

Credit Card Utilization and Consumption over the Life Cycle and Business Cycle

Scott L. Fulford and Scott Schuh

Abstract:

The revolving credit available to consumers changes substantially over the business cycle, life cycle, and for individuals. We show that debt changes at the same time as credit, so credit utilization is remarkably stable. From ages 20–40, for example, credit card limits grow by more than 700 percent, and yet utilization holds steadily at around 50 percent. We estimate a structural model of life-cycle consumption and credit use in which credit cards can be used for payments, precautionary smoothing, and life-cycle smoothing, uniting their monetary and revolving credit functions. Our estimates predict stable utilization closely matching the individual, life-cycle, and business-cycle relationships between credit and debt. The preference heterogeneity implied by the different uses of credit cards drives our results. The revealed preference that some people with credit cards borrow at high interest, while others do not, suggests that around half the population is living nearly hand to mouth.

JEL Classifications: D14, D15, E21, E27

Keywords: Credit cards, life cycle, consumption, saving, precaution, buffer stock

Scott L. Fulford is an economist at the Consumer Financial Protection Bureau. He conducted some of this work while he was on the faculty at Boston College and a visiting scholar at the Consumer Payments Research Center in the Research Department at the Federal Reserve Bank of Boston. Scott Schuh is the Director of the Consumer Payments Research Center and a senior economist and policy advisor. Their email addresses are scott.fulford@cfpb.gov and Scott.Schuh@bos.frb.org, respectively.

This paper, which may be revised, is available on the website of the Federal Reserve Bank of Boston at <http://www.bostonfed.org/economic/wp/index.htm>.

The views expressed in this paper are the authors' and do not necessarily reflect the official position of the Federal Reserve Bank of Boston, the Federal Reserve System, or the Consumer Financial Protection Bureau. To ensure appropriate use of the data, Equifax required that we pre-clear results that used Equifax data before making them public.

We thank David Zhang for his excellent research assistance. This paper has benefited from the comments of participants at the 2015 Canadian Economics Association, the 2015 Boulder Summer Conference on Consumer Financial Decision Making, the 2015 FDIC Consumer Research Symposium, the 2016 NBER Summer Institute Aggregate Implications of Micro Behavior workshop, the 2017 QSPS workshop at Utah State University, the 2017 CEPR Workshop on Household Finance, and the 2017 CESifo Venice Summer Institute, and seminars at the Federal Reserve Bank of Boston, the Bank of Canada, the Consumer Financial Protection Bureau, and the Federal Reserve Board. We thank Chris Carroll, Éva Nagypál, Robert Townsend, Robert Triest, John Sabelhaus, and Fabio Schiantarelli for substantive comments and suggestions.

This version: September 2017

1 Introduction

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

Underneath the dramatic cyclical changes in credit and debt, even larger changes occur over the life cycle and for individuals. Using a large panel from the credit bureau Equifax and collected by the Federal Reserve Bank of New York, we show that average credit card limits increase by more than 700 percent from ages 20–40 and continue to increase after age 40, although at a slightly slower rate (see Figure 2). Because many households hold little or no liquid assets, these increases in credit are one of the largest sources of “savings” early in life. Despite the massive increases in credit with age, debt increases at almost the same rate, and so the fraction of credit used declines very slowly over the life cycle. Average utilization is from 40 percent to 50 percent of available credit until age 50. Individuals also face substantial credit limit volatility—several times greater than income volatility (Fulford 2015)—but we show individual credit utilization is extremely persistent, with shocks dying out almost completely after about two years. Changes in credit and debt are intimately linked over time.

This paper uses this link to study savings, debt, and consumption decisions. Credit cards combine three central aspects of individual decision-making. As precautionary liquidity, credit cards can help people smooth over shocks. By revolving debt over the short and long term, credit cards are a way of allocating life-cycle consumption. And as a payment mechanism, spending on

credit cards forms part of consumer expenditures.¹ Credit cards are the most widespread form of unsecured consumer credit, particularly early in the life cycle, and their credit limits are directly observable, unlike most other forms of consumer credit. High-frequency long-term observation of credit card debt and credit at the individual level is thus a powerful vehicle for understanding not just consumer finance and liquidity constraints, but consumption behavior more generally.²

To understand what the tight link between credit and debt tells us, we incorporate all three aspects of credit cards into a structural model of life-cycle consumption and savings. Our model allows for saving at a low rate of interest and borrowing at a high rate, the large life-cycle variation in credit, and the life-cycle variation in income with uninsured income shocks previously studied by Gourinchas and Parker (2002) and Cagetti (2003). Within the model, we allow the consumer to endogenously decide how much of current consumption to pay for with a credit card. Using new data from the Federal Reserve Bank of Boston's Diary of Consumer Payment Choice, we estimate that non-revolvers would be willing to pay 0.319 percent of their consumption to continue using credit cards. In aggregate, given the current payments infrastructure, rewards, and prices, our calculations suggest that the value to consumers of using credit cards for payments is around \$40 billion a year.

We then embed the value of payment choice in the life-cycle model and estimate the preferences necessary to match the life-cycle profile of consumption or credit card debt using the Method of Simulated Moments (McFadden 1989). To match the life cycle of both consumption and debt, we allow for populations with different preferences in addition to the heterogeneous-agent approach (Aiyagari 1994, Deaton 1991) of many individuals with the same preferences but distinct shocks.

¹In this way, credit cards are similar to debit cards or checks, which may take several days to clear and require an intermediary who promises to pay the merchant first and collect from the consumer later. This payments aspect of credit cards, which involves the inter-relationship between credit and liquidity, has been studied recently by Telyukova and Wright (2008) and Telyukova (2013).

² Except for some work on mortgages (Iacoviello and Pavan 2013), this paper appears to be the first to study the life cycle of credit limits. 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 and introduce endogenous borrowing constraints as an extension. 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 and examines the distributional consequences of changing default policy.

The estimates suggest that more than half the population must be very impatient and care little about risk to hold the amount of revolving debt we observe. The model successfully predicts the slow decline in overall credit card utilization.

The key revealed preference that gives the basic intuition and identification for our results is the different uses for credit cards. Some people, typically called convenience users, use their credit cards only for payments. They have the option to revolve debt and yet rarely, if ever, do. They must be willing to save to have a buffer of wealth so that they rarely need to borrow because of a shocks, and so they must discount the future around the return on liquid savings. Others exercise the option and revolve debt at 14 percent or higher interest for long periods and so must discount the future around the rate of borrowing. The rest of the model machinery of heterogeneous agents over the life cycle is then necessary to account for how individual shocks and the life cycle change decisions. Even patient people borrow when times are sufficiently bad, and young people may want to consume more now because their incomes will be higher in the future.

While the heterogeneity among individuals over the life cycle matters, the most important heterogeneity is revealed by the different uses for credit cards that separate preferences. Our results thus hearken to the older heterogeneous approach in Campbell and Mankiw (1989) and Campbell and Mankiw (1990), who estimate that the relationship between aggregate income and consumption can be explained by dividing the population into two representative consumers, one living hand to mouth and the other saving for the future. Indeed, our estimate of the share of impatient, nearly hand-to-mouth consumers is close to the estimates by Campbell and Mankiw (1990). Similarly, heterogeneous preferences seem necessary to match wealth inequality (Krusell and Smith 1998) or the average marginal propensity to consume (Carroll et al. 2017). At the individual level, building on Gross and Souleles (2002), recent estimates of the response of debt to changes in credit have suggested substantial heterogeneity depending on credit utilization and age (Agarwal et al. 2015, Aydin 2015, Fulford and Schuh 2015). The debt response to credit is closely linked to the marginal propensity to consume (Fulford and Schuh 2015). Our structural estimates capture the rich heterogeneity of use necessary to make sense of these results, and in doing so they closely

match the individual dynamics we estimate from the credit bureau data.

Using our structural estimates, we examine the relative importance of consumer credit for the business cycle and counter-cyclical policy. We simulate an unexpected decline in consumer credit of the same size that occurred in 2008–2009, affecting people at different ages and across the liquidity distribution. By itself, the decline in consumer credit explains one quarter of the decline in personal consumption over the period, although it cannot explain continuing weakness because our estimates suggest the adjustment to lower credit limits is rapid.

One of the central concerns for counter-cyclical fiscal policy is how much households respond to temporary increases in income from, for example, tax rebates (Parker et al. 2013). Kaplan and Violante (2014) summarize the literature and suggest that households consume approximately 25 percent of rebates within a quarter. Because standard models with one asset and no preference heterogeneity have trouble explaining this large response, Kaplan and Violante (2014) build and calibrate a model with an illiquid asset that endogenously generates a large hand-to-mouth population. Our approach is different, but complementary, since we model savings and debt with similar liquidity but different prices.³ The revealed preference of being willing to borrow then suggests a substantial portion of the population has a high marginal propensity to consume. Our simulated consumption response to a small unexpected cash rebate is about 28 percent, driven mostly by the impatient population, a result consistent with recent estimates by Parker (2017). Yet because so much of the available liquidity of U.S. households comes from credit, the simulated consumption response to an unexpected increase in credit is nearly as large as a cash rebate.

Allowing for heterogeneous uses for credit suggests an explanation for the hump shape of life-cycle consumption (Attanasio et al. 1999) that is subtly different from the combination of precaution and life-cycle savings suggested by Gourinchas and Parker (2002). While all agents have life-cycle considerations and their own idiosyncratic shocks, our estimates suggest that the

³The approaches also work along different parts of the income/wealth distribution. Kaplan et al. (2014) show that there are a large number of wealthy hand-to-mouth households who are illiquid-asset rich but cash poor. Revolving credit card debt suggests a high degree of impatience and corresponding low liquid savings on average. While both groups have low liquid assets, the Kaplan and Violante (2014) consumers have invested in illiquid assets, and so the reason for having a high marginal propensity to consume differs, as does how long a household spends living close to hand to mouth.

impatient population is impatient enough that it closely resembles the buffer-stock population in Carroll (1997) over the entire life cycle, with consumption and debt closely following income. The patient population looks much like a liquidity-constrained life cycle/permanent income hypothesis consumer. The average of these two populations has a distinct hump shape of consumption formed mostly by the income profile of the impatient population. Consistent with Gourinchas and Parker (2002), even our patient population is highly liquidity constrained early in life. We show that the low credit limits in early life have particularly negative consequences for welfare, comparable to very large changes in the interest rate. Approaches that do not take into account the large life-cycle variation in credit are missing something important.

2 Credit card use

Both credit and debt change substantially over the business cycle, the life cycle, and for individuals in the short term. This section briefly discusses the context of consumer credit in the United States, introduces our main data sources, and presents some non-parametric and reduced-form results. Fulford and Schuh (2015) provide additional descriptive statistics, including additional evidence on the distribution of credit and on credit card holding by age. In the next section, we turn to a model that helps make sense of these observations.

2.1 The data

The Equifax/Federal Reserve Bank of New York Consumer Credit Panel (CCP) contains a quarterly 5 percent sample of all accounts reported to the credit-reporting agency Equifax starting in 1999. We use only a 0.1 percent sample for analytical tractability for much of the analysis. Once an individual consumer's 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: all credit-cards, auto, mortgage, and student-loan debt, as well as some other, smaller categories.⁴ While the CCP

⁴Lee 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 U.S 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

gives a detailed panel on credit and debt, its coverage of other variables is extremely limited. It contains birth year and geography, but not income, sex, or other demographics. One reason to move to a structural model is to leverage the long, detailed panel on the credit and debt side of the balance sheet to learn about other decisions. An important advantage of the CCP over other data sources used by Gross and Souleles (2002), for example, is that it includes all the credit cards held by an individual. Throughout, we combine all credit cards, giving the complete credit and debt picture. Importantly, we cannot directly distinguish between revolving debt and debt from new charges that will be paid off. Both are credit card debt, and accounting for these different uses is another important reason for introducing the structural model in the next section.

Our analysis is limited to the potential or actual credit-card-using population of the United States because credit card use is what gives us insight into behavior. More than 70 percent of the U.S. population has a credit card at any given time, and a larger fraction has a credit card at some point, because gaining and losing access is common (Fulford 2015). We limit the sample from the credit bureau to include only accounts that have a birth year and that had an open credit card account at some point from 1999–2015. A sizable fraction of accounts represents fragmentary files, typically from incorrect or incomplete reporting to Equifax.⁵

Our analysis is focused primarily on credit card use rather than whether someone has a credit card. The likelihood of credit card possession increases for people when they are in their 20s, but then it quickly stabilizes. We show the age and year distribution of having a positive limit or debt in Figure A-1 in the appendix. Depending on the analysis, we also limit the sample to those with current open accounts, debt, or limits.⁶

has a credit bureau account, although around 11 percent lack credit bureau accounts. See Brevoort et al. (2015) for an examination of these “credit invisibles.”

⁵The accounts are based on Social Security numbers, and so reporting an incorrect Social Security number, for example, can create a fragmentary account that is not associated with other debts. Typically these accounts do not have credit cards, lack a birth year, and are recorded only for a few quarters. Twenty-six percent of accounts lack an age, and of these only 14 percent have an open credit card account at any time.

⁶The CCP reports only the aggregate limit for cards that are updated in a given quarter. Cards with current debt are updated, but accounts with no debt and no new charges may not be. To deal with this problem, we follow Fulford (2015) and create an implied aggregate limit by taking the average limit of reported cards times the total number of open cards. This method is exact if cards that have not been updated have the same limit as updated cards. Estimating the difference based on changes as new cards are reported and the limit changes, Fulford (2015) finds that non-updated cards typically have larger limits, and so the overall limit is an underestimate for some consumers with unused lines.

To estimate our payments model, we also use data from the Federal Reserve Bank of Boston’s Diary of Consumer Payment Choice, which asks a nationally representative sample of consumers to record all of their expenditures and how they paid for them over a three-day period (Schuh 2017, Schuh and Stavins 2017). This rich data source allows us to understand how the payments behavior of revolvers and convenience users differs. In addition, we estimate life-cycle profiles of consumption from the Consumer Expenditure Survey (CE).

2.2 Credit and debt over the business cycle

Since 2000, overall credit limits and debt have varied tremendously. Figure 1 shows how the average U.S. consumer’s credit card limit and debt have varied from 2000–2014. Although the Equifax data set starts in 1999, we exclude the first three quarters of that year, because the limits initially are not comparable (see Avery et al. (2004) for a discussion of the initial reporting problems). From 2000–2008, the average credit card limit increased by approximately 40 percent, from around \$10,000 to a peak of \$14,000. During 2009, overall limits collapsed rapidly before recovering slightly in 2012. Credit card debt shows a similar variation over time. From 2000–2008, 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 during 2009 and 2010.⁷

Utilization is much less volatile than credit or debt. The thick line in the middle of Figure 1 shows credit utilization, the average fraction of available credit used. Because the scale on the left axis of the figure is in logarithms for credit and debt, a 1 percentage point change in utilization on the right axis has the same vertical distance as a 1 percent change in credit or debt. The similar scales mean that we can directly compare the relative changes over time in limits, debt, and credit

For consumers who use much of their credit and so may actually be bound by the limit, the limit is accurate because all their cards are updated.

⁷The fall in debt is not because of charge-offs in which the bank writes off the debt from its books as unrecoverable. The consumer still owes the charged-off debt. Banks may eventually sell charged-off debt to a collection agency, in which case it may no longer appear as credit card debt within credit bureau accounts. Charge-offs are not large enough to explain the fall in debt, although they did increase in 2009. The average charge-off rate from 2000–2007 was 4.35, increasing to 5.03 in 2008 and to 6.52 in 2009, before declining again to 4.9 in 2010 and 3.54 in 2011, and averaging 2.41 since then. See <https://www.federalreserve.gov/releases/chargeoff/delallsa.htm> for charge-off rates for credit cards.

utilization. Credit and debt vary together in ways that produce extremely stable utilization that has no obvious relationship with the overall business cycle. The next two sections examine how the decisions made by individuals combine to form this aggregate relationship.

2.3 Credit and debt over the life cycle

We next examine how credit, debt, and utilization evolve over the life cycle. Figure 2 shows the credit card limit and debt in the top panel and credit utilization in the bottom panel. Each line is for an age cohort that we follow over the entire time possible. The figure therefore makes no assumptions about cohort, age, or time effects. Credit limits increase very rapidly early in life, rising by around 400 percent from age 20–30, and continue to increase after age 30, although less rapidly. Life-cycle variation dominates everything else in Figure 2; while there is clearly some common variation over the business cycle, cohorts move nearly in line with age. We show a more formal decomposition into age and year effects in Figure A-3 in the appendix.⁸ Despite the very large variation over the business cycle evident in Figure 1, changes over the life cycle are an order of magnitude greater.

The bottom panel of Figure 2 shows the average credit card utilization—credit card debt divided by the credit limit—for each cohort. Consumers with zero debt have zero credit utilization, and so they are included in utilization but are excluded from mean debt, which includes only positive values.⁹ Credit utilization falls slowly from ages 20–80. 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

⁸Estimating a simple model that separates the variation between age and year allows us to make the importance of life-cycle variation even clearer. Figure A-3 in the appendix shows the age and year effects from estimating a simple regression of the form:

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

where $\ln D_{it}$ is either log debt, log credit limits, or utilization, and allows these to vary between age effects θ_a and year effects θ_t but imposes common cohort effects. The excluded group is age 20 and year 2000, so each panel in Figure A-3 starts at zero at age 20 and year 2000. The estimated effect is in log units, and so the scale of the figure suggests that variation over the life cycle in credit is around nine ($e^{2.5}/e^{0.3}$) times larger than over time, even with a massive credit contraction.

⁹The calculations in Figure 2 are the average of log limits and log debts to match later analysis and so exclude zeros except for utilization. Figure A-1 in the appendix shows the fraction in each cohort who have positive credit and debt. Including the zeros would lower the average credit limit and debt, but it actually makes the life-cycle variation larger.

credit. Credit utilization does not fall to below 20 percent until around age 70.

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

2.4 The reduced form evolution of individual utilization

The previous two sections show that credit utilization is remarkably stable despite very large changes in credit and debt over the life cycle and business cycle. The aggregate data could be hiding substantial individual volatility in utilization, but this section shows that utilization for an individual rapidly reverts to the mean. While individuals have different credit utilization ratios that represent their own steady state, they return rapidly to their own typical ratio. Credit utilization is best characterized by fixed heterogeneity across individuals and relatively small transitory deviations for an individual over time. We present non-parametric results in Appendix A and Appendix Figure A-4 and reach almost identical conclusions to the parametric estimates. The non-parametric results suggest that the simple linear dynamic reduced-form model we employ is surprisingly accurate. Fulford and Schuh (2015) give additional variations for utilization and show results on how debt and credit co-evolve, rather than fixing the relationship by combining them into utilization. Relatively little is lost by simplifying only to utilization. Moreover, in a Granger Causality sense, the direction of causality moves primarily from changes in credit to change in debt.

