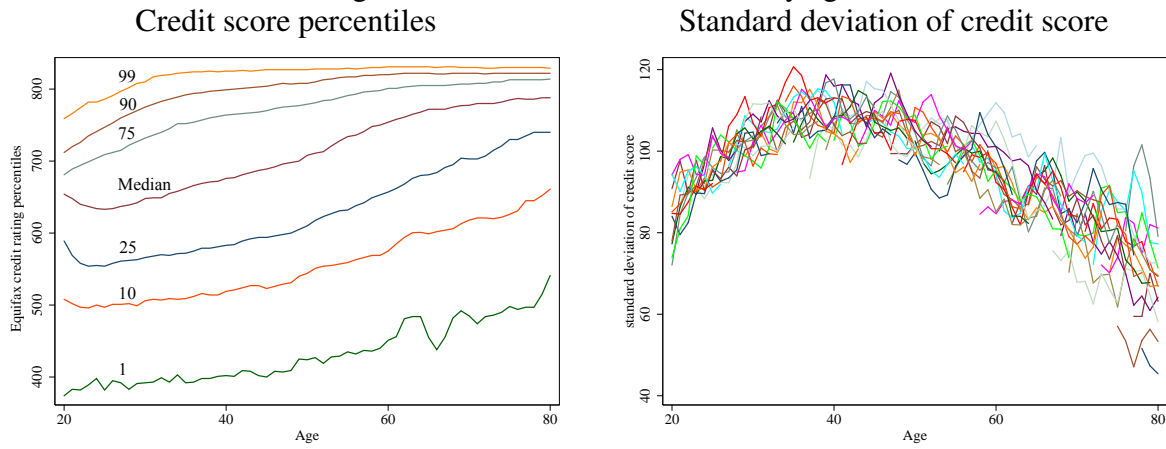
Notes: The left axis shows the average credit card limits (top line) and debt (bottom line). Note the log scale. The right axis shows mean credit utilization (middle line) defined as the credit card debt/credit card limit if the limit is greater than zero. Source: Authors’ calculations from Equifax/NY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

Figure 3: Credit score by cohort and age



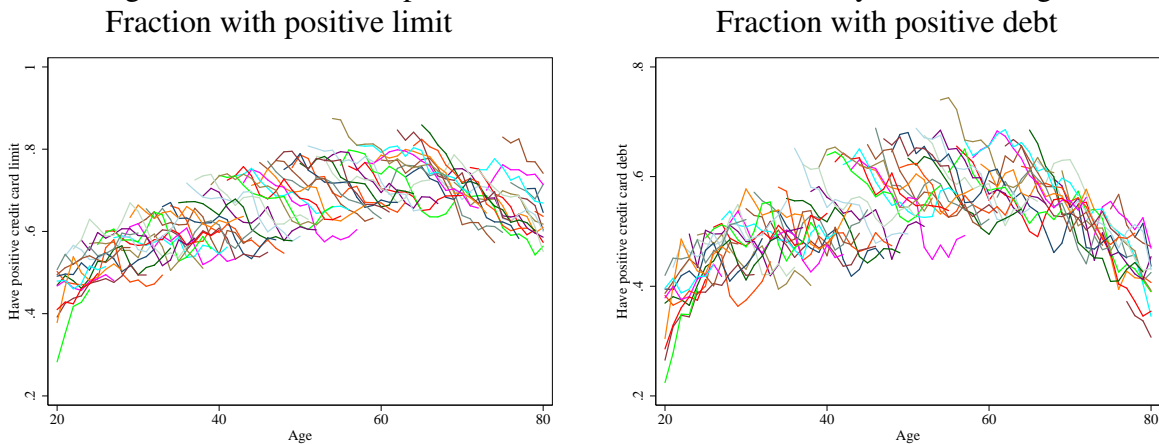
Notes: Each line represents the average credit score of one birth cohort from 1999 to 2014. Source: Authors’ calculations from Equifax/NY Fed CCP.

Figure 4: Credit score distribution by age



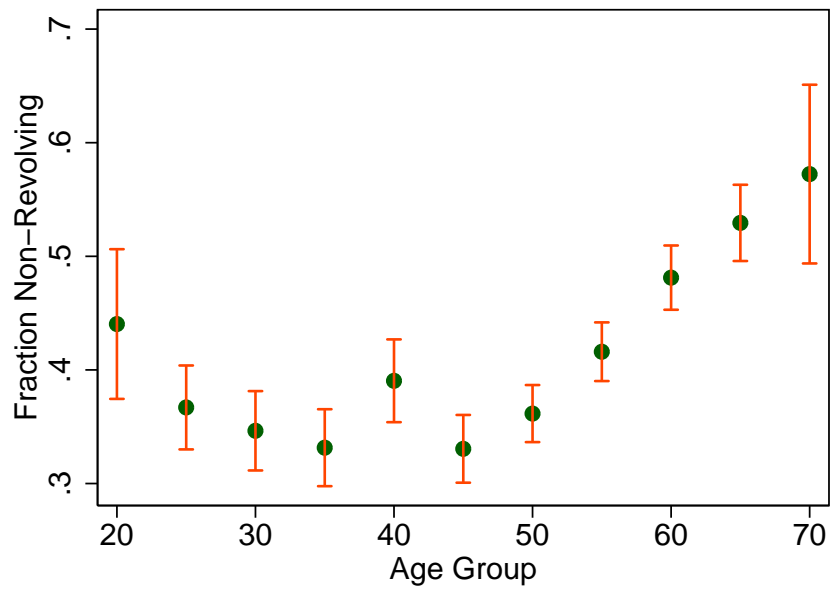
Notes: Left panel: The lines are the credit scores of the 1st, 10th, 25th, 50th, 75th, 90th, and 99th percentiles of credit score at each age over 1999-2014. The lines are in order from 1 at the bottom to 99 at the top. Right panel: Each line represents the standard deviation of the credit score of one birth cohort, 1999–2014. Source: Authors’ calculations from Equifax/NY Fed CCP.

