12-20-2020

Credit Cards, Credit Utilization, and Consumption

Scott Fulford
Consumer Financial Protection Bureau, scott.fulford@cfpb.gov

Scott Schuh
West Virginia University, scott.schuh@mail.wvu.edu

Credit Cards, Credit Utilization, and Consumption*

Scott L. Fulford[†] and Scott Schuh[‡]

December 2020

Abstract

Credit bureau data show remarkably stable consumer utilization of unsecured debt over the business cycle, life cycle, and individually quarter-to-quarter, despite massive variation in available credit. To explain these new findings, we propose a life-cycle consumption model with heterogeneous preferences, endogenous payment choice, and the option to revolve debt for consumption smoothing. Using diary data to identify payment use, the estimated model matches consumption and credit use at every frequency and suggests that around half the population has an endogenously high marginal propensity to consume. The results suggest understanding credit availability and heterogeneous use may lead to richer counter-cyclical policies.

Keywords: Credit cards; life cycle; consumption; saving; precaution; buffer-stock; payments
JEL Codes: D14, D15, E21, E27

*The views expressed in this paper are the authors' and do not necessarily reflect the official position of the Consumer Financial Protection Bureau 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 CFPB Office of Research 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 Consumer Financial Protection Bureau, 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: Consumer Financial Protection Bureau; email: scott.fulford@cfpb.gov. I conducted some of this work while I was on the faculty at Boston College and a visiting scholar at the Consumer Payments Research Center at the Federal Reserve Bank of Boston. I would like to thank the Bank and the Center for their knowledge and help.

‡Scott Schuh: West Virginia University; 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

A striking feature of the 2008 Financial Crisis and Great Recession was remarkable stability in average consumer utilization of unsecured revolving credit (the ratio of total debt to credit limit). Using comprehensive U.S. credit bureau data, the left panel of Figure 1 shows the average credit card limit fell about 40 percent from September 2008 to March 2010, nearly $1 trillion in aggregate. During this period, Americans also reduced their aggregate credit card debt by a similar percentage so average credit utilization was nearly constant during that tumultuous time, as it was from 2000–2015. In aggregate, this debt reduction was approximately double the value of the tax rebates from the Economic Stimulus Act (Parker et al. 2013); individually, average debt fell more than $1,000 dollars per cardholder.

This paper shows that remarkably stable *individual* credit utilization explains the aggregate stability. Beneath the dramatic cyclical changes in credit and debt, we show that even larger changes occur over the individual life cycles. Average credit card limits increase more than 700 percent from ages 20–40 and continue to increase after age 40, albeit more slowly (see Figure 2). These massive increases in credit with age are matched by increases in debt at almost the same rate, so average credit utilization is surprisingly stable over the life-cycle, declining only very slowly with age. Consistent with prior work (Gross and Souleles 2002, Agarwal et al. 2017, Fulford and Schuh 2015, Aydin 2015), we find that increases in credit pass through to increases in debt at the individual level, especially for individuals close to their limits. In addition, we show that: (1) utilization is largely fixed at the individual level despite credit volatility being several times greater than income volatility (Fulford 2015); and (2) individually persistent utilization is highly heterogeneous, with a large portion of the population persistently using most of its available credit while others use almost none.

To explain the expanded facts about credit card borrowing, we build and estimate a new model that allows for preference heterogeneity, payment choice, and credit cards used for long-term borrowing in a life-cycle model with uninsured shocks. The econometric estimates reveal a clear and significant distinction between two types of consumers. About half the population must have a high discount rate (about 11 percent) and low relative risk aversion to be willing to hold the amount of credit card debt observed. This impatient population has a high marginal propensity to consume, so increases in credit lead directly to increases in debt and a stable, but high, utilization. In contrast, the rest of the population has a "standard" discount rate (about 4 percent) and relative risk aversion and uses their credit cards only for payments. They have low and stable utilization because both their credit limits and their expenditure on credit cards are tied to their income, so move at about the same time. The econometric results also offer the first estimates of the direct utility value of credit cards as a means of payment. In the model, consumers endogenously decide how much of current consumption to pay for with a credit card. We estimate that consumers would be willing to pay about 0.3 percent of their consumption (around \$40 billion a year) to continue using credit cards given the current structure of U.S. credit card payment networks, interchange fees, rewards, and prices.

Existing models of consumption and saving are not well-suited to explain individually stable yet heterogeneous utilization of credit, so we add two key components to an otherwise standard model of consumption over the life-cycle model with uninsured shocks (Gourinchas and Parker 2002, Cagetti 2003) that allows for credit cards used for long-term borrowing (Laibson et al. 2003), default (Livshits et al. 2007, Chatterjee et al. 2007, Athreya 2008), and variable credit limits (Fulford 2015). First, we build on the growing literature that argues that heterogeneous preferences, not just heterogeneous agents, are necessary to understand consumption behavior. While we are not the first to introduce heterogeneous preferences, our identification of them is unique. Specifically, we infer impatience (high discount rate) from the clear revealed preference that about half the population with a credit card is willing to borrow at high rates of interest—averaging 14 percent—for extended periods of time while the other half has the option to borrow, but does not. To understand the full distribution of wealth and welfare, we need to understand not just why a few people accumulate a great deal, but why so many are willing to borrow at high interest rates.

Our second innovation is to model endogenous payment choices between settling consumption with liquid assets ("cash") or credit card liabilities and heterogeneous management of credit card debt month to month. No existing models incorporate both payment use of credit cards and longer-term revolving use, even though 90 percent of the population with a credit card uses it for payments in a given month. Explicit treatment of payment choices enables the model to distinguish heterogeneous credit card use: (1) convenience use, where consumers pay off all debt each month and incur zero interest; and (2) revolving use, where consumers exercise their option to roll over unpaid debt at high interest. Heterogeneous credit use is essential to *identifying*, rather than assuming, heterogeneous preferences among sub-populations of consumers as well as the heterogeneous persistent utilization. Our approach circumvents a limitation of credit bureau data, which only measures total credit card debt, so includes both current charges and debt revolved from the prior period.[1] We identify heterogeneous credit use from novel micro data on credit card spending (the Diary of Consumer Payment Choice). Revolvers have a credit card expenditure share 3 percentage points lower than convenience users because high interest on purchases accumulates immediately whereas convenience users benefit from free float. Under simple assumptions, differential payment use of revolvers and non-revolvers identifies the values they put on using credit cards for payments, including rewards like ”cash back.”

