Credit Cards and Consumption*

Scott L. Fulford† and Scott Schuh‡

June 2018

Abstract

Using credit bureau data, we show that the revolving credit available to consumers fluctuates substantially over the business cycle, the life cycle, and for individuals. Yet revolving debt changes proportionally at about the same time, so credit utilization (debt/limit) is remarkably stable. To understand this stability, we estimate a structural model of life-cycle consumption with preference heterogeneity and expenditures funded by liquid assets or unsecured credit with default. The model includes joint roles for credit cards as a means of payment and as a source of short and long-term smoothing. We identify the value of credit cards (about 0.3 percent of expenditures) from the lower share of credit card payments by consumers who revolve debt. Our estimates suggest that around half the population has an endogenously high marginal propensity to consume, which helps explain stable credit utilization and the puzzlingly high reaction to liquidity increases such as tax rebates.

Keywords: Credit cards; life cycle; consumption; saving; precaution; buffer-stock

*The views expressed in this paper are the authors' and do not necessarily reflect the official position of the Bureau of Consumer Financial Protection, or the United States. To ensure appropriate use of the data, Equifax required that we pre-clear results that used Equifax data before making them public. The BCFP reviewed the paper prior to making it public and the Federal Reserve Bank of Boston reviewed previous versions. 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, the 2017 CESifo Venice Summer Institute, 2018 Cambridge University Workshop on Social Insurance and Inequality, and seminars at the Federal Reserve Banks of Boston, Atlanta, and Richmond, the Bank of Canada, the Bureau of Consumer Financial Protection, and the Federal Reserve Board. We thank Chris Carroll, Eva Nagypal, Jonathan Parker, Robert Townsend, Robert Triest, and John Sabelhaus for substantive comments and suggestions.

†Scott Fulford: Bureau of Consumer Financial Protection; email: scott.fulford@cfpb.gov. I conducted some of this work while I was on the faculty at Boston College and a visiting scholar at the Consumer Payments Research Center at the Federal Reserve Bank of Boston. I would like to thank the Bank and the Center for their knowledge and help.

‡Scott Schuh: University of West Virginia email: scott.schuh@mail.wvu.edu. I conducted some of this work while Director of the Consumer Payments Research Center at the Federal Reserve Bank of Boston.

1 Introduction

Over the year and half from September 2008 to March 2010, a combination of financial crisis, recession, and regulatory pressures led banks to reduce the credit card limits of millions of people in the United States and to be less willing to extend new credit. Aggregate credit card limits fell by nearly a trillion dollars and the average limit fell by about 40 percent (see Figure 1). At the same time, Americans reduced their credit card debt by a similar amount. As a consequence, 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.

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.

Why are changes in credit and debt so intimately linked at both the micro and macro levels? 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 means of payment, spending on credit cards forms part of consumer expenditures.[1] We estimate a structural model of

[1]This payments aspect of credit cards, which involves the inter-relationship between credit and liquidity, has been

life-cycle consumption and savings that incorporates the payments, precautionary smoothing, and life-cycle smoothing aspects of credit cards. Our model allows for the large life-cycle variation in credit that we show is important in the data, saving at a low rate of interest and borrowing at a high rate, and the life-cycle variation in income with uninsured income shocks previously studied by Gourinchas and Parker (2002) and Cagetti (2003). To capture the unsecured nature of credit card debt, we also introduce the ability to default (Livshits et al. 2007, Chatterjee et al. 2007, Athreya 2008), so that the interest rate schedule faced by consumers is endogenous to past behavior. Crucially, 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. We estimate the preferences necessary to match the payments behavior of credit card users and the life-cycle profiles of consumption, credit card debt, and bankruptcy using the Method of Simulated Moments (McFadden 1989). Our work appears to be the first to study credit limits and utilization over the life cycle, although models with endogenous credit constraints (Lawrence 1995, Cocco et al. 2005, Lopes 2008, Athreya 2008) typically imply increasing credit limits with age as lenders gain more information.

We estimate that about half the population must have a high discount rate and low risk aversion to explain the amount and profile of credit card debt that we observe. This population has a high marginal propensity to consume and is living close to hand-to-mouth for most of the life-cycle, so increases in credit lead directly to increases in debt, explaining most of the close link. The key revealed preference that gives the basic intuition for our results is the different uses of credit cards. About half of credit card holders 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 shock, and so they must discount the future around the return on liquid savings. The other half 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 make

studied recently by Telyukova and Wright (2008) and Telyukova (2013).

sure we properly account for payments use of credit cards, how individual shocks and the life cycle affect consumption decisions, and the ability to default. 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.[2]

With a large population with a high marginal propensity to consume, our estimated model explains the smooth utilization at both the micro and macro levels. In sample, it simultaneously fits the life-cycle paths of debt, consumption, and default that it is estimated to match. Out of sample, the estimates predict the slow decline in utilization over the life cycle and the smooth utilization over the business cycle. At the individual level, the estimated model matches the rapid return to individual specific credit utilization that we document in the credit bureau data. In doing so, our estimates suggest a puzzle: Because the gap between the borrowing and saving interest rate is so large, it is difficult to explain why people stop revolving debt as they age with standard life-cycle approaches.

We also provide the first estimates of the value of credit cards as a means of payment. Embedded in 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 about 0.3 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.

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

[2]The impatient population explains most of why credit utilization is stable, but the rest of the model matters as well, because all uses of credit cards push for credit and debt to be closely linked and need to be properly accounted for. 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 so 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 so increasing their debts at the same time. Finally, since 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 increase their debts at nearly the same pace as their limits.

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 estimate preferences in a model where savings and debt offer 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 23 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.

Our results suggest that while the heterogeneity among individuals over the life cycle matters, the most important heterogeneity is revealed by the different uses of credit cards that separate preferences. Our results thus hearken back 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), the average marginal propensity to consume (Carroll et al. 2017), and the persistence of financial distress among a small population over the life cycle (Athreya 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 de-

[3]The 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.

pending 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.