Figure 5: Fraction with positive credit card limit and debt by cohort and age



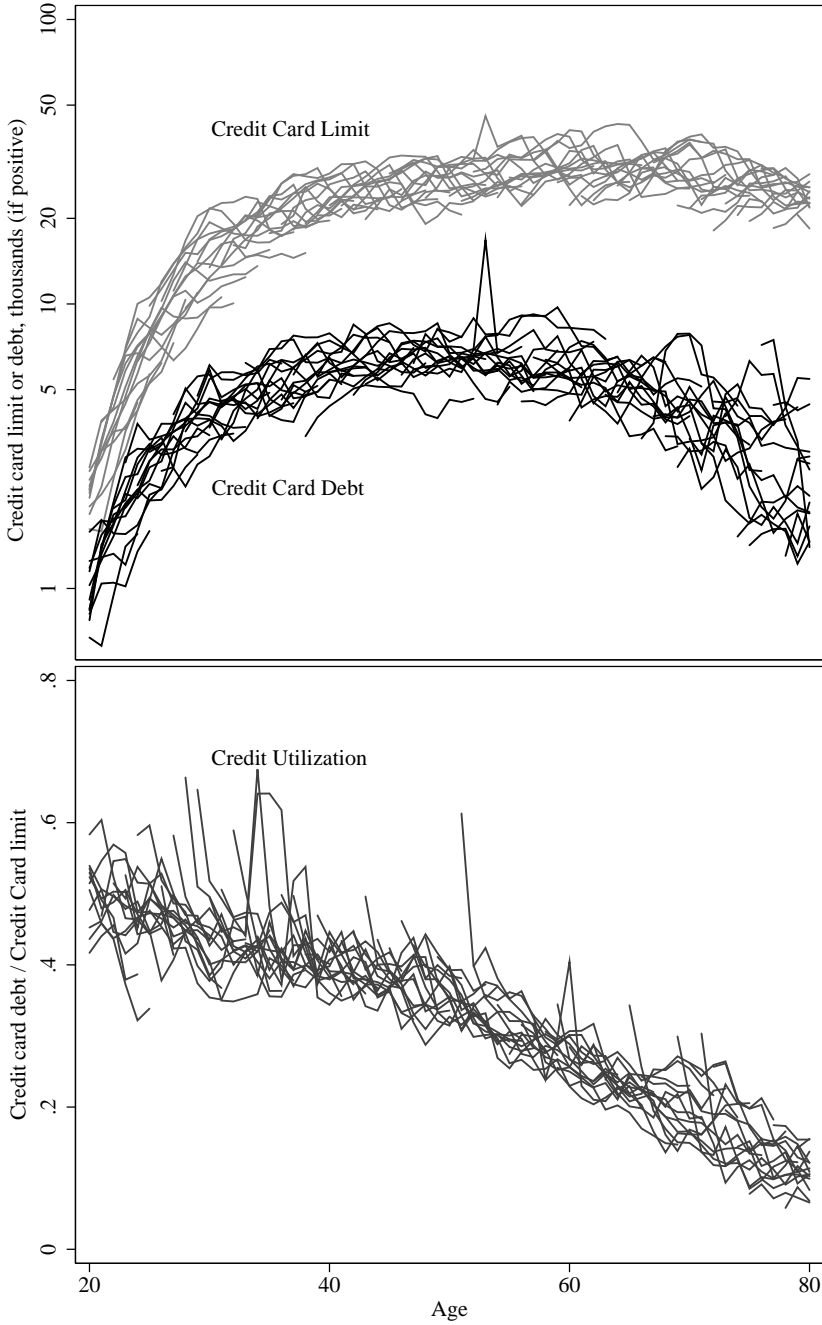
Notes: Each line represents the fraction with positive credit card limits or debt of one birth cohort, 1999–2014. Source: Authors’ calculations from Equifax/NY Fed CCP.