The estimated model explains smooth utilization at the micro and macro levels and simultaneously fits the life-cycle paths of debt, consumption, and default. Within sample, the estimates provide several new insights: (1) Because many households hold little or no liquid assets, increases in credit are one of the largest sources of “savings” early in life. Existing consumption models that do not incorporate the life-cycle changes in credit that we document are overlooking an important form of early-life liquidity. (2) Heterogeneous credit use explains the humped shape of life-cycle consumption (Attanasio et al. 1999) in a subtly different way than the combination of precaution and life-cycle savings in Gourinchas and Parker (2002). Our estimates suggest the hump comes primarily from the impatient population whose life-cycle consumption follows its income. (3) Incorporating increasing credit also reinforces the hypothesis that default must be substantially driven by shocks outside of the consumer’s control rather than strategic default (Livshits et al.

[1] Fulford and Schuh (2015) show how to use a simpler consumption model to form econometric predictions of the the unobserved revolving and convenience components in credit bureau data.

2007). Because credit limits are increasing over the life cycle, the incentive to voluntarily run up a large balance and default is increasing, while default after an expenditure shock is decreasing because credit is more available. Thus, if voluntary default is important, the frequency of default should be increasing over the life cycle rather than decreasing after age 30 as observed in credit bureau data.

Out of sample, the simulated model replicates important facets of consumption and credit use. It matches the qualitative smoothness of credit utilization during the Financial Crisis and Great Recession. At the micro-level, simulations produce estimates of the individual relationship between credit and debt that closely match reduced-form estimates from credit bureau data (Fulford and Schuh 2015). Combining the estimated moments for payment choice and life-cycle consumption with out-of-sample predictions at the micro and macro levels, the estimated model thus fits observed data at every frequency: very short-term (payment choice), quarterly (consumption and savings), life-cycle (accumulation of assets and liabilities), and even the business cycle (aggregate changes during the Great Recession).

The relative success of the estimated model suggests implications for policy. The simulated consumption response to a small unexpected cash rebate is about 23 percent within a quarter—very close to estimates based on tax rebates (Parker et al. 2013). Such a strong response is puzzling (Kaplan and Violante 2014). In our approach, it is driven by an large impatient population which is consistent with recent estimates of heterogeneous responses by Parker (2017). Perhaps even more important for policy, the simulated consumption response to an unexpected increase in credit is nearly as large as a cash rebate because so much of household liquidity comes from credit. The observed decline in credit during 2008-2009 may account for one-quarter of the fall in consumption during the Great Recession. Moreover, we show that the more that declines in credit are concentrated among the high-utilization population—those who are often the highest risk for banks—the larger is the consumption response. This result provides additional support for the emphasis in Guerrieri and Lorenzoni (2017) on the relationship between consumer credit, precaution, and the macroeconomy. Policy makers may benefit from considering the importance of consumer credit supply and heterogeneity of its use as a complement to existing conventional counter-cyclical policies.

Our work complements the literature finding preference heterogeneity necessary to explain aggregate consumption and wealth. It extends and enriches the basic idea of Campbell and Mankiw (1989, 1990), who explained aggregate income and consumption with two representative consumers of similar populations that we identify as revolvers and convenience users. Heterogeneous preferences also seem necessary to match wealth inequality (Krusell and Smith 1998, Iacoviello 2008, De Nardi and Fella 2017); portfolio choice over the life-cycle (Calvet et al. 2019); the average marginal propensity to consume (Carroll et al. 2017); heterogeneous marginal propensities to consume (Parker 2017, Gelman 2020); persistent financial distress (Athreya et al. 2017); housing finance and bequests (Favilukis et al. 2017); simultaneous holdings of liquid assets and credit card debt (Gorbachev and Luengo-Prado forthcoming); and experimentally elicited preferences (for example, Andreoni and Sprenger (2012)). Our econometric estimates correspond well with the largest two groups of calibrated heterogeneous agents assumed by Aguiar et al. (2020) to explain hand-to-mouth behavior. Closest to our work is Athreya et al. (2017) who use use a related structural model applied to credit bureau data to study the less than 10 percent of consumers in persistent distress. Much as we use preference heterogeneity to explain persistent high and low utilization, they use preference heterogeneity to explain persistent financial distress. Although not our focus, these consumers are part of our broader class of impatient consumers who revolve high-interest debt and sometimes default.

Our contribution distinguishes itself from the literature several ways. The model permits feasible estimation of *all* parameters, rather than full or partial calibration, including the heterogeneous discount factors. Unlike much prior research, we use credit bureau data that encompasses all credit card debt for the universe of consumers. We appear 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. In contrast to our focus on the heterogeneity and stability of credit use, Laibson et al. (2003) explain aggregate life-cycle accumulation of both credit card debt and illiquid saving by calibrating a model in which agents have hyperbolic preferences. Our approach models the impatient consumers as time-consistent, but our results do not rule out that some consumers could have hyperbolic preferences (Laibson et al. 2003, 2007, Meier and Sprenger 2010) and our model could be extended to include them. Yet because half of consumers do not revolve, at least some consumers must have standard exponential discounting to fit the aggregate data.

Our approach bridges the monetary and payment choice literatures, and ties them closer to the consumption literature. The parsimonious payment specification captures many different reasons consumers might choose one payment means over another including: non-pecuniary preferences and pecuniary credit card rewards (Koulayev et al. 2016, Wakamori and Welte 2017), an ordering of accounts by interest rate cost (Alvarez and Lippi 2017), and the costs of non-acceptance of the preferred method at the point of sale (Telyukova and Wright 2008, Telyukova 2013). Our model allows within-period co-holding of credit card debt and liquid savings because both are used for payments within period (Telyukova and Wright 2008), but for tractability we abstract from cross-period co-holding puzzles (Gross and Souleles 2002, Fulford 2015, Gorbachev and Luengo-Prado forthcoming) and long-term illiquid asset and debt puzzles (Laibson et al. 2003). Because it bridges the payment and consumption literatures, our structural model could be used to evaluate credit card use and consumer decision making for daily individual transactions from linked-account data.[2]

2 Data

This section briefly introduces our main data sources. Fulford and Schuh (2015) provide a more detailed discussion and additional descriptive statistics, including additional evidence on the distribution of credit and on credit card holding by age.

