

DO LIQUIDITY CONSTRAINTS AND INTEREST RATES MATTER FOR CONSUMER BEHAVIOR? EVIDENCE FROM CREDIT CARD DATA*

DAVID B. GROSS AND NICHOLAS S. SOULELES

This paper utilizes a unique data set of credit card accounts to analyze how people respond to credit supply. Increases in credit limits generate an immediate and significant rise in debt, counter to the Permanent-Income Hypothesis. The “MPC out of liquidity” is largest for people starting near their limit, consistent with binding liquidity constraints. However, the MPC is significant even for people starting well below their limit, consistent with precautionary models. Nonetheless, there are other results that conventional models cannot easily explain, for example, why so many people are borrowing on their credit cards, and simultaneously holding low yielding assets. The long-run elasticity of debt to the interest rate is approximately -1.3 , less than half of which represents balance-shifting across cards.

I. INTRODUCTION

The canonical Permanent-Income Hypothesis (PIH) assumes that consumers have certainty-equivalent preferences and do not face any liquidity constraints. Under these assumptions the marginal propensity to consume (MPC) out of liquid wealth depends on model parameters, but generally averages less than 0.1. The MPC out of predictable income or “liquidity” (e.g., increases in credit limits), which do not entail wealth effects, should be zero. The leading alternative view of the world is that liquidity constraints are pervasive. Even when they do not currently bind, they can be reinforced by precautionary motives concerning the possibility that they bind in the future. Under this view the MPC out of liquidity can equal one over a range of levels for “cash-on-hand,” defined to include available credit [Deaton 1991; Carroll 1992; Ludvigson 1999].

* A previous version of this paper circulated under the title “Consumer Response to Changes in Credit Supply.” We would like to thank the editors and two anonymous referees, Lawrence Ausubel, Paul Calem, Simon Gilchrist, Kathleen Johnson, Anil Kashyap, Anthony Santomero, and seminar participants at Yale University, the Universities of Maryland, Wisconsin, Michigan, and Chicago, the NBER monetary economics meeting, the Chicago, Kansas City, Philadelphia, Richmond, and New York Federal Reserve Banks, the ASSA meetings, Carnegie Mellon University, Columbia University, the NBER Summer Institute, and various workshops at the University of Chicago Graduate School of Business and The Wharton School of the University of Pennsylvania. We are grateful to the Wharton Financial Institutions Center and several credit card issuers for numerous discussions and assistance in acquiring the data. All remaining errors are our own.

© 2002 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology.

The Quarterly Journal of Economics, February 2002

Liquidity constraints have implications in many areas of economics, in addition to consumption theory. For instance, they can amplify the effects of fiscal policy and other business cycle shocks [Hubbard and Judd 1986]. While they increase saving, they can retard investment in human capital and thus slow growth [Jappelli and Pagano 1999]. Nonetheless, there is still no agreement about the quantitative importance of liquidity constraints and precautionary motives in practice [Browning and Lusardi 1996]. Although some micro studies have found evidence that liquidity constraints distort consumption (e.g., Hall and Mishkin [1982] and Zeldes [1989]), other studies have not (e.g., Altonji and Siow [1987] and Runkle [1991]). Part of this disagreement is due to the difficulty identifying which households in the data are in fact constrained. Most studies split the sample on the basis of net worth, but this conflates credit demand and supply. The fact that someone has low or negative net worth does not imply that he cannot increase his borrowing [Jappelli 1990].

Under the PIH the elasticities of consumption and saving to interest rates also depend on model parameters, such as the intertemporal elasticity of substitution. Like liquidity constraints, these elasticities have wide-ranging implications, for instance for monetary policy, business cycles [King, Plosser, and Rebelo 1988], and tax incentives for saving. Most studies have found small effects of interest rates on consumption and saving (e.g., Hall [1988]). However, it remains unclear whether interest-rate elasticities are truly small, or whether these findings are spurious, due, for instance, to measurement problems like the difficulty of observing household-specific interest rates [Browning and Lusardi 1996; Mishkin 1995].¹

To test whether liquidity constraints and interest rates really matter in practice, this paper uses a unique new data set containing a panel of thousands of individual credit card accounts from several different card issuers. The data set is of very high quality. It includes essentially everything that the issuers know about their accounts, including information from people's credit applications, monthly statements, and credit bureau reports. In particular, it separately records credit limits and credit balances, allowing us to distinguish credit supply and demand, as well as

1. Ausubel [1991]'s "underestimation hypothesis" tries to explain why consumers might be insensitive to credit card rates in particular, by assuming that consumers underestimate the probability they will borrow on their cards or that the costs of switching to other cards are large. (See also Calem and Mester [1995].)

account-specific interest rates. These data allow us to analyze the response of debt to changes in credit limits and thereby estimate the MPC out of liquidity, both on average and across different types of consumers. The analysis generates clean tests distinguishing the PIH, liquidity constraints, precautionary saving, and behavioral models of consumption, which are outlined in the next section. We also estimate the sensitivity of debt to interest rates, and the interaction of interest rates with liquidity constraints.

Credit cards play an important role in consumer finances, so they are a good place to look for the effects of liquidity constraints and interest rates. About 20 percent of aggregate personal consumption is already being purchased using credit cards [Chimerine 1997] and with the growth of e-commerce this fraction is likely to grow. Moreover, for most households credit cards, in particular bankcards (i.e., Visa, Mastercard, Discover, and Optima cards), represent the leading source of unsecured credit. About two-thirds of households have at least one bankcard, and of these households *at least* 56 percent—a remarkably large fraction—are borrowing on their bankcards, that is, paying interest, not just transacting [1995 Survey of Consumer Finances (SCF)].² Conditional on borrowing, the median bankcard account is borrowing over \$2000, with about another \$5000 of balances on other cards (authors' calculations; see also Yoo [1998]). These are large magnitudes in the context of typical household balance sheets. They are also large in the aggregate: total credit card borrowing amounts to about \$500B [Board of Governors of the Federal Reserve System 1998].

To preview the results, we find that increases in credit limits generate an immediate and significant rise in debt, counter to the

2. This figure probably understates the fraction of households borrowing on their bankcards, because it is computed as the fraction of SCF households with positive debt on their bankcards, but SCF households substantially underreport their bankcard debt. This underreporting can be seen by comparison with our account data (described below) or the aggregate data on revolving consumer credit collected by the Board of Governors of the Federal Reserve System, which is mostly credit card debt (including debt on retail-store cards). The 1995 SCF records average credit card debt (including retail-store debt) of around \$2000 per household with credit cards. By contrast, allocating the approximately \$400B of aggregate revolving debt in 1995 (adjusted for transactions balances) evenly across the three-quarters of households with credit cards (including retail-store cards) yields over \$5000 of debt per household in 1995. In our account data about 52 percent of all bankcard accounts revolve, but this includes accounts that are not currently being used to borrow or transact. Since the average cardholder has multiple cards, this understates the fraction of *individuals* borrowing. Conditional on an account currently being used, about 70 percent of accounts borrow.

PIH. The average “MPC out of liquidity” ($d\text{Debt}/d\text{Limit}$) ranges between 10–14 percent. The MPC is much larger for people starting near their credit limit, providing concrete evidence that liquidity constraints are often binding. However, the MPC is significant even for people starting well below their limit. We show that this response is consistent with buffer-stock models of precautionary saving. Nonetheless, there are other results that conventional models cannot easily explain, such as the fact that many credit card borrowers simultaneously hold low-yielding assets. Unlike most other studies, we also find strong effects from changes in account-specific interest rates. The average long-run elasticity of debt to the interest rate is approximately -1.3 . Less than half of this elasticity represents balance-shifting across cards, with most reflecting net changes in total borrowing. The elasticity is larger for decreases in interest rates than for increases, which can explain the widespread use of temporary promotional rates. The elasticity is smaller for people starting near their credit limits, again consistent with liquidity constraints.

Section II outlines our framework for testing for liquidity constraints and other alternatives to the PIH, and for controlling for the endogeneity of credit supply. Section III describes the data, and Section IV develops the econometric methodology. Sections V and VI estimate the average response of debt to changes in credit supply, credit limits and interest rates, respectively. Section VII analyzes the heterogeneity in people’s responses to credit, focusing on the ability of liquidity constraints and other models to explain the heterogeneity. Section VIII concludes, and is followed by a Data Appendix.

II. FRAMEWORK

1. Liquidity Constraints and Other Models

Because bankcards are the marginal source of credit for most households, they can be used to measure the pervasiveness of liquidity constraints. Considering the usefulness of bankcards for borrowing and transacting, it is plausible that many of the one-third of households without bankcards are liquidity constrained. Jappelli, Pischke, and Souleles [1998] found that their consumption is disproportionately sensitive to their income; i.e., people with bankcards are better able to smooth their consumption than people without bankcards. Of the two-thirds of households with

bankcards, the over 56 percent who are borrowing and so are paying high interest rates (averaging around 16 percent) might also be considered liquidity constrained, lacking access to cheaper credit. Combined with the households lacking bankcards, they bring the overall fraction of potentially constrained households to over $\frac{2}{3}$ ($\approx \frac{1}{3} + \frac{2}{3} * .56$). However, this calculation relies on the weaker notion of liquidity constraints as a wedge between borrowing and lending rates. Below we show that household balance sheets are complicated, simultaneously including both assets and liabilities with many different interest rates.

Instead, this paper mostly focuses on the stricter notion of liquidity constraints as quantity constraints. We find that about 14 percent of bankcard accounts have a utilization rate, defined as the balance divided by the credit limit, above 90 percent, leaving less than 10 percent of the credit line free. For these accounts liquidity constraints are arguably binding. However, focusing on just currently binding constraints can understate the full impact of liquidity constraints. In models with precautionary motives liquidity constraints can have large effects on the level of consumption and welfare even when they do not currently bind, as long as there is a possibility that they bind in the future.