Changes in credit come from both the supply and demand side. Card-offering banks cancel cards for their own balance sheet and business reasons—as happened during the crisis—and based on changes in cardholder credit worthiness. In addition, individual account holders often cancel credit cards or apply for new credit (Fulford 2015).

Table 1 shows how utilization this quarter relates to utilization in the previous quarter. For

simplicity, we estimate AR(1) regressions of the form:

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

where $v_{it} = D_{it}/B_{it}$ is the credit utilization given the credit limit B_{it} and the current debt D_{it} , conditional on the credit limit $B_{it} > 0$, and age (θ_a) and quarter (θ_t) effects that allow utilization to vary systematically by age and year. Column 1 does not include fixed effects and so assumes a common intercept. Column 2 includes quarter and age effects, while the other columns include individual fixed effects, quarter effects, and age effects.¹⁰

Without fixed effects, credit utilization is very persistent and returns to a non-zero steady state of approximately 40 percent utilization ($\alpha/(1 - \beta) = 0.38$). Note that this utilization is close to the average in Figure 1, as it should be because both are estimated from the same data, and the non-parametric conditional expectation function shown in Appendix Figure A-4 is nearly linear. Including age and year effects in column 2 barely changes the persistence.

The next column shows how credit utilization varies around an individual-specific mean. Nearly half of the overall variance in utilization comes from these fixed effects. In other words, about half of the distribution comes 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. After a 10 percentage point increase in utilization, 6.47 percentage points remain in one quarter, 1.7 percentage points in a year, and fewer than 0.3 percentage points after two years.

The estimates in Table 1 indicate that while there are deviations from the long-term mean for individuals, these dissipate quickly and are almost entirely gone within two years. The slow decline of utilization with age and the quick return to individual credit utilization suggest that the

¹⁰The combined age, year, and individual fixed effects in equation (2) are not fully identified. As in the age-cohort-period problem, it is impossible to fully identify all effects because there can be an observationally equivalent trend in any one of the age, time, or individual effects. The size of the data set means that rather than estimating individual coefficients—sometimes referred to as nuisance parameters—we instead must use the within transformation. To implement the additional necessary restriction, we follow 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. This restriction is innocuous in the sense that there can still be a trend with age because individuals who 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.

pass-through from an increase in the credit card limit to an increase in credit card debt is large and occurs relatively rapidly. In the next section, we describe a model that helps explain this tight link.

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

We have demonstrated that there is a strong tendency for individual debt and credit to change at the same time, with credit utilization falling only slowly over the life cycle. To explain these observations, this section describes a life-cycle consumption model that is similar to those of Gourinchas and Parker (2002) and Cagetti (2003) but includes the addition of a payment choice, the ability to borrow, and changing credit over the life cycle.

To keep the model numerically tractable and thus able to be estimated, we focus on unsecured credit card debt of individual consumers and do not directly model the endogenous decision to take on non-credit card debt or interactions within households. While these other elements likely affect credit card decisions to some extent, data limitations and numerical complexity make them difficult to address directly, although we can deal with some indirectly.¹¹

3.1 The decision problem

From any age t , a consumer seeks to maximize her utility for remaining life given current resources and expected future income. With additively separable preferences, the consumer at age t with cash

¹¹Most other forms of household debt, such as mortgages, home equity, and auto loans, are secured directly against a household asset, and so their main influence on credit card decisions is how they affect liquidity. The model allows for asset accumulation and income from illiquid assets in late life, but it does not directly model an endogenous liquidity decision as in Kaplan and Violante (2014) or Kaboski and Townsend (2011). In diagnostic regressions in Fulford and Schuh (2015), we have found that the reduced-form relationship between credit card limits and debts explored in Section 2.4 does not seem to change based on whether someone has a mortgage. Student loans are generally taken out before our youngest age of decision-making and so they act mainly to modify disposable income. Households may provide insurance across members (Blundell et al. 2008) and across generations. We observe individual accounts, not households, in the credit bureau data and so cannot directly observe all relevant household interactions, such as household formation, and both members of joint credit card accounts. Within the model, the existence of within-household or intergenerational insurance could be handled indirectly by modifying the uninsurable-income process to allow for a degree of co-insurance.

at hand W_t and current credit limit B_t maximizes the discounted value of expected future utility:

$$\begin{aligned} \max_{\{X_s, \pi_s\}_{s=t}^T} & \left\{ E \left[\sum_{s=t}^T \beta^{s-t} u(C_s) + \beta^{T+1} S(A_T) \right] \right\} \text{ subject to} \\ & C_s = \nu_s X_s \\ & X_s \leq W_s \\ & W_s = R(A_{s-1})A_{s-1} + Y_s + B_s \\ & A_{s-1} = W_{s-1} - B_{s-1} - X_{s-1} \\ & \nu_s = \nu(\pi_s; A_{s-1}), \end{aligned} \tag{3}$$

where she gets period utility $u(\cdot)$ from consumption C_t , which she gets by making expenditures X_t . The decision at t depends on what she expects her future decisions and utility to be at ages $s \geq t$. Within each period she decides what portion of expenditures to fund using credit versus liquid funds. Making payments from different sources of funds comes at a price that drives a small wedge ν_t between expenditures and consumption, the evolution of which we explain below. Expenditures are limited by the available liquidity W_t , which is the sum of assets left at the end of the previous period A_{t-1} (which may be positive or negative), income this period Y_t , and the credit limit this period B_t . Borrowers face a higher interest rate than savers. If the assets A_{t-1} at the end of the period are positive, her assets grow at the return on savings; if assets are negative, she is revolving debt, and her debt grows at the rate for borrowers:

$$R(A_{t-1}) = \begin{cases} R & \text{if } A_{t-1} \geq 0 \\ R_B & \text{if } A_{t-1} < 0, \end{cases}$$

with $R_B \geq R$. The consumer discounts the future with a fixed discounted factor β and so has time-consistent preferences. We therefore drop the distinction between age t and future ages $s \geq t$ for clarity. Most of the elements in this problem are standard. We focus on the nonstandard ones first.

The payments wedge between expenditures and consumption Credit card debt includes unpaid revolving debt from a previous period as well as all new charges. Even if the consumer intends to pay back the new charges by the next bill, convenience debt from new charges is still debt and is reported to credit bureaus as debt. To understand credit card debt, we must account for this convenience use as well as the revolving-debt use of credit cards. Doing so requires us to model why a consumer might use a credit card for some purchases and not others. Using a credit card implies that the consumer finds this way of accessing liquid funds more valuable than other possible ways for making those purchases. Removing this option would come at a cost that we measure. Yet consumers do not use credit cards to pay for all expenditures, and so credit cards must not be usable or the costs of using them must be greater than other methods for some expenditures. We model this within-period decision of what portion of expenditures to pay for using credit cards in a simple way that allows us to estimate it with observable behavior and embed it in the consumption model.¹²

A consumer has two choices for converting liquid funds into consumption. She can use a credit card or some other option that, for simplicity, we will call cash. The consumer must pay a cost to use each method, although we can measure the costs only relative to each other. Each fraction of expenditures $\pi \in [0, 1]$ has a value $N(\pi)$ of using a credit card relative to all other payment methods, so that if $N(\pi) > 0$, using a credit card is less costly than other methods. By making the value relative to other means, we effectively normalize the cost of using cash to zero. Thus we ask whether, for that fraction of expenditures, using a credit card is less costly than cash. The normalization is key to our identification approach, which can identify the value of credit cards only relative to other choices, not in absolute terms. The normalization is innocuous in the consumption model because it affects the marginal value of expenditures in all periods. By indexing the value using the fraction of expenditures, we rule out the possibility that the size of expenditures affects the costs of paying for them. This simplification is important for fitting the within-period payment

¹²Doing so necessarily abstracts from some important monetary concerns around acceptance and general equilibrium. In particular, we do not model firm decisions, but instead assume that the consumer takes all prices and options as given and must make choices given these options. The goal is to write a model that allows us to estimate the consumer's willingness to pay to use credit cards for payments over other means.

decision into the consumption decision.

We next put a simple functional form on $N(\pi)$, which allows us to directly identify willingness-to-pay given observable behavior. We order expenditures so that the value of using a credit card at $\pi = 0$ is the largest and $\pi = 1$ the smallest. With this order, we assume that the relative value of using a credit card is falling at a linear rate with the fraction of expenditures:

$$N(\pi) = \nu_0 - v_1\pi.$$

For the first fraction of expenditures, consumers are willing to pay ν_0 to use a credit card instead of cash. For expenditures for which $N(\pi) \geq 0$, the consumer prefers using a credit card. When $N(\pi) < 0$, she prefers cash because it is less costly. By ordering the costs and assuming a continuous and strictly monotonically decreasing function, we have simplified the consumer's decision from which option to use for every iota of expenditures to finding the optimal fraction of expenditures π^* , where $N(\pi^*) = 0$. The consumer uses a credit card only for the fraction of expenditures for which she gets positive value, relative to other payment methods.

Consumers who revolved debt the previous period have to immediately pay interest on new payments, while convenience users do not. The cost of using a card therefore depends on the borrowing decision in the previous period, creating a feedback from the asset-accumulation decision to the payment decision. Revolving makes consumption slightly more costly, and so the payment decision influences the consumption decision. If expenditures are spread evenly over the month, then a revolver will pay additional interest of $(r_B/12)/2$ on her credit card expenditure that month.¹³ Assuming the loss of float is the only factor explaining different usage, the cost function for revolvers shifts down by $(r_B/24)$.

Figure 3 illustrates these two cost functions and why these simple assumptions help us find

¹³This formula comes from the way that annual credit card rates are reported and interest charged. The interest rate on debt is $r^B = R^B - 1$. The Annual Percentage Rate, or APR, is not a compound rate, and so it is appropriate to divide it by 12 to find the rate of interest. The financing charge on a credit card is calculated based on the average daily balance within a month, and so the financing charge on consumption spread evenly throughout a month is half the interest rate. Note that while the APR is not a compound rate, interest charges not paid off each month will compound in both reality and in our model.

the payments wedge. As the fraction spent on a credit card increases, the value of paying for the next bit of expenditures declines. Eventually, expenditures on a credit card are less valuable than expenditures with cash, and so there is an optimum π^C . Because revolvers start at a lower initial value, their optimum π^R is lower, a prediction we see in the data and will discuss more when we estimate this model in Section 4. Figure 3 also makes clear the identification strategy. With estimates of π^C , π^R , and r_B , it is possible to solve for the two parameters ν_0 and ν_1 and find the area of the wedge for convenience users and revolvers. The area is the sum of the benefits of using a credit card to access funds instead of using cash when a credit card is a better choice. Because the consumer has a choice of how to access funds, and can always choose the other option, the relative cost for the rest of expenditures is zero. The wedge therefore takes on two values:

$$\nu_t = \max_{\pi_t} \nu(\pi_t, A_{t-1}) = \begin{cases} \nu^C = 1 + (\pi^C \nu_0)/2 & \text{if not revolving } (A_{t-1} \geq 0) \\ \nu^R = 1 + (\pi^R(\nu_0 - r_B/24))/2 & \text{if revolving } (A_{t-1} < 0), \end{cases}$$

where π^C and π^R are the optimum fraction for revolvers and convenience users. Appendix C goes through the algebra of exact expressions for π^C and π^R given ν_0 and ν_1 , and it shows how to calculate standard errors given estimates of π^C and π^R using the delta method.

Except when including durable goods, it is often convenient to set expenditures equal to consumption so that $\nu_t \equiv 1$. When there are costs to access funds, however, doing so no longer makes sense. Some payment means have direct costs. For example, obtaining a cashier's check from a bank requires a fee and the time to obtain the check. The consumption paid for with a cashier's check is therefore less than the expenditures. Obtaining cash may require indirect costs and direct costs from ATM fees. To the individual consumer, other payment mechanisms may actually offer benefits. If a credit card offers cash back, for example, then the cost of consumption may be less than the amount spent. Similarly, credit cards sometimes offer insurance on some purchases. Our simple model combines all these costs and benefits into a single value for each iota of expenditures.

To understand why we need to model the payments use of credit cards, consider what the

model says we will see for convenience use and revolving debt. The observed credit card debt at age t in the credit bureau data includes both new charges and previous debt for revolvers, but only convenience debt from charges in the past month for convenience users:

$$D_t = \begin{cases} \pi^C X_t & \text{if not revolving so } A_{t-1} \geq 0 \\ \pi^R X_t + A_{t-1} & \text{if revolving so } A_{t-1} < 0. \end{cases}$$

Debt evolves differently because for revolvers it includes the stock of previous debt, while for convenience users it is only the flow of expenditures.

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

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

$$P_{t+1} = G_{t+1}P_tM_{t+1},$$

where G_{t+1} is the known life-cycle income growth rate from period to period, and the “permanent” or random-walk shocks M_{t+1} are independently and identically distributed as lognormal with mean one: $\ln M_{t+1} \sim N(-\sigma_M^2/2, \sigma_M^2)$. The transitory shocks are similarly distributed lognormally with mean one and variance parameter σ_U^2 . We allow for a temporary low income U_L from unemployment or other shocks with probability p_L each period, and we adjust the shocks so that the mean is always one.¹⁴ The structure of the shocks ensures that the expected income next period is always $E_t[Y_{t+1}] = G_{t+1}P_t$, because the mean of both transitory and permanent shocks is one.

The credit limit Life-cycle variation in credit limits is proportionally several times larger than life-cycle variation in income (compare Figure 2 to Appendix Figure A-5), and the dispersion of credit limits across individuals of the same age is also large (Appendix Figure A-2). We allow for

¹⁴Low-income shocks, in addition to lognormal shocks, may matter for precautionary reasons by putting additional probability on very bad outcomes. Formally, the transitory shocks are distributed as: $U_{t+1} = U_L$ with probability p_L and $\tilde{U}_t(1 - U_L p_L)/(1 - p_L)$ with probability $1 - p_L$, where \tilde{U} is i.i.d. lognormally distributed with mean one: $\ln \tilde{U}_{t+1} \sim N(-\sigma_U^2/2, \sigma_U^2)$ and U_L is unemployment income as a fraction of permanent income.

dispersion across consumers by assuming that the credit limit B_t is an age-dependent proportion of permanent income:

$$B_t = b_t P_t,$$

where $b_t \geq 0$ is the age-varying fraction of this amount that can be borrowed, which is set outside the control of the consumer. This approach means that across consumers, B_t will be in proportion to income P_t , but it allows credit to follow an average path over the life cycle that is different from income.¹⁵

Iso-elastic preferences and normalization. We assume that period utility displays Constant Relative Risk Aversion (CRRA):

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

With CRRA preferences, it is possible to normalize the problem in terms of permanent income P_t at any given age. Using lower case to represent the normalized value, we denote $c_t = C_t/P_t$, $w_t = W_t/P_t$, and $a_t = A_t/P_t$. Appendix B.2 discusses how to rewrite the consumer's problem recursively in terms of the normalized state variable w_t and thus write the solution of the consumer's normalized recursive problem as an age-specific expenditure/consumption function $c_t(w_t, a_{t-1})$.

The beginning and end of life Several important decision parameters affect initial distributions and decisions late in life. We assume the initial distribution of the wealth/permanent-income ratio is lognormal with variance that matches the variance of permanent income shocks and mean λ_0 .

The consumer lives for T periods, where T is a random number that we match to actual life tables,

¹⁵The consumer's problem as written, with W_t as a sufficient period budget constraint, implies that a consumer must immediately repay all debt over her limit if her credit limit falls. To see this, consider what happens if $B_{t-1} > 0$ and the consumer borrows, leaving negative assets at the end of period $A_{t-1} < 0$. If $B_t = 0$, then assets at the end of period t must be weakly positive ($A_t \geq 0$), and so all debt has been repaid within a single period. A cut in credit limits implies an immediate repayment of debt in excess of the limit. This debt repayment when credit is cut below debt does not match credit card contracts, which do not require immediate and complete payment following a fall in credit (Fulford 2015). Instead, credit card borrowers can pay off their debt under the same terms; they just cannot add to it. However, allowing for such behavior means that there must be an additional continuous state variable, because W_t and B_t no longer fully describe the consumer's problem. This adds substantially to the numerical complexity of the solution through the curse of dimensionality.

and we assume she dies with certainty at age \tilde{T} . At death, she receives a final utility $S(\cdot)$ from leftover positive resources. In our base estimations, we set the bequest motive to allow for an annuity to heirs. Appendix B.1 discusses the specific function.¹⁶

Late in life, consumers may face income and expenses different from those they face during working years. Labor income may drop, but consumers may start claiming illiquid retirement benefits such as pensions and Social Security, and they may derive income from other illiquid assets such as housing. They may also face an increase in necessary expenses from additional medical care or other needs. We summarize all of these changes by assuming that income starting at T^{Ret} is a fraction λ_1 of pre-retirement permanent income ($\lambda_1 P_{T^{Ret}-1}$). Allowing for a fall in outside disposable income is a flexible way of combining the many late-in-life changes that consumers may want to plan for during working years, including possibly the acquisition of illiquid assets for retirement. Consumers still earn the return on their liquid assets accumulated before T^{Ret} , but they face no income volatility and continue to consume optimally given their income and expected longevity.