Our main data source is the Equifax/Federal Reserve Bank of New York Consumer Credit Panel (CCP) which contains a quarterly 5 percent sample of all accounts reported to the credit-reporting

[2]See related work by Gelman et al. (2014), Baker (2018), and Kuchler and Pagel (2018) among others. With more data on individual payment choices, the model could implement the integrated financial accounting framework proposed by Samphantharak et al. (2018) that measures exact cash flows by linking household balance sheets with income statements at the level of individual transactions.

agency Equifax starting in 1999. 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.[3] 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. We limit the sample to include only accounts that have a birth year and that had an open credit card account at some point from 2000–2015. 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. Throughout, we combine all credit cards, giving the complete credit and debt picture.[4]

The main advantage of the CCP is that it provides a uniquely long panel at the individual level with administrative data. But it has two disadvantages that we design our structural model to overcome. First, credit bureau data do not contain information on liquid assets. A key modeling insight is that the credit limit is a liquid asset in a buffer-stock sense (Fulford 2013), so how consumers choose to manage their credit card debt while limits change gives insight into their overall decisions. Second, credit bureau data do not directly distinguish between revolving debt and convenience debt. An important contribution of our approach is that we model both convenience use and revolving use so we can account for the different incentives and dynamics of different kinds of uses observed in the data in our estimation. A separate paper (Fulford and Schuh 2015) takes a non-structural approach to estimate the probability of revolving using a finite mixture model. As a practical matter, because payment credit card use is a monthly flow of less than one fifth (see

[3]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. We use only a 0.1 percent sample for analytical tractability for much of the analysis.

[4]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.

Table 2 discussed in section 5) of one twelfth of yearly expenditures, almost all observed credit card debt is held by consumers who are revolving. But to correctly model the dynamics of debt and heterogeneous uses, we need to also model the convenience credit card debt of consumers with typically low utilization.

We use several other data sources. To estimate our payments model, we use data from the Federal Reserve Bank of Atlanta's Diary of Consumer Payment Choice (DCPC), 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 and Stavins 2017, Schuh 2018). In addition, we estimate life-cycle profiles for: 1) consumption from from the Consumer Expenditure Survey (CE); 2) fraction revolving from the Federal Reserve Bank of Atlanta's Survey of Consumer Payment Choice (SCPC); and 3) bankruptcy rates from the Consumer Financial Protection Bureau's Consumer Credit Panel, which is derived from a 1 in 48 sample of credit bureau data.

3 Credit card use

This section presents the empirical evidence on credit card use.

3.1 Credit and debt over the business cycle

Figure 1 shows how the average U.S. consumer's credit card limit and debt varied significantly from 2000–2014. 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.[5]

[5]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. Note that our econometric estimation captures defaults.

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

Observed from credit bureau | Model prediction given fall in credit limits

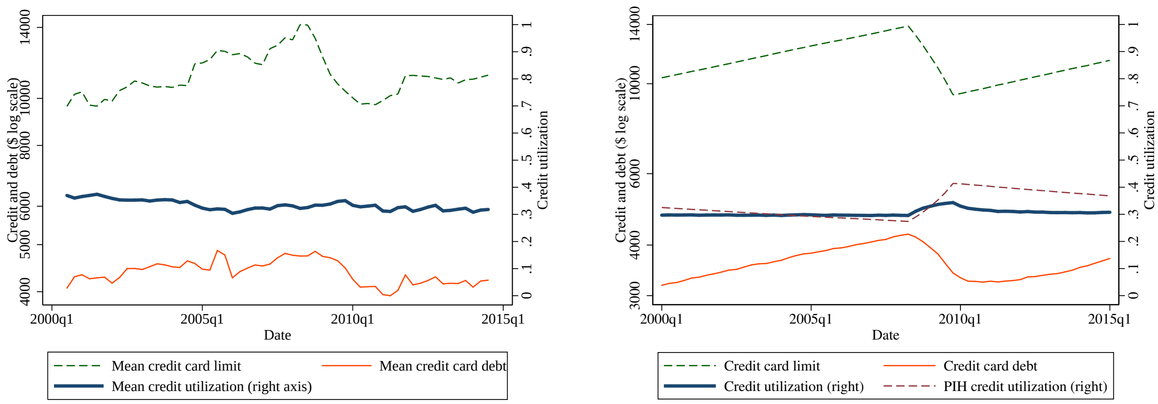

Notes: The left panel shows observed limits, debts, and utilization from credit bureau data (see Section 3 for details). The right panel shows model predictions given an unexpected fall in credit (see section 6 for details). For both panels, the left axis shows the average credit card limits (top line) 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.

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.

3.2 Credit and debt over the life cycle

Figure 2 shows how credit card limits, debt and utilization evolve over the life cycle. In the left panels, 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