Figure 6: Fraction of convenience users of credit cards, conditional on use



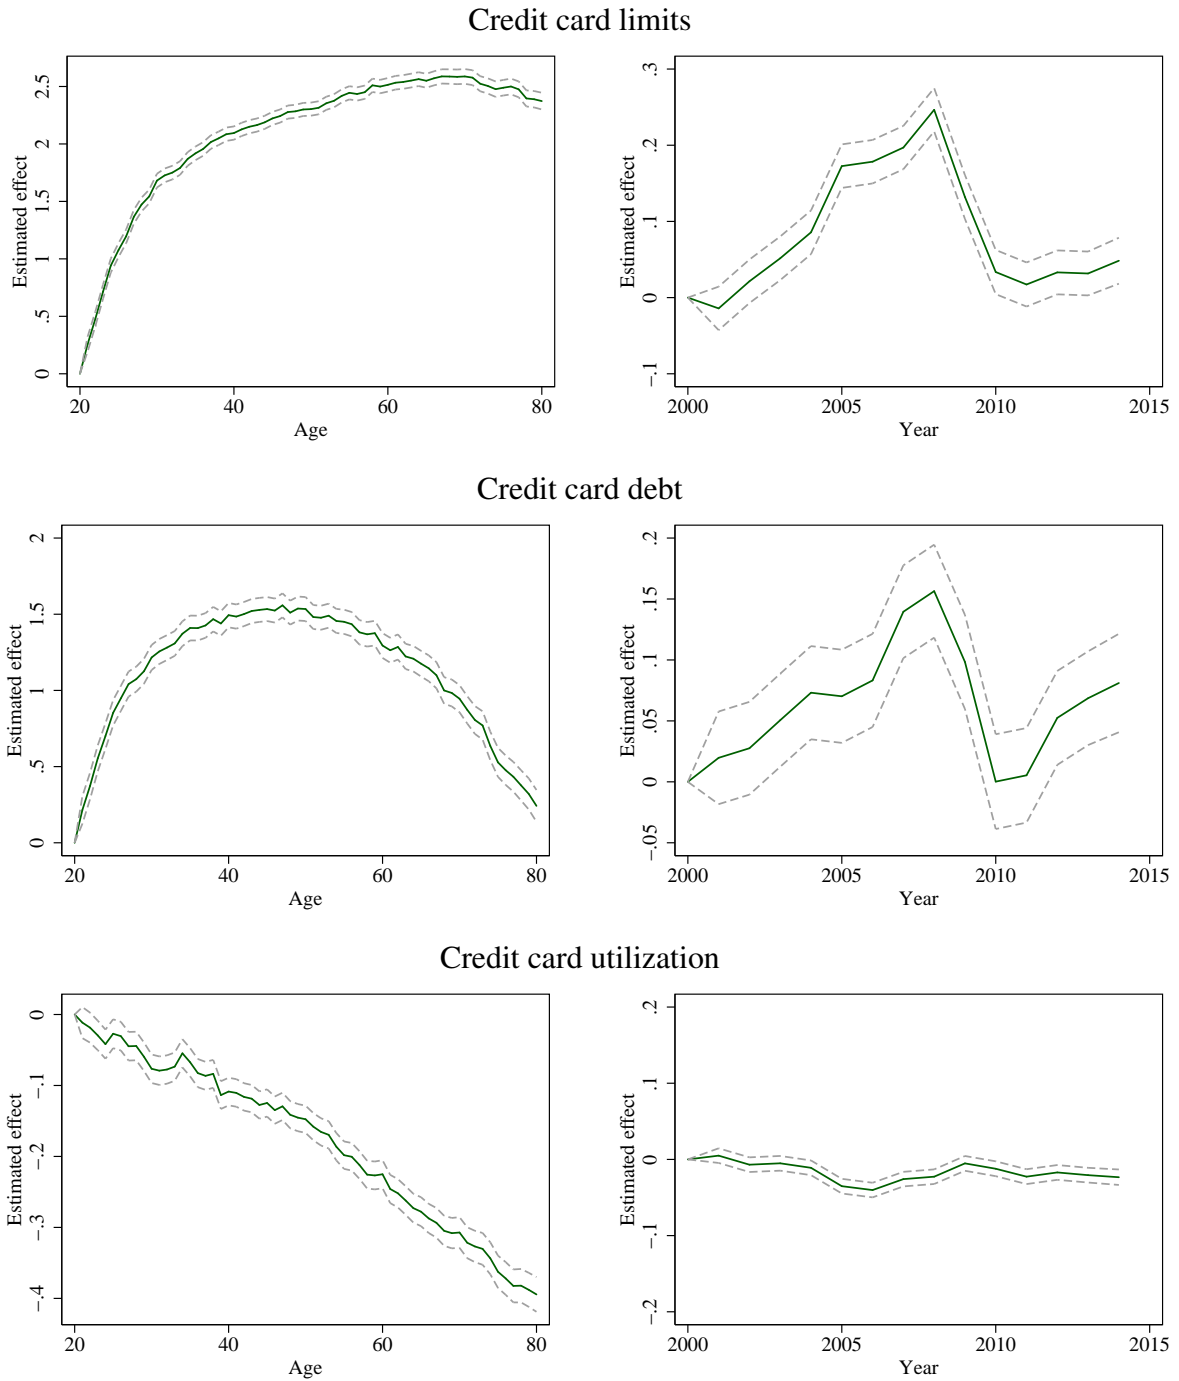
Notes: Each dot shows the fraction in that five year age group (labeled by the youngest member) who did not revolve credit card debt from month to month, 'conditional on using their credit card. 95 percent confidence intervals in bars. Source: Authors' calculations from the 2008–2014 Surveys of Consumer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

Figure 7: Credit card limits, debt, and credit utilization



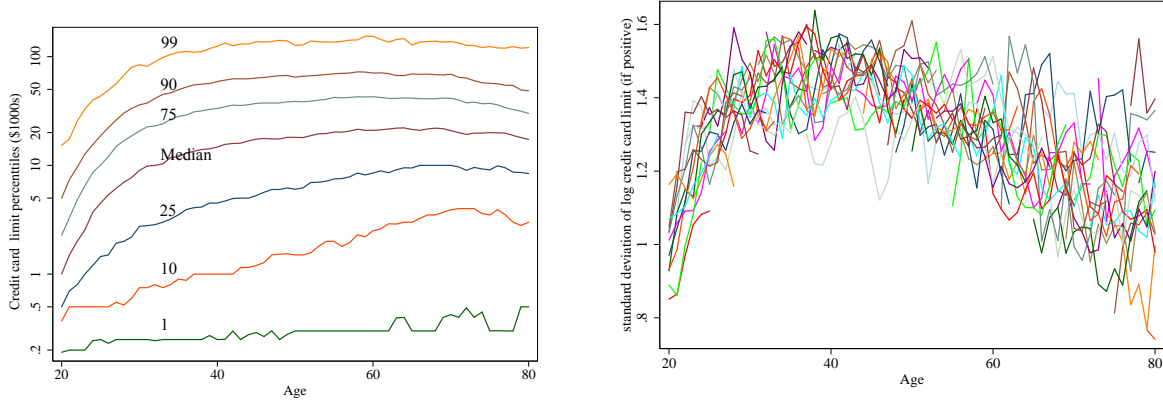
Notes: Each line represents the average credit card limit, debt, and utilization of one birth cohort, 1999–2014. Source: Authors' calculations from Equifax/NY Fed CCP.