The most direct way to test whether liquidity constraints matter is to see whether changes in liquidity have real effects. Accordingly, this paper estimates the MPC out of credit card liquidity, $d\text{Debt}/d\text{Limit}$. When someone's credit limit (credit line) increases, what fraction of that extra liquidity does she use to borrow and spend? Under the PIH, which assumes no liquidity constraints, the answer should be zero. Alternatively, if liquidity constraints exist but only matter when currently binding, the MPC should be positive and large only for people with very high utilization rates. However, if the possibility of future liquidity constraints also matters, for precautionary reasons, the MPC would be positive even at lower levels of utilization. Consider someone with a credit limit of \$10,000 and a balance of \$2500. Does raising the limit to say \$11,000 have any effect on debt?

We also estimate the elasticity of debt to interest rates, which is of independent interest. When someone's credit card rate rises, by how much does she change her debt? We highlight the interaction of interest rates with liquidity constraints. In the absence of constraints, people with substantial debt might be more sensitive to interest rates than people with little or no debt. But binding liquidity constraints would make people less sensitive,

ceteris paribus. Hence we also test for liquidity constraints by examining how interest rate elasticities vary with utilization rates and demographic characteristics.

2. *The Credit-supply Function*

Endogeneity is a generic problem in studies of the effects of credit supply, including monetary policy [Christiano, Eichenbaum, and Evans 1996; Kashyap and Stein 1994]. In our case there could be a problem if credit card issuers increase credit supply when they expect credit demand to rise. Then part of the observed response in debt could be the result of a demand shock, not just a response to supply. However, our data allow us to go further than most previous studies to address the endogeneity of both credit limits and interest rates. We follow two general strategies.

First, we use an unusually rich set of control variables to capture the endogenous part of credit-supply changes. Monthly time dummies control for seasonality and all other aggregate economic effects. First-differencing and fixed individual effects accommodate all persistent characteristics of the account-holders. Indicator variables for whether the line change was requested by the account-holder help control for changes in individual credit demand. Most notably, we also control for the credit-risk scores, which are the issuers' own summary statistics for the default risk and profitability of each account. With millions of accounts to manage, the issuers have largely automated their decision-making, relying very heavily on the scores in deciding credit policy for each account, including credit-limit and interest-rate policy [Moore 1996]. The issuers estimate the scores using all the information at their disposal, both in-house (the "internal" scores) and at the credit bureaus (the "external" scores). Thus, the scores essentially summarize the fundamental characteristics of each account.³ We also control for the most important of these characteristics directly, including account debt levels and account age (time since booking).⁴ Because the data include the account

3. While the scoring functions are proprietary, they are known to heavily reflect a person's past borrowing and delinquency history.

4. While we considered many other account characteristics as well, the reported results focus on these variables based on results in a companion paper. Gross and Souleles [2002] used the same data source to estimate hazard models of default, including bankruptcy. They empirically evaluated the implications of the various account characteristics that the issuers track for account risk and profitability. Apart from the scores, debt levels and being a newly opened account

characteristics tracked by the issuers themselves, this strategy of controlling should be a compelling check against endogeneity.

Second, we also use instrumental variables (IV) to isolate exogenous changes in credit supply. In particular, we exploit exogenous "timing rules" built into the credit-supply functions. For instance, many issuers will not consider (or are less likely to consider) an account for a line change if it has been less than six months or less than one year since the last line change. Hence for a given account the probability of a change is exogenously higher in certain months than in others. Accordingly, we instrument for line changes with dummy variables for the number of months since the latest line change, controlling for account fundamentals. Consider, for example, two accounts opened at the same time that currently have the same credit scores but are on different timing cycles for exogenous reasons. Suppose that one account had its latest line increase twelve months ago and the other had it eleven months ago. Because of the timing rules, the first account is more likely to have its line go up this month, even though there is no fundamental difference between the accounts. Our IV estimates compare the resulting debt of the first account with the debt of the second account. As a result, the estimates exploit only exogenous variation in the timing of a line change. In addition, we can condition on the total number of line increases that each account receives during the sample period. This controls for all the factors justifying the increases in the first place. The remaining variation is purely in the high-frequency timing of the line increases, as induced by the exogenous timing rules, conditional on the increases taking place. Such IV specifications are an unusually powerful response to endogeneity.

We handle interest rates analogously. In particular, credit card issuers also use exogenous timing rules for changing interest rates, which we exploit as instruments. For instance, many issuers offer low promotional rates that expire after a prespecified period like six months or twelve months. Also, for contractual reasons interest rates on fixed rate cards can often be changed only periodically.

Credit supply still varies across accounts with the same fundamental characteristics and timing cycles. Although the de-

played the most significant roles, although not nearly so significant as the scores themselves.

tails of the credit-supply functions are proprietary, and vary across issuers, there are a number of factors generating this remaining variation. First, there is a time-series dimension to credit supply. Issuers can vary the total amount of credit they allocate to their credit card portfolios over time. Such variation is likely endogenous, because aggregate shocks influence both the supply and demand for credit. We use the month dummies to control for this endogeneity.

Second, cross-sectionally, credit-supply policies are functions of both fundamental account characteristics and numerous institutional constraints, in addition to the timing rules, which are arguably exogenous. At a given issuer, different accounts can be considered for line changes at different times, by different managerial divisions, whose operating procedures vary somewhat, even though they all draw upon the same fundamental account information. For example, as a result of mergers, portfolio acquisitions, or other historical reasons, an issuer's accounts might have been divided into different subportfolios, across multiple computer systems. Because of computing and other constraints, in any given month only the accounts in some subportfolios on some systems might be considered for a line change; the other accounts would be considered in subsequent months. Also, some accounts might have been booked under a myriad of different marketing programs, and so be managed by different people changing credit supply at different times. Some of these programs involve randomized experiments, but even the others are unlikely to confound our analysis. Consider, for example, credit cards that were designed to attract Chicago Bears football fans (perhaps by including a picture of the Bears' logo on the face of the cards). Knowing that the person who opened the account is a football fan does not help predict subsequent high-frequency changes in his demand for credit that the issuers use to change credit limits. Hence such institutional variation is exogenous for our purposes, even before instrumenting. Also, first-differencing and individual account dummies directly control for differences across marketing programs and other subportfolios, as well as any other persistent characteristics of the account and account-holder.

Background econometric analysis supports this description of the credit-supply function. Account fundamentals like the scores, debt levels, and account age are statistically very significant in explaining line changes (especially the scores, which have

chi-squared statistics above 1500 in probit models of whether the line changes). Nonetheless, the fundamentals explain a relatively small part of the high-frequency variation in credit limits exploited here. For instance, they account for only about 5 percent of the probability of a line change in any given month, according to the pseudo R^2 from a probit model. Adding as explanatory variables our instruments for the number of months since the latest line change raises this figure to 17 percent. These results confirm the importance of institutional constraints, of which our instruments are but one example. That is, even though our fundamental control variables do not explain all the variation in credit supply, the remaining variation can reasonably be taken as exogenous. Further, our instruments should be powerful guards against any remaining endogeneity. Similarly for interest rates, the fundamentals alone are very significant but explain only about 5 percent of the monthly probability of a rate change; but adding the instruments raises this figure to 33 percent.⁵

III. DATA DESCRIPTION

The data used in this paper are proprietary, coming from the account archives at several anonymous issuers of bankcards. Each issuer provided a representative sample of all their personal bankcard accounts open as of 1995.⁶ Because the issuers include some of the largest credit card companies in the United States, the data should be generally representative of credit cards in the United States in 1995. For computational tractability, this paper uses a large randomized subset of the original data consisting of about twenty-four thousand accounts. These accounts are followed monthly for different periods of time, depending on the issuer, or until they attrite; but on average for just over 24

5. Adding other available account characteristics does not substantially increase these pseudo R^2 figures. Adding month dummies to control for the time-series variation raises them to 27 percent for credit limits and 48 percent for interest rates. Hence this analysis has accounted for about a quarter to a half of the credit-supply functions.

Industry sources confirmed this analysis anecdotally. One even pointed out its implication for monetary policy lags. Although it would be optimal to adjust credit supply simultaneously for all accounts immediately after a change in monetary policy, in practice institutional constraints significantly delay the full adjustment. These constraints can also help explain the "stickiness" of credit card interest rates pointed out by Ausubel [1991].

6. Some accounts may in fact be used for business purposes; however, they are the personal liability of the individual borrower.

months. Thus, the sample is unbalanced, but in total extends from January 1995 through January 1998.

The unit of observation is a credit card account, which generates a single monthly statement no matter how many people use the account, even if multiple cards have been issued for the account. There are four general types of information available. First, the bulk of the data consists of the main information listed on the account's monthly billing statement, including statement totals such as balances, payments, and interest charges, as well as the credit limit and interest rate. The majority of issuers archived this information by the monthly billing cycle for each account, while other issuers stored the data by calendar month. Second, there are also data that the issuers obtain from the credit bureaus, usually at a monthly or quarterly frequency, such as the external credit scores and balances on other credit cards. Third, issuers track administrative data related to each account, such as the internal credit scores, information about previous changes in the credit limit and interest rates, and the date the account opened. Finally, there are some limited demographic data from the credit application, such as the age and income of the account-holder.

Different credit card issuers track somewhat different sets of variables depending on their source. To protect the identity of the accounts and the issuers, the data from different issuers were pooled together, with great care taken to define variables consistently across issuers. The reported results will focus on variables common to multiple issuers. They have been confirmed separately for each issuer to insure that pooling is not driving the conclusions. Table I provides summary statistics for the main variables used in the analysis. A Data Appendix further describes the data.