An important question in understanding bankruptcy in the United States is the importance of liabilities and expenses at least partially outside the consumer's control, such as medical debt. Livshits et al. (2007) discuss the evidence for expense shocks and highlight the importance of these shocks for explaining the rate of bankruptcy and for conducting welfare analysis.[4] Similarly, Chatterjee et al. (2007) conclude that such shocks are important for explaining the frequency of default. Our estimates support the view that expenditure shocks must be important for understanding bankruptcy. After increasing early in life, the fraction of people with a bankruptcy on their credit record is declining after age 30. Because credit limits are increasing over the life cycle, the incentive to voluntarily run up a large balance and default is increasing as well, suggesting that if voluntary default is important, the frequency of default should be increasing over the life-cycle. On the other hand, default caused by unexpected expense shocks is decreasing over the life cycle exactly because credit limits are increasing, giving even impatient consumers a greater buffer. Our estimates thus suggest that, given the large increase in credit limits we document, it is difficult to reconcile the life cycle pattern of bankruptcy without expenditure shocks being the main driver of default.

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 hump comes mainly from the average of two populations: one impatient enough that consumption largely follows income over the life-cycle, closely resembling the buffer-stock population in Car-

[4] In the working paper version, the authors similarly conclude that reduced incidence of these shocks in Germany compared to the United States is necessary to explain the differences in bankruptcy rates.

roll (1997), and the other patient population with flat or growing consumption. Consistent with Gourinchas and Parker (2002), even our patient population is highly liquidity constrained early in life. Approaches to life-cycle savings and consumption insurance that do not take into account the large life-cycle variation in credit are missing an important component.

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.[5] While the CCP 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

[5]Lee and van der Klaauw (2010) provide additional details on the sampling methodology and how closely the overall sample corresponds to the demographic characteristics of the overall 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 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."

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.[6]

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.[7]

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

[6]The 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.

[7]The 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. 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.

behavior of revolvers and convenience users differs. In addition, we estimate life-cycle profiles of consumption from the Consumer Expenditure Survey (CE) and bankruptcy rates from the Bureau of Consumer Finance's Consumer Credit Panel which is derived from credit bureau data.

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.[8]

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 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.

[8]The 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 and it generally still appears on the credit record. 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.

Figure 1: Credit card limits, debt, and utilization: 2000–2015

Observed limits, debts, and utilization from credit bureau

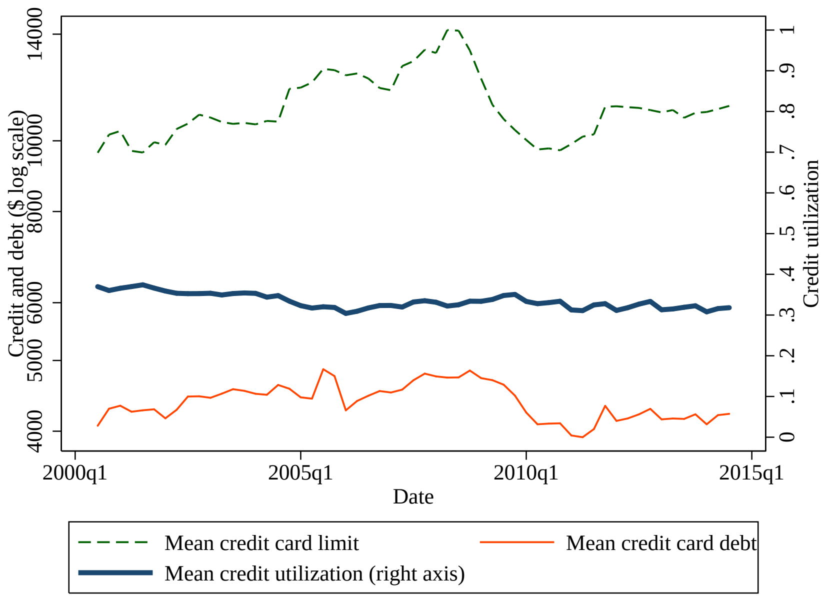

Model prediction given fall in credit limits

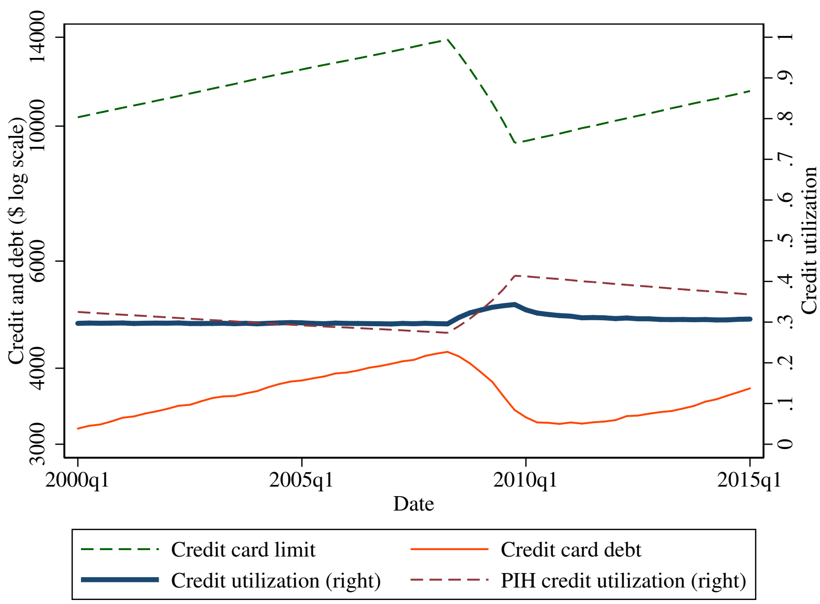

Notes: The top panel shows observed limits, debts, and utilization from credit bureau data (see Section 2 for details). The bottom panel shows model predictions given an unexpected fall in credit (see section 5 for details). For both panels, the left axis shows the average credit card limits (top line) 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.

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

2.4 The reduced form evolution of individual utilization

This section shows that utilization for an individual rapidly reverts to an individual specific mean. Credit utilization is therefore best characterized by fixed heterogeneity across individuals and relatively small transitory deviations for an individual over time. We present parametric results in here and non-parametric results in Appendix A and Appendix Figure A-4 that reach almost identical

[9]The 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-1in 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.