Figure 8: Credit card limits, debt, and utilization: age and year effects

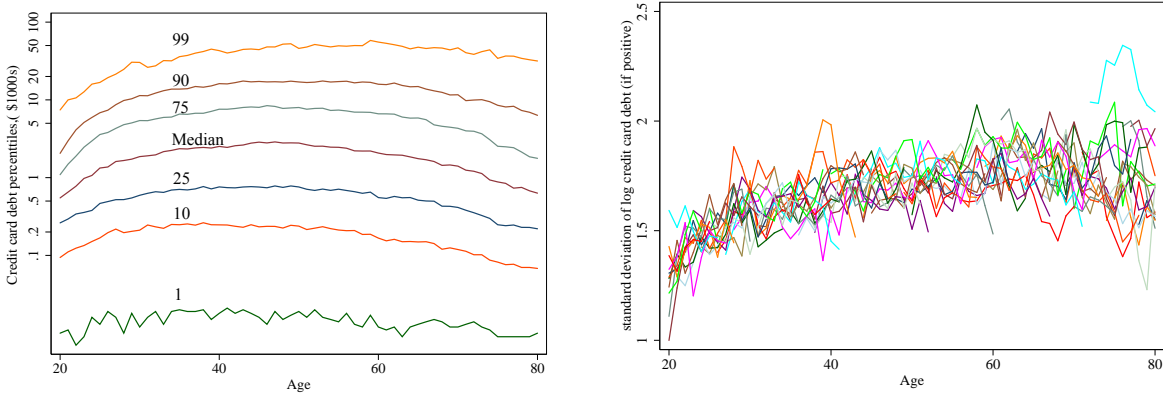


Notes: Each line shows the estimated age or year effects from equation (1). Note the different scales. Source: Authors' calculations from Equifax/NY Fed CCP.

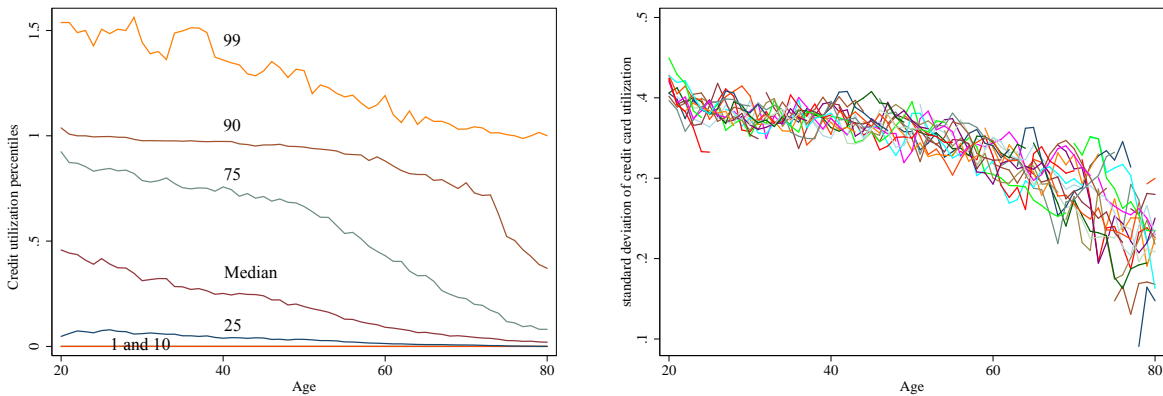
Figure 9: Credit card limit, debt, and credit utilization distributions by age
 (A) Credit Card Limits