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

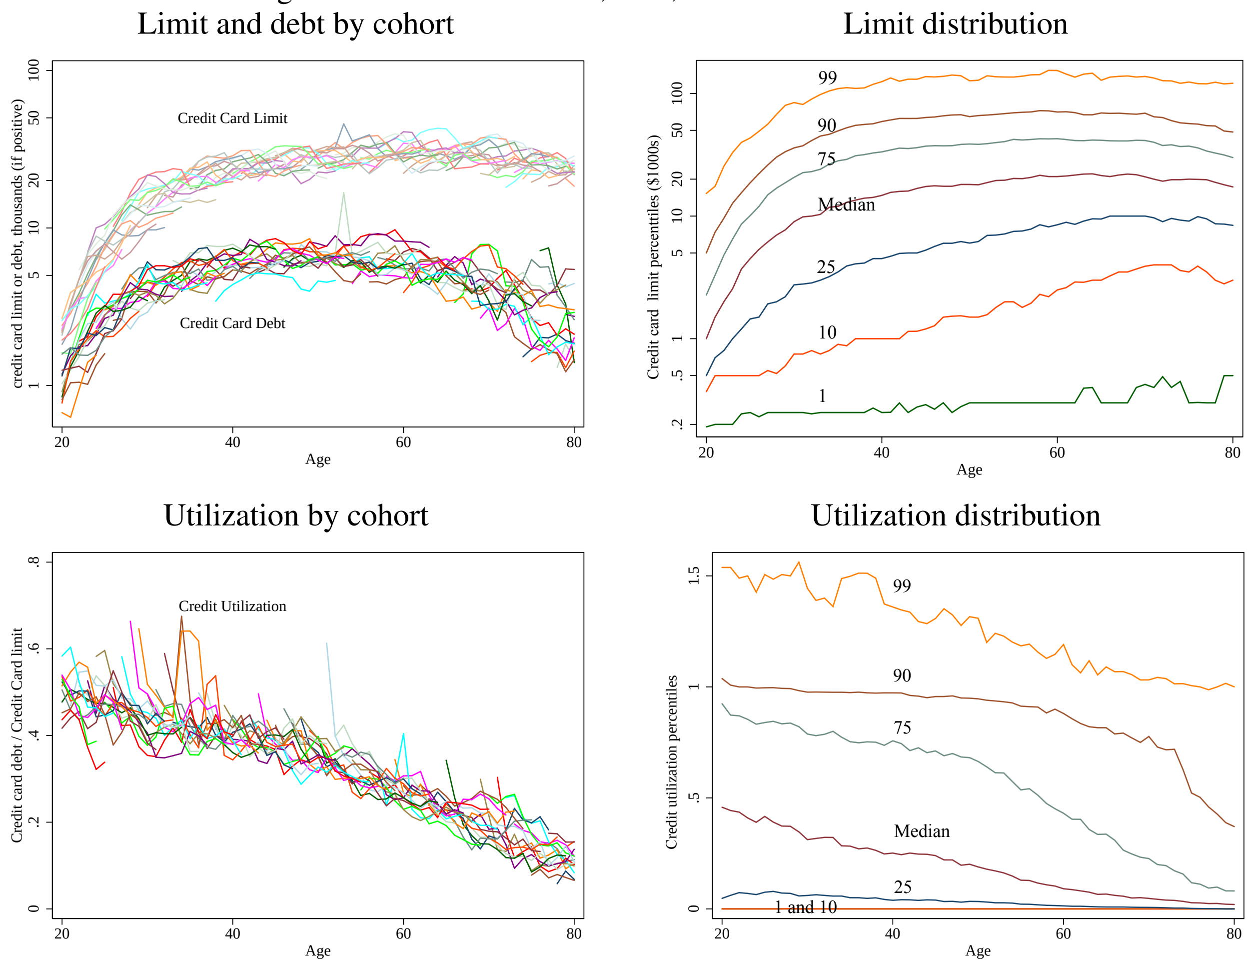

Notes: The left panels show individual birth cohorts. Each line represents the average credit card limit (conditional on being positive, log scale), debt (conditional on being positive, log scale), and utilization (conditional on having a limit, right panel) of one birth year cohort from 1999–2014. The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

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 two panels of Figure 2 show the credit card utilization—credit card debt divided by the credit limit—for each cohort and the distribution of utilization. Consumers with zero debt have zero credit utilization, so they are included in the calculation of utilization but are excluded from mean debt, which includes only positive values.[6] 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 20 percent until around age 70. Moreover, there is substantial and persistent heterogeneity of utilization. More than 10 percent of the population is nearly at its credit limit even past age 70.

3.3 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 here and non-parametric results in Appendix A and Appendix Figure A-4. 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.

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.[7]

[6] 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-1 in the appendix shows the fraction in each cohort who have positive credit and debt. Including the zeros would lower the average credit limit and debt, but makes the life-cycle variation larger.

[7] The combined age, year, and individual fixed effects in equation (1) are not fully identified. 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.

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.

Without fixed effects, credit utilization is very persistent in column 1. Including age and year effects in column 2 barely changes the persistence. The third 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.

Moreover, this individual persistent utilization is highly heterogeneous. As Figure 2 shows, for most of the life cycle, the 25th percentile is using less than 10 percent of available credit, while the 75th percentile is using more than 80 percent. Following people quarterly for 15 years, people who are using more than 60 percent of their credit on average spend 80 percent of the time using more than have 60 percent; people using less than 30 percent, spend more than 90 percent of time using less than 30 percent. In the next section, we describe a model that helps explain this persistent yet heterogeneous utilization.

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

To explain the observations in the previous section, this section describes a life-cycle consumption model that builds on the models in 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.[8]

4.1 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

[8]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). Fulford and Schuh (2015) show that the reduced-form relationship between credit card limits and debts explored in Section 3.3 does not seem to change based on whether someone has a mortgage. Households may provide insurance across members (Blundell et al. 2008) and across generations.

utility:

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

$$\begin{array}{ll}
C_{is} = \nu_{is}(1 - f_{is}\phi_s^c)X_{is} & \text{(Consumption from expenditures)} \\
X_{is} \leq W_{is} & \text{(Expenditures limited by liquidity)} \\
W_{is} = R_{i,s}A_{i,s-1} + Y_{is} + B_{is} - K_{is} & \text{(Evolution of liquidity)} \\
A_{i,s-1} = W_{i,s-1} - B_{is-1} - X_{is-1} & \text{(Relationship between liquidity and assets)} \\
\nu_{is} = \nu(\pi_{is}; A_{i,s-1}) & \text{(Payment decision)} \\
f_{is} = f(F_{is}, W_{is}) & \text{(Default decision)} \\
F_{is} = H(F_{i,s-1}, f_{i,s-1}) & \text{(Evolution of default state)}
\end{array}$$

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. We assume that period utility displays Constant Relative Risk Aversion (CRRA) and allow the risk aversion parameter γ_j to vary across types. Appendix B.2 discusses how to rewrite the consumer's problem recursively in terms of the normalized state variable w_t and thus write the solution of the consumer's normalized recursive problem as an age-specific expenditure/consumption function $x_t(w_{it}, a_{i,t-1}, F_{it})$.

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 that may be paid off. To understand credit card debt, we must account for this payment or "convenience" use as well as the revolving-debt use of credit cards. We model the 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.

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 pays a cost or receives some possibly non-pecuniary benefit when using each method. 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 continuous and strictly monotonically decreasing function, we have simplified the consumer's decision from which option to use for every iota of expenditures to finding the optimal fraction of expenditures π^* to use a credit card for, 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. 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.[9] Assuming the loss of float is the only factor explaining different usage, the cost function for revolvers shifts down by $(R_B - 1)/24$.