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

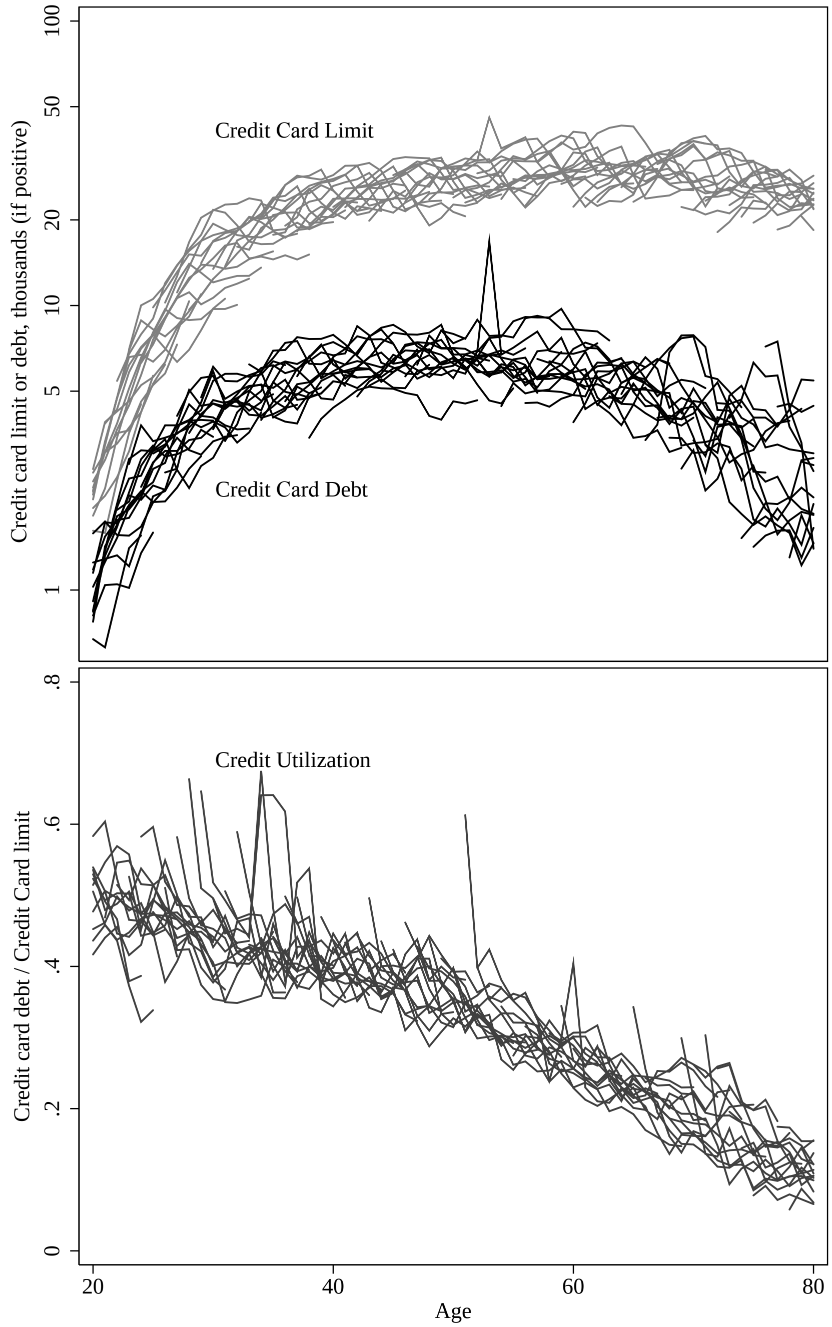

Notes: Each line represents the average credit card limit (conditional on being positive, log scale), debt (conditional on being positive, log scale), and utilization (conditional on having a limit, bottom panel) of one birth year cohort from 1999–2014. Source: Author's calculations from Equifax/NY Fed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

conclusions.[10]

Table 1 shows how utilization this quarter relates to utilization in the previous quarter. For simplicity, we estimate AR(1) regressions of the form:

$$\upsilon_{it} = \theta_t + \theta_a + \alpha_i + \beta\upsilon_{it-1} + \epsilon_{it}, \tag{1}$$

where υ_{it} is the credit utilization, conditional on a positive credit limit, 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.[11]

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.

[10] 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.

[11] The combined age, year, and individual fixed effects in equation (1) 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.

Table 1: Credit utilization

	Equifax/NY Fed CCP			Model
	Credit utilization$_t$			
Credit utilization$_{t-1}$	0.874***	0.868***	0.647***	0.699***
	(0.000876)	(0.000892)	(0.00131)	(0.000492)
Constant	0.0479***			
	(0.000461)			
Observations	347,642	347,642	347,642	2,168,011
R-squared	0.741	0.743	0.429	0.491
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	0.217

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.

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 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 at a higher interest rate, the choice to default on debt, expenditure shocks, and changing credit over the life cycle. Although we describe the decision making for a particular consumer, in the estimation we allow for multiple populations of consumers with distinct preferences.

To keep the model numerically tractable and thus able to be estimated, we make a number of modeling decisions that simplify the full richness of the decision environment—particularly of the payment choice and default—but allow us to capture the important dimensions of the problem. 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.[12]

3.1 The decision problem

From any age t, a consumer indexed by i seeks to maximize her utility for remaining life given current resources and expected future income. Consumers may belong to a population with distinct preferences which we denote with j. With additively separable preferences, the consumer with liquid funds W_{it} and current credit limit B_{it} maximizes the discounted value of expected future

[12]Most 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.

utility:

$$\max_{\{X_s,\pi_s,f_s\}_{s=t}^{T}} \left\{ E\left[\sum_{s=t}^{T} \beta_j^{s-t} u(C_{is}) + \beta_j^{T+1} S(A_{iT}) \right] \right\} \text{ subject to}$$

$$C_{is} = \nu_{is}(1 - f_{is}\phi_s^c)X_{is} \quad \text{(Consumption from expenditures)}$$

$$X_{is} \leq W_{is} \quad \text{(Expenditures limited by liquidity)}$$

$$W_{is} = R_{i,s}A_{i,s-1} + Y_{is} + B_{is} - K_{is} \quad \text{(Evolution of liquidity)}$$

$$A_{i,s-1} = W_{i,s-1} - B_{is-1} - X_{is-1} \quad \text{(Relationship between liquidity and assets)}$$

$$\nu_{is} = \nu(\pi_{is}; A_{i,s-1}) \quad \text{(Payment decision)}$$

$$f_{is} = f(F_{is}, W_{is}) \quad \text{(Default decision)}$$

$$F_{is} = H(F_{i,s-1}, f_{i,s-1}) \quad \text{(Evolution of default state)}$$

where she gets period utility $u(\cdot)$ from consumption C_{is}, which she gets by making expenditures X_{is} adjusted for the payment choice and default. The decision at t depends on what she expects her future decisions and utility to be at ages $s \geq t$. The consumer discounts the future with a fixed discounted factor β_j and so has time-consistent preferences. We therefore drop the distinction between age t and future ages $s \geq t$ for clarity. The discount factor is fixed for the individual consumer, but may vary across consumers in different groups j and we will estimate the importance of this variation.