(B) Credit Card Debt

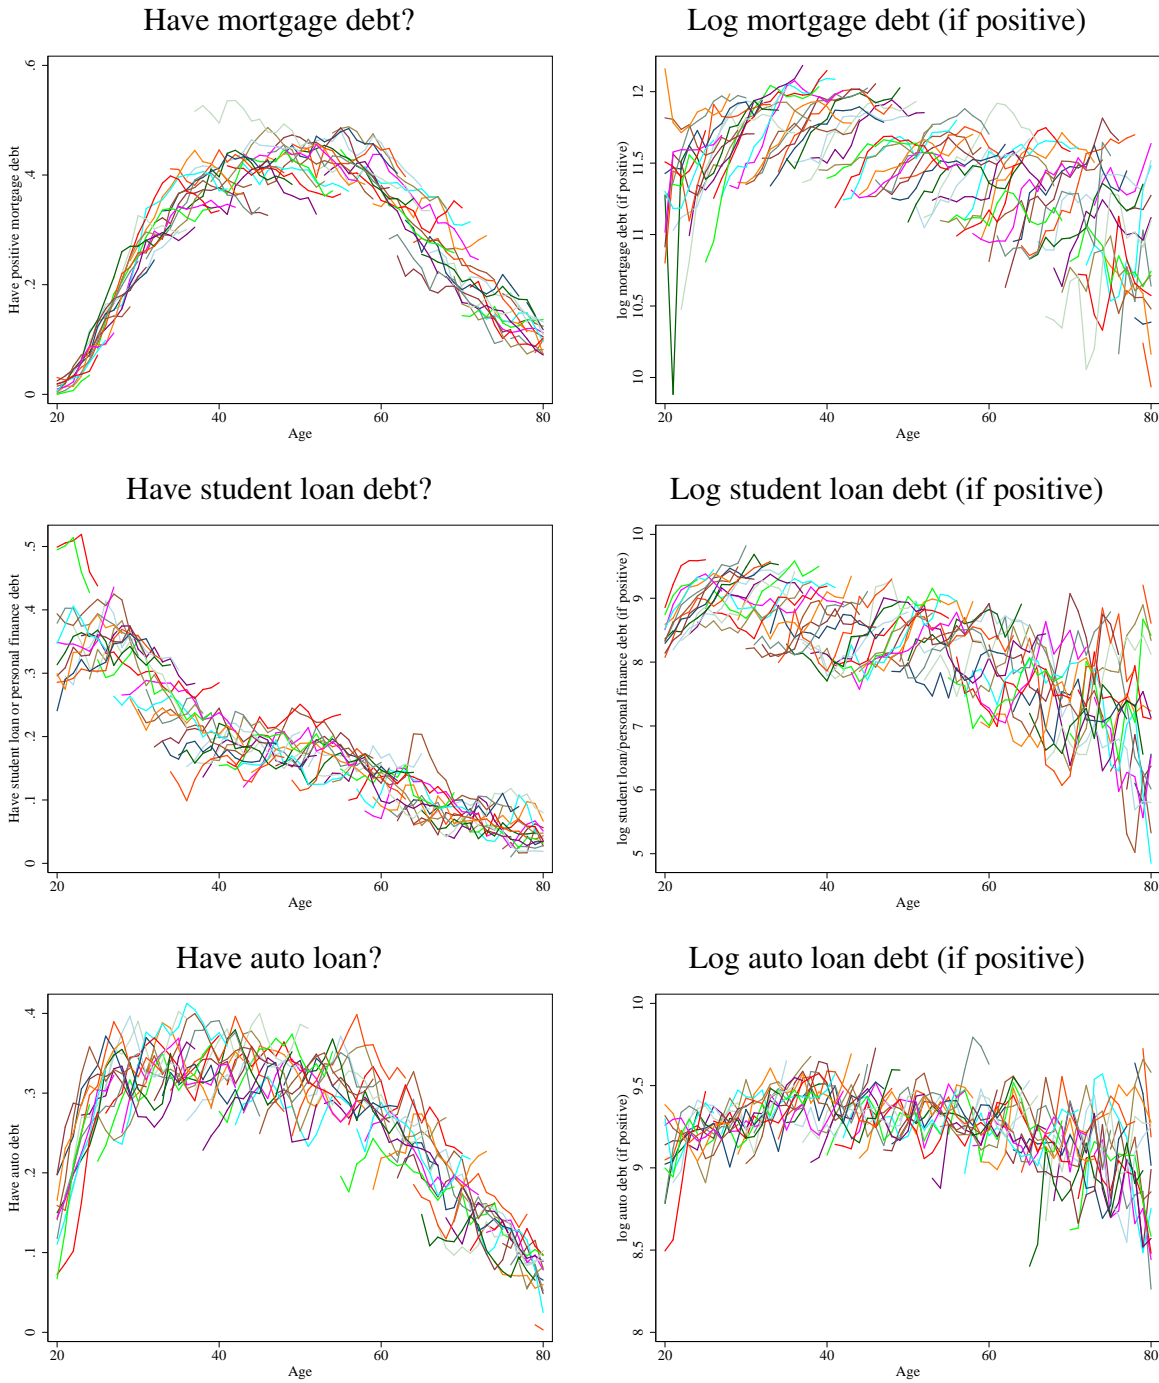


(C) Credit utilization



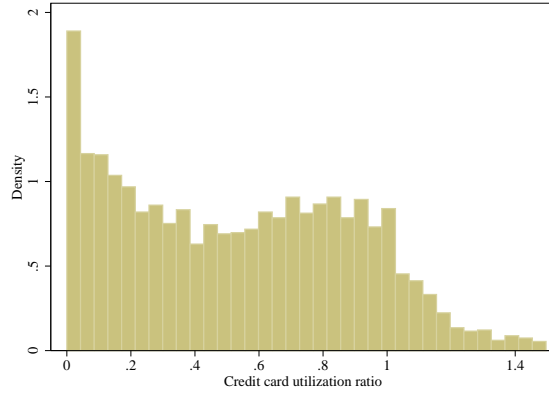
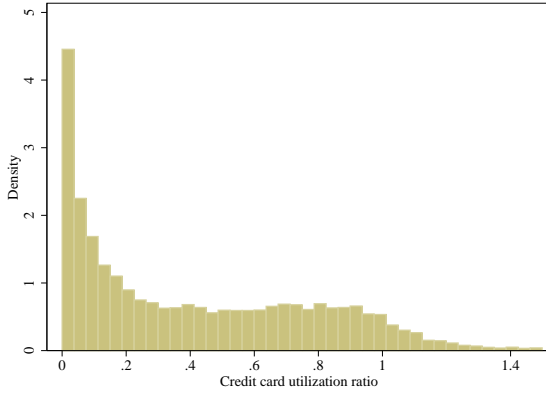
Notes: Each line is the percentile of credit limit at that age, conditional on having a positive credit limit on a log scale. For example, the 90th percentile line shows that 10 percent of the population (with a positive credit limit) have a limit larger than that line. Source: Authors' calculations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

Figure 10: Mortgage, student loan, and auto loan debt over the life cycle

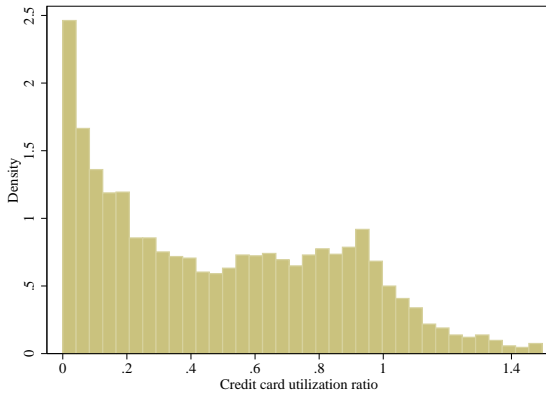


Notes: Each line represents the average for one birth cohort, 1999–2014. Source: Authors' calculations from Equifax/NY Fed CCP.

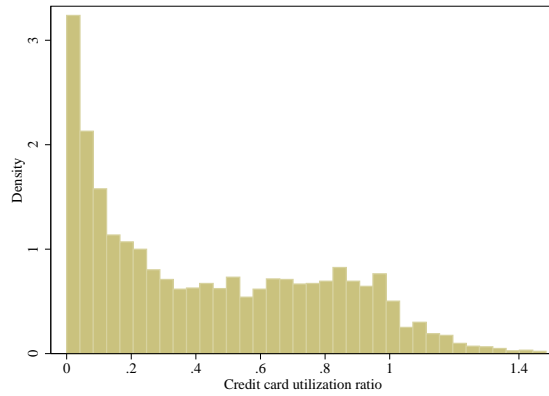
Figure 11: Credit card utilization distribution by age
 All ages Age 20–30