Credit cards are used for both transactions and borrowing purposes. To distinguish the two uses, the key dependent variable is interest-incurring debt. For the issuers that provided balance information by billing cycle, debt is defined to be actual balances rolled over into the next month. This equals the closing balance in the cycle net of the subsequent payment. For issuers that provided balance information by calendar month, debt is defined as the balance at the end of the month net of the subsequent payment, *if* interest payments were subsequently incurred. The difference in this figure from one calendar month to the next pro-

TABLE I
SAMPLE STATISTICS

Variable	Mean	Median
debt	1349	70
debt debt > 0	2809	2120
credit limit	6207	5000
Δ credit limit	76.8	0
Δ credit limit Δ credit limit \neq 0	1985	1000
interest rate	16.6	17.2
debt-weighted	17.0	17.7
Δ interest rate	.036	0
abs(Δ interest rate) Δ interest rate \neq 0	.914	.250
# observations	231644	

The unit of observation is a credit card account. Debt is interest-incurring balances, as defined in the text. The differenced variables represent monthly changes. Debt and credit limits are measured in current dollars (1995–1998). The sample is that used for Table II, except for the interest rate statistics which correspond to Table III (# obs = 185,151).

vides a consistent measure of the change in account liabilities between the same point in two consecutive billing cycles.⁷

These data have a number of unique advantages compared with traditional household data sets like the SCF or the Panel Study of Income Dynamics. First, the large sample with little measurement error should provide much more power than usual to identify the effects of credit supply. Second, because each account is observed over many months, it is possible to study dynamics and control for fixed effects. Third, unlike most studies, credit demand can be distinguished from credit supply. Not only do we separately observe credit limits and balances, we can also control for the endogeneity of credit supply changes.

Using account data does, however, entail a number of limitations. First, there is little information about some potentially important variables like household assets or employment status. However, the issuers also lack access to this information, so its

7. The results below were checked separately across both types of balance information. Also, for a subset of accounts we acquired balance information by both cycle and month. For these accounts the results were insensitive to the distinction. For instance, the estimated “impulse responses” (as in Figures I and II below) retained their original shape, and their long-run effects b_{Tot} changed by only about 3 percent in magnitude.

absence will not affect our identification strategy. Second, the main unit of analysis in the data is a credit card account, not an individual or a household. We partially circumvent this limitation by using data from the credit bureaus, which cover all sources of credit used by the account-holder, in particular other credit card balances. (The credit bureau data on balances do not distinguish transactions balances and debt, however.)

IV. ECONOMETRIC METHODOLOGY

The results in this paper can be interpreted as an event study of how credit card debt responds to changes in credit supply, both in quantities and prices, over a period of about a year. Let $D_{i,t}$ be the amount of debt held by account i at the end of month t , and let $L_{i,t}$ be the account's credit limit (the line) as of the start of month t . A traditional event study might focus on a month t in which the limit changes, and regress $D_{i,t+s} - D_{i,t-1}$ on $L_{i,t} - L_{i,t-1}$ for various $s = 0, 1, 2, \dots$. The resulting coefficients would give the response of debt, per dollar of line increase, over horizons of length s . However, since a given account can experience multiple line changes and especially multiple interest rate changes, a more flexible approach is required. The natural extension in this setting is to estimate a distributed lag model, which controls for the effects of multiple changes within the estimation window. The main specification identifying the effect of changes in credit limits is equation (1):

$$(1) \quad \Delta D_{i,t} = \alpha' \mathbf{time}_t + \beta_0 \Delta L_{i,t} + \beta_1 \Delta L_{i,t-1} + \beta_2 \Delta L_{i,t-2} \\ + \dots \beta_{12} \Delta L_{i,t-12} + \gamma' \mathbf{X}_{i,t} + \varepsilon_{i,t},$$

where $\Delta D_{i,t} \equiv D_{i,t} - D_{i,t-1}$. The vector \mathbf{time} represents a complete set of month dummies (a different dummy for each month of the year in the sample), which controls for all aggregate effects, including seasonality and the business cycle, trends in debt over time, and aggregate credit supply (both liquidity and interest rates). Sometimes additional control variables \mathbf{X} will be added, often with twelve lags corresponding to the twelve lags for ΔL .

The coefficient β_0 measures the contemporaneous increase in debt in response to a line increase, per dollar of line increase. The marginal coefficients $\beta_1, \beta_2, \dots, \beta_{12}$ measure the *additional* in-

creases in debt one month after the line increase, two months later, . . . , and twelve months later, respectively. Consequently, $b_k \equiv \sum_{j=0}^k \beta_j$ gives the *cumulative* increase in debt after k months, $k = 0-12$. Twelve lags were sufficient for b_k to converge. Attention will therefore focus on $b_{Tot} \equiv b_{12}$, the long-run, total effect on debt. b_{Tot} can be interpreted as the fraction of a line increase that is borrowed, or the MPC out of liquidity (dD/dL). This interpretation is analogous to the MPC out of predictable income estimated in Campbell and Mankiw [1990].⁸

The dependent variable is the change in debt between months $t - 1$ and t . Using differences controls for individual effects in the level of debt or the credit line. In the event study interpretation, i 's debt after an increase in his credit line is being compared with his debt before the line increase. This identifies the within variation in his debt due to within variation in his line. Various extensions will consider additional controls for the endogeneity of line changes. For instance, account dummies will sometimes be added to equation (1), controlling for individual effects in the *growth* in debt as well. We will also distinguish people who requested their line increase, add the credit scores and other fundamental account characteristics as control variables, and use instruments exploiting the exogenous timing rules built into the credit supply function. Further extensions will test whether the MPC out of liquidity b_{Tot} differs across various groups of accounts. Indicator variables for these groups will be added to equation (1), both directly and interacted with all lags of ΔL .

About 4 percent of credit lines change in any given month. Unlike interest rates which both increase and decrease, issuers have been reluctant to decrease lines, so ΔL is generally nonnegative.⁹ People who do not receive a line increase in a given month remain in the sample that month with the corresponding ΔL equal to zero. These people serve in a sense as a control group, in that the estimated coefficients will pick up the effects of line changes relative to the debt of this group.

8. They use aggregate data. Using micro data, Souleles [1999, 2002] estimates analogous MPCs out of predetermined tax refunds and cuts in payroll taxes (see also Parker [1999]). Ludvigson [1999] relates changes in aggregate consumption to changes in aggregate consumer credit.

9. Thus, line changes are essentially permanent. Apart from the promotional rates discussed below, interest rates are also very persistent, with an AR(1) coefficient above 0.9.

The specification for changes in interest rates is analogous, except that only nine lags were needed for convergence:¹⁰

$$(2) \quad \Delta D_{i,t} = \alpha' \mathbf{time}_t + \beta_0 \Delta r_{it} + \beta_1 \Delta r_{i,t-1} + \beta_2 \Delta r_{i,t-2} \\ + \dots + \beta_9 \Delta r_{i,t-9} + \gamma' \mathbf{X}_{it} + \varepsilon_{i,t},$$

where $r_{i,t}$ is account i 's interest rate as of the start of month t . The coefficients β_j identify the within variation in debt due to within variation in interest rates. The cumulative coefficient $b_{Tot} \equiv b_9$ ($\equiv \sum_{j=0}^9 \beta_j$) gives the total change in debt from a one percentage point increase in the interest rate, i.e., the long-run derivative (dD/dr). Again, these coefficients can be motivated as the product of an event study regressing $D_{i,t+s} - D_{i,t-1}$ on $r_{i,t} - r_{i,t-1}$ for an interest rate change at t . Alternative specifications will also be considered. The controls for the endogeneity of interest rates are analogous to those for the credit limits. Some extensions will jointly estimate the effects of interest rates and credit limits, but the results are similar to those found estimating equations (2) and (1) individually.

While the month dummies in equation (2) partial out aggregate interest rates, there remains substantial idiosyncratic variation in the account-specific interest rates $r_{i,t}$. About 20 percent of interest rates change in any given sample month, both up and down.

Equations (1) and (2) are estimated by OLS, unless stated otherwise. The standard errors allow for heteroskedasticity across accounts as well as serial correlation within accounts. The analysis will begin by examining the debt held on the credit card accounts in the main sample. This examination is of interest for understanding how people use individual credit instruments and the industrial organization of the credit card industry. For the macroeconomic implications of credit, we also want to examine the effect of credit supply on the total amount of debt held by the account-holder. To this end, we will also use the credit bureau data on the balances held on other credit cards owned by account-holder i . In extensions we will replace the dependent variable in equations (1) and (2) with the change in these other balances. If the resulting long-run effects b_{Tot} have the opposite sign as those

10. In both equations (1) and (2) we started with twelve lags of the regressors ΔL and Δr , and removed insignificant lags so long as the adjusted R^2 did not decrease. Graphs of the cumulative coefficients $\{b_k | k = 0-12\}$ (the impulse responses) visually confirmed the convergence (see Figures I and II for examples).

TABLE II
THE RESPONSE OF DEBT TO INCREASES IN THE CREDIT LIMIT

Row		b_{Tot}	<i>S.e.</i>	# obs
(1)	average MPC (dD/dL)	0.126	0.021	231644
(2)	automatic dL	0.096	0.016	145429
	manual dL	1.673	0.773	
(3)	fixed account effects	0.138	0.029	145429
(4)	credit scores	0.085	0.021	127431
(5)	scores, debt, account age	0.122	0.022	131636
(6)	IV	0.111	0.018	145396
(7)	IV: scores, account age	0.083	0.030	131603
(8)	IV: $\#(dL \neq 0)$	0.104	0.020	145396
(9)	$d(\text{interest rates})$	0.122	0.024	137279
(10)	interest rates	0.122	0.024	137279
(11)	balances on other cards	0.026	0.030	130486

A distributed lag model (equation (1)) was used to estimate the dynamic response of credit card debt ΔD to changes in credit limits ΔL including twelve lags. b_{Tot} ($=b_{12}$) gives the long-run, cumulative change in debt as a fraction of the change in the line (the MPC out of liquidity), dD/dL . All regressions include a full set of month dummies. The standard errors allow for heteroskedasticity across accounts as well as serial correlation within accounts. Sample sizes vary with missing variables. Row (2) distinguishes manual line changes (requested by the consumer) from automatic changes (initiated by the issuer), and includes separate intercepts for each case. Rows (3)–(8) include the controls for the manual line changes; the reported b_{Tot} is for the automatic changes. Row (3) includes a fixed effect by account. Row (4) includes as controls a cubic polynomial in the normalized credit scores, both internal and external, interacted by issuer dummies; and then twelve lags of all these terms (i.e., 2 scores \times 3 polynomial terms \times 13 lags \times issuer dummies). Row (5) includes cubic polynomials in the two scores, debt and account age; all from month $t - 1$ and all interacted by issuer dummies. Rows (6)–(8) instrument for ΔL with indicator variables for the number of months since the latest change in line. Row (7) includes as controls cubic polynomials in the two scores and account age, from month $t - 1$. Row (8) includes dummy variables for the total number of line changes each account received during the sample period. Row (9) includes as controls the change in account interest rate with twelve lags. Row (10) includes instead the level of the interest rate with twelve lags. In row (11) the dependent variable is balances on *other* credit cards held by the account-holder.

for the cards in the main sample, that will provide evidence of “balance shifting,” i.e., of people shifting balances across cards in response to changes in credit supply.