Beyond expenditures, the consumers faces two additional decisions each period: how to pay for her expenditures and whether to default. 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_{it} between expenditures and consumption, the evolution of which we explain below. Expenditures are limited by the available liquidity W_{it}, which is the sum of assets left at the end of the previous period $A_{i,t-1}$ (which may be positive or negative) earning total return R_{it} which depends on the default status and assets in the previous period, income this period Y_{it}, and the credit limit this period B_{it}, minus an expenditure shock K_{it}. The consumer may choose to default, indicated by the binary variable f_{it} and enter the default state F_{it}, or be forced to default if the expenditure shock pushes liquidity below zero. Defaulting reduces expenditures in the current period and puts the consumer in the default state which has costs in future periods, but removes all debt. We discuss the consumption and credit implications of default below. Many of the elements in this problem are standard. We focus on the nonstandard ones.

Rate of return on assets Borrowers face a higher interest rate than savers, and those in default face an even higher interest rate. If the assets $A_{i,t-1}$ 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 or defaulted borrowers if she has a bankruptcy on her credit record:

$$R_{it} = R(A_{i,t-1}, F_{i,t-t}) = \begin{cases} R & \text{if } A_{i,t-1} \geq 0 \\ R_B & \text{if } A_{i,t-1} < 0 \\ R_D & \text{if } A_{i,t-1} < 0 \text{ and in default } (F_{i,t-1} = 1), \end{cases}$$

with $R_D \geq R_B \geq R$.

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.[13]

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 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 contin-

[13]Doing 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.

uous 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 $((R_B - 1)/12)/2$ on her credit card expenditure that month.[14] Assuming the loss of float is the only factor explaining different usage, the cost function for revolvers shifts down by $(R_B - 1)/24$.

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

[14]This formula comes from the way that annual credit card rates are reported and interest charged. The interest rate on debt is $R_B - 1$. 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.

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

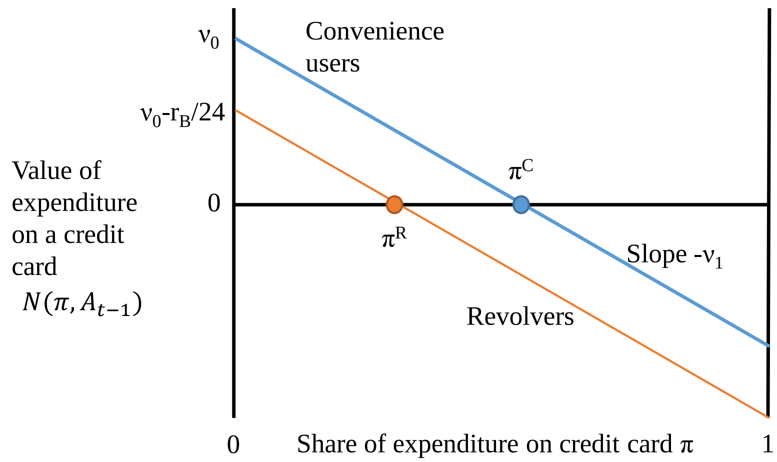

Notes: This figure shows the value or cost of expenditure on a credit card at each expenditure share π relative to cash. The top line is for convenience users who put an optimal share π^c of consumption on a credit card. The bottom line for revolvers is shifted down by the amount $-r_B/24$, because revolvers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

relative cost for the rest of expenditures is zero. The wedge therefore takes on two values:

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

where π^C and π^R are the optimum fraction for revolvers and convenience users. Appendix D 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.

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:

$$D_{i,t} = \begin{cases} \pi^C X_{it} & \text{if not revolving so } A_{t-1} \geq 0) \\ \pi^R X_{i,t} + A_{i,t-1} & \text{if revolving so } A_{t-1} < 0). \end{cases}$$

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 and expenditure shocks Income or disposable income follows a random walk with drift:

$$Y_{i,t+1} = P_{i,t+1}(U_{i,t+1} - F_{i,t+1}\phi^y_{t+t})$$
$$P_{i,t+1} = G^j_{t+1} P_{it} M_{i,t+1},$$

where G^j_{t+1} is the known life-cycle income growth rate from period to period for population j. $F_{i,t+1}\phi^y_{it+t}$ is an income cost of being in the default state $F_{i,t+1} = 1$ discussed more below. The "permanent" or random-walk shocks $M_{i,t+1}$ are independently and identically distributed as lognormal with mean one: $\ln M_{i,.t+1} \sim N(-\sigma^2_M/2, \sigma^2_M)$. The transitory shocks are similarly distributed lognormally with mean one and variance parameter σ^2_U. We allow for a temporary low income U_L from unemployment or other shocks with probability p_L each period.[15] The structure of the shocks ensures that the expected income next period is always $E_t[Y_{i,t+1}] = G^j_{t+1}P_{it}$ when not defaulting, because the mean of both transitory and permanent shocks is one.

A consumer also faces expenditure shocks $K_{i,t}$ which are either 0 or a multiple of permanent income, $kP_{i,t}$ with probability p^k. These expenditure shocks represent expenditures the consumer is required to make, but derives no utility from. Thus, while they do not count as consumption for utility purposes, they are expenditures for accounting purposes, and we include them when we compare model expenditures to actual consumer expenditures.

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-6), and the dispersion of

[15] Low-income shocks, in addition to lognormal shocks, may matter for precautionary reasons by putting additional probability on very bad outcomes. We introduce low-income shocks in such a way that $E_t[U_{i,t+1}] = 1$. Formally, the transitory shocks are distributed as: $U_{i.t+1} = U_L$ 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: $\ln \tilde{U}_{i,t+1} \sim N(-\sigma^2_U/2, \sigma^2_U)$ and U_L is unemployment income as a fraction of permanent income.

credit limits across individuals of the same age is also large (Appendix Figure A-2). We allow for life-cycle growth and dispersion across consumers by assuming that the credit limit B_{it} is an age-dependent multiple of permanent income:

$$B_{it} = b_t P_{it} b_f^{F_{it}},$$

where $b_t \geq 0$ is the age-varying fraction of permanent income that can be borrowed, which is set outside the control of the consumer. By defaulting and entering the default state (described in greater detail below) so that $F_{it} = 1$ the amount the consumer can borrow is reduced by b_f. This approach means that across consumers, B_{it} will be in proportion to income P_{it}, but it allows credit to follow an average path over the life cycle that is different from income and affected by consumer decisions.[16]

The decision to default The consumer may voluntarily decide to default $(f_{it} = 1)$ and enter the default state $(F_{it} = 1)$. Alternatively, if the expenditure shock is sufficient to push $W_{it} \leq 0$, the consumer is forced into involuntary default.