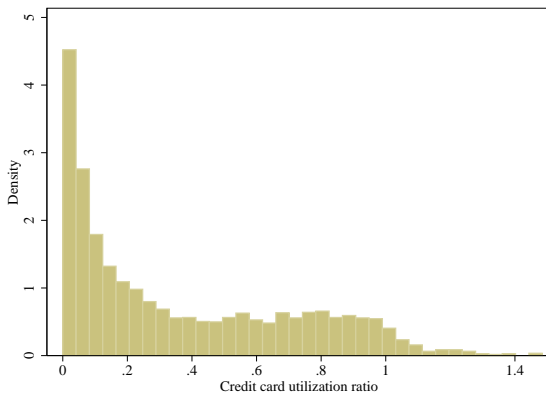
Age 30–40



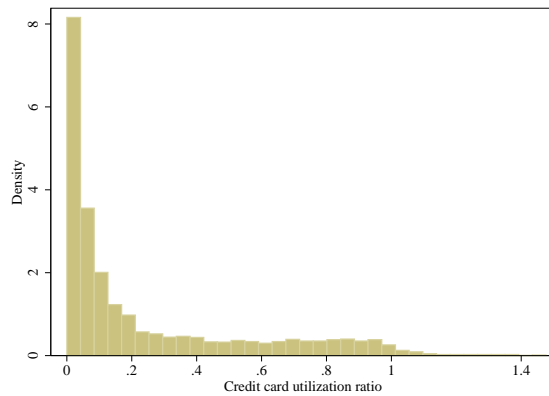
Age 40–50



Age 50–60

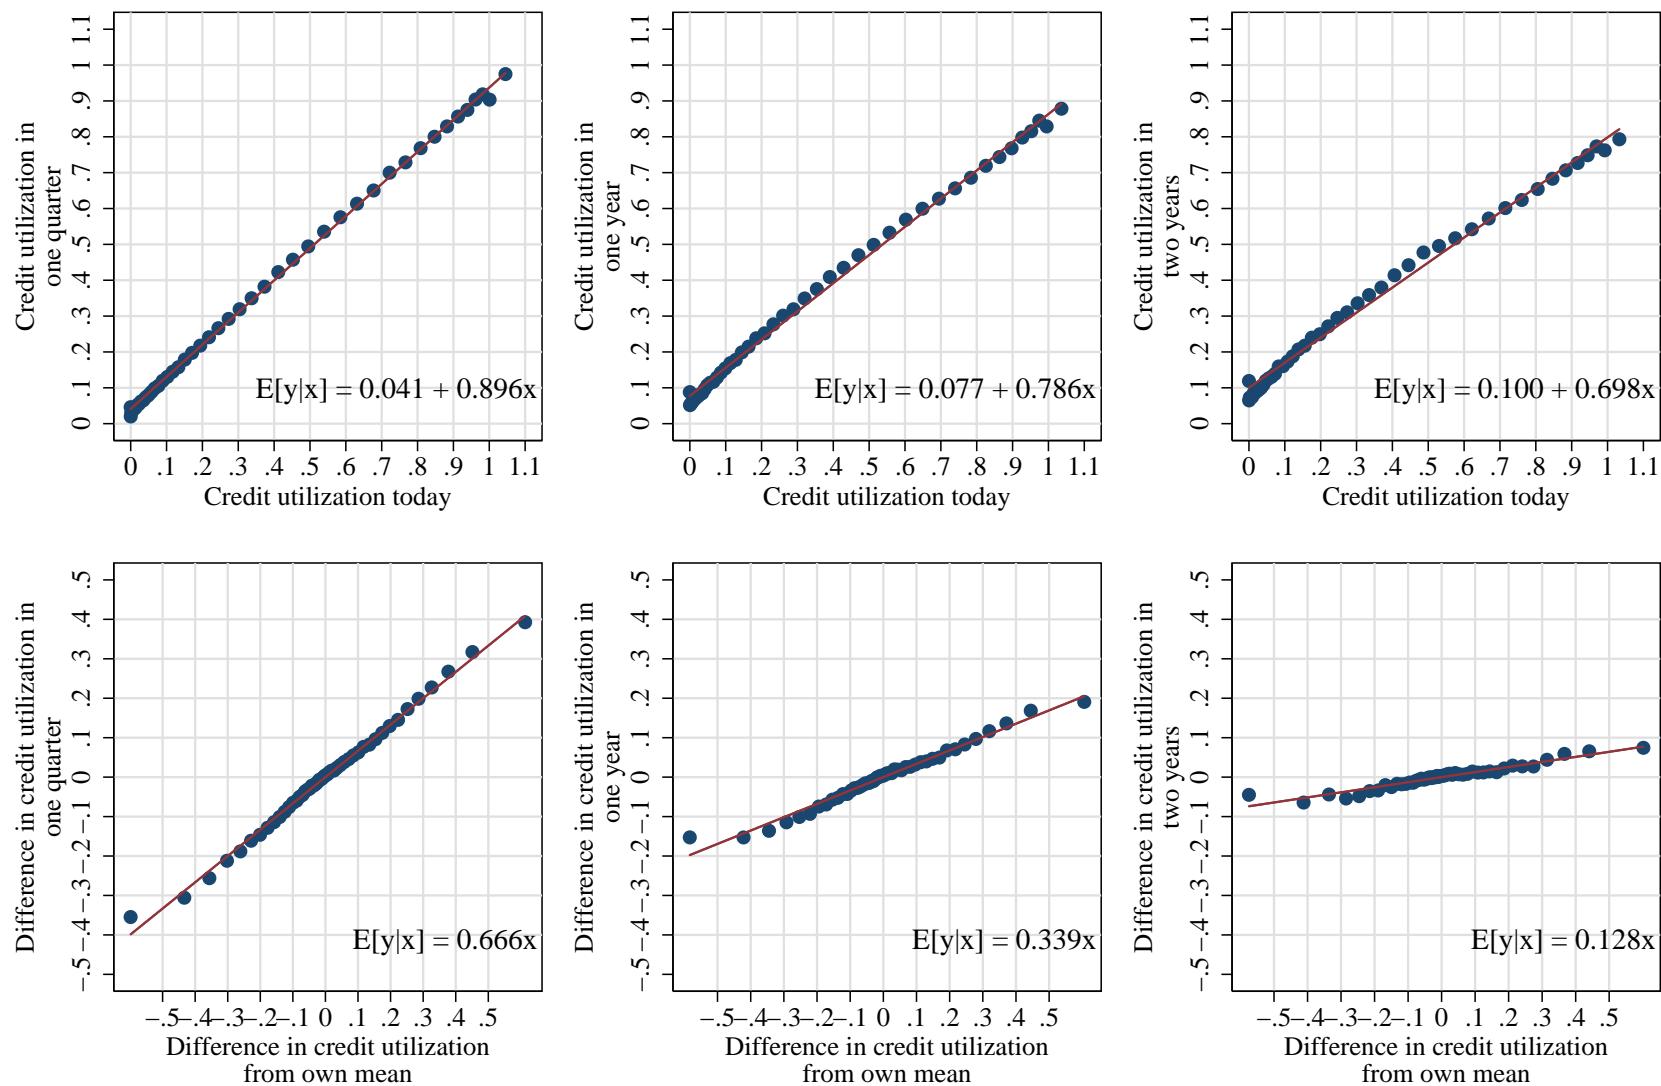


Age 60–80



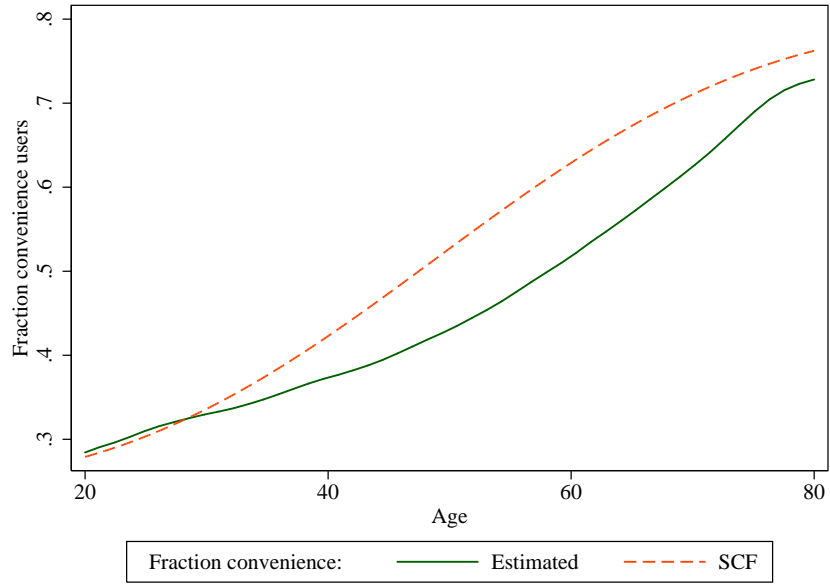
Notes: Each panel shows the histogram for select age groups of credit utilization (credit card debt/credit card limit if the the limit is positive). The histograms exclude utilizations greater than 1.5 Source: Authors' calculations from Equifax/NY Fed CCP.

Figure 12: Changes in credit utilization in one quarter, one year, and two years