V. RESULTS: CHANGES IN CREDIT LIMITS

This section estimates the average response of debt to credit line increases, beginning with the credit card accounts in the main sample. The analysis of heterogeneity appears in Section VII, after the results for interest rates in Section VI.

Table II records the long-run response of debt b_{Tot} , which can be interpreted as the MPC out of liquidity. As a starting point, row (1) uses the entire sample. The estimated b_{Tot} is relatively large at just under 0.13 and quite significant (t -ratio of 6). Over the year following a line increase, each extra thousand dollars of

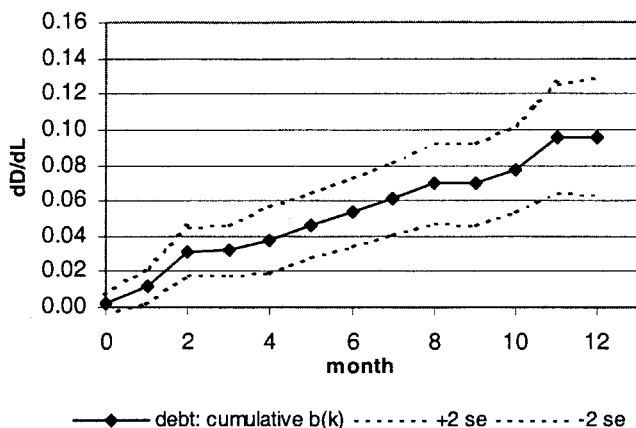


FIGURE I

The Cumulative Response of Debt to "Automatic" Increases in the Credit Line, per Dollar of Extra Line (Table II, Row (2))

liquidity generates on average a \$130 increase in debt, an MPC of 13 percent. In absolute terms (calculated by replacing the regressors ΔL in equation (1) with indicator variables $I(\Delta L > 0)$), this corresponds to a \$350 average increase in debt per line increase. Hence liquidity matters, counter to the PIH.

For brevity the intermediate cumulative coefficients $\{b_k | k = 0-11\}$ are not reported but can be found in the working paper version of this paper. When graphed, they give the "impulse response" to liquidity. Figure I provides an example of the impulse response, for a smaller but preferred sample used below. Debt rises sharply and significantly over the first two months after a line increase, and then smoothly asymptotes to b_{Tot} . Such a clean result is rare in micro data. Presumably, the difference is due to the large sample with little measurement error. The time dummies in equation (1) are also jointly significant (not reported).

The remaining rows of Table II show the results of various extensions.¹¹ Rows (2) to (8) address the potential endogeneity of credit limits, employing a number of strategies to ensure that it is not driving the results. First, recall that the time dummies con-

11. Since different issuers collect somewhat different data, we are unable to apply all of the extensions simultaneously and still pool across a sufficiently large number of accounts and issuers. We therefore apply the extensions in groups. We have confirmed in unreported regressions that our conclusions are robust to applying the extensions simultaneously.

trol for all aggregate effects. For instance, they control for the fact that issuers might offer additional credit before Christmas knowing that on average debt will subsequently rise. But the time dummies do not control for the analogous idiosyncratic situation in which, when someone expects to make a big purchase in the near future, he calls his issuer and requests a line increase. In this situation there would be reverse causality: the expected future purchase would be responsible for the line increase.

Fortunately, for a large subset of our data, we know whether the line increase was requested by the account-holder or not. The requested changes, which the issuers call “manual” changes, account for about 10 percent of the total number of line changes; the remaining 90 percent, called “automatic” changes, are initiated by the issuers. A dummy variable was created to identify manual line changes, and this dummy and its twelve lags were interacted with the regressors ΔL and its lags in equation (1). Row (2) reports the resulting b_{Tot} for both manual and automatic line increases. Not surprisingly, debt rises much more after a manual increase, significantly so. b_{Tot} in this case is over 100 percent of the extra line, although the standard errors are larger because of the smaller number of manual changes. More interestingly, debt also responds to automatic line changes. Their MPC b_{Tot} remains significant and large at about 0.10 (\$290 in absolute terms). Figure I shows the impulse response for automatic changes, which remains relatively smooth despite the smaller sample. Therefore, even though in the full sample we do not always know which line changes were automatic and which were manual, we do know that the small number of manual changes are not driving the results. Nonetheless, we will often add the interaction terms for manual line changes as controls in subsequent extensions when the sample size permits, including in rows (3)–(8) which focus on endogeneity. In these cases we report b_{Tot} for the uninteracted ΔL terms, which represents the MPC for automatic line changes. In other specifications that do not explicitly include the controls for manual changes, we have checked that their omission does not change our conclusions.

Second, issuers try to identify the types of potential customers who will borrow and pay interest but not default. Row (3) adds fixed account effects to equation (1), which control for all persistent characteristics of the account and account-holder, including the subportfolio to which the account belongs. These are fixed effects in the *change* in debt, which should be a powerful control

for endogeneity. Nonetheless, the MPC b_{Tot} remains significant and actually rises slightly to 0.14.

Third, we control directly for the credit-risk scores, which summarize the fundamental account characteristics affecting credit policy for each account. We use both the internal and external scores. Each score was first normalized relative to the average score for the corresponding issuer. A third-order polynomial in these normalized scores was interacted with dummy variables for the issuers, to allow each issuer to use its scores nonparametrically and differently. These interacted variables were added to equation (1), along with twelve lags to control for past line changes.¹² The results are in row (4). While the scores are significant (not reported), they reduce b_{Tot} only slightly to .085, and it remains significant.¹³

We can also directly control for the account fundamentals. Row (5) includes third-order polynomials in the two scores and in account debt and account age, the other variables highlighted by Gross and Souleles [2002].¹⁴ b_{Tot} remains significant and similar in magnitude to the preceding estimates. The results are also similar on controlling for other salient account characteristics.¹⁵

12. We also interacted these variables with time dummies, to allow each issuer to use the scores differently over time. Because of the large number of resulting regressors, we did this in smaller test samples. This time interaction had little effect on the main results in the test samples.

13. In this and other extensions, the sample size can vary with missing variables. For instance, the scores are occasionally missing or changed to a flag value (e.g., if there were not enough data on record at the bureaus to have computed a score).

14. Again these variables are interacted with issuer dummies. To conserve memory, the reported specification includes these three characteristics (normalized scores, debt, age) only from month $t - 1$ in equation (1), without lags. The results are robust to adding lags. We also added account age and debt from month $t - 1$ to the full set of scores with twelve lags used in row (4), but again the results are similar.

15. The results are also robust to controlling for past debt levels D with twelve lags, or for twelve lags of the dependent variable ΔD . There is little information available to issuers that is useful in predicting future high-frequency changes in an account's demand for credit, other than the account's previous debt dynamics. These control variables summarize this information (which also appears in the scores), including account-specific seasonality, and so are powerful controls. We also added six *future* values of the regressors ΔL to equation (1), i.e., $\Delta L_{i,t+1}$ to $\Delta L_{i,t+6}$. Their coefficients were statistically insignificant, both individually and jointly, and small in magnitude. This suggests that the high-frequency variation in lines exploited here is not an endogenous response to preceding debt dynamics. This implication can also be illustrated using the event-study specification. Let b_s be the coefficient from regressing $(D_{i,t+s} - D_{i,t-1})$ on $(L_{i,t} - L_{i,t-1})$, for a line change at t . $s < 0$ represents the "pre-event" window. The estimated b_s for $s = -6$ to -1 are insignificant and small, especially conditioning on the credit scores. Thus, debt is not significantly changing in advance of the line increases, only afterwards.

This robustness suggests that the variation driving the results is indeed exogenous, the product of nonfundamental institutional constraints.

Fourth, to control for any remaining endogeneity, we exploit the exogenous timing rules described in Section II. We instrument using dummy variables for the number of months since each account's latest line change. These dummies and their lags were interacted with issuer dummies, to allow each issuer to use different timing rules.¹⁶ As already noted, in the first-stage these instruments are quite significant (p -value $< .0001$), with an R^2 of about 0.03, which is relatively large in the context of related micro studies; they also pass the overidentification test. The results are in row (6). Again, b_{Tot} does not change very much, rising slightly to 0.11. Reassuringly, the standard errors remain small, and the underlying impulse response remains smooth.

To bolster these results, we add some control variables (from month $t - 1$) to equation (1), in addition to instrumenting for the line changes. Row (7) again adds the third-order polynomials in the two credit scores, which for the issuers summarize the history of account behavior. This specification compares accounts whose past behavior was similar, but some are getting line increases in a given month, and others are not, because they happen to be on different timing cycles. To ensure that our instruments are not biased against younger accounts, we also include a third-order polynomial in account age. There is no reason to believe that the correlation of the remaining variation in credit line policy with the instruments is endogenous. These controls decrease b_{Tot} slightly to 0.08, but it remains quite significant. Row (8) instead adds dummy variables for the total number of line increases an account receives during the entire sample period. These variables control for the factors justifying the line increases. The remaining variation is purely in the high-frequency timing of the increases. Nonetheless, b_{Tot} remains significant and large.

We conclude that endogeneity is not driving the results. They remain remarkably consistent across a large number of specifications, with richer controls and instruments than available in most previous studies of credit supply. There appears to be a

16. If the issuer's variable that directly records the month of the latest line change is missing, we recreate it when possible. We can compute the number of months since an account's most recent line change within the sample period, for all but its first line change within the sample.

causal link from liquidity to debt, with an average MPC between 10–14 percent depending on the specification.