Defaulting has a series of consequences. Involuntary defaulters consume the consumption minimum $c^{min} P_{it}$. In the period of default for voluntary defaulters, expenditure is all of available liquidity ($X_{it} = W_{it}$), but the consumption value of this expenditure is reduced by $(1 - \phi_t^c)$. We think of this reduction as capturing three costs: a non-pecuniary cost of default; pecuniary default penalties that apply during the period of default; and the possible ability of card issuers to limit default exposure by reducing credit limits proactively. After defaulting, the consumer enters the

[16]The 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 $B_{i,t-1} > 0$ and the consumer borrows, leavings negative assets at the end of period $A_{i,t-1} < 0$. If $B_{it} = 0$, then assets at the end of period t must be weakly positive ($A_{it} \geq 0$), 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.

next period with no debt ($A_{i,t+1} = 0$).

Having entered the default state, the consumer faces a modified consumption problem of someone with a bankruptcy on her credit record. Her credit limits is a fraction b_f of non-defaulted credit limits. Her cost of borrowing is higher. To reflect possible wage garnishment or the effect default may have on available employment, the income process is reduced by a multiple of the default debt $\phi^y_{it} = \phi(R^B - 1)b_t P_{it}$ in every period. Formulated this way, the cost of default is increasing with the credit limit, so that as credit limits increase with age, so does the cost of default. Because the credit limit is increasing over the life cycle, the consumption value of maxing out credit cards is also increasing, so the incentive to default is increasing. The current period and future costs of default are conceptually distinct, but difficult to distinguish empirically, so we link them and set $\phi^c_t = \phi^y_t$ so that only one parameter governs the total cost of default.

To keep the state space tractable, we model the evolution of the default state $F_{i,t} = H(F_{i,t-1}, f_{i,t-1})$ as an absorbing Markov process: A consumer in default in the previous period ($F_{i,t-1} = 0$) stays in default with probability p^F, and exits default with probability $1 - p^F$. The consumer is in default with certainty if she defaulted in the previous period ($f_{i,t-1} = 1$).

Given the costs and benefits of default, consumers must decide whether to default. Only consumers not currently in default may decide to default. Because default is a discrete decision, consumers decide whether the value of current and expected future utility from defaulting is greater than defaulting:

$$f_{it} = f(F_{i,t}, W_{i.t}) = \begin{cases} 1 & \text{if } V^{\text{Default}}(W_{i,t}) > V^{\text{Not Default}}(W_{i,t}) \text{ and not in default } (F_{it} = 0) \\ 0 & \text{else.} \end{cases}$$

Following Chatterjee et al. (2007), we can simplify this decision into finding the crossing point, if it exists, of the two value functions, so characterize the decision as finding the liquidity below which default occurs: W^{Default}_t.

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

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

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.2discusses 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 $x_t(w_{it}, a_{i,t-1}, F_{it})$.

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^j that may be different for consumers in different populations j. The consumer lives for T periods, where T is a random number that we match to actual life tables, 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.[17]

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}

[17] Recent work has disagreed over the importance of a bequest motive as opposed to other possible motives for keeping assets late in life, such as long-term care and medical needs (De Nardi et al. 2010). Since we focus primarily on 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 4.4.

is a fraction λ_1^j of pre-retirement permanent income ($\lambda_1^j P_{i,T^{Ret}-1}$). 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 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.[18]

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

[18]This 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.

substantially reduces the computation costs for the expenditure problem. 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. 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. Numerically, we solve the expenditure/payments problem each period, then in a separate value function iteration problem let the consumer decide whether to default. The value function iteration is computationally costly and suffers from the curse of dimensionality: the default decision and default state add a continuous decision at all liquidity levels that must be solved for all combinations of endogenous payment states and ages.

Figure 4 illustrates some of the complexities of the decision problem. Along the x-axis is the ratio of liquidity 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 liquidity 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 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 liquidity between these two points, the

Figure 4: Expenditure functions over the life cycle with borrowing

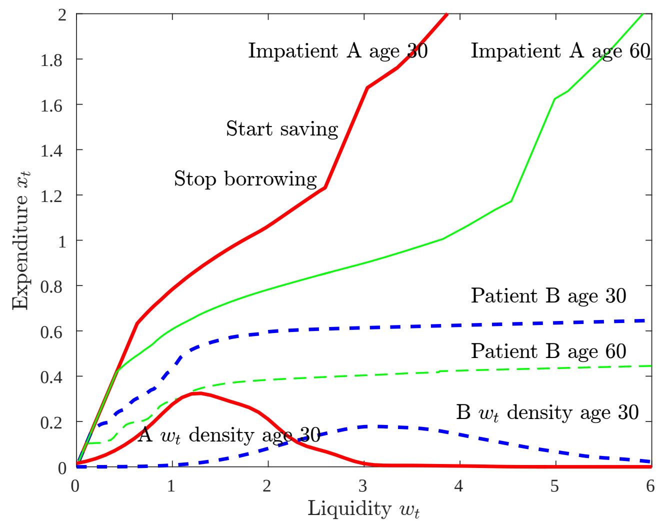

Notes: This figure uses the estimates in Table 3 column 1 at age 30 and age 60 to show the quarterly expenditure function for impatient (A) consumers and patient (B) consumers. Liquidity w_t is a multiple of quarterly permanent income P_t and includes available credit. The densities for liquidity are 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.

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 discus how to allow numerically for these inflection points so that the consumption function is suitably kinky. Figure 4 is based on the estimates in the next section which suggest a cost of default parameter high enough that voluntary default is never optimal. With a lower cost of default, the decision becomes even more complex as is illustrated by appendix Figure A-5. With a voluntary default, there is a discrete jump in consumption at the optimal default liquidity; below the default point, consumers spend all available liquidity and suffer the costs of default, above the default point consumers leave some liquidity for the next period.

4 Estimation

This section describes how we estimate the structural model using life-cycle profiles of consumption, debt, and default. 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 other observable parameters at this stage. 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, consumption, and default 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 suggest 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. [19] For each sub-population, the entire decision

[19] Together f^A and ζ^A directly determine ζ^B. For 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 $f^B = 1 - f^A$.