Model frequency We model all decisions as being made quarterly and adjust the discount rates and interest rates accordingly, although we report the yearly equivalent for straightforward comparison to other work. Quarterly decision-making is approximately four times more computationally intensive than yearly. Because of data and computational constraints, much of the structural consumption literature has been limited to examining decisions made at a yearly frequency. Yet consumption decisions must be made more frequently than yearly. If smoothing within the year is perfect, then the frequency should not matter. However, the logic of the model and the data suggest that people do occasionally hit their budget constraint, which implies that ignoring decisions made within the year may miss important facets of consumer behavior. In addition, if we want to

¹⁶Recent work has disagreed over the importance of a bequest motive as opposed to other possible motives for keeping assets late in life, such as long-term care and medical needs (De Nardi et al. 2010). Since we focus primarily on debt, our model and estimates are not well situated to distinguish between motives. While the exact form of the bequest motive or another motive for keeping assets late in life is not important, removing it entirely is consequential. Because the likelihood of dying is increasing with age, people with no bequest motive are effectively getting more impatient. Therefore, they should not decrease the amount of debt they hold as much as the data shows they do. We discuss the effects of alternate formulations of the bequest motive more in Section 4.5.

understand whether the model can match the quarterly dynamics of individual and aggregate credit utilization, it must have at least a quarterly frequency. We adjust convenience credit card debt appropriately so that it represents only one month of expenditure when we estimate the model.¹⁷

3.2 Numerical solution

For a given set of parameters, we find a numerical approximation of the consumer's problem by writing the problem recursively and proceed through backward recursion from the end of life. We briefly discuss some of the unique characteristics of the problem here and give a more detailed discussion in Appendix B.3. We follow the method of endogenous gridpoints (Carroll 2006), which substantially reduces the computation costs. The payments problem can be solved separately from the decision problem in each period, which makes the model numerically tractable. However, the payments problem depends on whether the consumer was borrowing in the previous period, so A_{t-1} is a state variable. The problem depends only on whether there are separate expenditure/consumption functions for revolvers and convenience users: $c_t(w_t, A_{t-1}^R)$. Moreover, consumers take into account the loss of float on new credit card debt when making decisions about whether to leave debt for the next period. Losing the float makes the decision to borrow slightly more expensive.

Figure 4 illustrates some of the complexities of the decision problem. Along the x-axis is the ratio of cash at hand to permanent income w_t . Normalizing this way is useful numerically and because it allows us to compare the decisions of someone earning \$20,000 to those of someone earning \$200,000 in terms of their relative liquidity. Because credit limits also scale with permanent income, only age, previous borrowing, and the current cash-at-hand ratio enter the consumption decision. The consumption functions then tell how much a consumer at that age with those preferences will consume at each liquidity. There are three kinks in the consumption function, which are

¹⁷This adjustment represents a subtle but important point for matching the model to the data. The CCP is a quarterly snapshot of total reported debt at the end of a quarter. Some of the debt was revolved from the previous month—a stock—while other debt is new from the previous month, and represents a monthly flow, since the debt will be paid off before the consumer revolves it. The consumption in the model is all consumption from the previous quarter and so would give convenience consumption three times too large if it were not adjusted to a monthly frequency.

most visible for the impatient 30-year-olds. First, the consumption function has an inflection point where the consumer goes from leaving nothing for the next to period to leaving some liquidity by not borrowing up to her credit limit. When the consumer's liquid resources are below this point, the Euler equation is instead an inequality, because she would like to spend more today but cannot (Deaton 1991). The second two inflection points arise because the interest-rate differential means there are two solutions to the Euler equation for leaving zero assets. One, the limit with assets approaching zero from below, uses the borrowing rate R^B , and the other uses the savings rate R . The economic intuition is that leaving zero assets for the next period is optimal at a high borrowing rate well before it is optimal at a low savings rate. For cash at hand between these two points, the consumer has a marginal propensity to consume of one because the return on savings is not high enough to induce her to save, but the cost of borrowing is sufficient to keep her from borrowing, and so additional resources go straight to consumption. In Appendix B.3, we discuss how to allow numerically for these inflection points so that the consumption function is suitably kinky.

4 Estimation

This section describes how we estimate the structural model using life-cycle profiles of consumption and debt. The estimation works in two stages: First, we estimate the payments value of credit cards for revolvers ν^R and convenience users ν^C in Section 4.1. The structure of the payments problem means it can be estimated separately. We also estimate the processes for life-cycle credit limit and income in a separate first stage described in Section 4.2. Second, we estimate the parameters of the model that minimize the difference between the life-cycle profiles the model produces and the life-cycle profiles of debt and consumption we observe in the data.

We allow for preference heterogeneity by introducing two sub-populations with different preferences and overall income. Of course, additional preference heterogeneity is possible, but our results show that this is the minimum heterogeneity necessary, and we prefer this parsimonious form because it makes obvious the contribution of different populations while not adding too much complexity to the computational problem. Moreover, it is not clear that more preference

heterogeneity is identified without additional assumptions or data. We estimate differences in the income-generating process between the two populations to allow for correlation between preferences and income.

There are thus three forms of heterogeneity in the estimated model: (1) *life cycle*, as people make different decisions at different ages; (2) *heterogeneous agents*, as people are hit with different shocks and so have different assets and incomes and make different decisions based on their current wealth; and (3) *population-level preference and income heterogeneity*, as distinct sub-groups that have different preferences and different income processes react differently to shocks.

To combine groups we estimate the share of group A (f^A) and the multiple of the average permanent income earned by group A (ζ^A). We constrain the population average income of the two groups to match the empirical income profile so that if population A has a higher income, then population B must have a lower income. Together f^A and ζ^A directly determine ζ^B .¹⁸ For each sub-population, the entire decision is described by four parameters: the discount rate β , the coefficient of relative risk aversion γ , the initial wealth-to-income ratio λ_0 , and the fraction of permanent labor income expected from illiquid assets such as housing, pensions, or Social Security in late life λ_1 . We therefore estimate 10 parameters jointly in the second stage: $\theta = \{\gamma^A, \beta^A, \lambda_0^A, \lambda_1^A, \gamma^B, \beta^B, \lambda_0^B, \lambda_1^B, f^A, \zeta^A\}$.

We estimate the parameters of the nonlinear model using the Method of Simulated Moments (MSM) of McFadden (1989). For a given set of parameters $\theta \in \Theta$ and first-stage parameters χ such as the interest rates, payments parameters, and income process estimated separately, we numerically find consumption/expenditure functions at each age. These same θ and χ determine the initial distribution of assets, income, and credit limits across consumers, and how these processes evolve stochastically. For each consumer, we draw from the initial distribution, then for each period we draw from the income-shock distribution. Then the consumer chooses her consumption, and her assets or debt accumulates for the next period. This process proceeds until the final period, generating for a large number of simulated consumers their own idiosyncratic paths of expenditure, assets,

¹⁸For the average income of the combined populations to equal the average observed income $f^A \zeta^A + f^B \zeta^B = 1$, which implies that $\zeta^B = (1 - f^A \zeta^A)/(1 - f^A)$, since $f^B = 1 - f^A$.

and debts at every quarter over their entire life cycle. Combining the simulated consumers, a given set of model parameters generates a life-cycle distribution of consumption, debt, and savings.

The estimation then finds the parameters θ that produce a life-cycle evolution of average simulated consumption and debt that best matches their quarterly empirical counterparts from ages 24–74. We describe the sources and construction of the empirical moments in more detail in Section 4.3. Each profile is quarterly, so there are $T = 204$ quarters. More formally, for a given $\theta \in \Theta$, and first stage parameters χ estimated below, let $g_t(\theta; \chi)$ be the difference between an empirical moment and a simulated moment for each of $2T$ total moments. The MSM then seeks to minimize the weighted square of these differences:

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

where $g(\theta; \chi) = (g_1(\theta; \chi), \dots, g_{2T}(\theta; \chi))$, and W is a $(2T) \times (2T)$ weighting matrix. We generally use a weighting matrix proportional to the inverse variance of the empirical moments, which is the optimal weighting matrix with no first-stage correction. It gives more weight to better-estimated moments. We also show results using the “optimal” weighting matrix, which takes the estimated $\hat{\theta}$ using our standard weights and calculates the optimal weights, taking into account the impact of the first-stage estimates. We adjust the variance-covariance matrix of the estimates of θ for the first-stage estimates, following Laibson et al. (2007), who improve on the work of Gourinchas and Parker (2002) by allowing for the empirical moments to have different numbers of observations.

4.1 Estimation and identification of the payments model

Because of the structure of the consumer’s problem, whether the consumer was revolving as of the previous period is the only way the consumption decision influences the payment decision. We can thus find the solution to the payments problem first and then allow the solution to the payments problem to influence the consumption problem. Table 2 shows the fraction of all expenditures over a three-day period that the nationally representative sample of consumers from the Diary of

Consumer Payment Choice puts on a credit card. The average consumer pays for 17.2 percent of expenditure with a credit card. Revolvers pay for slightly less at 15.6 percent, and convenience users pay for slightly more at 18.2 percent.¹⁹

The difference between revolvers and convenience users then exactly identifies the payment model, as Figure 3 illustrates. We show the algebra for the identification of the payment parameters ν_0 and ν_1 and the delta method to calculate their standard errors in Appendix C. Table 2 shows the estimated coefficients with an interest rate on borrowing of 14.11 percent adjusted for inflation of 2.15 percent (see discussion in Section 4.2 for sources).

The model then directly gives the convenience value of credit cards. For a real borrowing rate of close to 12 percent, the value of using a credit card for payments over other methods is worth 0.319 percent of expenditures to convenience users and 0.235 percent to revolvers, although with fairly wide standard errors. The implied aggregate value of using credit cards for payments is around \$40 billion a year.²⁰ As a comparison, the fees that banks charge merchants for processing credit cards are roughly \$60 billion per year.²¹ The value of the intercept ν_0 suggests that for the most valuable purchases, using a credit card has a value of 4.1 percent of all expenditures for these purchases. For comparison, if all convenience consumers received the equivalent of 1 percent cash back on their purchases with credit cards, the implied consumer surplus would be 0.182 percent of consumption. This calculation likely overstates the direct value of rewards because not all cards offer rewards, but it suggests that about half of the convenience value from credit cards comes from direct rewards or other card benefits, and the other half comes from their value as a convenient payment mechanism.

¹⁹Credit card use is fairly stable with age, although with wide standard errors (Fulford and Schuh 2015). Interestingly, both revolvers and convenience users over 65 tend to spend more on a credit card.

²⁰Personal consumption expenditures were \$12.3 trillion in 2015, according to the BEA. If half of the population is revolving, then $12283 * (0.319/100 + 0.234/100)/2 = 36.6$ billion. Note that this calculation is an estimate of the consumer surplus of credit cards as a payment mechanism over other means, given the current payments ecosystem, and so does not directly calculate welfare. For example, the calculation does not take into account the costs of operating the payments system or the producer surplus from additional sales made because some purchases are more convenient, or the gains to the processors, network operators, and banks.

²¹The total value of credit card payments was \$3.16 trillion in 2015 (see the 2016 Federal Reserve Payments Study <https://www.federalreserve.gov/newsevents/press/other/2016-payments-study-20161222.pdf>). The percentage charged to merchants varies from approximately 0.75 percent to 4 percent, but appears to average around 2 percent. Fee revenue is therefore around \$60 billion, most of which is accounted for by the interchange fees shared by banks after payouts to card networks, processors, and other parties.

4.2 First-stage estimates

This section describes the sources and estimates from other data sets that identify the ancillary parameters of the model. We estimate a fifth-order polynomial of the average life cycle of income to find income growth (G_s) at each age using after-tax income per adult household member from the Consumer Expenditure Survey from 2000–2015. The raw data and the fitted lines are in Appendix Figure A-5. Similarly, we take a fifth-order polynomial estimate of the total credit limit per account from the Equifax/NY Fed CCP to form B_s .²²

While average income follows the observed life-cycle path, individual incomes vary based on their idiosyncratic shocks. We use the estimates of the annual income process from Gourinchas and Parker (2002), which are updates of Carroll and Samwick (1997), calculated from the Panel Study of Income Dynamics. We adjust these volatilities for quarterly dynamics so that four quarterly shocks combine to produce the same variance as one yearly shock. The quarterly transitory variance is approximately four times the annual variance because quarterly shocks average out, while the quarterly permanent variance is approximately one-fourth the yearly variance because permanent shocks stack. We estimate the probability of low income p_L based on the average monthly unemployment rate from 2000–2015 of 6.3 percent.

We observe two interest rates directly, although there is likely greater heterogeneity in interest rates than we incorporate in the model. We set the interest rate on debt $R_b - 1 = 14.11$ percent based on the average revolving interest rate over the period, adjusted slightly for the lack of explicit default risk in the model.²³ We would like to capture the returns that people expect to receive on their savings, but the appropriate rate of return is not obvious because there is only one riskless asset. We therefore set the return on savings at 5.4 percent, which is the average return on an

²²Not smoothing these two budget constraints makes little difference to the overall estimates, but it introduces distracting jumps in life-cycle consumption and debt as consumers respond to sudden changes in the budget constraint driven by jumps in income or credit that disappear. B_s is proportional to permanent income for an individual consumer, and the problem is set up so that the average permanent income across all consumers is the average income, allowing us to back out b_s .

²³The Federal Reserve series G19 (Commercial Bank Interest Rate on Credit Card Plans NSA) average over the period is 14.73 percent (the average credit card interest rate reported in the SCF is 14.22 percent). Based on calculating the risk of default from the PSID, Edelberg (2006) calculates the zero bankruptcy risk rate would be 0.62 percentage points lower, a smaller adjustment than in Angeletos et al. (2001), who adjust for default by 2 percentage points.

all-bond portfolio from 1926–2015 as calculated by the mutual fund company Vanguard. We adjust both borrowing and saving prices for the geometric average inflation rate from 2000–2015 of 2.15 percent. In the expected growth over the life cycle, we also include expected real aggregate growth of 1.5 percent, the average compounded rate from 1947–2015 from the Bureau of Economic Analysis (2009 chained dollars GDP per capita).²⁴

We consider economically active life to last for 51 years (204 quarters) from age 24, when most people have finished schooling, through age 74, when differential death rates and other end-of-life concerns dominate. While the Equifax/NY Fed credit data have many observations even for older ages, the Consumer Expenditure Survey (CE) becomes increasingly sparse and, for privacy reasons, topcodes ages above 80, with the top age varying by year of the survey. Before age 94, individuals have a probability of dying and leaving a bequest at each age. We set the probability of death to match the age structure of the population in 2010.²⁵

4.3 The empirical life-cycle moments

We estimate the model to provide the best fit to two life-cycle profiles: (1) the observed credit card debt over the life cycle from the Equifax/NY Fed CCP described in Section 2, and (2) the observed household consumption over the life cycle from the CE from 2000–2014. Because our observed credit data are for individuals rather than households, we adjust household consumption by dividing by the number of adults in the household. We allow for some unobserved taste changes over the life cycle by adjusting consumption for the number of children in the household.²⁶

Using two different life-cycle profiles—one a flow of consumption, the other a stock of debt—presents some unique challenges relative to other work. We want the model to fit both, and so we want the

²⁴While each of these parameters is volatile, and different agents may experience different prices, there is no sampling variance about them, and so we do not adjust the MSM variance-covariance matrix for them.

²⁵See https://www.cdc.gov/nchs/data/nvsr/nvsr63/nvsr63_07.pdf, accessed 8 August 2017.

²⁶Formally, we estimate:

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

and then calculate average household consumption per adult at each age after removing the effect of children at the individual level. Removing the implied consumption effect of children has a surprisingly small effect. Figure A-5 in the appendix shows the unadjusted and adjusted consumption. Children slightly raise expenditures per adult household member from ages 35–45, but the adjustment is small.

estimates to give approximately equal weight to departures from debt and consumption. The flow of quarterly consumption is larger than the stock of credit card debt. By dividing each moment by the average over the entire life cycle, we define our moments in such a way that they are scaled comparably.²⁷ For example, for the debt moments at each age t we define:

$$g_t^D(\theta; \chi) = \frac{1}{\bar{D}} \left((1/J_t) \sum_{j=1}^{J_t} D_{j,t} - (1/K) \sum_{k=1}^K \hat{D}_{k,t}(\theta; \chi) \right),$$

where $\hat{D}_{k,t}(\theta; \chi)$ is the simulated debt of person k at age t given parameters θ averaged over the K people we simulate, and $D_{j,t}$ is the debt of person j in the CCP at age t , averaged over the J_t people in the CCP sample. A second potential problem is that the two sets of moments may be estimated with different degrees of precision, and the standard estimation weights will attempt to fit one set of moments perfectly at the expense of large departures from the other. Our consumption and CCP sample sizes are relatively similar, however, and so we avoid this problem.²⁸ Appendix B.4 discusses the construction of the variance-covariance matrix of the combined moments.

4.4 Estimation and identification of the life-cycle model

Using the first-stage estimates of the payments problem and the other parameters, we next estimate the full life-cycle model and then discuss the variation that helps identify the different parameters. Table 3 shows the model estimates, while Figure 5 shows how debt and consumption vary over the life cycle in the model and empirical moments. Because the scales of the two top panels of Figure 5 are in logs, the estimation approximately finds the parameters so that the weighted sum of the squared differences between the predicted consumption and debt lines is as small as possible. It

²⁷An alternative we have also tried is to define $g_t^D(\theta; \chi)$ as the difference in the log of the mean debt and log of mean simulated debt. Since we have to adjust the variance matrix for estimating logs instead of levels, doing so introduces a different sort of complication. The results using this approach are similar.

²⁸This problem is not hypothetical. We used a 0.1 percent sample from the CCP for computational reasons, but we could have calculated the full 5 percent sample, which would have been much more precisely estimated than the CE. With a standard weighting matrix that is the inverse of the variance-covariance matrix of the moments, the consumption moments would have had very little weight. One possible solution is to divide each moment block of the weighting matrix by its trace, which normalizes the overall weight of the blocks relative to each other, but it preserves the relative weighting within each block.

is clear that, given the constraints of the life-cycle optimization model, the model estimates can successfully capture the life-cycle profiles of debt and consumption.

To do so, the model suggests that about 64 percent of the population (f^A) must be fairly impatient (β^A) and not care very much about risks (γ^A). This portion of the population, which the figure and tables call population A, has already acquired some debt (λ_0^A) by age 24 and has substantial revolving debt throughout the life cycle. To match the amount of debt and consumption, the estimates suggest that this population has an income slightly higher than the average (ζ^A), although with a large standard error that does not reject average income.²⁹ Because individual credit limits are proportional to income, the members of this group cannot be too poor, otherwise they would not be able to hold and make payments on their debts. Because the discount rate is high and risk aversion is low, most of this population lives essentially hand to mouth over the entire life cycle, relying on credit for all of their smoothing. This population's average utilization is high through much of the life cycle (see the third panel in Figure 5).

The estimates suggest that the other portion of the population must be relatively patient and risk averse. Population B is too patient to ever want to hold much debt and has not acquired much debt by age 24 in any case (λ_0^B). So consumers in population B rarely borrow except in their 20s, when some have enough shocks to want to borrow for a brief time. Their credit card debt is thus almost entirely from convenience use.³⁰ Because this population expects to receive little income after expenses (λ_1^B) in late life and is relatively patient, this population spends early life accumulating savings for late life. Consumption increases early in the life cycle as income and savings increase, but it becomes relatively flat afterward as this population smooths consumption over the rest of the life cycle.