[9] 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.

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

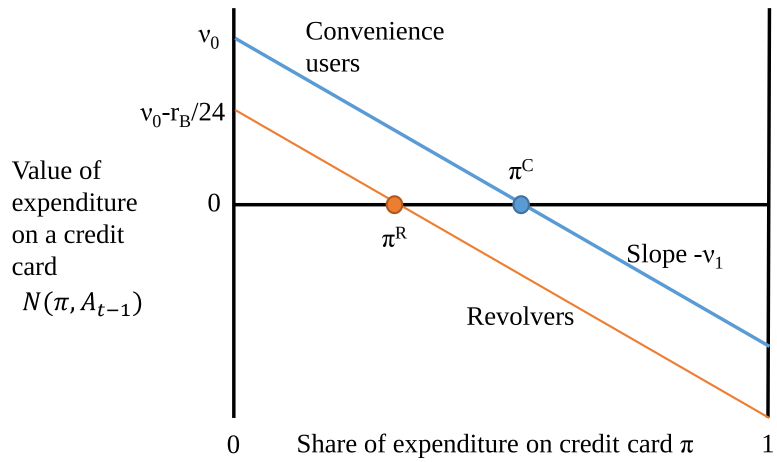

Notes: This figure shows the value or cost of expenditure on a credit card at each expenditure share π relative to cash. The top line is for convenience users who put an optimal share π^c of consumption on a credit card. The bottom line for revolvers is shifted down by the amount $-r_B/24$, because revolvers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

Figure 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 5. 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, ν^C, and revolvers, ν^R. 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. Appendix D goes through the algebra of exact expressions for ν^C, ν^R, 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 time 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.[10] 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.

[10]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.

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 credit limits across individuals of the same age is also large (see 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 and b_f is the fraction that the consumer can borrow in the default state ($F_{it} = 1$). 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.[11]

Some consumers may have the ability to increase their credit limits relatively quickly by applying for more credit. For these consumers, the observed limit is only a short term constraint. Yet the evidence in section 3.3 makes clear that credit and debt co-move and many consumers are close to their limits, so credit limits appear binding (either directly or in expectation) on average. Because low credit utilization is one of the factors that allows consumers to successfully apply for more credit, the consumers for whom the limit is not strict are the same consumers who are far from the the limit binding. Put a different way, the more credit is used, the harder it is to increase the limit, so the the limit is only soft when a consumer is far from the limit and generally using credit solely for convenience. For this reason, although we model the limit as strict for everyone for parsimony, little is lost through this simplification.

[11] 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.

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 $(1 - \phi^c)$. 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 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 $\phi^y_{it} = \phi^y(R^B - 1)b_t P_{it}$ 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 $\phi^c = \phi^y$ 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_t^{Default}.

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.[12]

Late in life, consumers may face income and expenses different from those they face during working years. Labor income may drop, but consumers may start claiming illiquid retirement benefits such as pensions and Social Security, and they may derive income from other illiquid assets such as housing. They may also face an increase in necessary expenses from additional medical care or other needs. We summarize all of these changes by assuming that income starting at T^{Ret} is a fraction λ_1^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

[12]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 5.4.

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 to match the data 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 yet helps to capture the within year consequences of hitting a budget constraint. Because of data and computational constraints, much of the structural consumption literature has been limited to examining decisions made at a yearly frequency. We adjust convenience credit card debt appropriately so that it represents only one month of expenditure when we estimate the model.

4.2 The consumer's decision

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 give a detailed discussion in Appendix B.3. We follow the method of endogenous gridpoints (Carroll 2006), which 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.

Figure 4 illustrates some of the complexities of the decision problem. The consumption functions then show how much a consumer at that age with those preferences will consume at each liquidity. Because credit limits also scale with permanent income, only age, default status, previous borrowing, and the current liquidity ratio enter the consumption decision. 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 as examined by 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

Figure 4: Expenditure functions over the life cycle with borrowing

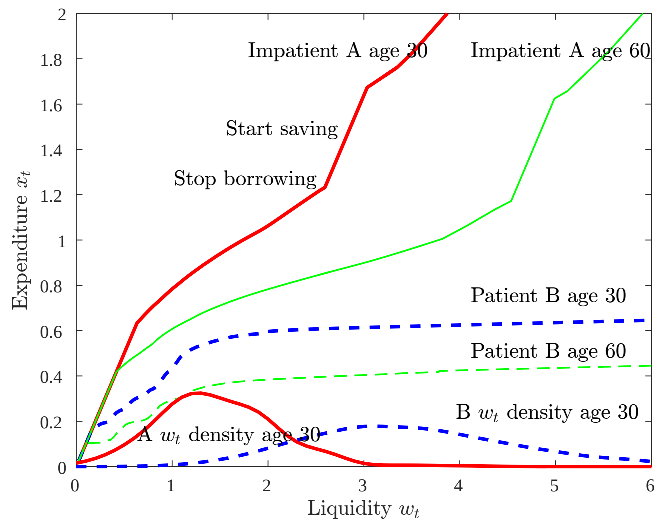

Notes: This figure uses the estimates in Table 3 column 1 at age 30 and age 60 to show the quarterly expenditure function for impatient (A) consumers and patient (B) consumers. Liquidity w_t is a multiple of quarterly permanent income P_t and includes available credit. The densities for liquidity are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

approaching zero from below, uses the borrowing rate R^B, and the other uses the savings rate R. 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.[13]

5 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

[13]The standard method of endogenous gridpoints breaks down when there is a discontinuity in the value function at the default point. We therefore use a modified version that forms endogenous gridpoints on either side of the discontinuity and then enters a successive approximation around the discontinuity from above and below to choose gridpoints that capture the discontinuity correctly.

of credit cards for revolvers ν^R and convenience users ν^C in Section 5.1. 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 arbitrary 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.[14] For each sub-population, the entire decision is described by four parameters: the discount rate β, the coefficient of relative risk aversion γ, the initial wealth-to-income ratio λ_0, and the fraction of permanent labor income expected from illiquid assets such as housing, pensions, or Social Security in late life λ_1. Finally, we estimate the probability (p^k) and cost (k) of expenditure shocks. We show that the default cost parameter

[14]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 $f^B = 1 - f^A$.