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. Finally, we estimate the probability (p^k) and cost (k) of expenditure shocks. We show that the default cost parameter ϕ is identified only up to an inequality, so jointly estimate 12 parameters in the second stage: $\theta = \{\gamma^A, \beta^A, \lambda_0^A, \lambda_1^A, \gamma^B, \beta^B, \lambda_0^B, \lambda_1^B, f^A, \zeta^A, p^k, k\}$.

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, whether to default or be forced into default, 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, debt, and default 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, savings, and default.

The estimation then finds the parameters θ that produce a life-cycle evolution of average simulated consumption, debt that best matches their empirical counterparts from ages 24–74. We describe the sources and construction of the empirical moments in more detail in Section 4.2. Each profile is annual, so there are $T = 51$ quarters.[20] 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 $3T$ total moments. The MSM then seeks to minimize the

[20]The model is quarterly, so we aggregate appropriately to match the data by taking the average model debt for a given age. The annual nature of the empirical age profiles is driven by the data source. The Equifax/NY Fed CCP, for example, reports only the year of birth from which we calculate age.

weighted square of these differences:

$$\min_{\theta \in \Theta} g(\theta; \chi)' W g(\theta; \chi), \tag{2}$$

where $g(\theta; \chi) = (g_1(\theta; \chi), \ldots, g_{3T}(\theta; \chi))$, and W is a $(3T) \times (3T)$ weighting matrix. Our standard weighting matrix is block proportional to the inverse variance of the empirical moments (the optimal weighting matrix with no first-stage correction). Because our life-cycle moments come from surveys and administrative data, they are estimated with very different levels of precision and the estimates tend to attempt to fit only the administrative data. We therefore weight each life-cycle moment block so that the each bock receives the same weight, but within each life-cycle moment better-estimated moments receive more weight.[21] 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

[21] Starting from V_M, the $(3T)\times(3T)$ block-diagonal variance covariance matrix of the moments, we form $\tilde{W} = V_M^{-1}$ which is also block-diagonal. For each block of $\tilde{W}$ we take a vector ι of ones of size T and calculate the weight function for each block of life-cycle moments $w_1 = \iota' W_1 \iota$ which is the impact on the objective function if that all of the moments $(g_1(\theta; \chi), \ldots, g_T(\theta; \chi))$ in that life-cycle block were equal to one. We then form W by dividing each life-cycle block by its scalar weight.

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.

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.[22]

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 D. 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 Appendix C.1 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.[23] As a comparison, the fees that banks charge merchants for processing credit cards

[22]Credit 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.

[23]Personal 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,

are roughly \$60 billion per year.[24] 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.

4.2 The empirical life-cycle moments and first stage moments

We estimate the model to provide the best fit to three life-cycle profiles: (1) log mean credit card debt over the life cycle from the Equifax/NY Fed CCP described in Section 2, (2) log mean household consumption over the life cycle from the CE from 2000–2014,[25] (3) the fraction of consumers with a credit card line charged off in bankruptcy from the CFPB Consumer Credit Panel which is derived from credit bureau data, and which, unlike the Equifax/NY Fed CCP, shows individual credit card lines.[26] We show each of these moments in Figure 5 together with their

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.

[24]The 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.

[25]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. Formally, 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-6 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.

[26]We use bankruptcy as the appropriate empirical comparison because we model default as wiping away debts, but there are many forms of default not directly coming from bankruptcy. At any age, the fraction of consumers with a credit card line marked as charged off by the issuer (including for bankruptcy) is approximately double the fraction with a line charged off for bankruptcy. Only in bankruptcy is the debt actually removed for the consumer allowing a clean start, a charge off simple means that the bank has marked the debt on its books as uncollectable for regulatory

estimates from the model, and have already discussed the debt profile in Section 2.3. Consumption follows the characteristic hump shape (Gourinchas and Parker 2002, Attanasio et al. 1999). Most types of bankruptcy take seven years to leave the credit record and people need time to accumulate debt and bad shocks, so bankruptcy is increasing early in the life-cycle, before declining. Appendix C.2 discusses the construction of the variance-covariance matrix of the combined moments.

First-stage estimates We briefly describe the sources and estimates from other data sets that identify the ancillary parameters of the model. We provide greater detail in Appendix C.1. We estimate the average life cycle of income growth (G_t) using the Consumer Expenditure Survey to match consumption, adjusting for the aggregate growth. While average income follows the observed life-cycle path, individual incomes vary based on their idiosyncratic shocks. We use the estimates of these shocks 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 and allow for unemployment shocks. We estimate the total credit limit for a consumer from the Equifax/NY Fed CCP to form B_t. We estimate the interest rate on debt $R_b - 1 = 14.11$ percent based on the average revolving interest rate over the period. From the SCF, we estimate that those with a bankruptcy pay 1.92 percentage points more in interest on their credit card debt and have only 42 percent of the credit limit (b_f). We estimate the return on savings for an all-bond portfolio. We adjust both borrowing and saving prices for the geometric average inflation rate from 2000–2015 of 2.15 percent.

4.3 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, consumption, and bankruptcy

purposes. The bank may continue to try to collect the debt or sell it to a firm specializing in collection. See Athreya et al. (2017) for a model that allows both non-payment default and bankruptcy.

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

To do so, the model suggests that about two thirds 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 a income about average (ζ^A).[27] Because individual credit limits are proportional to income, the members of this group cannot be too poor on average, 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.[28] 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.

[27] Depending on the particular weights, some estimates suggest an impatient income higher than average. In 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.

[28] The 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.