Although interest rates are analyzed in the next section, it is important to ensure that they are not confounding the results here for credit limits. Row (9) adds to equation (1) the change in interest rates (Δr) with twelve lags. The reported b_{Tot} for credit limits does not materially change.¹⁷ In analogy with an Euler equation (e.g., Campbell and Mankiw [1990]), row (10) adds instead the *level* of interest rates with twelve lags (r , not Δr). b_{Tot} remains the same as in the previous row.

Finally, to determine the effect of line increases on the aggregate economy, it matters whether *net* debt is increasing, or whether people are simply shifting balances away from their other credit cards to their card in the main sample with the new, larger credit line. Row (11) shows the results using the change in *other* credit card balances held by the account-holder as the dependent variable in equation (1). A negative value for b_{Tot} would represent balance-shifting, an offsetting decline in balances on other cards. The estimated b_{Tot} , however, is insignificant and positive. Thus, there is no evidence of offset from the other cards, so the previous results represent a net increase in debt, and hence have aggregate implications.¹⁸

VI. RESULTS: CHANGES IN INTEREST RATES

Table III estimates the effect of interest rates on debt, using equation (2), again beginning with the credit cards in the main sample. In row (1) the long-run effect b_{Tot} ($\equiv b_9$) is about $-\$110$, a large and significant coefficient (with t -ratio about 13). Each percentage point increase in the interest rate leads to a $\$110$ decline in debt on average, within nine months. Thus, higher interest rates lead to substantially less borrowing—people are in fact sensitive to interest rates. The impulse response, graphed in Figure II, is again quite smooth and significant. It shows that debt responds immediately, declining by about $\$70$ in the first two months after a rate change. These estimates correspond to short-

17. This follows from the fact that line changes and rate changes are not disproportionately correlated. This fact also supports the conclusion above regarding endogeneity. Unconstrained, optimal credit policy would determine credit lines and interest rates jointly, presumably resulting in greater correlation.

18. Furthermore, adding other balances directly as a control variable in equation (1) has little effect on the previous results.

TABLE III
THE RESPONSE OF DEBT TO CHANGES IN INTEREST RATES

Row		b_{Tot}	S.e.	# obs
(1)	interest sensitivity (dD/dr)	-112.6	8.4	185151
(2)	normalize by line: $d(D/L)/dr$	-0.016	0.001	185151
(3)	fixed account effects	-132.4	7.9	185151
(4)	credit scores	-109.4	11.9	107542
(5)	scores, debt, account age	-117.1	9.5	173406
(6)	IV	-79.1	27.2	147158
(7)	IV: scores, account age	-156.4	32.1	138709
(8)	IV: $\#(dr \neq 0)$	-80.8	26.8	147158
(9)	$d(\text{credit limit})$	-111.8	-8.3	185099
(10)	increase r	-90.1	9.2	185151
	decrease r	-317.7	63.7	
(11)	balances on other cards	37.8	20.2	142334

A distributed lag model (equation (2)) was used to estimate the dynamic response of credit card debt ΔD to changes in credit card interest rates Δr including nine lags. b_{Tot} ($=b_9$) gives the long-run, cumulative change in debt per percentage point increase in rates, dD/dr . All regressions include a full set of month dummies. The standard errors allow for heteroskedasticity across accounts as well as serial correlation within accounts. Sample sizes vary with missing variables. In row (2) only, the dependent variable is the change in the ratio of debt to the credit limit. Row (3) includes a fixed effect by account. Row (4) includes as controls a cubic polynomial in the normalized credit scores, both internal and external, interacted by issuer dummies; and then nine lags of all these terms (i.e., 2 scores \times 3 polynomial terms \times 10 lags \times issuer dummies). Row (5) includes cubic polynomials in the two scores, debt and account age; all from month $t - 1$ and all interacted by issuer dummies. Rows (6)–(8) instrument for Δr with indicator variables for the number of months since the latest change in rates. Row (7) includes as controls cubic polynomials in the two scores and account age, from month $t - 1$. Row (8) includes dummy variables for the total number of rate changes each account received during the sample period. Row (9) includes as controls the change in account credit limit with nine lags. Row (10) distinguishes the response to increases versus decreases in interest rates, including separate intercepts for each case. In row (11) the dependent variable is balances on *other* credit cards held by the account-holder.

run and long-run elasticities of about -0.8 and -1.3 , respectively. They are larger and more significant than in most other micro studies of interest rates. Again the difference is probably due to the large sample and reduced measurement error in our data, both in debt and in interest rates, which are the account-specific marginal cost of credit.¹⁹

The other rows of Table III display various extensions. First, row (2) controls for cross-sectional differences in credit limits by normalizing debt by the credit limit. The dependent variable is the change in this debt ratio ($\Delta(D/L)_{i,t}$). The impulse response remains qualitatively similar to that in Figure II, with each percentage point increase in interest rates reducing the normalized debt ratio by about 1.7 percentage points. Since the average

19. The results using the percentage change in the interest rate, $\Delta r/r$, are qualitatively similar to those reported.

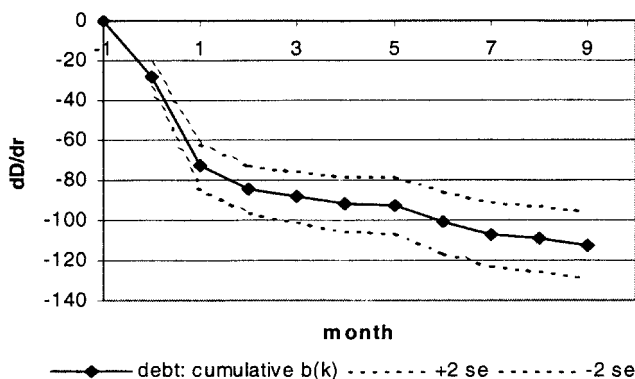


FIGURE II

The Cumulative Response of Debt to Increases in the Interest Rate, per Percentage Point (Table III, Row (1))

credit limit is about \$6200, this is consistent with the results in row (1). The next few rows address the endogeneity of interest rates, using specifications analogous to those above for credit limits. In rows (3)–(5), adding to equation (2) fixed effects or the fundamental account attributes (the scores, plus account age and debt) does not materially change the results. This again suggests that the variation being exploited is exogenous.

As with credit limits, the issuers also use exogenous timing rules for interest rates, for instance, changing rates after fixed periods like six months or twelve months. Row (6) therefore instruments with dummy variables for the number of months since the latest rate change. The first-stage R^2 is about 0.02. (Again the instruments are very significant, with p -value $< .0001$, and pass the overidentification test.) The resulting b_{Tot} is somewhat smaller in magnitude than before, at about $-\$80$, but still quite significant. In row (7), adding to the IV specification third-order polynomials in the credit scores and account age increases the magnitude of b_{Tot} . Thus, even controlling for account fundamentals the exogenous timing rules affect debt. Row (8) instead controls for the number of interest rate changes each account receives during the sample period. b_{Tot} remains quite significant. This coefficient reflects purely exogenous variation in the timing of rate changes, even controlling for the fact that rates changed. Row (9) confirms that these results are not due to changes in credit limits. Overall, the results are again remark-

ably consistent across specifications. We conclude that they are not driven by endogeneity.

Row (10) distinguishes the effects of increases versus decreases in interest rates, by interacting the regressors Δr and its lags in equation (2) with indicator variables for their signs. The resulting long-run effect for increases in rates is $b_{Tot}^{\Delta r > 0} = -\90 , for decreases in rates $b_{Tot}^{\Delta r < 0} = -\320 . That is, debt declines in response to increases in borrowing rates, and rises even more strongly in response to decreases in borrowing rates (with $b_{Tot}^{\Delta r < 0}$ significantly larger in magnitude than $b_{Tot}^{\Delta r > 0}$). Such an asymmetry implies that temporary interest rate reductions can ratchet up the amount of debt. This implication can explain the widespread use of promotional offers like teaser rates.²⁰

These results also suggest the possibility that when interest rates change people are transferring balances between their card in the main sample and their other cards. Row (11) uses other card balances as the dependent variable in equation (2). A positive b_{Tot} would represent balance-shifting away from a card whose interest rate has increased. In fact, the estimated b_{Tot} is almost +\$40, and significant at the 6 percent level. For each percentage point rise in the interest rate on the card in the main sample, while its debt declines by about \$110, other card balances rise by about \$40 on average, for a net decline of about \$70. Therefore, there appears to be some balance-shifting between cards. People's sensitivity to interest rates is large enough to overcome some switching costs. However, the balance-shifting is only partial; there is a net decline in total borrowing. The net long-run elasticity is still large at about -0.85.

VII. HETEROGENEITY AND INTERPRETATION

The results so far do not distinguish whether all people are responding equally to credit supply, or whether the response is limited to part of the sample. We now investigate the heterogeneity in people's responses. The goal is to identify what models of consumption explain these responses and credit card usage more generally.

Table IV starts with the leading alternative to the PIH,

20. Debt is more sensitive to large changes in rates, resulting from promotional offers, than small changes, but even the response to small changes is significant (see the working paper).