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). Appendix C gives additional details. Briefly, given set of parameters $\theta \in \Theta$ and first-stage parameters χ, we solve the consumer's problem and then simulate the life-cycle decisions for a large population of consumers. We then minimize the weighted sum of square differences between the empirical and simulated life-cycle moments for the population. Our standard weighting matrix is block proportional to the inverse variance of the empirical moments (the optimal weighting matrix with no first-stage correction). We also show results using the "optimal" weighting matrix, 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.

5.1 Estimation and identification of the payments model

Because of the structure of the consumer's problem, whether the consumer was revolving as of the previous period is the only way the consumption decision influences the payment decision. We can thus find the solution to the payments problem first and then allow the solution to the payments problem to influence the consumption problem. Table 2 shows the fraction of all expenditures over a three-day period that the nationally representative sample of consumers from the Diary of Consumer Payment Choice puts on a credit card. The average consumer pays for 17.2 percent of expenditure with a credit card. Revolvers pay for slightly less at 15.6 percent, and convenience users pay for slightly more at 18.2 percent.[15]

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.2 for sources).

[15]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.

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.

The model 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.[16] As a comparison, the fees that banks charge merchants for processing credit cards are roughly $60 billion per year.[17] 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.

[16] 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, 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.

[17] 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.

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.

5.2 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 3, (2) log mean household consumption over the life cycle from the CE from 2000–2014,[18] (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.[19] We show each of these moments in Figure 5 together with their estimates from the model, and have already discussed the debt profile in Section 3.2. Consumption follows the characteristic hump shape (Gourinchas and Parker 2002, Attanasio et al. 1999). Bankruptcy is increasing early in the life-cycle, before declining. Appendix C.3 discusses the construction of the variance-covariance matrix of the combined moments.

We briefly describe the sources and estimates from other data sets that identify the ancillary parameters of the model, providing greater detail in in Appendix C.2. We estimate the average life cycle of income growth (G_t) using the Consumer Expenditure Survey to match our consumption data, adjusting for aggregate growth. s. We use the estimates of income shocks from Gourinchas and Parker (2002), which are updates of Carroll and Samwick (1997), calculated from the Panel

[18] 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.

[19] 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 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.

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. For the other parameters and prices, 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 interest rates for the geometric average inflation rate from 2000–2015 of 2.15 percent.

5.3 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. Because this is a nonlinear model, all moments are typically used to identify all parameters. Appendix C.5 provides a discussion of how different sources of variation help identify the parameters. Table 3 shows the model estimates, while Figure 5 shows how debt, consumption, 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 an income about average (ζ^A).[20] Because individual

[20]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.

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.