Table 3: Model estimates

	Standard Weights	Optimal Weights	Endogenous payments	Low bequest
Population A				
CRRA γ^A	0.067	0.121	0.067	0.066
	(0.018)	(0.008)	(0.018)	(0.014)
Discount β^A	0.892	0.887	0.892	0.893
	(0.001)	(0.003)	(0.001)	(0.001)
Initial wealth λ_0^A	0.516	0.481	0.516	0.516
	(0.087)	(0.220)	(0.114)	(0.072)
Late life inc. λ_1^A	0.727	0.719	0.727	0.731
	(0.033)	(1.766)	(0.049)	(0.033)
Population B				
CRRA γ^A	2.023	1.975	2.023	2.023
	(1.007)	(42.111)	(0.946)	(1.134)
Discount β^B	0.963	0.963	0.963	0.962
	(0.018)	(0.350)	(0.020)	(0.020)
Initial wealth λ_0^B	1.728	1.658	1.728	1.728
	(2.245)	(90.275)	(2.267)	(1.826)
Late life inc. λ_1^B	0.212	0.200	0.212	0.212
	(0.260)	(8.444)	(0.251)	(0.180)
Share A f^A	0.669	0.648	0.669	0.669
	(0.008)	(0.005)	(0.008)	(0.013)
Inc. mult. A ζ^A	0.991	0.971	0.991	0.995
	(0.099)	(0.506)	(0.137)	(0.099)
Prob. of exp. shock	0.040	0.031	0.040	0.040
	(0.005)	(0.008)	(0.005)	(0.005)
Size of exp. Shock	0.660	0.532	0.660	0.659
	(0.146)	(0.073)	(0.117)	(0.135)
SSR (g'g)	0.3484	1.5633	0.3618	0.4558
J-stat	4.46E+08	1.37E+09	3.62E+08	5.08E+08
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. Endogenous payments makes the consumer's aware that revolving affects the value of credit cards for payments. The standard default cost parameter $\phi_f = 7$ bequest parameter is 1. Low bequest reduces the bequest motive. For a description of each moment and the estimation method, see the beginning of section 4. The Sum of Squared Residuals ($g'g$) is the unweighted sum of the difference between the model and actual moments, so is the closest to the sum of differences in Figure 5.

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

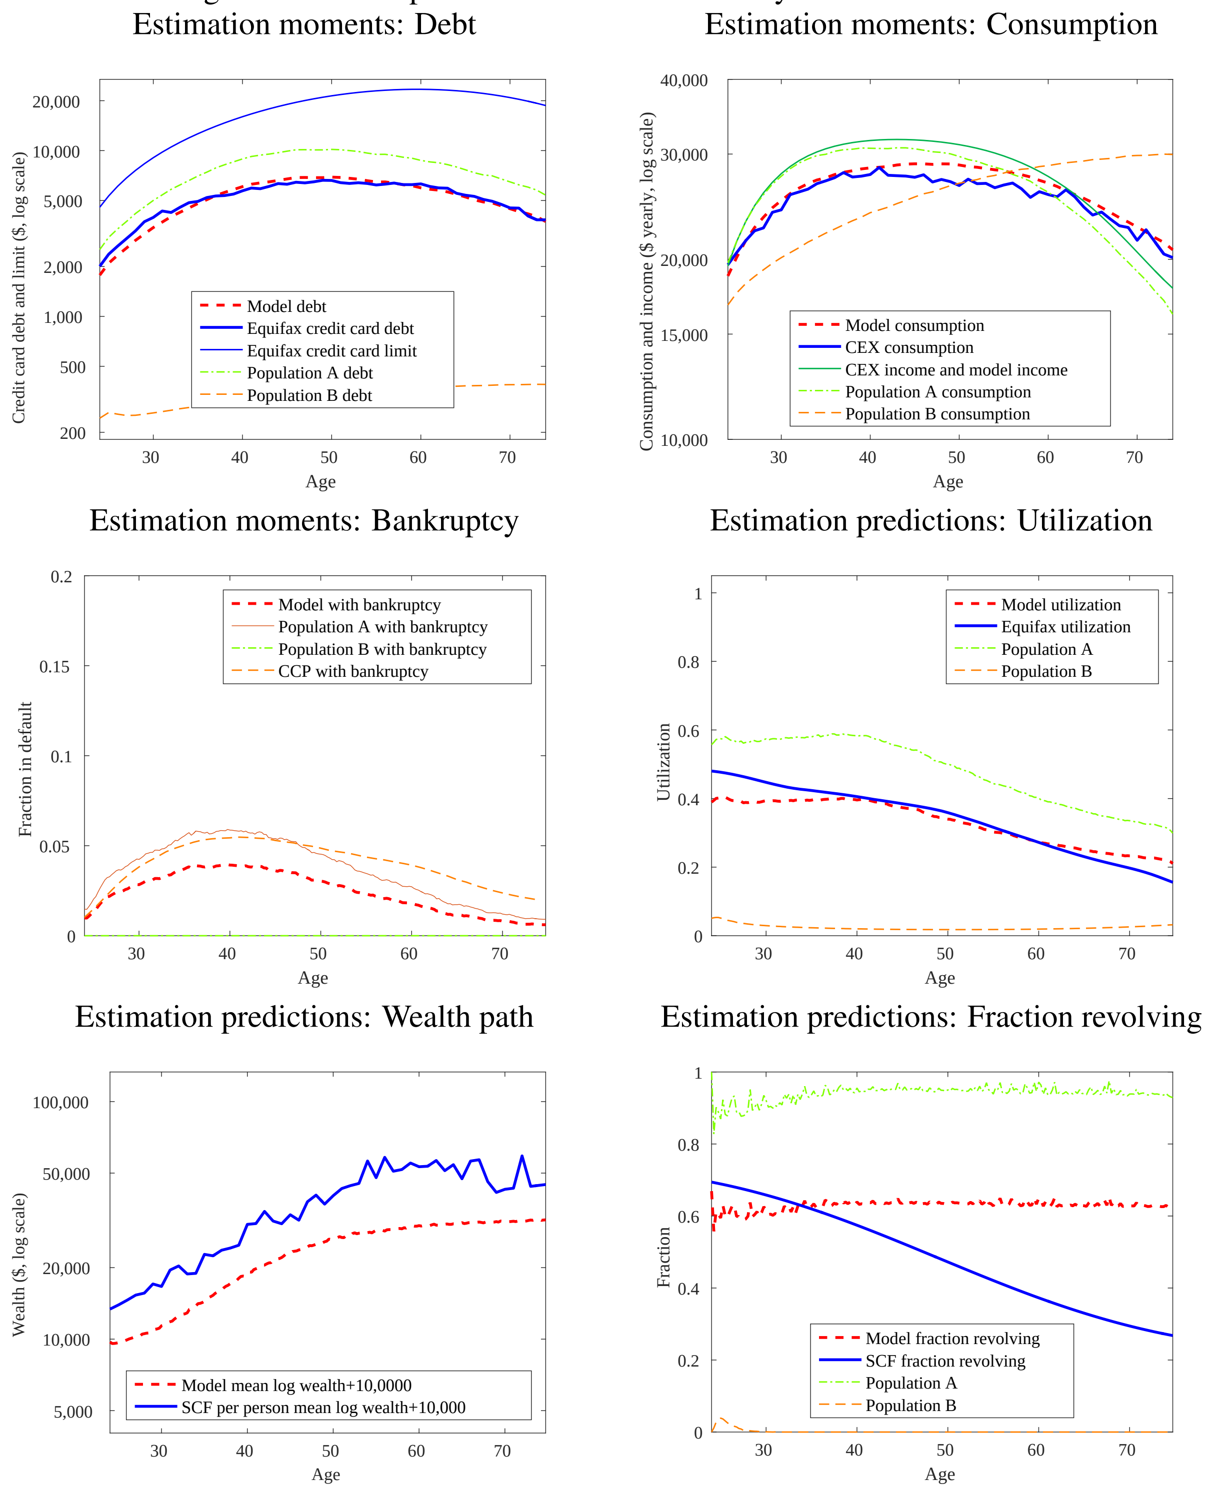

Notes: Life-cycle paths from simulated population using the estimates in column 1 of Table 3.