Because this is a nonlinear model, all moments are typically used to identify all parameters, but it is useful to understand how different sources of variation identify the parameters. Both the

²⁹In comparisons using the SCF, we found that the median income of revolvers was larger than the median income of convenience users, while the mean income of convenience users was larger. The slightly higher income of revolvers is likely also a result of using CE income, which may miss high incomes, with administrative debt data in the CCP, which includes large debt values.

³⁰The added debt from convenience use of credit cards is one month's worth of consumption (one-third of quarterly consumption) times the estimated rate of consumption on a credit card for a convenience user from Table 2.

consumption and debt that we observe over the life cycle are population averages, so the model is identified from the average of the two model populations. The share of population A (f^A) and its relative income (ζ^A) change the mix of the two populations. For the model to produce as much debt as in the data, a large portion of the population (f^A) must be relatively impatient and not overly concerned about debt. This population's impatience (β^A) is mostly pinned down by the borrowing rate (R^B) to make its members willing to hold debt. If the population is too patient, it will not accumulate enough debt. If it is too impatient, it will acquire too little debt. It must have enough income to support the amount of debt it holds, helping to identify ζ^a . To get an average consumption profile in which consumption is below income for much of the life cycle therefore requires the other portion of the population to be relatively patient, with its discount rate (β^B) close to the savings rate R .³¹

While the levels of consumption and debt come from the average, the life-cycle profiles are largely determined by only one of the populations. Because the patient population carries almost no debt—the flow of debt from payments is relatively small compared to the stock of revolving debt—the profile of credit card debt largely identifies the preferences of the impatient population, their initial wealth, and their expected residual income late in life. Given this population's impatience, consumption must closely follow income. The hump shape of debt comes from increases in credit limits early in life, which allow this population to increase its debts, and the fall in income after age 50, which makes carrying as much debt less affordable. This population's risk aversion (γ^B) is identified by how much credit it keeps as a buffer.

The more patient and risk-averse population carries little revolving debt, so all of its debt comes from the convenience use as a share of consumption. The impatient population A has a strong hump in consumption as it follows income. For the average consumption profile to be below average income, the patient population B must have a relatively flat consumption profile

³¹Since we include expected aggregate growth and adjust for inflation, β^A and β^B are more closely pinned down relative to $R^B - \text{Inflation} + \text{Real Aggregate Growth}$. We thank Chris Carroll for pointing out that even if we remove trends from life-cycle profiles, the economic decision of the agent includes expected aggregate growth, and so we need to include it to correctly model their decisions. Aggregate growth implies everyone expects to have more income next period and so should be more impatient.

without a downturn late in life. Its preference for risk (γ^B) and expected late-life income after expenses (ζ^B) are determined by this shape, with its discount rate (β^B) pinned down by the rate of return on savings. Its risk aversion determines the size of the buffer of savings it builds up early in life, and so the initial level and slope of consumption over the life cycle help identify γ^B . The risk aversion and initial wealth (λ^B) are not particularly well identified by the life-cycle moments, and their standard errors are relatively large.

The remaining four panels of Figure 5 show model predictions for other life-cycle paths. The model captures the slow fall in credit utilization over the life cycle. The fall comes primarily from revolvers using less of their credit as their limits increase and, secondarily, from incomes decreasing and making debts less affordable. To examine the evolution of wealth, which may be negative, we take the log of wealth after giving everyone \$10,000, which allows us to consider the full distribution in a single graph. The model estimates predict less wealth accumulation over the life cycle than estimates from the Survey of Consumer Finances, but it predicts a similar trend increase and flattening after age 55. The model captures the level of the variance of credit card debt reasonably well, although it does not predict the shape very well. The variance in debt is largely coming from the combination of preference heterogeneity and permanent-income heterogeneity. Because convenience and revolving use both count as debt, there is a big gulf between the average debt of convenience users and revolvers, and income heterogeneity then adds to this variance.³² The model was not estimated to match these profiles, and so its ability to successfully predict something close to their level and evolution suggests that the model is capturing important facets of life-cycle decision-making.

³²We have also estimated the variance in the change in debt from quarter to quarter, which controls for the permanent income and preference heterogeneity. While the cross-section variance of debts is largely a reflection of persistent differences, the change in debt captures changes for the same individual over time. The variance in the change in log debt from quarter to quarter in the Equifax/NY Fed data is approximately 1.5 and is nearly constant over the life cycle. The variance of the change of log debt from our standard estimates is around 0.5, with little life-cycle variation. Changes in debt come from short-term income shocks, which cause people to want to spend more or less, and long-term income shocks that change the credit limit. Since our estimates do not include credit limit volatility apart from income volatility, and Fulford (2015), using the Equifax/NY Fed data, shows that credit-limit volatility is about four times greater than income volatility, our model has too little credit-limit volatility. We also do not include shocks to marginal utility other than income, and these are likely to be important. For example, spending for a vacation or a durable good may increase convenience or revolving debt, but it is a lumpy consumption decision, not an income shock.

The heterogeneity in preferences is key to the model's ability to capture, even approximately, more than one life-cycle profile. Gourinchas and Parker (2002) estimate parameters to match the consumption profile and under-predict wealth accumulation, while Cagetti (2003) estimates parameters to match the wealth profile but needs such a high degree of risk aversion that it is difficult to capture the consumption profile. Allowing for multiple populations with different preference lets us come closer to both profiles, even if our primary purpose is to explain debt.

By allowing for preference heterogeneity that is a natural outcome of the multiple uses of credit cards, our estimates also suggest a mechanism for the life-cycle hump of consumption that is somewhat different from the one in Gourinchas and Parker (2002). The estimates suggest that the hump is caused by the combination of a very impatient population consuming almost all of its income and a more patient population with a relatively flat consumption profile consuming less than its income. The combined population consumes less than its income on average and has the distinct hump shape in consumption. The hump in our model comes primarily from population heterogeneity, rather than from the combination of precaution early in life and retirement accumulation late in life for a single preference-homogeneous population suggested by Gourinchas and Parker (2002). While differences in income in the short and long term are still important, the major source of heterogeneity in outcomes such as debt and consumption come from differences in preferences. In many ways, this explanation fits with the approach of Campbell and Mankiw (1989) and Campbell and Mankiw (1990), who suggest there might be a fraction of the population living hand to mouth. Similarly, while Krusell and Smith (1998) do not have a life cycle component, income differences are not sufficient to explain wealth inequality in their estimates; instead, differences in discount rates are the major driver of wealth inequality.

One continuing puzzle is that the model predicts the fraction revolving will be approximately constant over the life cycle, while the surveys suggest it should decline over the life cycle (see Figure 5).³³ The average fraction revolving over all ages in the model is approximately correct,

³³Since the SCF sometimes has trouble with credit limits (Zinman 2009) and is at a household rather than individual level, there is reason to question whether the SCF fraction revolving is the best benchmark. Estimates from the Diary of Consumer Payments Choice in Fulford and Schuh (2015) suggest that the fraction revolving is approximately constant until age 50, but then declines steadily. While this profile is closer to the model prediction, the model is still not

but because the impatient population A is always in debt and the patient population B almost never revolves, the fraction revolving does not change much over the life cycle. It is possible that additional heterogeneity across the population could replicate this behavior by allowing part of the population to be on the edge between revolving and not, with the balance shifting over the life cycle. The problem with this approach is that the difference between the savings interest rate and the borrowing rate is so large that it is for only very particular preferences that someone who is not willing to borrow late in life would be willing to borrow early in life. Such particular preferences are sensitive to changes in interest rates, which suggests that credit use is sensitive to interest rates. The stability of utilization in Figure 1 despite changes in interest rates over the period suggests that the answer likely lies elsewhere. Perhaps reductions in income volatility as people age (Sabelhaus and Song 2010) might explain it, although other risks seem to be increasing late in life (De Nardi et al. 2010). Instead, there may be some sort of preference change over the life cycle, such as increased patience or financial literacy.

4.5 Robustness and variations

In this section, we examine the robustness of the estimates to changes in weighting matrices, starting points of the estimation, and model choices. Table 3 shows the over-identification statistic for each estimation, which always decisively rejects the hypothesis that the model is not over-identified. The choice of weighting matrix is therefore not innocuous; because the model is over-identified, different weighting matrices will give statistically different results, so the best estimate we present should be viewed as one of many possible estimates. In this section, we characterize how the estimation would change based on alternative choices and whether the changes affect our conclusions. Doing so also provides additional evidence about how different parameters are identified. While particular parameters are sensitive to estimation and model choices, our overall conclusions are not.

The second column of Table 3 shows estimates that use the two-stage “optimal” weighting

predicting a movement out of revolving late in life.

matrix, which first estimates the parameters using our standard weighting matrix and then uses those estimates to calculate the weights that asymptotically minimize the variance of the estimator. The estimates are broadly similar; the impatient population is more risk averse but less patient, and so it carries more debt and is a smaller share of the population. The “optimal” weight matrix puts weight on different moments, so its predicted life-cycle profiles are somewhat different (see Appendix Figure A-6) and do not fit the debt profile as well as our standard weighting matrix does.

Our overall conclusions are also robust to alternative starting points for estimation. Our numerical procedure for finding the minimum of equation (4) proceeds by using numerical derivatives calculated from a starting θ_0 to move to a local minimum where the derivatives in all dimensions are zero to within a small tolerance. This procedure is guaranteed to find only a local minimum, however. We therefore start the procedure with θ_0 in a grid that covers the 10 dimensional parameter space. Not all starting points produce the same estimate of θ , indicating that the objective function in equation (4) has multiple local minima. The procedure converged to our best θ^* from a wide range of starting θ_0 , and so θ^* is a candidate for the global minimum. We discuss other local minima in Appendix B.5. The overall conclusion holds for all local minima: Around half of the population must be fairly impatient and close to risk neutral.

In the last two columns of Table 3, we examine how changing the model changes estimates. Our baseline estimates allow consumers to take into account the effect their consumption decisions will have on their payments decisions, at the cost of substantial numerical complexity. In column 3, we remove this feedback. The payments decision still affects consumption, but the consumer does not take into account this future value when making consumption decisions. Removing endogenous payments choice leaves the estimates almost exactly the same. Because so few people switch from revolving to convenience use, the value an individual gets from credit card consumption this period is almost always the same as next period. Since the value of consumption on a credit card does not affect the marginal utility tradeoff between today and the future, it does not affect the decision. Including convenience use as part of credit card debt is necessary, however, because the debts we observe in the credit bureau data include both revolving and convenience debts. Allowing

consumers to take into account the impact of consumption choices on payment choices in the future does not appear to be particularly important for their consumption decisions.

We do not estimate directly the strength of the bequest motive since it is not well identified. Our bequest function, described in greater detail in Appendix B.1, gives people the discounted utility from their heirs consuming the annuity value of assets at death as well as the heirs' own income. The strength of the bequest motive is determined by how much more income the heirs have compared to the individual; as the heirs' income increases, the marginal value of leaving anything to them diminishes. Our baseline estimates assume heirs have twice the permanent income upon death of the individual. The last column assumes heirs instead have 10 times the permanent income. The estimates are similar, suggesting that our estimates are robust to other assumptions about bequests, and that, given our approach and data, the bequest motive is not well identified.

5 Model predictions and welfare analysis

In this section, we take the estimated model and ask how well it predicts phenomena outside the life cycle. These results provide both an out-of-sample examination of how good the model estimates are and whether the model can successfully explain other phenomena that we did not estimate it explicitly to explain. We then examine the welfare impact of changing some parameters of the model.

5.1 Aggregate credit-fall experiment

We simulate a large population with an age profile matching the population from age 24–74 and a credit drop of the same size as the one that occurred over 2008–2009. In addition to life-cycle income growth and individual income volatility, aggregate income grows at a constant rate of 1.5 percent per year, just as the consumers in the model assume. We also adjust the dollar values for the average inflation rate. Finally, to mimic the fall in credit limits that started in the final quarter of 2008 and continued through 2009, we introduce a fall in credit of 35 percent for one-sixth of the population for six quarters. This experiment is the simplest way to produce the approximately 35

percent drop in credit limits spread over more than a year that is evident in Figure 1, but it is not a full replication of the changing environment. In particular, it does not include a fall in income or a possible decline in expectations of future income growth.

The individual dynamics of credit utilization from the simulations closely match the dynamics from the credit bureau data. Table 1 shows that once we control for fixed unobserved heterogeneity with fixed effects, shocks to utilization disappear quickly, with 64.7 percent of a shock surviving each quarter (the third column). The last column performs exactly the same regression on the simulated data. The simulated consumers experience the large unexpected fall in credit in 2009 and the expected increase over the life cycle, but the only unexpected credit volatility that they face comes because credit is proportional to volatile permanent income. Because volatility in income is much less than volatility in credit (Fulford 2015), the consumers in the model face less credit volatility than actual consumers do over the time period. Nonetheless, their average response to changes in credit limits is very close to that of actual consumers; the estimated model captures the dynamics of credit utilization closely, with 72.3 percent of a shock persisting to the next quarter.

We finally examine the business-cycle dynamics with which we started the paper. The bottom panel of Figure 1 shows the aggregate response of the simulated consumers to the 35 percent fall in credit introduced over six quarters. Credit continues to increase over the entire period at the same 1.5 percent rate as income, plus 2.1 percent for average inflation, partly counteracting the large fall. Model credit growth is slightly slower than actual credit growth over the period, suggesting that pegging credit to income does not fully capture the aggregate growth. Since consumers expect credit growth, their debt grows at the same time, and credit utilization is stable despite the growth before and after the crisis, just as in the data. The model successfully predicts about the same credit utilization as in the data.

During the crisis, debt quickly adjusts to the fall in credit, so utilization is much smoother than either credit or debt, although not as smooth as the data. As the individual dynamics show, while shocks at the individual level disappear quickly in both the model and data, it still takes several quarters for consumers to fully adjust their debt and savings to a 35 percent fall in credit. The

excessive smoothness of utilization in the credit bureau data suggests that there must be additional features of the period not captured by the simple simulated shock, because the simulated data closely matches the individual responses to a fall in credit. The simulation does not match the distribution of the fall in credit, which was initially concentrated among those consumers with high credit scores, while later declines occurred among those with lower credit scores (Fulford 2015). Even without these features, our model produces a notably smoother path than a simple version of the Life Cycle/Permanent Income Hypothesis (LC/PIH) would suggest.³⁴

How important was the fall in credit for consumption? Our model makes clear a causal connection between the fall in credit limits and the fall in debt through a reduction in consumption. From the second quarter of 2008 to the final quarter of 2009, real consumption per person fell 9.2 percent relative to the trend from 2000–2008. Our simulations suggest that the fall in credit limits over the same period was responsible for a fall in consumption of 2.3 percent relative to trend, or about one-quarter of the fall. The fall in consumption from the simulations quickly rebounds, however, as consumers rebuild their liquidity, and so a fall in credit does not explain the continuing weakness in consumption after 2009.³⁵

5.2 Implications for stimulus policy

We next briefly examine the welfare and policy consequences of modifying the economic environment, given our structural estimates. A caveat with all of these results is that our model is in partial equilibrium, and so it does not reflect any changes in prices that might result in general equilibrium. Guerrieri and Lorenzoni (2017) examine some of the complications and amplifications that

³⁴Constructing the path of the LCH/PIH is not entirely trivial or without assumptions. By definition, in the PIH, liquidity constraints can never bind, otherwise a precautionary motive arises (Carroll and Kimball 2001). Counterfactually to the results in this paper, credit limits cannot matter for the PIH. We construct the PIH line in Figure 1 by taking the 2008Q1 debt as the optimal distribution. Since we do not vary the age structure of the population or the growth rate, that amount of debt, adjusted for inflation, is the correct amount of debt for the entire period. Utilization is therefore falling until 2008, as limits increase, and then increases proportionally to the fall in credit limits.

³⁵Figure A-7 in the appendix shows the relative paths of consumption from our simulations and real personal consumption per person from the BEA. The fall in consumption in the data, relative to trend, continues even after the credit contraction stops, while a fall in consumption caused by consumer credit produces a V-shaped path. Consumption from the simulations is actually higher after several years, because debt is lower, and so interest payments decline. This path is general following credit changes in precautionary models (Fulford 2013).

occur in a buffer-stock model in general equilibrium.

The ability to temporarily boost consumption is an important tool for counter-cyclical policy. One way to provide such a boost is with direct cash infusions through tax rebates (Parker et al. 2013). For such a policy to be effective as a stimulus, individuals must increase spending soon after the rebate. Kaplan and Violante (2014) summarize the literature and suggest that the additional non-durable consumption within a quarter is around 25 percent of the rebate. Yet standard models, even with income uncertainty, predict very small responses. Figure 4 illustrates why. Our patient population B has preferences that look similar to standard assumptions based on calibration or estimation that attempts to match the level of wealth. The distribution of liquidity for our patient population at age 30 puts almost no one at a steep part of the consumption function, even this early in the life cycle, and so rebates have a small impact.

Our population estimates produce responses to temporary payments that are similar to empirical estimates, because our estimates suggest that a large portion of the population has a strong marginal propensity to consume. Using the estimates from column 1 in Table 3, we simulate the population response to a temporary, unexpected cash gift of 5 percent of permanent income distributed evenly over age groups. The results are in Table 4. On average, 28 percent of the gift is consumed within a quarter, driven by a strong consumption response by the impatient population A. In Figure 4, the mass of this population is generally along a high marginal propensity to consume part of the consumption function. Our results thus provide an alternate but complementary explanation to Kaplan and Violante (2014) for why the consumption response to rebates is so large.

Both the reduced-form estimates from the credit bureau data and the structural estimates suggest that changes in consumer credit produce large consumption responses. An alternate way to increase liquidity is to increase credit rather than income. When we increase the credit limits of the population by 5 percent in Table 4, we get consumption effects that are almost as large as direct cash infusions, again driven mostly by our impatient population. While the structural model allows us to increase credit in a way that is uncorrelated with anything else, our reduced-form estimates from the credit bureau data give nearly the same estimates in response to an increase in credit that

reduces utilization (see Table 1).

5.3 Welfare analysis

We next examine the welfare consequences of changes in the economic environment. Because our population has heterogeneous preferences, how to conduct welfare analysis is not obvious. We avoid the normative issue of whose preferences to use by asking how changes affect the welfare of each population according to its own preferences. For each possible change, we recalculate the optimal responses using the estimated parameters in column 1 of Table 3 with the change in environment. Then we ask how much extra consumption we would have to give individuals in each period to make them indifferent at the start of life between the new environment and our standard one.