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

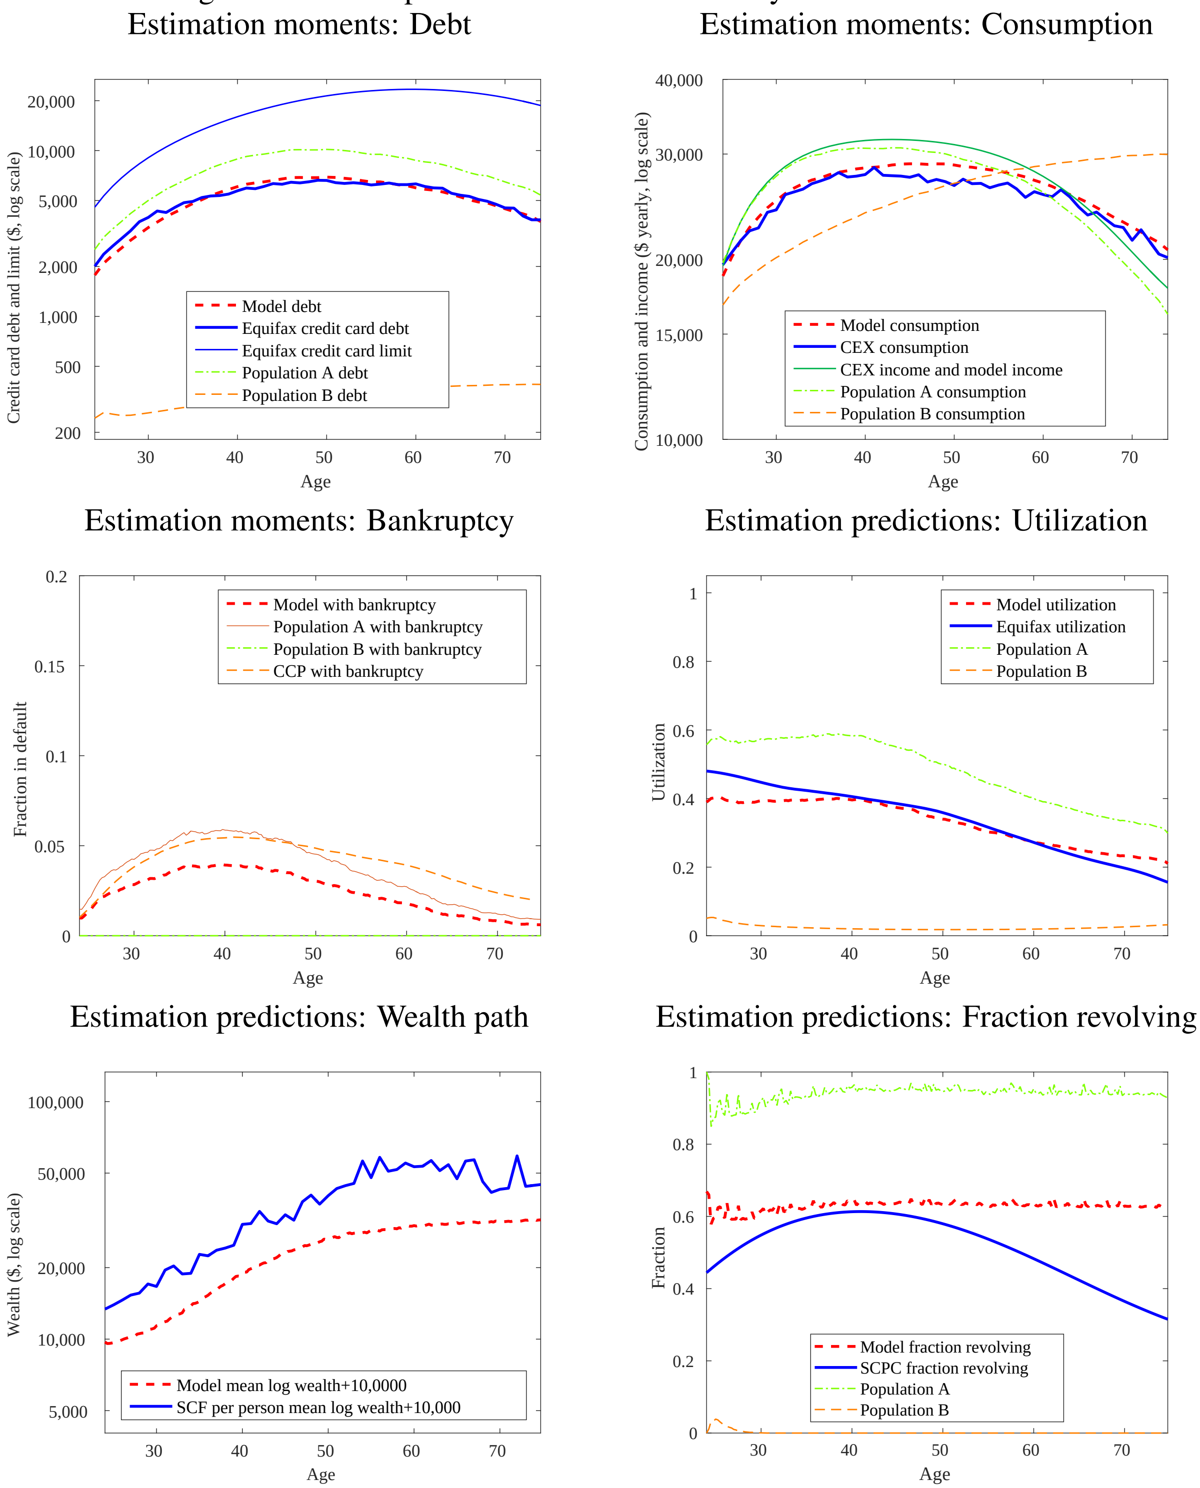

Notes: Life-cycle paths from simulated population using the estimates in column 1 of Table 3.

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 fourth 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.[21] 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.

Our model allows both strategic and involuntary default, but the data suggest that only involuntary default is important in this framework. The cost of default is identified only up to an inequality. The framework suggests default cost parameters allowing strategic default give a much worse fit than default costs parameters in which all default is involuntary. As Appendix Figure A-8 shows, holding other parameters fixed at their values in column 1, changing the default cost parameter does not improves the fit, unless it is below a threshold for strategic default, when the model fit rapidly deteriorates.

Rising credit limits over the the life cycle explain why the data push against strategic default. Appendix Figure A-9 shows default over the life-cycle with the standard default cost parameter and one below the strategic default threshold. As credit limits increase and the remaining expected life decreases, the gains from default get ever larger. At approximately age 35, much of the impatient population finds it better to default and the default rate changes precipitously. Increasing the cost

[21]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.

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, strategic default is not useful for explaining the fraction in default within our simple default framework. Note that the model ties the default cost to the credit limit, so the costs of default are increasing over the life-cycle. Not doing so makes default even more likely at older ages. Whether a richer model could allow for strategic default and increasing credit limits over the life cycle we leave for future research.

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

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.[22] 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.

[22]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.

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.

5.4 Robustness and variations

In this section, we briefly examine the robustness of the estimates to changes in weighting matrices, starting points of the estimation, and model choices. Appendix C.6 offers additional details. The general conclusion is that while particular parameters are sensitive to estimation and model choices, our overall conclusions are not. Our overall conclusions are also robust to alternative starting points for estimation. Table 3 shows the over-identification statistic for each estimation, which always decisively rejects the hypothesis that the model is not over-identified.[23] 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.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 last two columns of Table 3 examine how changing the model changes estimates. Allowing consumers to take into account how their their payments decisions will affect their consumption decisions does not not appear to be important in column 3. Similarly, substantially increasing the bequest motive in column 4 barely changes the estimates. Appendix C.6 provides additional

[23] 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.

discussion.

6 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. After showing it has substantial success out of sample, we explore the estimated model's implications for stimulus policy.

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 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 right 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 spread evenly among the population., Even so, our model produces a notably smoother path than a simple version of the Life Cycle/Permanent Income Hypothesis (LC/PIH) would suggest.[24]

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. Figure 6 shows the relative paths of consumption from our simulations and detrended real personal consumption per person from the BEA. 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. The sim-

[24]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). 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.

Figure 6: Consumption over the business cycle

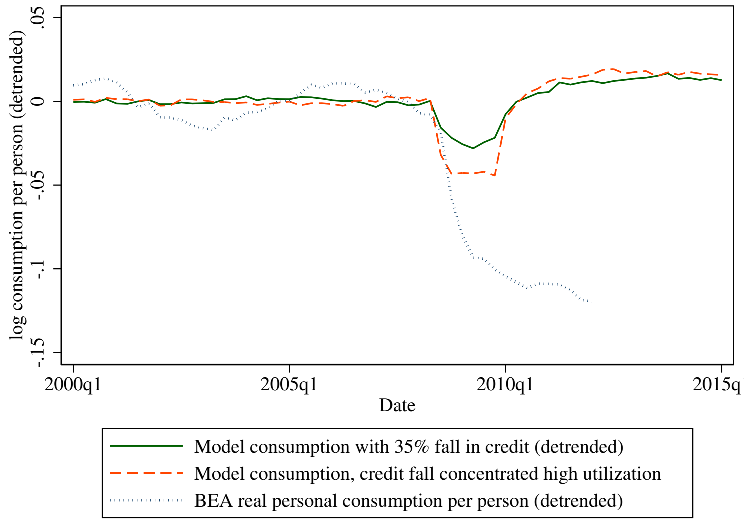

Notes: This figure shows personal consumption from the BEA and average consumption from simulations with a 35 percent fall in available credit starting in 2008q3. Each series is detrended using the 2000–2008 period. When the credit fall is concentrated among the high utilization, only consumers in this population have a fall in credit, but the aggregate fall in credit credit is held constant. The BEA consumption series continues to fall after 2012 relative to the 2000-2008 trend. We omit the continuing fall to focus on the impact of consumer credit changes.