The cost of default is identified only up to an inequality, so it is not estimated directly. As Appendix Figure A-8 shows, holding other parameters fixed at their values in column 1, changing the bankruptcy cost parameter does not improves the fit until it goes below a threshold, after which the model fit rapidly deteriorates. For estimates larger than the threshold, no one voluntarily defaults. To see why the data push this conclusion so strongly, Appendix Figure A-9 shows bankruptcy over the life-cycle with the standard bankruptcy cost parameter and one below the threshold. As credit limits increase and the remaining expected life decreases, the gains from bankruptcy get ever larger. At approximately age 35, much of the impatient population finds it better to default and the bankruptcy rate changes precipitously. Increasing the cost of default increases the age at which it becomes optimal for much of the impatient population to voluntarily default, but does not change the rapid shift to default. Since the observed fraction of the population with a bankruptcy on record is declining over the life-cycle as shown in Figure 5, voluntary default is not useful for explaining the fraction in default. The estimates thus reject cost-of-default parameters in which there is substantial voluntary default. Note that the model ties the default cost to the credit limit, so the costs of default are increasing over the life-cycle, but even this approach suggests that older people with higher credit limits should be the ones defaulting voluntarily.

Expenditure shocks, on the other hand, are useful for explaining the default over the life cycle. Early in life, consumers hit by an expenditure shock are likely to go bankrupt. Because bankruptcy stays on the record for seven years, the fraction in bankruptcy is increasing. But as credit limits increase, expenditure shocks are less likely to push someone into bankruptcy, and the fraction in bankruptcy starts to decline.

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 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 $(\zeta^{A)})$ 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.[29]

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 or increasing consumption profile without a downturn late in life. Its preference for risk (γ^B) and expected late-life income after expenses $\zeta^B)$ 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

[29]Since we include expected aggregate growth and adjust for inflation, β^A and β^B are more closely pinned down relative to $R^B -$ Inflation $+$ Real Aggregage 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.

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 three 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. [30] 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.

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.

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 bottom right panel of Figure 5).[31] The average fraction revolving over all ages in the model is approximately correct, but because the impatient population A is always in debt and the patient

[30] We have also estimated the variance of credit card debt and the variance in the change in debt from quarter to quarter, which controls for the permanent income and preference heterogeneity. The model captures the level of the variance of credit card debt reasonably well, although it does not predict the shape very well. Our simulations of the variance of the change are somewhat lower than the empirical counterparts because the only change in credit limits comes from changes in permanent income. 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.

[31] Since 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 predicting a movement out of revolving late in life.

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.

4.4 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.[32] 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. Our standard weight matrix gives equal weight to all three blocks of life-cycle moments. 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.

[32]The over-identification statistic is large because the debt and bankruptcy moments are estimated very precisely from the administrative data. The over-identification statistic rejects that the model can simultaneously fit all moments because the debt and default moments are so precisely estimated that even small departures from exact fit leads to a rejection of the hypothesis.

The second column of Table 3 shows estimates that use the two-stage "optimal" weighting 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-7) 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 (2) 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 only guaranteed to find a local minimum, however. We therefore start the procedure with θ_0 at random points in a grid that covers the 12 dimensional parameter space. Not all starting points produce the same estimate of θ, indicating that the objective function in equation (2) 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 mimum. We discuss other local minima in Appendix C.3. The overall conclusion holds for all local minima: Around half of the population must be fairly impatient and have low risk aversion.

In the last two columns of Table 3, we examine how changing the model changes estimates. Our baseline estimates do not allow consumers to take into account the effect their consumption decisions will have on their payments decisions. In column 3, we allow for this feedback, at the cost of substantial additional computation time. Allowing this feedback 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 the same permanent income upon death as the individual. The last column assumes heirs instead have five 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 policy implications

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.

5.1 How well does the model predict individual and aggregate utilization?

The model was estimated to match the life cycle. We first ask how well it predicts the out of sample behavior of the individual and business cycle relationship between credit and debt we estimated in Section 2. 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 over 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 in the credit bureau data, 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 69.9 percent of a shock persisting to the next quarter compared to 64.7 percent in column 3.

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. In addition, 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.[33]

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 2.5 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, so a fall in credit does not explain the continuing weakness in consumption after 2009.[34]

5.2 Implications for stimulus policy

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

[33]Constructing 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.

[34]Figure A-10 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).

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, 23 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 and holds relatively little wealth. 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).

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	0.226***	0.270***	0.0904***			
	(0.0250)	(0.0334)	(0.0333)			
Permanenent credit limit increase				0.296***	0.340***	0.162***
				(0.0248)	(0.0330)	(0.0337)
Observations	533,288	329,560	203,728	533,288	329,560	203,728
R-squared	0.000	0.001	0.000	0.001	0.001	0.000
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.

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, default, and payments. 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. Out of sample, it 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 at a high rate of interest early life to start saving at a low rate 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).

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.

Athreya, Kartik B. 2008. "Default, Insurance, and Debt over the Life-Cycle." *Journal of Monetary Economics* 55(4): 752–774.

Athreya, Kartik, Jose Mustre del Rio, and Juan M. Sanchez. 2017. "The Persistence of Financial Distress." Federal Reserve Bank of Richmond Working Paper 17-14.

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., 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.

Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull. 2007. "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default." *Econometrica* 75(6): 1525–1589.

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 Microeconometric Approach to Development Policy*. Baltimore: Johns Hopkins University Press.

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.

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.

Livshits, Igor, James MacGee, and Michèle Tertilt. 2007. "Consumer Bankruptcy: A Fresh Start." *American Economic Review* 97(1): 402–418.

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.

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.