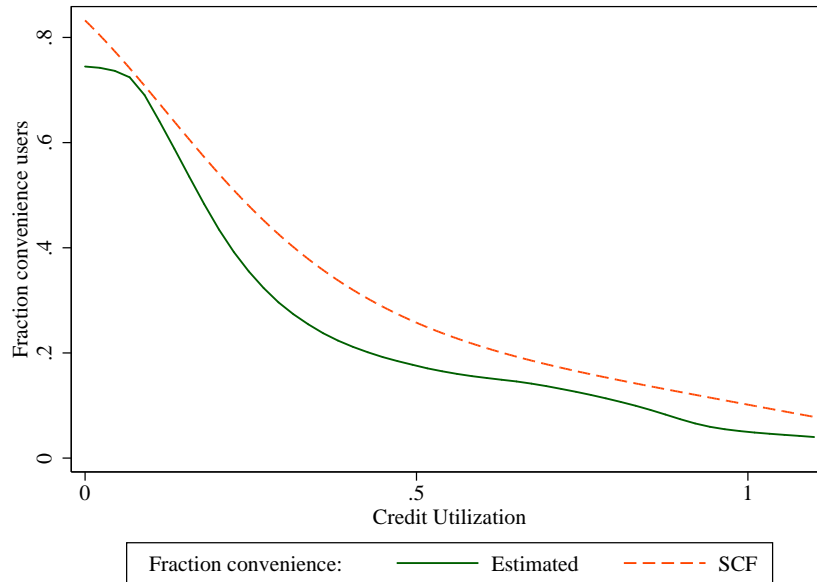


Notes: Each point in the top row shows the credit utilization in the future conditional on having a credit utilization on shown on the x-axis today. The points represent an even distribution of the sample (see figure 11). The bottom row shows the conditional relationship between deviations from the individual 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 13: Fraction convenience from estimates and SCF
(A) Fraction convenience over age