TABLE IV
UTILIZATION RATES AND LIQUIDITY CONSTRAINTS

Row	b_{Tot}	<i>S.e.</i>	# obs
A. Credit limit changes			
(1) dD/dL			143511
utilization < .50	0.068	0.018	
utilization .50-.90	0.158	0.060	
utilization > .90	0.452	0.125	
(2) $d(\text{utilization})/dL$			143511
utilization < .50	-0.012	0.004	
utilization .50-.90	0.013	0.015	
utilization > .90	0.012	0.026	
B. Interest rate changes			
(3) dD/dr			182317
increase r			
utilization < .50	-65.0	11.4	
utilization .50-.90	-129.5	20.5	
utilization > .90	-85.5	21.4	
decrease r			
utilization < .50	-296.9	99.7	
utilization .50-.90	-429.3	120.2	
utilization > .90	-151.9	63.7	

This table contrasts the response of debt to credit supply (b_{Tot}) across accounts starting with different initial utilization rates (balances divided by the credit limit). In Panel A, for changes in credit limits (equation (1)), utilization is taken from month $t - 13$; in Panel B, for changes in interest rates (equation (2)), it is taken from month $t - 10$. Row (3) also distinguishes increases versus decreases in interest rates, for each utilization group. In rows (1) and (3) the dependent variable is the change in debt ΔD . In row (2) it is the change in utilization rate. Panel A includes controls for manual credit limit changes; the reported b_{Tot} is for the automatic changes. All regressions include separate intercepts for each utilization group, and a full set of month dummies. The standard errors allow for heteroskedasticity across accounts as well as serial correlation within accounts.

liquidity constraints. Panel A analyzes the effects of credit line changes. It includes the controls for manual changes to ensure that they are not driving the results; the reported b_{Tot} is for the automatic changes. Equation (1) is reestimated distinguishing accounts according to their initial utilization rate (balances divided by the line).²¹ To avoid endogeneity, this rate is taken from month $t - 13$, the beginning of the distributed lag horizon in equation (1). Indicator variables were created for accounts with initial utilization above 0.90 and for accounts with utilization below 0.50. The omitted category is initial utilization between

21. Because the credit limit constrains the magnitude of total balances, not just debt, we include transactions balances as well as debt in the numerator of utilization. The conclusions are similar using just debt in the numerator.

0.50 and 0.90. These indicators were interacted with all lags of ΔL in equation (1). The results appear in row (1). The people starting at over 90 percent utilization are the ones most likely to be facing binding liquidity constraints. In fact, their MPC out of liquidity $b_{Tot}^{util>.9}$ is almost 50 percent and quite significant. This coefficient is significantly larger than $b_{Tot}^{util.5-.9}$ and $b_{Tot}^{util<.5}$ for the lower utilization groups. Hence the MPC varies with utilization. The sharp response by the high utilization group suggests that their liquidity constraints are indeed binding. This is concrete evidence that liquidity constraints matter.

However, even the people starting at lower levels of utilization significantly respond to liquidity. $b_{Tot}^{util.5-.9}$ is almost 0.16, and even $b_{Tot}^{util<.5}$ for people with most of their line free is 0.07. Both coefficients are statistically significant (although not significantly different from each other).²² In absolute terms, these coefficients correspond to average increases in debt of about \$500 and \$170 per line increase, substantial amounts. This evidence is inconsistent with the view that only currently binding liquidity constraints matter.

The evidence, however, is potentially consistent with the interaction of precautionary motives and the possibility that liquidity constraints bind in the future, as in buffer-stock models [Deaton 1991; Carroll 1992; Ludvigson 1999]. In the current context, these models generate an optimal amount of credit card line to keep free, as a precautionary buffer against future adverse shocks, and hence pin down the utilization rate. Changes in credit limits can then generate changes in borrowing even at utilization rates well below 100 percent, as in row (1).

To further evaluate the buffer-stock model, we consider a simple, homothetic version (as in Ludvigson [1999]), in which the optimal buffer of unused credit can be a constant fraction of the credit line, generating a constant target utilization rate. To test this possibility, in row (2) the dependent variable in equation (1) is replaced by the change in the utilization rate ($\Delta utilization_{i,t}$). The coefficients b_{Tot} now represent the long-run, cumulative response of utilization after a line increase, separately by initial utilization rate as before. The estimated b_{Tot} are quite small, within a percentage point or so of zero. Hence, for each utilization

22. The impulse response coefficients $\{b_k\}$ jointly differ across the two lower utilization groups. And both of their impulse responses differ from that for the high utilization group. Thus, there are differences in the paths of debt, not only in the long-run effects b_{Tot} .

group, in the long run utilization returns back near its initial level (that is, about 90 percent for someone starting at 90 percent utilization, about 50 percent for someone starting at 50 percent utilization, etc.).²³ These dynamics are consistent with people having constant target utilization rates, at least on average. Of course, such targeting could be due to behavioral rules-of-thumb, but as noted it is also compatible with conventional precautionary motives.²⁴ Under either explanation, the important point is that additional liquidity increases debt even at low utilization rates, leading to larger aggregate effects than if only currently binding constraints mattered.

Row (3) tests whether liquidity-constrained people are less sensitive to changes in interest rates. Equation (2) is estimated distinguishing increases and decreases in interest rates (as in row (10) of Table III), both interacted with the three initial utilization groups as before (with utilization taken from month $t - 10$). This yields six different coefficients b_{Tot} for the long-run effects of interest rates. The estimated coefficients vary nonmonotonically with utilization, for both increases and decreases in rates. For decreases in rates, accounts with utilization between 0.5–0.9 are the most sensitive to interest rates, accounts with utilization above 0.9 are the least sensitive (with the difference between them statistically significant). The credit limit automatically constrains high utilization people from substantially increasing their debt in response to lower interest rates, but it does not prevent people from decreasing their debt in response to higher rates. For increases in rates, the estimated $b_{Tot}^{\Delta r > 0}$ is again smaller for accounts with utilization above 0.9 than for accounts with utilization between 0.5–0.9, although not significantly so (only at the 13 percent level). This nonmonotonicity of interest-rate elasticities is again consistent with liquidity constraints.

Table V examines other dimensions of heterogeneity in credit card usage. To begin with, column (1) shows how utilization rates vary across accounts of different characteristics, reporting the fraction of accounts with high utilization rates, above 0.90. Each

23. The underlying impulse responses are illustrated in the working paper. To summarize them, in the month of the line increase utilization rates sharply decline, because the line in the denominator of utilization has risen. But thereafter the cumulative coefficients $\{b_k\}$ smoothly rise, back up near zero within five months or so.

24. Indeed, rule-of-thumb targets for utilization rates might serve to approximate optimal buffer-stock policies for liquidity. Allen and Carroll [2000] consider how close rule-of-thumb consumption behavior comes to optimal behavior.

TABLE V
 HETEROGENEITY IN CREDIT-CARD USAGE

Row	(1) % with utilization > .90	(2) Credit limit MPCs $\Delta D/\Delta L$		(3) Interest rate sensitivities $\Delta(D/L)/\Delta r$	
		b_{Tot}	S.e.	b_{Tot}	S.e.
(1) Credit limit					
small ($\leq \$5000$)	16.9	0.081	0.035	-0.016	0.002
large ($> \$5000$)	9.2	0.100	0.023	-0.017	0.002
<i>p</i> -value		0.675		0.512	
N		145429		185151	
(2) Credit score					
low ($\leq .5$)	25.8	0.125	0.026	-0.016	0.002
high ($> .5$)	1.8	0.070	0.020	-0.020	0.002
<i>p</i> -value		0.066		0.059	
N		132121		169234	
(3) Income					
low ($\leq \$31000$)	14.0	0.122	0.049	-0.015	0.002
high ($> \$31000$)	12.6	0.114	0.019	-0.018	0.009
<i>p</i> -value		0.876		0.103	
N		116542		141675	
(4) Age of account-holder					
young (≤ 43)	17.7	0.123	0.033	-0.020	0.002
old (> 43)	11.5	0.068	0.027	-0.017	0.002
<i>p</i> -value		0.182		0.196	
N		91832		132243	
(5) Age of account (years since booking)					
new (≤ 4)	14.7	0.103	0.018	-0.017	0.001
old (> 4)	12.9	0.067	0.029	-0.015	0.002
<i>p</i> -value		0.259		0.365	
N		145429		185151	
sample average	13.8				

This table contrasts the response of debt to credit supply (b_{Tot}) across accounts of different characteristics. In column (2) the dependent variable is the change in debt ΔD , and the independent variables are changes in the credit limit ΔL (equation (1)). In column (3) the dependent variable is the change in debt normalized by the line and the independent variables are changes in interest rates Δr (equation (2)). Each row in column (2) and in column (3) represents a separate regression. Column (1) reports the fraction of accounts with utilization rates above 90 percent. Each row splits the sample at (approximately) the median of the characteristics in that row. For the time-varying characteristics in rows (1) and (2), the characteristics are taken from month $t - 13$ for credit limits (column (2)), month $t - 10$ for interest rates (column (3)). For the characteristics in rows (3)–(5), which come from cross-sectional application data, their 1995 values are used. Row (2) uses the normalized internal credit score. Income is real application income (1982–1984\$), extrapolated to 1995 assuming a real growth rate of 2.5 percent per year. The reported *p*-values are for tests that b_{Tot} is the same across both groups in the row. Except for row (3), column (2) includes controls for manual credit limit changes; the reported b_{Tot} is for the automatic changes. See text for further details. All regressions include separate intercepts for each characteristic group, and a full set of month dummies. Sample sizes vary with missing variables. The standard errors allow for heteroskedasticity across accounts as well as serial correlation within accounts.

row splits the sample at the median of the corresponding characteristic, including the real income and age of the account-holder derived from the credit application data. For the overall sample, almost 14 percent of accounts have high utilization. But for young account-holders, in row (4), this fraction rises to about 18 percent, which is consistent with other evidence that the young are disproportionately liquidity constrained [Jappelli 1990]. In row (2) people with low credit scores are much more likely to have high utilization rates than people with high scores, 26 percent versus 2 percent.²⁵ This result is also consistent with liquidity constraints, because lenders are reluctant to extend credit to people with low scores. However, reverse causation is likely also to be at play, because high debt levels can lead to lower scores. Low income and newer accounts (time since booking) have slightly higher utilization rates than average (rows (3) and (5)).²⁶

The more direct test of liquidity constraints looks for high MPCs out of liquidity, as reported in column (2). Indicator variables for being above the median of each characteristic are interacted with all the lags of ΔL in equation (1). For the time-varying characteristics (credit limits and scores), the indicator is taken from month $t - 13$ ($t - 10$ for interest rates in column (3)). For the other characteristics derived from application data and hence essentially cross-sectional (income and age of the account-holder, age of the account), the indicator is based on their values in 1995, the start of the sample period.²⁷ Column (2) includes again the controls for manual line changes, except in row (3).²⁸ In row (1), the MPCs b_{Tot} are similar in magnitude for accounts with small and large credit limits, and not significantly different (p -value =

25. Table V uses the normalized internal scores; using the external scores yields similar conclusions.

26. These results apply to individual bankcard accounts. Utilization rates calculated from the 1995 SCF, which combines all bankcards a household owns, are smaller in magnitude because of the underreporting of debt discussed in footnote 2. Only 3 percent of SCF households with bankcards report combined utilization rates above 90 percent. Nonetheless, the cross-sectional effects of household age and income in the SCF are qualitatively the same as in Table V. Also, credit bureau sources report that during our sample period about 15 percent of *individuals* with bankcards had combined utilization rates above 90 percent, consistent with our account data.