Table 5 shows the results of the consumption-equivalent welfare comparison of changing the interest rate on debt and the credit limit. Reducing the nominal annual interest rate from 14.11 percent to 8 percent has the utility equivalent of increasing consumption in every period by 1.5 percent for population A, which is generally borrowing, but by only 0.04 percent for population B. Similarly, doubling the interest rate to 28 percent reduces utility by the consumption equivalent of 0.86 percent for A and by 0.13 for B. Most of the welfare loss from an increase in interest comes from the first increase; an increase in interest to 60 percent is only slightly worse than 28 percent.

Allowing credit limits to go immediately to their life-cycle maximum multiple of income so that agents can borrow approximately one year's worth of income even at the beginning of life increases welfare by the consumption equivalent of about 1 percent for both populations. Credit limits have strong welfare consequences, particularly for the young people who are initially quite credit constrained.

6 Conclusion

This paper uses the consumer's decision about how to use credit cards to provide a window into more general savings and consumption decisions. We show that credit changes are very large over

the business cycle, the life cycle, and for individuals. Changes in credit are therefore some of the largest changes in liquidity faced by households. On average, people react quickly to these credit changes, and so credit utilization is stable over the business cycle, life cycle, and for individuals.

We take the insight this tight link between credit and debt gives and estimate a model of life-cycle consumption, debt, and payments using credit cards. The model has a number of notable successes. It captures the hump shape of debt and consumption. It predicts the slow decline in utilization over the life cycle and the steady increase in wealth. It also predicts smooth utilization over the business cycle, and it closely matches the reduced-form relationship at the individual level between credit and debt that we estimate from the credit bureau data. The model also reveals an important puzzle: It predicts a relatively constant fraction borrowing over the life cycle, even as the amount of debt changes, while surveys suggest that the share should fall over the life cycle. It seems that life-cycle concerns are not sufficient to cause those willing to borrow in early life to stop borrowing late in life. Perhaps additional population level heterogeneity is needed, so that life-cycle concerns are enough to shift a portion of the population out of borrowing. Alternatively—and intuitively to anyone who is no longer a teenager—it seems likely that people learn and change their preferences as they age.

Many of our results come directly from the insight that not everyone who has a credit card uses it to borrow, while some people are willing to borrow at a high rate of interest. Borrowing implies the consumer places substantial weight on consumption today versus tomorrow because of shocks or impatience. Other people have a credit card and use it only to make payments, suggesting they place more equal weight on today and the future. This heterogeneity of use suggests that preference heterogeneity must be an important part of understanding consumption decisions, and that a large fraction of the population must have a relatively high marginal propensity to consume. The preference heterogeneity is key to the estimated model's ability to match the impact of a cash infusion (Kaplan and Violante 2014, Parker et al. 2013).

So why is credit utilization stable? Because all the uses for credit cards push toward stability, although in different ways. Payments use of credit cards is proportional to consumption and so

moves in the same way it does. Changes in permanent income that increase credit limits also increase consumption and thus payments use, keeping utilization stable for convenience users. When credit is useful as a buffer against shocks, an increase in credit effectively makes people more wealthy, allowing them to spend more in the short run (Fulford 2013) and increasing their debts at the same time. Finally, because credit limits increase faster than income early in life, consumers using credit cards to smooth over the life cycle are particularly constrained early on, and so they increase debt at nearly the same pace as their limits increase.

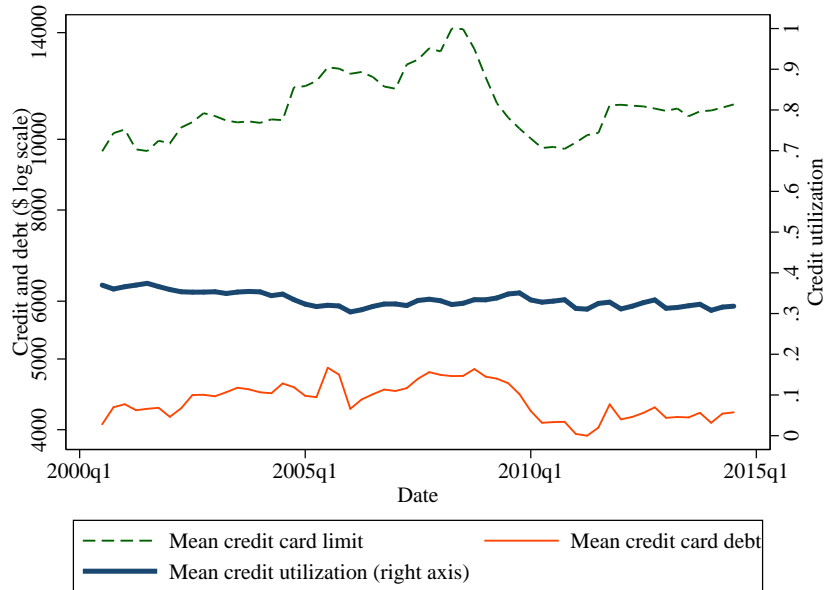
References

- 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." Working Paper 2617116, SSRN.
- Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics* 109(3): 659–84.
- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. 2001. "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *The Journal of Economic Perspectives* 15(3): 47–68.
- Athreya, Kartik B. 2008. "Default, Insurance, and Debt over the Life-Cycle." *Journal of Monetary Economics* 55(4): 752–774.
- Attanasio, Orazio P., James Banks, Costas Meghir, and Guglielmo Weber. 1999. "Humps and Bumps in Lifetime Consumption." *Journal of Business & Economic Statistics* 17(1): 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.
- Aydin, Deniz. 2015. "The Marginal Propensity to Consume out of Liquidity: Evidence from a Randomized Controlled Trial." 2015 Meeting Papers 270, SED.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98(5): 1887–1921.
- Brevoort, Kenneth P., Philipp Grimm, and Michelle Kambara. 2015. "Data Point: Credit Invisibles." Data Point 2015, Consumer Financial Protection Bureau Office of Research.
- Cagetti, Marco. 2003. "Wealth Accumulation over the Life Cycle and Precautionary Savings." *Journal of Business and Economic Statistics* 21(3): 339–353.
- Campbell, John Y., and N. Gregory Mankiw. 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." *NBER Macroeconomics Annual* 4: 185–216.
- Campbell, John Y., and N. Gregory Mankiw. 1990. "Permanent Income, Current Income, and Consumption." *Journal of Business & Economic Statistics* 8(3): 265–79.
- Carroll, Christopher D. 1997. "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112(1): 1–55.
- Carroll, Christopher D. 2006. "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems." *Economics Letters* 91(3): 312–320.
- Carroll, Christopher D. 2012. "Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems." Available: <http://www.econ2.jhu.edu/people/ccarroll/SolvingMicroDSOPs/>, accessed 3 June 2015.

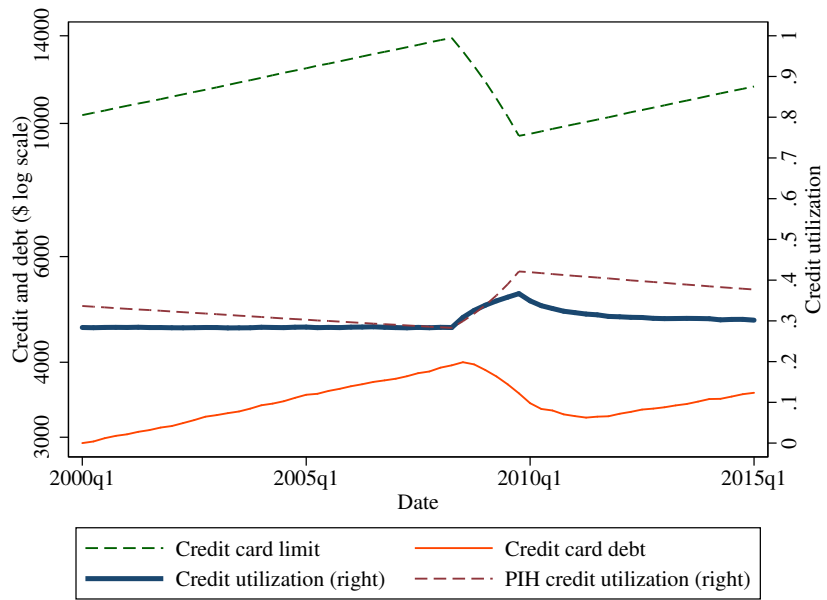
- Carroll, Christopher D., and Miles S. Kimball. 2001. "Liquidity Constraints and Precautionary Saving." Working Paper W8496, NBER.
- Carroll, Christopher D., and Andrew A. Samwick. 1997. "The Nature of Precautionary Wealth." *Journal of Monetary Economics* 40(1): 41–71.
- Carroll, Christopher, Jiri Slacalek, and Kiichi Tokuoka. 2017. "The Distribution of Wealth and the Marginal Propensity to Consume." *Quantitative Economics*: Forthcoming.
- Cocco, João F., Francisco J. Gomes, and Pascal J. Maenhout. 2005. "Consumption and Portfolio Choice over the Life Cycle." *Review of Financial Studies* 18(2): 491–533.
- De Nardi, Mariacristina, Eric French, and John B. Jones. 2010. "Why Do the Elderly Save? The Role of Medical Expenses." *Journal of Political Economy* 118(1): 39–75.
- Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica* 59(5): 1221–1248.
- Deaton, Angus. 1997. *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. Baltimore: Johns Hopkins University Press.
- Edelberg, Wendy. 2006. "Risk-Based Pricing of Interest Rates for Consumer Loans." *Journal of Monetary Economics* 53(8): 2283–2298.
- Fulford, Scott. 2013. "The Effects of Financial Development in the Short and Long Run: Theory and Evidence from India." *Journal of Development Economics* 104: 56–72.
- Fulford, Scott L. 2015. "How Important is Variability in Consumer Credit Limits?" *Journal of Monetary Economics* 72: 42–63.
- Fulford, Scott L., and Scott Schuh. 2015. "Consumer Revolving Credit and Debt over Life Cycle and Business Cycle." Research Department Working Papers 15–17, Federal Reserve Bank of Boston.
- 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.
- Guerrieri, Veronica, and Guido Lorenzoni. 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *The Quarterly Journal of Economics* 132(3): 1427–1467.
- Iacoviello, Matteo, and Marina Pavan. 2013. "Housing and Debt over the Life Cycle and over the Business Cycle." *Journal of Monetary Economics* 60(2): 221–238.
- Kaboski, Joseph P., and Robert M. Townsend. 2011. "A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative." *Econometrica* 79: 1357–1406.

- Kaplan, Greg, and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82(4): 1199–1239.
- Kaplan, Greg, Giovanni L. Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity*: 77–154.
- Krusell, Per, and Anthony A. Smith, Jr. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *The Journal of Political Economy* 106(5): 867–896.
- Laibson, David, Andrea Repetto, and Jeremy Tobacman. 2007. "Estimating Discount Functions with Consumption Choices over the Lifecycle." Working Paper 13314, NBER.
- 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.
- Lopes, Paula. 2008. "Credit Card Debt and Default over the Life Cycle." *Journal of Money, Credit and Banking* 40(4): 769–790.
- McFadden, Daniel. 1989. "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration." *Econometrica* 57(5): 995–1026.
- Parker, Jonathan A. 2017. "Why Don't Households Smooth Consumption? Evidence from a \$25 Million Experiment." *American Economic Journal: Macroeconomics* 9(4): 153–83.
- Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103(6): 2530–53.
- Sabelhaus, John, and Jae Song. 2010. "The Great Moderation in Micro Labor Earnings." *Journal of Monetary Economics* 57(4): 391–403.
- Schuh, Scott. 2017. "Measuring Consumer Expenditures with Payment Diaries." *Economic Inquiry*: Forthcoming.
- Schuh, Scott, and Joanna Stavins. 2017. "The 2012 Diary of Consumer Payment Choice." Research Data Report, Federal Reserve Bank of Boston.
- Stepner, Michael. 2013. "BINSCATTER: Stata Module to Generate Binned Scatterplots." Statistical Software Components, Boston College Department of Economics, October.
- Telyukova, Irina A. 2013. "Household Need for Liquidity and the Credit Card Debt Puzzle." *Review of Economic Studies* 80(3): 1148–1177.
- Telyukova, Irina A., and Randall Wright. 2008. "A Model of Money and Credit, with Application to the Credit Card Debt Puzzle." *Review of Economic Studies* 75: 629–647.
- Zinman, Jonathan. 2009. "Where is the Missing Credit Card Debt? Clues and Implications." *Review of Income and Wealth* 55(2): 249–265.

Figure 1: Credit card limits, debt, and utilization: 2000–2015
Observed limits, debts, and utilization from credit bureau

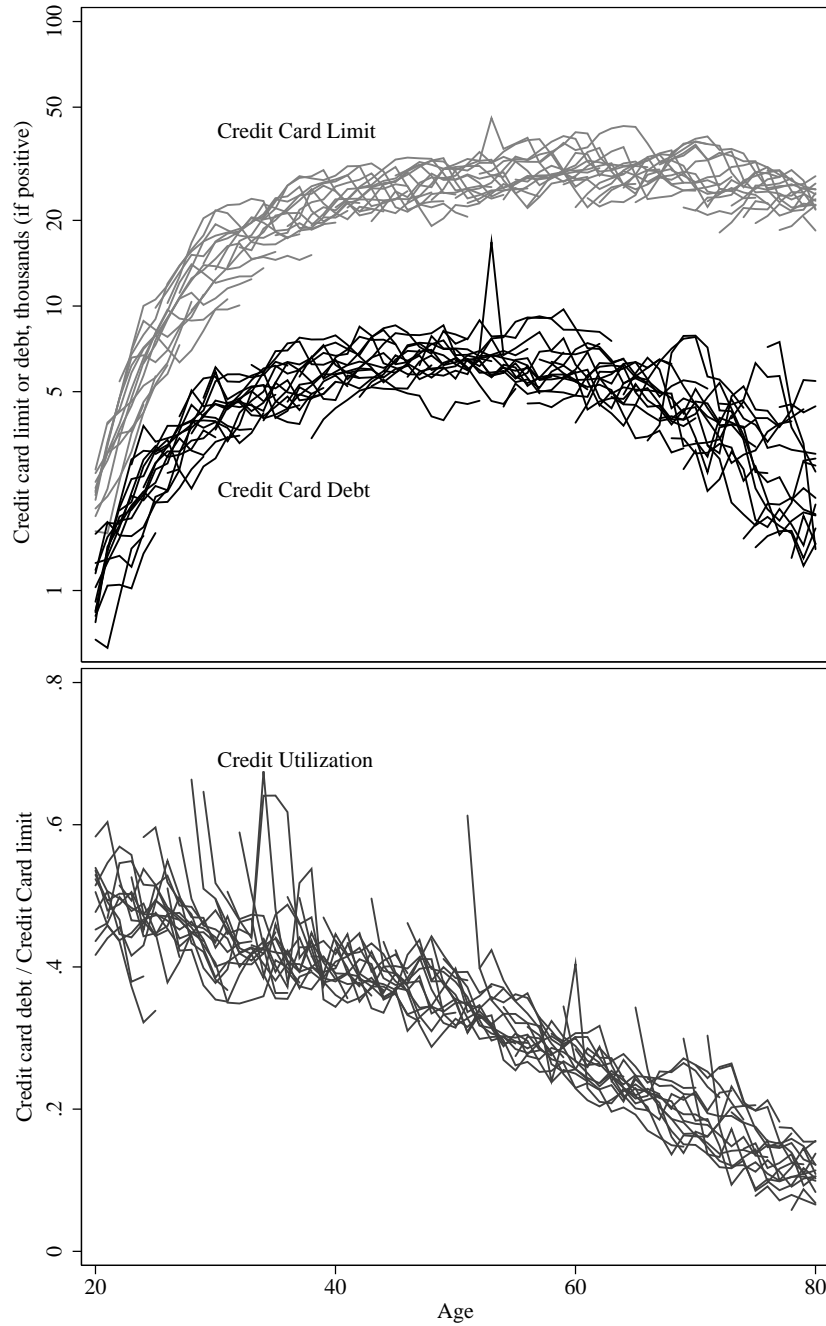


Model prediction given fall in credit limits



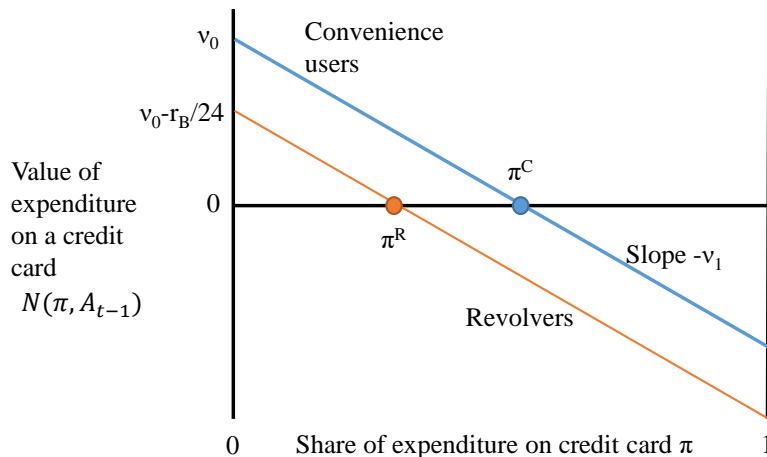
Notes: The top panel shows observed limits, debts, and utilization from credit bureau data (see Section 2 for details). The bottom panel shows model predictions given an unexpected fall in credit (see section 5 for details). For both panels, the left axis shows the average credit card limits (top line) and debt (bottom line). Note the log scale. The right axis shows mean credit utilization (middle line) defined as the credit card debt/credit card limit if the limit is greater than zero. Source: Authors’ calculations from Equifax/NY Fed CCP.

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



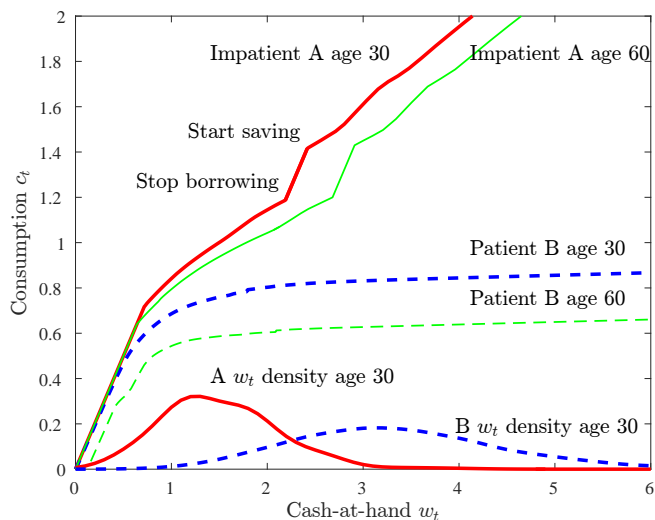
Notes: Each line represents the average credit card limit (conditional on being positive, log scale), debt (conditional on being positive, log scale), and utilization (conditional on having a limit, bottom panel) of one birth year cohort from 1999–2014. Source: Author’s calculations from Equifax/NY Fed CCP.