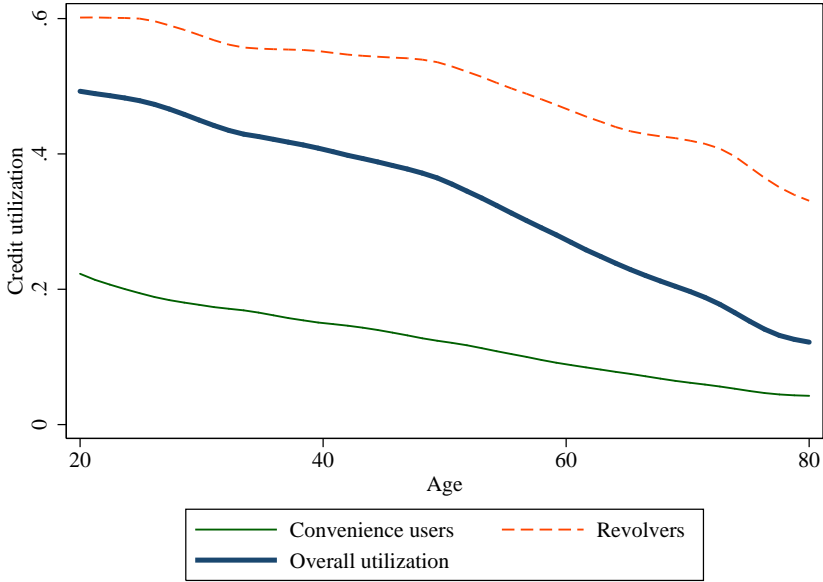


(B) Fraction convenience over utilization



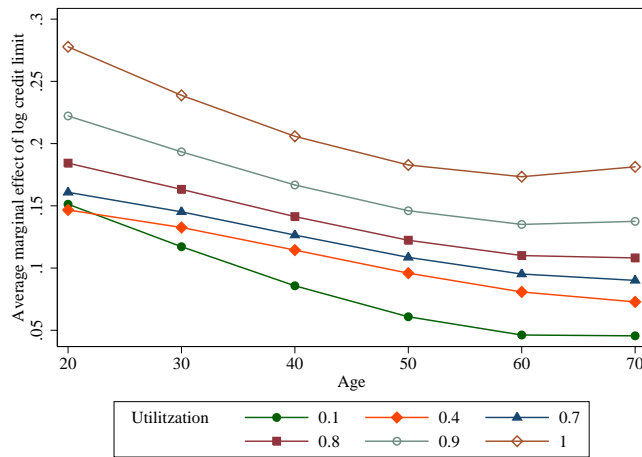
Notes: Source: Authors' calculations from Equifax/NY Fed CCP and SCF. The SCF estimates are based on a cubic logistic regression for age or utilization.

Figure 14: Credit utilization of revolvers and convenience users over the life cycle

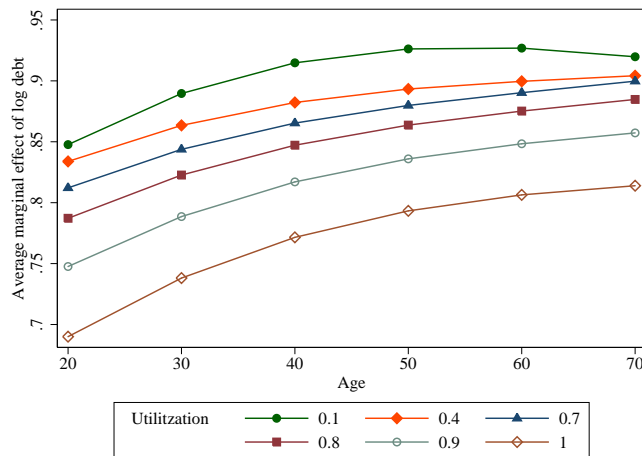


Notes: Authors' calculations from the Equifax/NY Fed CCP based on Finite Mixture Model.

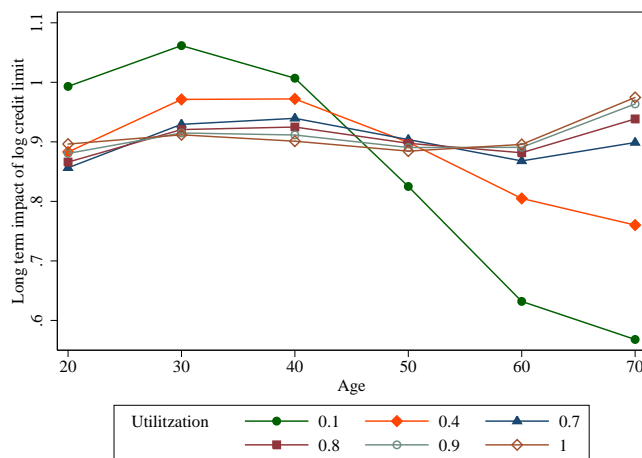
Figure 15: Average marginal effects of credit and previous debt on debt for revolvers
 (A) Marginal effect of log credit limit on log debt next quarter



(B) Marginal effect of log debt this quarter on log debt next quarter



(C) Long-term effect of change in log credit limit on log debt



Notes: For panel (C) each point is calculated using the marginal effect in (A) divided by one minus the marginal effect in (B). Source: Authors' calculations from Equifax/NY Fed CCP based on a finite mixture model separating convenience users from revolvers.

Table 1: Credit utilization

	Credit utilization _t					
	All	All	All	All	All	Revolver
Credit utilization _{t-1}	0.874*** (0.000876)	0.868*** (0.000892)	0.647*** (0.00131)	0.647*** (0.00139)	0.514*** (0.00441)	0.766*** (0.00125)
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	238,111
R-squared	0.741	0.743	0.429	0.444	0.431	0.616
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	
Frac. Variance from FE			0.477	0.467	0.498	

Notes: The sample always includes 0 credit utilization but excludes individual quarters where the utilization is undefined since the limit is zero or utilizations greater than 5 (a very small fraction, see figure 11). 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: Debt and credit changes

	Log Debt _t				Log Limit _t		
	All [1]	All [2]	All [3]	Revolvers [4]	All [5]	Convenience [6]	Revolvers [7]
Log Debt _{t-1}	0.505*** (0.00157)	0.758*** (0.00119)	0.604*** (0.00132)	0.861*** (0.000593)	0.00687*** (0.000561)	-0.0141*** (0.000580)	0.0240*** (0.000616)
Log Credit Limit _{t-1}	0.414*** (0.00262)	0.134*** (0.00148)	0.313*** (0.00243)	0.130*** (0.000787)	0.848*** (0.000933)	0.736*** (0.00114)	0.903*** (0.000807)
Observations	296,369	296,369	361,280	293,014	307,805	361,280	361,280
R-squared	0.432	0.667	0.610	0.926	0.778	0.412	0.412
Accounts	10,028		10,718	10,718	10,149	10,718	10,718
Fixed effects	Yes	No	Yes	Yes	Yes	Yes	Yes
Zero included	No	No	Yes	No	No	Yes	Yes
Age effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Long-term credit impact	0.862	0.875	0.879	0.990			
Credit salience σ	0.443	0.665	0.530	0.756			

Notes: The columns marked “Zero included” include zeros in the level regressions and give all individuals in all quarters \$10 in credit and debt before taking the log in the log regressions and so include individuals with either zero in credit or debt instead of dropping them when zeros not included. Source: Authors’ calculations from Equifax/NY Fed CCP.