27. That is, we computed account and account-holder age in 1995. For real income (1982–1984\$) in 1995, we assumed a real growth rate of 2.5 percent per year to allow for growth in income between the time of application and 1995. The results are similar using nearby growth rates.

28. To ensure sufficient pooling, since income is often missing. We have confirmed that the conclusions for income are the same including the manual controls.

.675). Thus, the response of debt to the credit limit is independent of the size of the limit, at least within the range in the sample. This result supports the specification in equation (1), which uses the absolute change in the limit as the independent variable. In row (2) the MPC for people with low credit scores is substantially larger than for people with high scores, significantly so at the 7 percent level. This bolsters the previous evidence that low credit scores reinforce liquidity constraints. Unless the scores are themselves functions of the MPC, which seems unlikely, this result is not due to reverse causation. Row (3) shows that low income people have a larger MPC on average, although not significantly so. In row (4) young people have a much larger MPC than older people, about double in magnitude although not significantly different. (The result is similar for the youngest quartile of account-holders.) This is again consistent with liquidity constraints. Similarly, in row (5) newer accounts have somewhat larger MPCs on average. Overall, these results suggest that liquidity constraints are pervasive and vary across the population, disproportionately affecting low income and, in particular, young people. They also highlight the importance of the credit scores in gaining access to credit. Yet the results are less significant than those in Table IV, which split the sample directly by utilization rates. There is a great amount of heterogeneity regarding liquidity constraints even *within* each of these characteristic groups of people.

Column (3) contrasts how the same groups respond to changes in interest rates. To control for cross-sectional variation in credit limits, the dependent variable is again normalized by the limit ($\Delta(D/L)_{i,t}$), as in row (2) of Table III. People with low scores and low income are less sensitive to interest rates, significantly at the 6 percent and 10 percent levels, respectively. This is again consistent with liquidity constraints. However, young people and newer accounts are slightly more sensitive to interest rates, although neither distinction is significant.

Turning to Table VI, the bottom two rows examine alternative measures of liquidity constraints, using data from the 1995 SCF. Overall (in column (1)), 67 percent of households have a bankcard, leaving 33 percent of households without access to this important source of credit and so possibly liquidity constrained. Access also varies substantially across demographic groups, as shown in columns (2)–(4). Low income, young, and low education households are much less likely to have bankcards. The bottom

TABLE VI
ASSET HOLDINGS OF BANKCARD BORROWERS

Variable	(1) All borrowers			(2) By income		
	Mean	Median	% positive	<\$15K	% positive \$15K–50K	>\$50K
Net worth	\$127980	\$61530	95	89	94	99
Net housing equity	40568	22700	69	54	63	86
Total assets	180657	112850	100	98	100	100
Financial assets	47346	11700	98	93	98	100
Nonretirement financial assets	29546	5450	98	93	98	100
Liquid assets	5536	1800	96	91	96	98
Liquid assets – one month's income	1671	–983	33	46	30	34
% with bankcards	67			32	71	94
% borrowing have bankcards	56			56	59	52
Variable	(3) By age			(4) By education		
	<35	% positive 35–62	>62	<HS	% positive HS	COL
Net worth	89	97	100	95	95	95
Net housing equity	43	79	80	77	69	67
Total assets	100	100	100	99	100	100
Financial assets	98	98	97	95	98	99
Nonretirement financial assets	98	98	97	95	98	99
Liquid assets	95	96	97	91	96	99
Liquid assets – one month's income	25	32	54	35	31	36
% with bankcards	59	72	64	35	65	88
% borrowing have bankcards	69	64	29	54	63	49

Source: 1995 SCF. N = 3258. Results are weighted to be representative of the United States. The bottom two rows give the fraction of households in each group (each column) that have bankcards, and the fraction that are borrowing on their bankcards conditional on having bankcards. Borrowing households are those rolling over positive bankcard balances after their latest payments. The columns labeled “% positive” give the fraction of households in the group whose holdings of the asset in that row are positive. Asset definitions generally follow standard SCF codebook suggestions, with the exception that net worth does not include credit card debt. Nonretirement financial assets exclude IRAs and defined-contribution pensions. Liquid assets include amounts in checking, saving, money-market, and brokerage call accounts. Income is gross total household income. Asset holdings are all conditional on having and borrowing on at least one bankcard. Age and education refer to the household head. Columns (2) and (3) contrast the bottom 25 percent, the middle 50 percent, and the top 25 percent of the income and age distributions, respectively.

row of Table VI reports the fraction of households borrowing (i.e., paying interest) on bankcards, conditional on having a bankcard. (See the Data Appendix for SCF data definitions.) On average 56 percent of bankcard-holding households are borrowing on their cards (column (1)), a remarkably large fraction, and probably a lower bound considering the measurement error in the SCF (see footnote 2). These households would be considered liquidity constrained under the weaker notion of an operative wedge between borrowing and lending rates. Credit card interest rates averaged over 16 percent during the sample period (Table I), much larger than typical saving rates.²⁹ The fraction of households borrowing also varies across demographic groups, as indicated in columns (2)–(4). The effect of income and education is not monotonic, with both very low and very high income and education households less likely to borrow than households in the middle. Nonetheless, it is surprising that 52 percent of high income households and 49 percent of high education households are borrowing on their credit cards. Young households are much more likely to borrow than old households, 69 percent versus 29 percent, and so again appear disproportionately liquidity constrained.

The results for equations (1) and (2) are arguably consistent with conventional models of liquidity constraints reinforced by precautionary motives. However, conventional models have more trouble explaining some other aspects of credit card use. Most saliently, why does such a large fraction of the population hold credit card debt? And why do so many of these borrowers simultaneously hold assets yielding low returns?

We begin with the first question. Conventional buffer-stock models calibrated using historical income profiles have difficulty rationalizing so much borrowing. Motivated in part by these results, Laibson, Repetto, and Tobacman [2000] show that allowing for hyperbolic discounting substantially increases the fraction of households that would borrow on their credit cards, although not quite enough to match the data, especially for high wealth households. Hyperbolic discounting also raises the average MPC.³⁰ Alternatively, perhaps people expect that their future

29. As Table I reports, the debt-weighted interest rate is actually a little larger than the average rate. This rules out the possibility that most of the credit card borrowing is done at low, promotional interest rates.

30. The consumption policy functions for the hyperbolic buffer-stock consumers are qualitatively similar to those for the exponential buffer-stock consumers. Hence the difference between them is primarily quantitative, and depends on the

income will rise faster than calibrated using historical data (the “new era” view), justifying more borrowing.

Turning to the second question above, the portfolio puzzle, Table VI describes the asset holdings of bankcard borrowers. Column (1) starts with the overall sample averages. Conditional on borrowing on their bankcards, 95 percent of households have positive net worth and so could have paid off some of their expensive credit card debt by drawing down various assets. For instance, almost 70 percent have positive housing equity, and so would be better off using lower cost home-equity debt (currently charging around 7–9 percent, not including tax deductions). Over 90 percent have positive holdings of financial assets, even excluding illiquid, tax-favored retirement assets. Most puzzling of all, over 90 percent of people with credit card debt have some very liquid assets in checking and savings accounts, which usually yield at most 1–2 percent. Granted, some transactions cannot easily be made using credit cards—although it behooves borrowers to seek out vendors accepting credit cards, and the ability to write checks against credit card accounts, or to take out cash advances, increases the scope of possible transactions. To accommodate cash transactions, Table VI allows for one month’s worth of gross total household income to be kept in liquid assets for transactions purposes, which is arguably generous. Yet one-third of credit card borrowers still have more than this much income in liquid assets. These assets could instead have been used to pay down their credit card debt. These results persist even for people with substantial debt. For over 10 percent of bankcard-holding households, bankcard debt amounts to more than one month’s income. Yet again about one-third of these high-debt households have over one month’s income in liquid assets.

Such behavior is puzzling, apparently inconsistent with no-arbitrage and thus inconsistent with any conventional model. Perhaps behavioral models of self-control or mental accounts might help explain it. For example, some people might need to undertake costly actions to limit their “impulse” spending (or

model calibration. Laibson, Repetto, and Tobacman [2000] carefully calibrate their models with realistic returns, if anything biased in favor of credit card borrowing. Related simulations in Angeletos et al. [2000] of conventional buffer-stock models generate an average MPC out of predictable income, which should be an upper bound to the MPC out of liquidity, of about 0.03–0.06, well below the 0.10–0.14 estimated in this paper. Changing the discounting from exponential to hyperbolic raises the MPC out of predictable income to about 0.17–0.27.

spending by their spouses). By not fully paying off credit card balances, people can reduce their liquidity and thereby help reduce the temptation of available credit. Laibson, Repetto, and Tobacman [2000] develop a more sophisticated model of self-control. They show that hyperbolic consumers are more likely than exponential consumers to simultaneously borrow on credit cards and invest in an *illiquid* asset of a given return. But hyperbolic discounting alone cannot justify simultaneously borrowing and holding *liquid* assets, which violates no-arbitrage. Alternatively, borrowers planning to file for bankruptcy have an incentive to hold some assets, up to the amounts protected by the bankruptcy exemption rules [Lehnert and Maki 2000]. However, such strategic bankruptcy planning should be undertaken only just before filing, to avoid paying unnecessary interest.

The remaining columns of Table VI show how these asset holdings vary with demographic characteristics. Each borrowing group holds substantial assets on average. Even for high income and high education households (columns (2) and (4)), over a third of borrowers simultaneously hold more than one month's income in liquid assets. Hence the puzzle of credit card borrowing is not limited to a single demographic group.