Figure 3: Value or cost of expenditure using a credit card, relative to other means



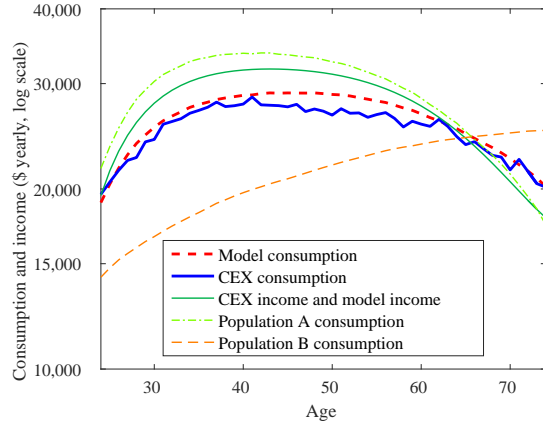
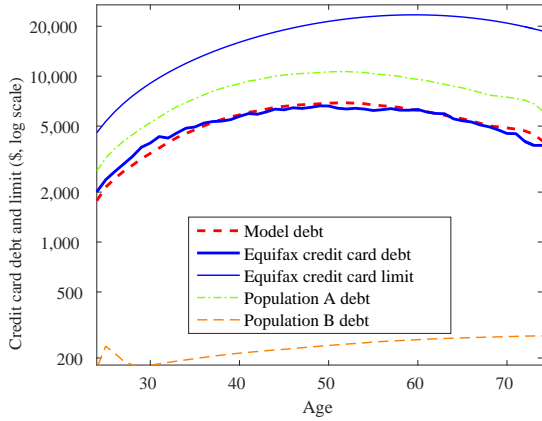
Notes: This figure shows the value or cost of expenditure on a credit card at each expenditure share π relative to cash. The top line is for convenience users who put an optimal share π^C of consumption on a credit card. The bottom line for revolvers is shifted down by the amount $-r_b/24$, because revolvers lose the float on payments made using credit cards and therefore put a smaller optimal share on their credit cards π^R .

Figure 4: Expenditure functions over the life cycle with borrowing

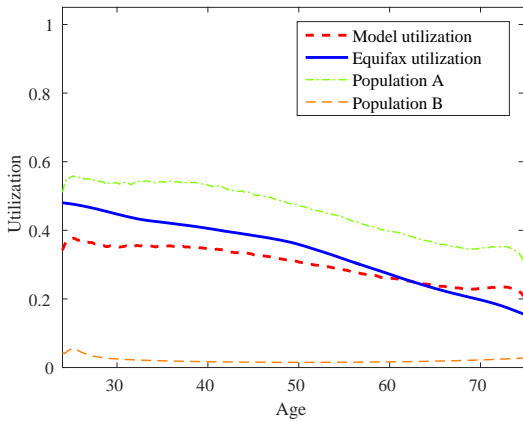


Notes: This figure uses the estimates in Table 3 column 1 at age 30 and age 60 to show the quarterly expenditure function for impatient (A) consumers and patient (B) consumers. Cash at hand w_t is the available liquidity (including credit) as a multiple of quarterly permanent income P_t . The densities for cash at hand are for age 30 and show where individuals are along their consumption functions. Because the rate of savings is lower than the rate of borrowing, the expenditure function has kink going from borrowing to saving nothing for the next period to actively saving.

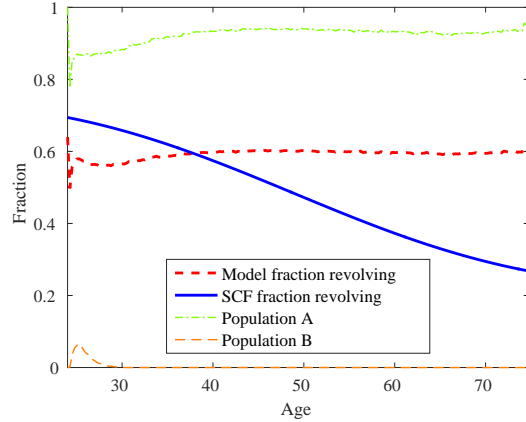
Figure 5: Consumption and debt over the life cycle: model estimates
 Estimation moments: Debt Estimation moments: Consumption



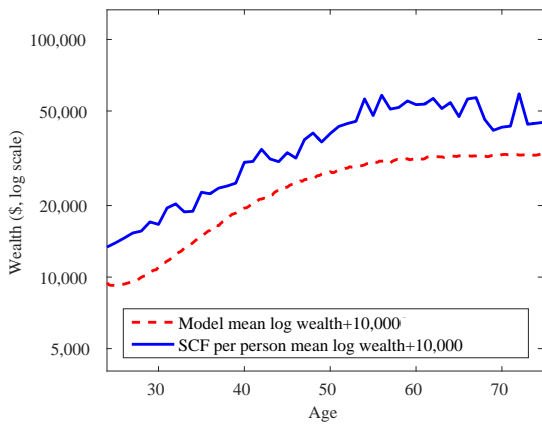
Estimation predictions: Utilization



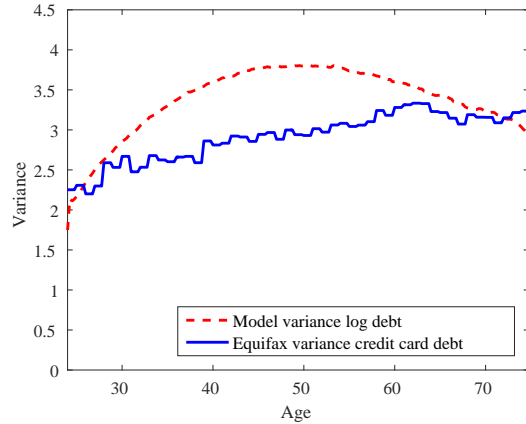
Estimation predictions: Fraction revolving



Estimation predictions: Wealth path



Estimation predictions: Variance of debt



Notes: Life-cycle paths from simulated population using the estimates in column 1 of Table 3.

Table 1: Credit utilization

	Equifax/NY Fed CCP			Model
	Credit utilization _t			
Credit utilization _{t-1}	0.874*** (0.000876)	0.868*** (0.000892)	0.647*** (0.00131)	0.723*** (0.000477)
Constant	0.0479*** (0.000461)			
Observations	347,642	347,642	347,642	2,168,011
R-squared	0.741	0.743	0.429	0.525
Fixed effects	No	No	Yes	Yes
Age and year effects	No	Yes	Yes	Yes
Number of accounts			10,451	46,607
Frac. Variance from FE			0.477	0.228

Notes: The sample includes zero credit utilization but excludes individual quarters where the utilization is undefined since the limit is zero and when utilization is greater than five (a very small fraction, see distributions of utilization in Fulford and Schuh (2015)). Source: Authors' calculations from Equifax/NY Fed CCP.

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

	Fraction on Credit card	Std. error	Std. dev.
All consumers	0.172	0.0082	0.310
All revolvers	0.156	0.0130	0.283
All convenience users	0.182	0.0105	0.324
Model Estimates			
Level ν_0	0.035	0.0216	
Slope ν_1	0.194	0.1259	
Implied value of credit card use (percent of consumption)			
Revolvers	0.235	0.1512	
Convenience users	0.319	0.0962	

Notes: Authors' calculations from the Federal Reserve Bank of Boston Diary of Consumer Payment Choice. The standard errors are calculated by bootstrapping.

Table 3: Model estimates

	Standard Weights	Optimal Weights	No end. payments	Low bequest
Population A				
CRRA γ^A	0.115 (0.034)	0.568 (0.523)	0.115 (0.034)	0.117 (0.038)
Discount β^A	0.890 (0.002)	0.801 (0.196)	0.890 (0.002)	0.891 (0.002)
Initial wealth λ_0^A	0.500 (0.046)	0.477 (0.206)	0.500 (0.045)	0.500 (0.046)
Late life inc. λ_1^A	0.599 (0.040)	0.496 0.31057	0.599 (0.040)	0.598 (0.039)
Population B				
CRRA γ^B	1.751 (1.060)	1.82511 0.67573	1.751 (1.059)	1.751 (0.923)
Discount β^B	0.964 (0.018)	0.932 (0.025)	0.964 (0.018)	0.965 (0.016)
Initial wealth λ_0^B	1.700 (1.171)	1.742 (0.247)	1.700 (1.170)	1.700 (1.103)
Late life inc. λ_1^B	0.051 (0.185)	0.442 (0.071)	0.051 (0.185)	0.052 (0.151)
Share A f^A	0.640 (0.005)	0.365 (0.003)	0.640 (0.005)	0.637 (0.009)
Inc. mult. A ζ^A	1.094 (0.099)	1.015 (0.343)	1.094 (0.098)	1.092 (0.101)
J-stat	528.79	1168.80	528.77	571.52
p-val	0	0	0	0
Weights	Standard	Optimal	Standard	Standard
Endogenous payments	Yes	Yes	No	Yes

Notes: Standard errors in parentheses. Optimal weights are the inverse of the variance of each individual moment. No endogenous payments removes the consumer's awareness that revolving affects the value of credit cards for payments. Low bequest reduces the bequest motive. For a description of each moment and the estimation method, see the beginning of section 4.

Table 4: Effects of temporary cash infusion or permanent credit increase

	Full pop.	Pop. A	Pop B.	Full pop.	Pop. A	Pop B.
	Δ Expenditure from previous quarter					
Transitory income increase	0.288*** (0.0153)	0.349*** (0.0220)	0.121*** (0.00680)			
Permanent credit limit increase				0.180*** (0.0152)	0.221*** (0.0215)	0.0524*** (0.00684)
Observations	533,288	341,544	191,744	533,288	341,544	191,744
R-squared	0.001	0.001	0.002	0.001	0.001	0.001
Age effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows the results of experiments using the estimates from column 1 in Table 3. We give a randomly selected portion of our simulated population a cash gift of 5 percent of permanent income or a 5 percent increase in individual credit limit. The regression is then $\Delta Cons_t = \alpha + f(age) + \beta Cash_t + \epsilon_t$ measuring how much of the increase in cash or credit limit is consumed within one quarter.

Table 5: Welfare comparisons: Equivalent gain or loss in consumption (percent)

	Population A	Population B
Reducing the interest rate to 8% annually	1.5017	0.0401
Increasing the interest rate to 28% annually	-0.8633	-0.0656
Increasing the interest rate to 60% annually	-1.0581	-0.1329
Constant credit limit at life-cycle maximum	0.928	0.815

Notes: Each cell shows the increase (or decrease) in consumption in every period—discounted back to the beginning of life—that makes the utility of that population the same, on average, as without the change. For example, if population A shows an increase of 1.1, then the change in environment has the same utility impact as a 1.1 increase in consumption in every period for every agent using the preferences of population A.

Credit Card Utilization and Consumption over the Life Cycle and Business Cycle

Scott L. Fulford and Scott Schuh

Appendix For Online Publication

A Changes in credit utilization: non-parametric evidence

Figure A-4 shows conditional mean scatter plots of credit utilization in one quarter against credit utilization in the next quarter, in the next year, and in two years. The top row shows the mean in the future, conditional only on having the utilization shown on the x-axis in that quarter. The bottom row instead takes the within transformation and allows for age and year effects. It therefore shows how far from the individual's average credit utilization she is in the next quarter, conditional on differing from her average utilization by the amount on the x-axis this quarter. In other words, if an individual is 10 percentage points above her typical utilization in one quarter, how far will she be on average in the next quarter, next year, and in two years? Each dot contains an equal portion of the sample. Figure A-4 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 trend to zero on average. Credit utilization this quarter is typically very close to credit utilization next quarter, because 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. On closer examination, 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 0.39.³⁶ 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.

On average, individuals do not trend to zero utilization or to using all 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 her credit in one quarter is using less than 90 percent of it in a year.

The second row of Figure A-4 allows individuals to return to their own mean and adds substantial nuance. Credit utilization is so persistent in the top row because 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 67 percent of a shock to utilization remains after one quarter, and 13 percent remains after two years.

Even if individuals return very rapidly to their own means, it is important to note that those means are not zero. Credit utilization is persistent in the top row of Figure A-4 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 debt by 33 percent of the increase in credit within one quarter, and by 87 percent of the increase in credit in two years.

³⁶Since the conditional expectation of utilization next quarter given this quarter is $u_{t+1} = 0.041 + 0.896u_t$, the steady-state utilization is $0.39 = 0.041 / (1 - 0.896)$.

B Additional Model Details

This section provides additional discussion of the model’s formulation and numerical solution and the construction of the variance covariance matrix of moments from credit bureau data.

B.1 Bequests

In our base estimation, we give consumers who die with positive resources some extra utility from bequests. Bequests end up being important only because they help explain our impatient group’s decrease in debt after 60, although the exact parameters are not well identified. Without any bequest motive, as the likelihood of dying increases, it becomes optimal to increase debt for impatient nearly risk-neutral consumers because they are effectively becoming more impatient. A bequest motive keeps the increasing probability of death from effectively translating into increased impatience, and so it allows debt to decline with income after age 60. This appendix outlines one flexible approach to including bequests.

When introducing bequests in a model with both debt and savings, it is difficult to value what happens when people die in debt. Unsecured consumer credit is taken out of any estate passed on to heirs, but is not directly passed on. Simply including bequests in the sub-utility function will therefore produce negative infinite value from leaving no bequest, which consumers will counterfactually act to avoid by never being in debt. Instead, we model the bequest motive as the consumer’s considering the marginal utility a bequest to her heirs will bring them on top of the heirs’ own incomes and any non-liquid bequest she may consider leaving. We model these non-liquid bequests and heirs’ income as a multiple ζ of permanent income on death, adding the annuity value of the assets left at death and taking the present value using the consumer’s preferences:

$$S(A_t) = \left(\sum_{s=0}^{\tilde{T}} \beta^s \frac{(\zeta P_t + r_B A_t)^{1-\gamma}}{1-\gamma} \right).$$

The parameter ζ determines the marginal utility of bequests and can be thought of as how much more or less income children have compared to their parents, and \tilde{T} is the same finite life as the

parents'. The reason to use finite rather than infinite heirs' lifetimes is to allow for possibly very patient parents with β close to or greater than one.

B.2 Recursive formulation and normalization

Rewritten in recursive form and normalized by permanent income P_t , the problem is equivalent to:

$$v_t(w_t, b_t, a_{t-1}) = \max_{x_t, \pi_t} \{u(\nu_t x_t) + E_t[\beta_{t+1}(G_{t+1}N_{t+1})^{1-\gamma} v_{t+1}(w_{t+1}, b_{t+1})]\} \text{ subject to}$$

$$x_t \leq w_t$$

$$w_{t+1} = R_{t+1}(I_t^R) a_t + U_{t+1} + b_{t+1}$$

$$a_t = w_t - b_t - x_t$$

$$\nu_t = \nu(\pi_t; a_{t-1}),$$

where $R_{t+1}(I_t^R) = R/(G_{t+1}N_{t+1})$ if $a_t \geq 0$ and $R_{t+1}(I_t^R) = R_B/(G_{t+1}N_{t+1})$ if $a_t < 0$. The expectation at t includes the possibility of death before T and the certainty of death at \tilde{T} , leaving a bequest worth $\beta_{t+1}s(a_t)$, where $s(\cdot)$ is the bequest function normalized by P_t . Note that if b_t is not stochastic and instead follows the average path of the credit limit to permanent income ratio, then b_t , like G_{t+1} , is not a part of the state space that differs for individuals and the decision simplifies slightly to $v_t(w_t, a_{t-1})$. Of course, credit limits and income growth still matter, but they do not vary individually and so show up in each consumer's expenditure function $x_t(w_t, a_{t-1})$. Because of the structure of the payment problem, $\nu(\pi_t; a_{t-1})$ takes on only two values for an optimizing consumer, ν^R for revolvers and ν^C for convenience users. The value expenditure function depends only on whether $a_{t-1} \geq 1$, substantially reducing the dimensionality of the problem.

B.3 Numerical solution

With the problem written recursively, we proceed through backward recursion to find a numerical approximation of the consumer's problem. Let I_t^R indicate revolving status: It is 1 if $a_{t-1} < 0$ and 0 else. For a given set of parameters, once $v_{T+1}(a_T, 0, I^R)$ is given, it is possible to find

an approximation of $v_T(w, b, I^R)$ and use the approximation of $v_T(w, b, I^R)$ to find $v_{T-1}(w, b, I^R)$. Note that in each case we find a separate function for revolvers and convenience users. The solution to each period's value function is an expenditure function $x_t(w, b, I^R)$. We follow several standard steps (see Carroll (2012) for a more in-depth discussion of many of these approaches). First, we discretize the lognormal shocks using a Gauss-Hermite quadrature, which turns the integration in the expectation function into a summation over discrete states. Because the income process is surely not exactly lognormal, there is no gain or loss in accuracy from doing so; we are simply replacing one approximation of shocks with another.

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

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

where, despite its subscript, $\nu_{t+t} = \nu(I_{t+1}^R)$ is determined entirely by the choice of whether to leave positive or negative assets for the next period. Given the next-period expenditure function, it is straightforward to find the optimal expenditure that leaves end-of-period assets a_t as:

$$x_t^a(a, b, I^R) = (1/\nu_t) \left(E_t \left[\beta \hat{\beta}_{t+1} R_{t+1}(a) (G_{t+1} N_{t+1})^{1-\gamma} (\nu_{t+1} x(R_{t+1}(a)a + U_{t+1} + b_{t+1}, b_{t+1}))^{-\gamma} \right] \right)^{-1/\gamma}. \quad (5)$$

For a vector of end-of-period assets \vec{a} , it is nearly costless to find the optimal consumption at a vector of endogenous points for cash at hand where $\vec{w} = \vec{a} + b + x_t^a(\vec{a}, b, I^R)$ is the amount at

which consuming $x_t^a(a, b, I^R)$ and leaving a for next period is optimal. We linearly interpolate between these points to find an approximation of the expenditure function. Note that the expenditure function is a function of whether the consumer is revolving by having negative assets last period, in addition to the current state of cash at hand and the credit limit. While revolving status is not a continuous state, the addition of another state variable complicates the solution because we must find the optimal expenditure for convenience users and revolvers, who find consuming less valuable because they pay for it in a slightly less convenient way. For the most part, someone who is not revolving this period will not be revolving next period, and so $v_t = v_{t+1}$, and the payment choice does not affect the expenditure decision directly. It does, however, make revolving somewhat more costly.