ulations based on our estimated model 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 is also quite rapid initially, matching the BEA series well as high utilization consumers are pushed to deleverage and convenience users reduce their consumption to build up their buffer. 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. Other features not captured by our estimated model of consumer decision making must be important.[25]

The policy implications are even more striking when we instead assume that the fall in credit was concentrated among the high-utilization consumers. These consumers are often the highest

[25]Note that consumption from the simulations is actually higher after several years, because debt is lower, so interest payments decline. This initial decrease followed by a higher steady state is a general feature of credit changes in precautionary models (Fulford 2013). Because credit card interest rates were relatively steady over the period (the lack of response of credit card rates was noted in earlier work by Ausubel (1991)), our estimated model captures the consumption response to the fall in credit, but misses production and savings responses which we are not modeling. Guerrieri and Lorenzoni (2017) examine the interaction between credit contractions and precautionary preferences in general equilibrium with plausible, but not estimated, preferences.

risk, so are the most likely to be targeted for limit cuts when banks want to reduce risk. If the fall in credit had been concentrated among these consumers, it would have explained nearly half of the consumption fall during the Great Recession, yet utilization would have been just as smooth. This heterogeneity in response is thus a central feature in understanding the impact of both monetary and fiscal policy, a topic we turn to next.

6.1 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 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. 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 sug-

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.

gest 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).

7 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, 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.

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. 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 is 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 data on so many dimensions, including the impact of a cash infusion (Kaplan and Violante 2014, Parker et al. 2013). The implications of the heterogeneity of credit use we document for counter-cyclical policy are also important. The more that banks reduce risk or are encouraged to reduce risk by not extending or reducing credit among high-utilization customers, the larger the consumption impact of a credit crunch. Conversely, credit increases or other cash infusions targeted to the highest utilization consumers have an especially large impact.

An important unanswered question for future research is the source and nature of the impatience found in the high discount rate population. Certainly, behavioral economic approaches such as quasi-hyperbolic discounting and present bias could be consistent with our finding of a high discount rates for some, though not all, consumers. Our results do not rule out such approaches, but also show they are not necessary to explain credit card use.

References

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. 2017. "Do Banks Pass through Credit Expansions to Consumers Who want to Borrow?" *The Quarterly Journal of Economics* 133(1): 129–190.

Aguiar, Mark A, Mark Bils, and Corina Boar. 2020. "Who Are the Hand-to-Mouth?" Working Paper 26643, National Bureau of Economic Research.

Alvarez, Fernando, and Francesco Lippi. 2017. "Cash burns: An inventory model with a cash-credit choice." *Journal of Monetary Economics* 90: 99 – 112.

Andreoni, James, and Charles Sprenger. 2012. "Estimating Time Preferences from Convex Budgets." *American Economic Review* 102(7): 3333–56.

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.

Ausubel, Lawrence M. 1991. "The Failure of Competition in the Credit Card Market." *The American Economic Review* 81(1): 50–81.

Aydin, Deniz. 2015. "The Marginal Propensity to Consume out of Liquidity: Evidence from a Randomized Controlled Trial." 2015 Meeting Papers 270, SED.

Baker, Scott R. 2018. "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data." *Journal of Political Economy* 126(4): 1504–1557.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98(5): 1887–1921.

Cagetti, Marco. 2003. "Wealth Accumulation over the Life Cycle and Precautionary Savings." *Journal of Business and Economic Statistics* 21(3): 339–353.

Calvet, Laurent E, John Y Campbell, Francisco J Gomes, and Paolo Sodini. 2019. "The cross-section of household preferences."Technical report, Working Paper, Harvard University.

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. 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* 8(3): 977–1020.

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, and Giulio Fella. 2017. "Saving and wealth inequality." *Review of Economic Dynamics* 26: 280–300.

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.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh. 2017. "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium." *Journal of Political Economy* 125(1): 140–223.

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.

Gelman, Michael. 2020. "What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs persistent characteristics." *Journal of Monetary Economics*.

Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. 2014. "Harnessing naturally occurring data to measure the response of spending to income." *Science* 345(6193): 212–215.

Gorbachev, Olga, and María José Luengo-Prado. forthcoming. "The Credit Card Debt Puzzle: The Role of Preferences, Credit Access Risk, and Financial Literacy." *The Review of Economics and Statistics*.

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

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

Guerrieri, Veronica, and Guido Lorenzoni. 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *The Quarterly Journal of Economics* 132(3): 1427–1467.

Iacoviello, Matteo. 2008. "Household Debt and Income Inequality, 1963-2003." *Journal of Money, Credit and Banking* 40(5): 929–965.

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.

Koulayev, Sergei, Marc Rysman, Scott Schuh, and Joanna Stavins. 2016. "Explaining adoption and use of payment instruments by US consumers." *The RAND Journal of Economics* 47(2): 293–325.

Krusell, Per, and Anthony A. Smith, Jr. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *The Journal of Political Economy* 106(5): 867–896.

Kuchler, Theresa, and Michaela Pagel. 2018. "Sticking to Your Plan: The Role of Present Bias for Credit Card Paydown." Working Paper 24881, National Bureau of Economic Research.

Laibson, David, Andrea Repetto, and Jeremy Tobacman. 2003. "A Debt Puzzle." In *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps.*: P.

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.

Meier, Stephan, and Charles Sprenger. 2010. "Present-Biased Preferences and Credit Card Borrowing." *American Economic Journal: Applied Economics* 2(1): 193–210.

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.

Samphantharak, Krislert, Scott Schuh, and Robert M. Townsend. 2018. "Integrated Household Surveys: An Assessment of U.S. Methods and an Innovation." *Economic Inquirty* 56(1): 50–80.

Schuh, Scott. 2018. "Measuring Consumer Expenditures with Payment Diaries." *Economic Inquiry* 56(1): 13–49.

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.

Wakamori, Naoki, and Angelika Welte. 2017. "Why do shoppers use cash? Evidence from shopping diary data." *Journal of Money, Credit and Banking* 49(1): 115–169.