VIII. CONCLUSION

This paper has analyzed how consumer debt responds to changes in credit supply. We utilized a unique new data set of credit card accounts to estimate both MPCs out of liquidity and interest-rate elasticities. We also evaluated the ability of different models of consumption and saving to rationalize these estimates, distinguishing the PIH, liquidity constraints, buffer-stock, and behavioral models. The data set separately records credit limits and balances, allowing us to distinguish credit supply and demand, as well as account-specific interest rates. It also provides an unusually rich set of controls and instruments for the endogeneity of credit supply. While measurement error might have thwarted previous studies of credit supply using traditional household data sets, it is much less a problem here. As a result, the "event studies" in this paper provide more powerful tests than usual of whether liquidity constraints and interest rates really matter in practice.

We found that increases in credit limits generate an immediate and significant rise in debt, counter to the PIH. The average

long-run MPC out of liquidity ranges between 10–14 percent. The MPC is much larger for people starting near their credit limit, providing concrete evidence that liquidity constraints are binding. However, even people starting well below their limit significantly respond. We showed that this response is consistent with buffer-stock models of precautionary saving. Nonetheless, there are other results that conventional models cannot easily explain, such as the fact that many credit card borrowers simultaneously hold low yielding assets. Behavioral models might help explain some of the remaining puzzles.

Unlike most other studies, we also found strong effects from changes in account-specific interest rates. The average long-run elasticity of debt to the interest rate is approximately -1.3 . Less than half of this elasticity represents balance-shifting across cards, with most reflecting net changes in total borrowing. The elasticity is larger than average for declines in interest rates, which can explain the widespread use of temporary promotional rates. The elasticity is smaller for people starting near their credit limits, again consistent with liquidity constraints.

These results help quantify the aggregate implications of changes in credit supply. To illustrate, suppose that every account exogenously received the average credit line increase of about \$2000 (the intensive margin). Given the multiplier for automatic line increases in Table II (row (2)), this would lead to about a \$200 increase in debt per bankcard. Since the average household holds over two bankcards, this comes to over \$40B in the aggregate. For comparison, this figure is about 10 percent of the size of the annual increase in aggregate consumption. Offering new credit cards (the extensive margin), especially to liquidity-constrained households, would increase the aggregate effect. Such large responses, however, might also change the equilibrium interest rate. We do not know the magnitude of the resulting general equilibrium effect, but Table III (row (1)) implies that for every one percentage point rise in credit card rates, debt will decline by about \$110 per card, potentially partly offsetting the above response. (Presumably there would be less balance-shifting than in row (11) of Table III if interest rates rise on all cards at the same time.) By contrast, after a monetary expansion lower interest rates would bolster the stimulative effects of larger credit lines. Overall, the results imply that the consumer plays a potentially important role in the transmission of monetary policy and other credit shocks.

We also documented the pervasiveness of liquidity constraints, under different definitions and across different demographic groups. About one-third of households lack bankcards, and well over half of the households with bankcards are rolling over debt and paying high interest rates. Almost 15 percent of bankcard accounts have utilization rates above 90 percent. Most tellingly, for these accounts the average MPC out of liquidity is almost 50 percent. Under most of these definitions, liquidity constraints were found to disproportionately affect young and low income people, and people with low credit scores.

DATA APPENDIX

The unit of analysis is an individual bankcard account. For most accounts, statement totals such as balances are recorded by the monthly billing cycle; for other accounts, they are recorded by calendar month. The credit bureaus store their information by individual borrower, not household.

The main variable definitions are as follows. Let $bal_{i,t}$ be account i 's balance at the end of period t (whether by billing cycle or calendar month). Then i 's debt at the end of period t is defined as $D_{i,t} \equiv bal_{i,t} - payment_{i,t+1}$, where $payment_{i,t+1}$ is i 's payment in period $t + 1$. When the data are tracked by cycle, this is the amount of debt rolled over into the next period. If the data are tracked by calendar month, $D_{i,t}$ is set to zero if no interest charges are incurred in period $t + 1$.³¹ $L_{i,t}$ is the credit limit as of the start of period t . $r_{i,t}$ is the interest rate (annual percentage rate) as of the start of period t , including any promotional rate underway. If there are separate credit limits and interest rates for retail purchases, cash advances, etc., we use the variables for retail purchases. (The credit limit for retail purchases is usually the overall credit limit for the account.) The utilization rate is the ratio of balances to the credit limit, $utilization_{i,t} \equiv bal_{i,t}/L_{i,t}$. A manual line increase is one requested by the consumer; an automatic line increase is initiated by the issuers' automated systems. The internal credit-risk scores are measures of account risk and

31. For balances tracked by calendar month, the main results use $D_{i,t}$ only if nonnegative, but are insensitive to this distinction. For a subsample of accounts we have balances by both calendar month and cycle. In this subsample, this definition of debt by calendar month yields quantitatively very similar results as those using actual debt by cycle. Alternative definitions, including retaining the negative values of $D_{i,t}$, or not subtracting $payment_{i,t+1}$ from $bal_{i,t}$, yield similar but slightly noisier results.

profitability based on account data from the issuer; the external scores are based on data from the credit bureaus across all sources of credit held by the account-holder. The credit bureau variable on other balances records total balances on other credit cards held by the account-holder, including transactions balances, not just debt.

The SCF data are weighted to be representative of the United States, and generally limited to bankcard debt (e.g., not including retail-store credit cards). Bankcard debt comes from variable No. 413: "After the last payment(s) (was/were) made on (this/these) [bankcard] account(s), roughly what was the balance still owed on (this/all these) account(s)?" A household is said to be borrowing on its bankcards if this variable is positive.

LEXECON, INC.

FINANCE DEPARTMENT, THE WHARTON SCHOOL, UNIVERSITY OF PENNSYLVANIA, AND NBER

REFERENCES

- Allen, Todd, and Christopher Carroll, "Learning about Consumption," mimeo, The Johns Hopkins University, 2000.
- Altonji, Joseph, and Aloysius Siow, "Testing the Response of Consumption to Income with (Noisy) Panel Data," *Quarterly Journal of Economics*, CII (1987), 293–328.
- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "The Hyperbolic Buffer-Stock Model: Calibration, Simulation, and Empirical Evaluation," mimeo, Harvard University, 2000.
- Ausubel, Lawrence, "The Failure of Competition in the Credit Card Market," *American Economic Review*, LXXXI (1991), 50–81.
- Browning, Martin, and Annamaria Lusardi, "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature*, XXXIV (1996), 1797–1855.
- Chimerine, Lawrence, "Americans in Debt: The Reality," Mastercard International, 1997.
- Calem, Paul, and Loretta Mester, "Consumer Behavior and the Stickiness of Credit Card Interest Rates," *American Economic Review*, LXXXV (1995), 1327–1336.
- Campbell, John, and N. Gregory Mankiw, "Permanent Income, Current Income, and Consumption," *Journal of Business and Economics Statistics*, VIII (1990), 265–279.
- Carroll, Christopher, "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity* (1992), 61–156.
- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans, "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds," *Review of Economics and Statistics*, LXXVIII (1996), 16–34.
- Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica*, LIX (1991), 1221–1248.
- Board of Governors of the Federal Reserve System, "G.19, Consumer Credit," 1998.
- Gross, David, and Nicholas Souleles, "An Empirical Analysis of Personal Bankruptcy and Delinquency," *Review of Financial Studies*, XC (2002), 319–347.
- Hall, Robert, "Intertemporal Substitution in Consumption," *Journal of Political Economy*, XCVI (1988), 339–357.
- Hall, Robert, and Frederic Mishkin, "The Sensitivity of Consumption to Transi-

- tory Income: Estimates from Panel Data on Households," *Econometrica*, L (1982), 461–481.
- Hubbard, R. Glenn, and Kenneth Judd, "Liquidity Constraints, Fiscal Policy, and Consumption," *Brookings Papers on Economic Activity* (1986), 1–59.
- Jappelli, Tullio, "Who Is Credit Constrained in the U. S. Economy?" *Quarterly Journal of Economics*, CV (1990), 219–234.
- Jappelli, Tullio, and Marco Pagano, "The Welfare Effects of Liquidity Constraints," *Oxford Economic Papers*, LX (1999), 410–430.
- Jappelli, Tullio, Stephen Pischke, and Nicholas Souleles, "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources," *Review of Economics and Statistics*, LXXX (1998), 251–262.
- Kashyap, Anil, and Jeremy Stein, "Monetary Policy and Bank Lending," in *Monetary Policy*, N. Gregory Mankiw, ed. (Chicago: University of Chicago Press, 1994).
- King, Robert G., Charles I. Plosser, and Sergio I. Rebelo, *Journal of Monetary Economics*, XXI (1988), 195–232.
- Laibson, David, Andrea Repetto, and Jeremy Tobacman, "A Debt Puzzle," mimeo, Harvard University, 2000.
- Lehnert, Andreas, and Dean Maki, "The Great American Debtor," mimeo, Board of Governors of the Federal Reserve System, 2000.
- Ludvigson, Sydney, "Consumption and Credit: A Model of Time-Varying Liquidity Constraints," *Review of Economics and Statistics*, LXXXI (1999), 434–447.
- Moore, Mary, "Credit Scoring's Uses Expand as It Gains Acceptance," *The American Banker* (1996), 4A.
- Mishkin, Frederic, "Symposium on the Monetary Transmission Mechanism," *Journal of Economic Perspectives*, IX (1995), 3–10.
- Parker, Jonathan, "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review*, LXXXIX (1999), 959–973.
- Runkle, David, "Liquidity Constraints and the Permanent Income Hypothesis: Evidence from Panel Data," *Journal of Monetary Economics*, XXVII (1991), 73–98.
- Souleles, Nicholas, "The Response of Household Consumption to Income Tax Refunds," *American Economic Review*, LXXXIX (1999), 947–958.
- , "Consumer Response to the Reagan Tax Cuts," *Journal of Public Economics* (2002), forthcoming.
- Yoo, Peter, "Still Charging: The Growth of Credit Card Debt between 1992 and 1995," *Federal Reserve Bank of St. Louis Review* (1998), 19–27.
- Zeldes, Stephen, "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, XCVII (1989), 305–346.