Because the consumer's problem includes an externally imposed credit limit as well as interest rates that differ depending on whether assets are positive or negative, there are several additional complications. The first is that the standard Euler equation does not hold when the consumer is against her credit limit, and so she spends all available resources because she would like to spend more today but cannot (Deaton 1991). This problem is relatively easy to deal with, however, by including the inflection point that is the last point at which the Euler equation holds. At this point the assets left for the next period are $-b_t$. For any cash at hand less than $w^* = x_t^a(-b, b, I^R)$, the consumer expends all cash at hand, so $x_t(w, b, I^R) = w$ if $w \leq w^*$. The second problem is that the interest-rate differential introduces a step in the consumption function, because there are two solutions to equation (5) for $a = 0$. One, the limit with assets approaching zero from below, uses the borrowing rate R^B , and the other uses the saving rate R . The economic intuition is that leaving zero assets for the next period is optimal at a high borrowing rate well before it is optimal at a low savings rate. For cash at hand between these two points, the consumer has a marginal propensity to consume of one since the return on savings is not high enough to induce her to save, but the cost of borrowing is sufficient to keep her from borrowing, and so additional resources go straight to consumption. To deal with this issue, the endogenous gridpoints include two points where $a = 0$: The first, $x_t^B = x_t^a(0, b, I^R; R^B)$, is the solution to equation (5) when $a = 0$ using

R^B and $w_t^B = 0 + b + x_t^B$; and the second is $x_t^F = x_t^a(0, b, I^R; R)$ and w_t^F . Between the points (w_t^B, x_t^B) and (w_t^F, x_t^F) , the consumer has a marginal propensity to consume of one.

Figure 4 in the main paper illustrates these kink points, labeling w_t^B as the point where consumers stop borrowing and w_t^F as the point where they start saving. Several points are worth discussing. First, the consumption function generally falls with age. This occurs as the consumer plans for retirement, when having accumulated a large amount of savings is valuable. Second, for low cash at hand below w^* , the marginal propensity to consume is one. Between w^* and w^B , the consumer is leaving debt for next period and so is paying a high interest rate R^B . Between w^B and w^F , the consumer does not want to borrow, but the return on savings is not high enough, so she leaves zero assets and has a marginal propensity to consume of one. This kink in the consumption function implies that there can be a positive fraction of consumers who hold exactly zero assets. The distance between w^B and w^F depends on the interest-rate differential, with a wider differential implying a larger distance.

B.4 Construction of variance-covariance matrix

The variance-covariance matrix for the combined moments is simply block diagonal, because they are sampled independently from a large population. The Consumer Expenditure Survey block of the variance-covariance matrix is simply diagonal, since the survey does not repeat the same households over multiple years.

Because we observe individuals over time in the Equifax/NY Fed CCP data, the credit bureau portion of the variance-covariance matrix has off-diagonal elements. We populate this matrix by estimating the co-variance of debts in the population at various lags. Since the data is quarterly, the full matrix is 204x204, and because our data agreement limited what we could make public (and the size of the data limited what we could practically calculate), in practice we do not estimate each element separately. Instead we estimate the covariance at each age from 24–74 at quarterly lags 1, 2, 3, 4, 8, 12, 16, 32, and 48 and assume that the co-variance at each age changes smoothly in between them. Since the data cover only 16 years, all covariances beyond 64 quarterly lags are

zero. The combination of estimating covariances from a sample, constructing intermediate covariances, and numerical precision leaves the resulting covariance matrix with minimum eigenvalues that are slightly negative, and so the matrix is not positive definite. We make a “ridge adjustment” by adding a small amount to the diagonal until all eigenvalues are greater than or equal to zero. This adjustment effectively increases the variance of our main moments, and so it is generally conservative although not entirely innocuous, since it changes the variances but not the covariances.

B.5 Characterization and discussion of local minima

While Table 3 presents the best estimate after starting the estimation at a grid of points, given the over-identification, it is useful to briefly characterize other possible minima. Figure A-8 shows the starting points and the distribution of optimal estimates for each parameter. There is a high correlation between many parameters, and many of the estimates come from local maxima that are much worse than our best estimate, and so these distributions do not represent a menu of equivalent alternative estimates. Nonetheless, it is clear that the starting point for θ_0 does not determine θ^* , and so the estimation is identified in a narrow numerical sense.

Among the estimates with objective functions close to that produced by θ^* , the only important variation is that the estimates suggest a somewhat more patient population B ($\beta^B \approx 0.962$) with a lower risk aversion ($\gamma^B \approx 1.6$) is nearly equivalent to our main estimates. Because the coefficient of relative risk aversion is the inverse of the intertemporal elasticity of substitution, γ and β have similar roles in utility. Loosely, β governs how much the consumer cares about expected marginal utility in the future, while γ shifts expected marginal utility by making bad states better or worse. Beyond the best estimates, there are larger differences in the parameters at the local maximum, but the fit was always substantially worse. There are some obvious tradeoffs between certain parameters. The impatient population could either be very impatient ($\beta^A \approx 0.83$) and have some risk aversion ($\gamma^A \approx 1$) or be less impatient ($\beta^A \approx 0.88$) and have almost no risk aversion ($\gamma^A \approx .5$ or less). The more patient population B faces a similar tradeoff between β^B and γ^B .

The proportion of population A (f^A) changes across the local minima in more substantial ways.

While allowing for preference heterogeneity is a step forward in our work, imposing only two populations is still a simplification, albeit one that is useful both expositionally and computationally. At each minimum, f^A is tightly estimated. However, the range of f^A in the local minima suggests that different weights would produce different estimates of f^A . The overall conclusion holds for all estimates: In order to match the amount of debt we see in the data, half or more of the population must be fairly impatient and close to risk neutral.

C Identification of the payments model

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

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

the average expenditures by revolvers on a credit card, and similarly π^C , the average for convenience users. We denote our estimates of the standard errors of these means as σ_R and σ_C . Then the intercept for the average convenience user is just $\pi^C = \nu_0/\nu_1$, and for a revolver it is $\pi^R = (\nu_0 - r^B/24)/\nu_1$, where r^B is the APR interest charged on payments, which have an average daily balance of half of the month's consumption. Solving for ν_0 and ν_1 gives:

$$\begin{aligned} \nu_1 &= \frac{r^B/24}{\pi^C - \pi^R} \\ \nu_0 &= \pi^C \nu_1 = \frac{(r^B/24)\pi^C}{\pi^C - \pi^R}. \end{aligned}$$

The presence of a difference of two random variables whose supports may overlap in the denominator of the transformed variables makes calculating their variances potentially tricky. Since $\pi^C - \pi^R$ may be close to zero, then ν_1 and ν_0 may be very large, which is a different way of saying that the model is not identified if there is not a difference in the average behavior of convenience users and revolvers. We calculate the standard errors of the transformed variables using the delta

method, which avoids this issue by examining only small changes around the optimum, and so it does not consider the highly nonlinear increase around $\pi^C - \pi^R = 0$. For small changes ϵ^C and ϵ^R around π^C and π^R :

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

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

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

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

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

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

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

and for revolvers it is:

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

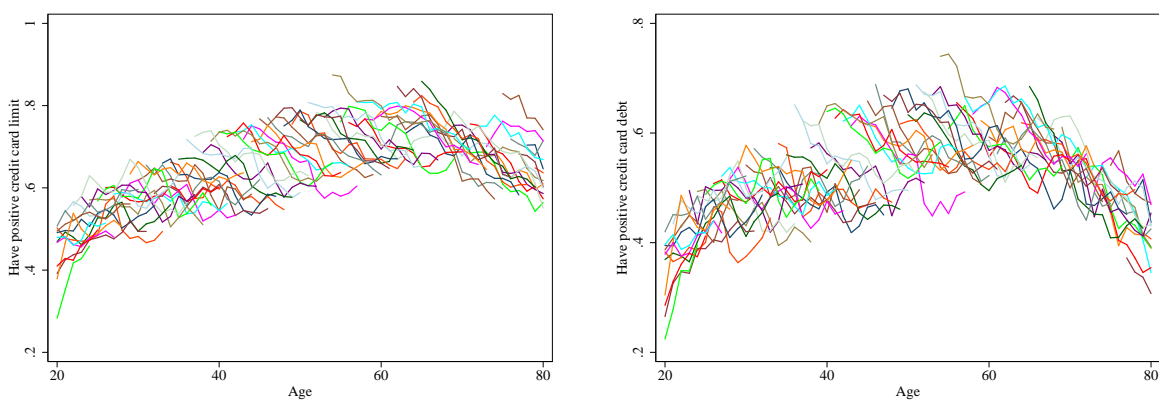
Taking an expansion around π^C and π^R yields:

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

and

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

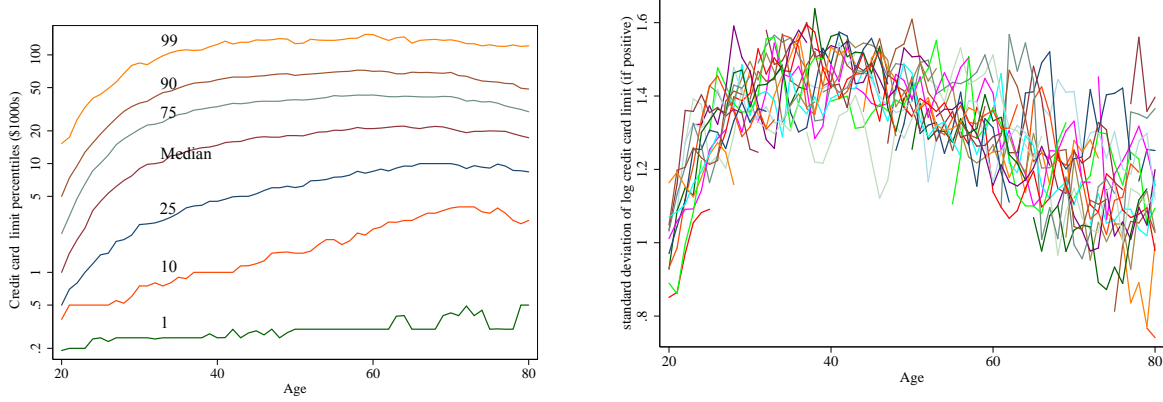
Figure A-1: Fraction with positive credit card limit and debt by cohort and age from CCP
(A) Fraction with positive limit (B) Fraction with positive debt



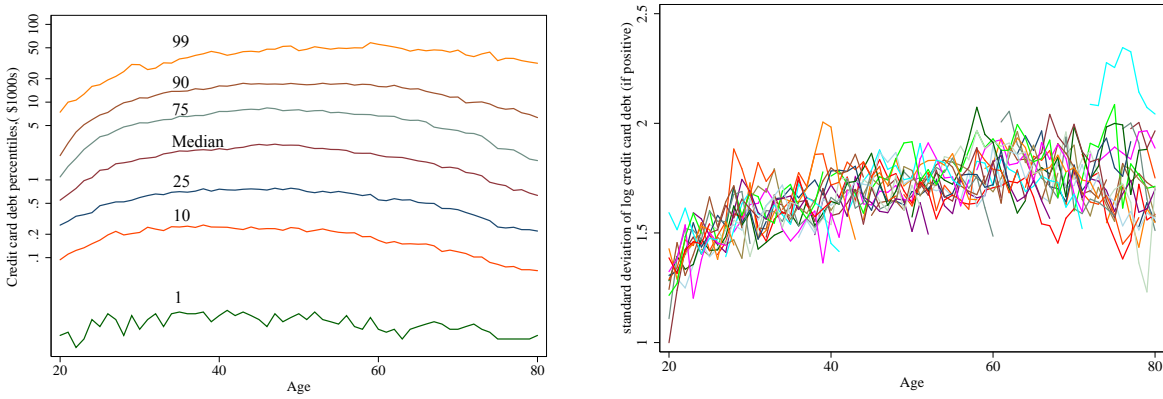
Notes: Each line represents the fraction with positive credit card limits or debt of one birth cohort, 1999–2014. Source: Authors' calculations from Equifax/NY Fed CCP.

Figure A-2: Credit card limit, debt, and credit utilization distributions and standard deviations by age

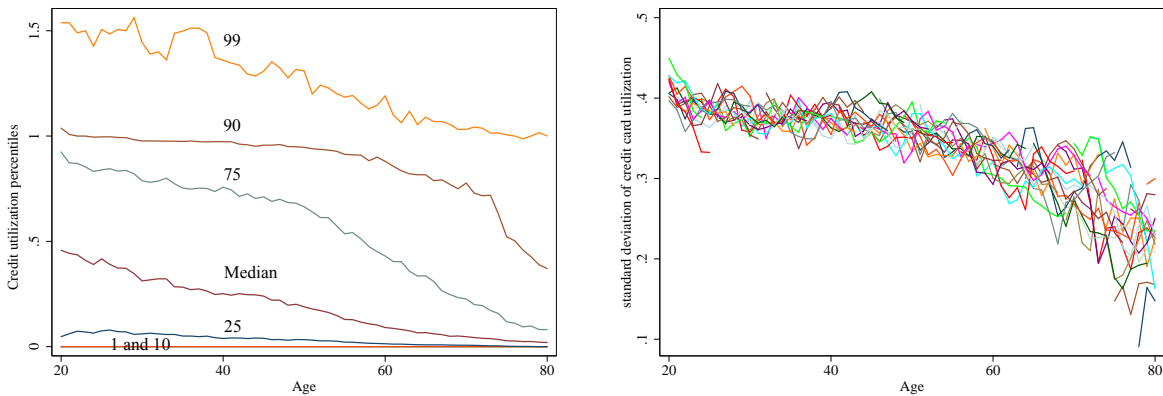
(A) Credit card limits



(B) Credit card debt

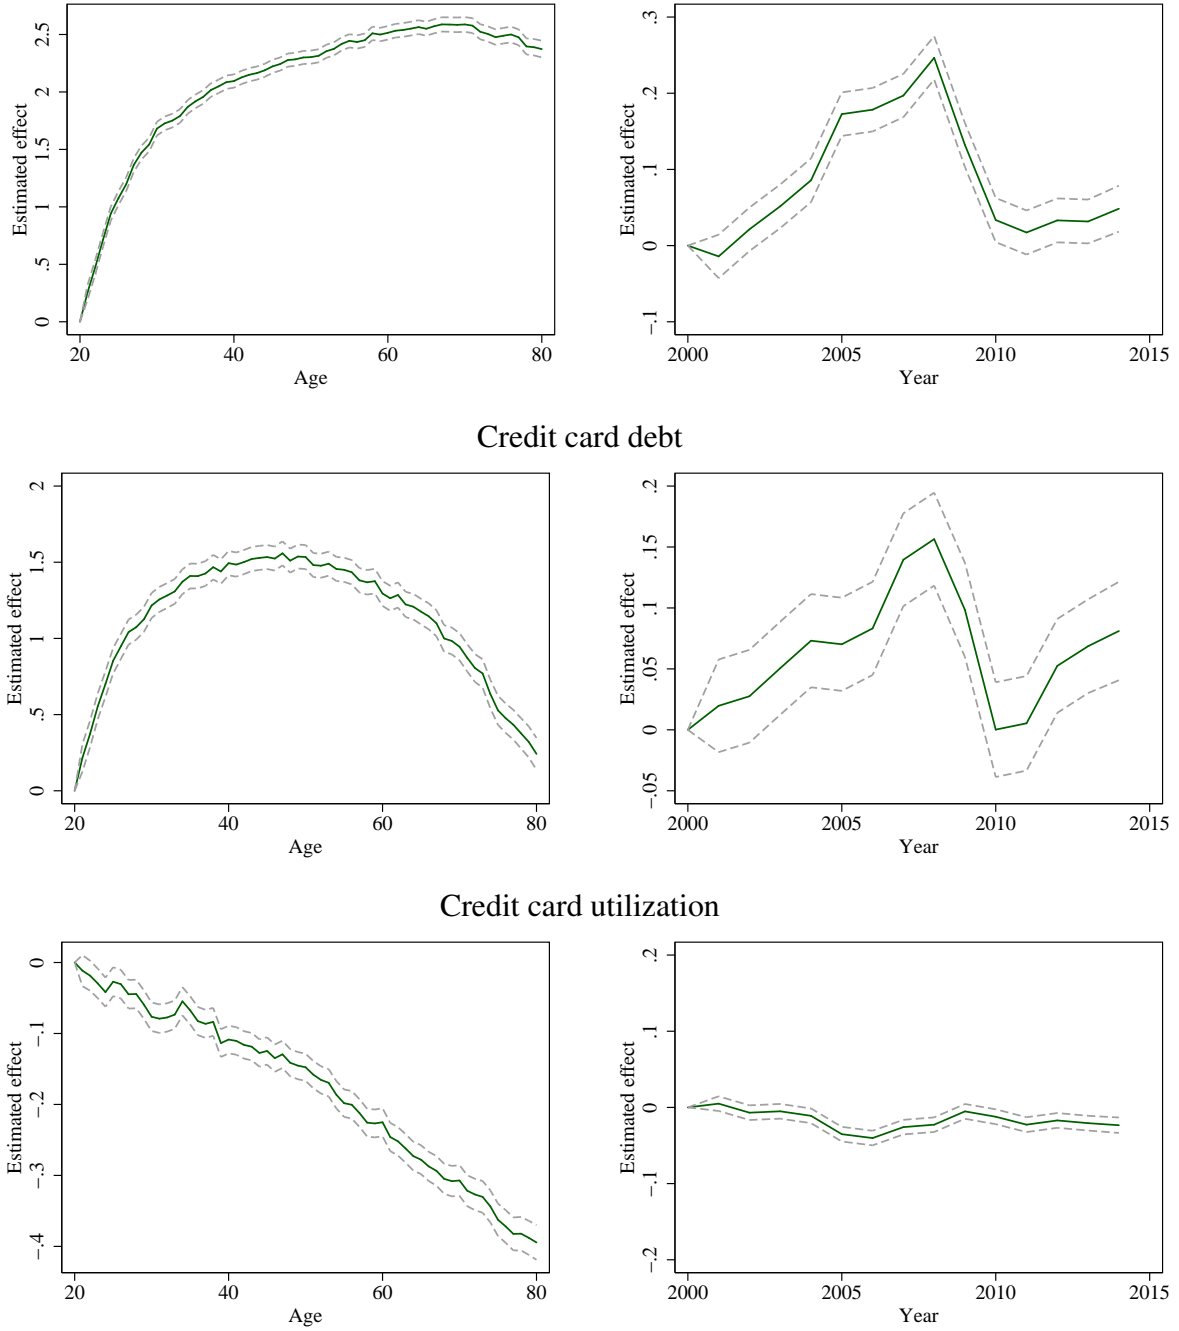


(C) Credit utilization



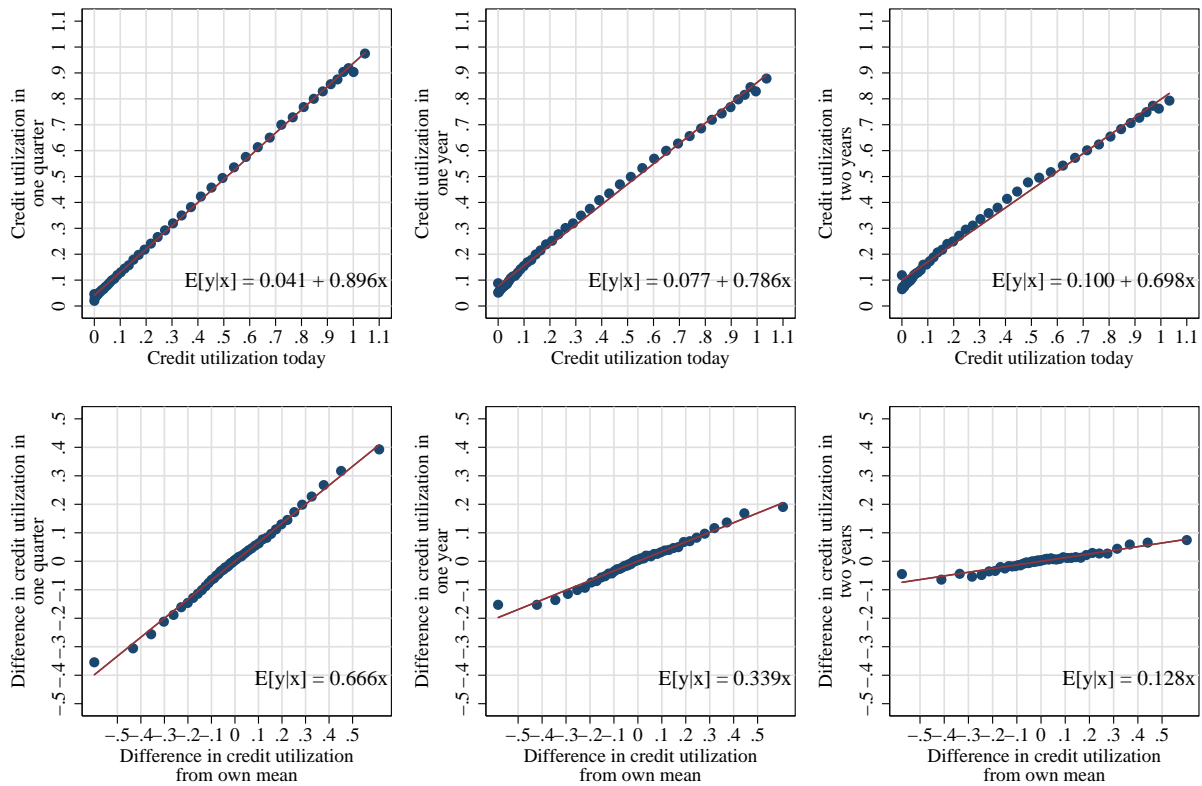
Notes: Each line is the percentile of credit limit at that age, conditional on having a positive credit limit on a log scale. For example, the 90th percentile line shows that 10 percent of the population (with a positive credit limit) has a limit larger than that line. Source: Authors' calculations from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

Figure A-3: Credit card limits, debt, and utilization: age and year effects



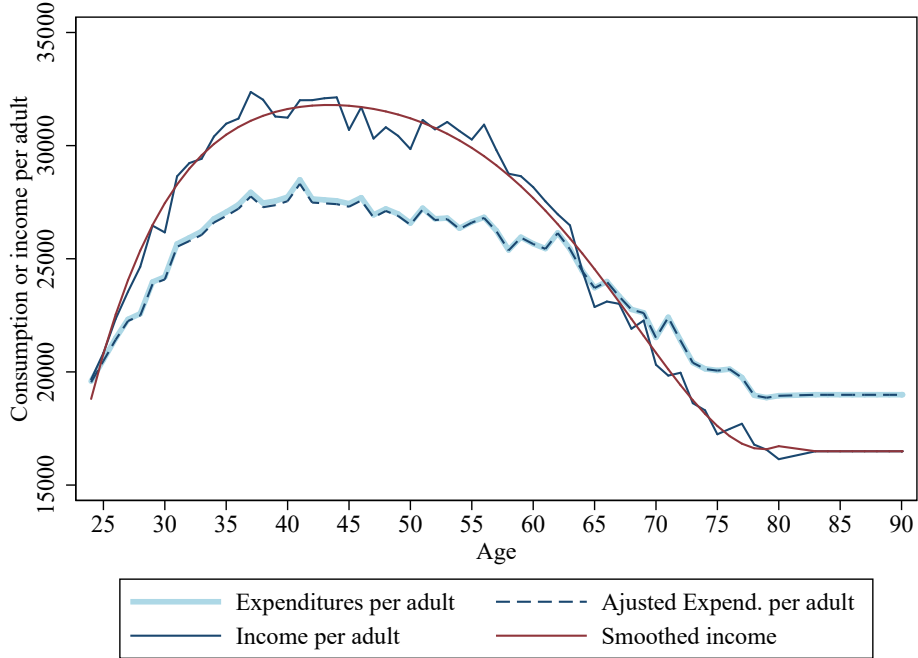
Notes: Each line shows the estimated age or year effects from equation (1). Note the different scales. Source: Authors' calculations from Equifax/NY Fed CCP.

Figure A-4: Changes in credit utilization in one quarter, one year, and two years



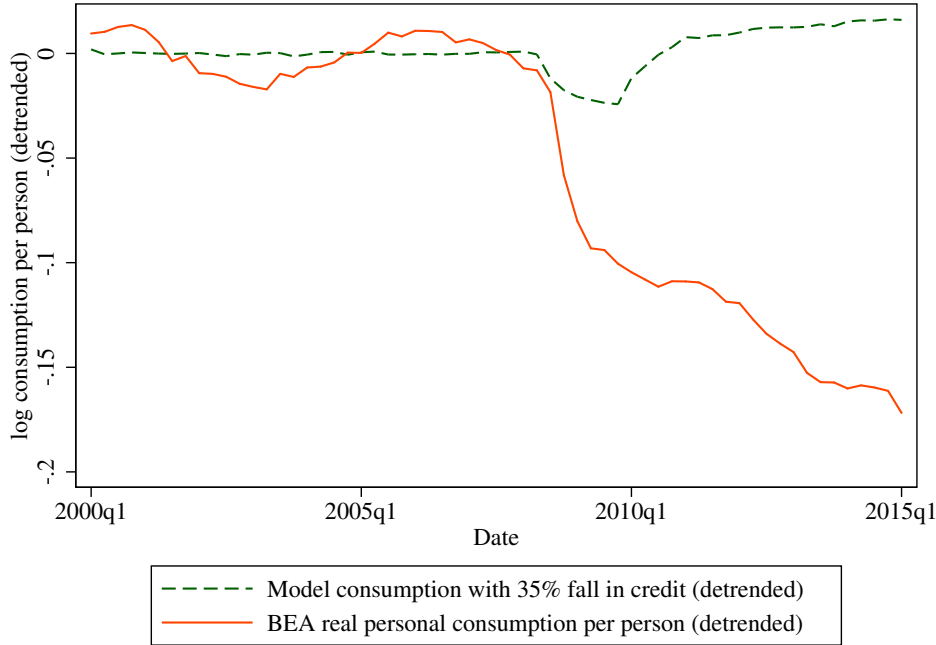
Notes: Each point in the top row shows the mean credit utilization in the future, conditional on being in the bin with a mean credit utilization on x-axis today. The bottom row shows the conditional relationship between deviations from the individual mean utilization over the entire sample, adjusting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

Figure A-5: Consumption and income over the life cycle from the Consumer Expenditure Survey



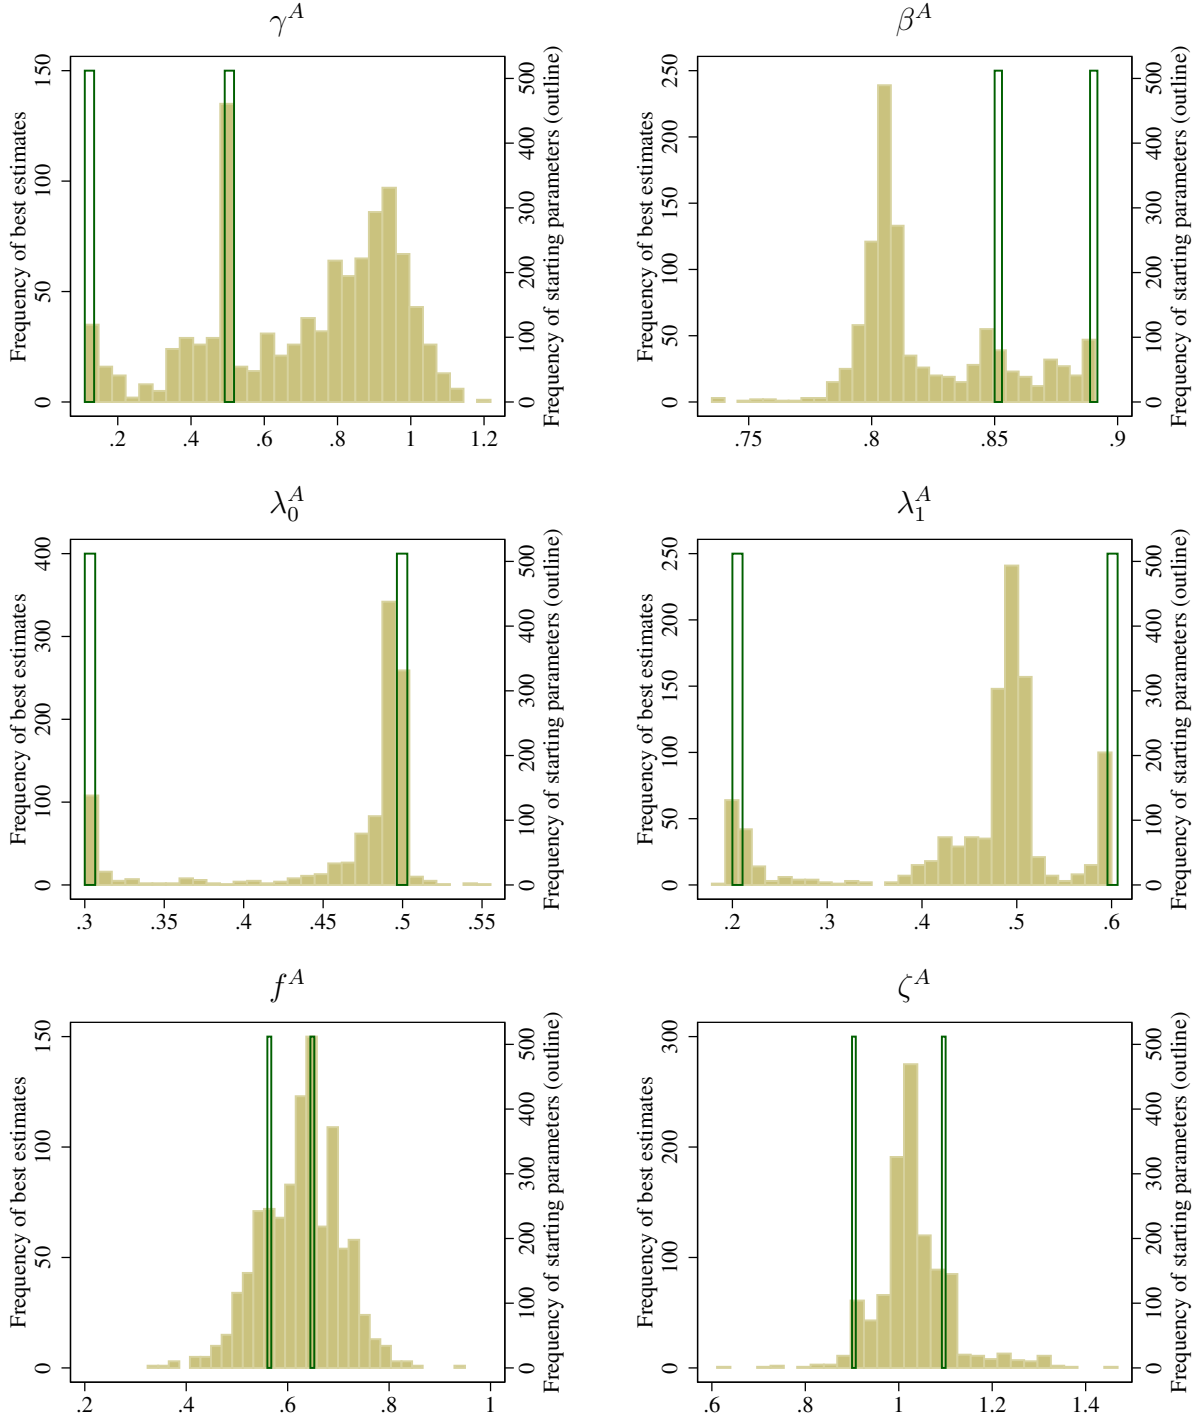
Notes: This figure shows the average consumption and income at each age from the CE, pooling all surveys from 2000–2014. Consumption is the total household expenditures divided by the number of adults. Adjusted consumption removes the estimated effect of children. Income is after-tax income, and its smoothed version is based on a quintic from ages 24–81. Since the survey pools income and consumption after age 81 (or 83 in later years), ages 81 and older are the average for this group.

Figure A-7: Consumption over the business cycle



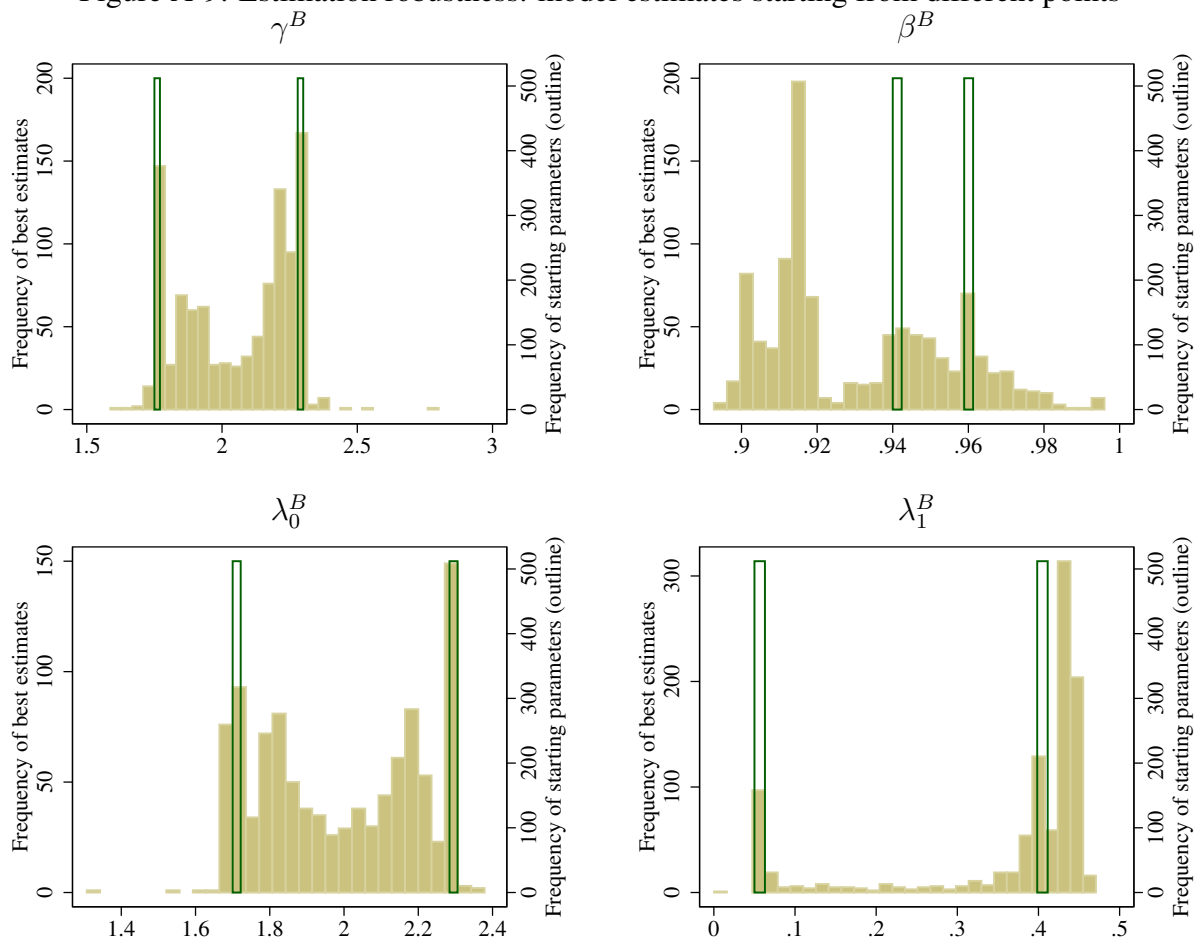
Notes: This figure shows personal consumption from the BEA and average consumption from simulations with a 35 percent fall in available credit starting in 2008q3. Each series is detrended using the 2000–2008 period.

Figure A-8: Estimation robustness: model estimates starting from different points



Notes: Shows the best estimates from starting the estimation optimization at the starting points given by the outline bars. Each starting point is crossed with every other starting point, resulting in $2^{10} = 1024$ starting points for the estimation. The optimization at the full grid of starting points took approximately six days to run using 16 processors at 2.7 Ghz.

Figure A-9: Estimation robustness: model estimates starting from different points



Notes: Continuation from previous figure.