CONSUMPTION AND CREDIT: A MODEL OF TIME-VARYING LIQUIDITY CONSTRAINTS

Sydney Ludvigson*

Abstract—This paper studies the optimal consumption behavior of individuals who face borrowing limitations that vary stochastically with their income. This framework is motivated by new empirical evidence that I document in U.S. aggregate data: predictable growth in consumer credit is significantly related to consumption growth, a finding that is inconsistent with existing models of consumer behavior. The time-varying liquidity constraint model considered here correctly predicts two key properties of the U.S. aggregate data: the correlation of consumption growth with predictable credit growth documented in this paper, and the well-known correlation between consumption growth and predictable income growth that has been documented extensively elsewhere.

I. Introduction

As a fraction of personal income, consumer credit has more than doubled during the postwar period, with particularly sharp and sustained increases occurring in the 1980s (see figure 1). This intensive use of consumer credit has recently lead to considerable speculation over a variety of ways in which consumer indebtedness may have an independent affect on real consumer expenditure. Despite the growing attention paid to ballooning consumer indebtedness, little formal research has been devoted to investigating the impact of credit aggregates on aggregate consumption. And, although the nature of consumer credit use has undoubtedly changed in recent years, any supposition that household borrowing is causally related to real activity raises both empirical and theoretical questions.

On the empirical side, what, if any, are the qualitative dynamic effects of consumer credit on consumption, and how should they be characterized? From a theoretical point of view, such effects would seem to be at odds with the certainty-equivalent, permanent income hypothesis (PIH) in which the demand for credit passively responds to the demand for consumption (Hall, 1978; Flavin, 1981). What kind of theoretical framework can be employed to think about how variation in consumer credit might instead influence consumption?

This paper has two objectives. One is to address the empirical question raised above by characterizing the time-series relationship between consumption and aggregate measures of consumer credit. It is no surprise that ex post consumer credit is correlated with consumption; it is a simple artifact of the consumer's budget constraint that, fixing income, debt will increase with consumption contemporaneously. If consumer credit plays an important role in generating or propagating fluctuations in real variables, something other than a simple ex post correlation should be evident in the data. By contrast, theory suggests that contemporaneous variables such as consumer credit should play no role in explaining the first difference in consumption: Hall (1978) concluded that a close approximation to the stochastic behavior of consumption under the permanent-income hypothesis (PIH) was that—conditional on lagged consumption—no other variable observed in earlier periods should have predictive power for current consumption.

Nevertheless, it is a well-known empirical fact that some contemporaneous variables do help forecast the change in consumption. Most notably, changes in aggregate consumption are significantly correlated with lagged, or predictable changes in income (Flavin, 1981; Campbell & Mankiw, 1989, 1991; Deaton, 1992; Attanasio & Weber, 1993). Thus, consumption is said to exhibit "excess sensitivity" to predictable income growth.

This paper takes an empirical approach that sets up the PIH as the null hypothesis, and tests whether predictable changes in consumer credit have any influence on the change in consumption that is independent of the variation in predictable income. A principal empirical finding of this paper is that forecastable, or ex ante, consumer credit growth has a significant influence on consumption that is independent of the variation in predictable income growth.[1]

The association of forecastable changes in credit with changes in consumption is not readily explained by existing models of consumer behavior. For example, the finding is inconsistent with popular formulations of the PIH, according to which the change in consumption is—at least to an approximation—unforecastable.[2] Hall and Mishkin (1982) and Campbell and Mankiw (1989) modify the PIH framework by suggesting that the data are generated by two types of consumers: those who consume their permanent income, and "rule-of-thumb consumers" who consume their current income. Although this simple paradigm can explain the

Received for publication October 27, 1997. Revision accepted for publication June 11, 1998.

*Federal Reserve Bank of New York.

I am grateful to Angus Deaton, Ben Bernanke, and John Y. Campbell, for their advice and many valuable discussions; to Phillippe Bacchetta, Christopher Carroll, Karen Dynan, Jeffrey Furher, Mark Gertler, Jordi Gali, Angelo Melino, Scott Schuh, and Joe Tracy for constructive suggestions; to Beethika Khan and Reagan Murray for excellent research assistance; and to the Alfred P. Sloan Foundation for financial support. The views expressed in the paper are those of the author and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the author.

[1] Since a previous version of this paper was circulated, Bacchetta and Gerlach (forthcoming) independently document similar findings using international data.

[2] One such approximation is that marginal utility is linear and the real interest rate constant and equal to the rate of time preference. Others assume that the log, rather than the level, of consumption follows a martingale (Mankiw, 1981; Hansen & Singleton, 1983; Campbell & Mankiw 1989, 1991), and that the variance of consumption growth is part of a constant drift. Zeldes (1989b) and Carroll (1997) suggest (but don't demonstrate) that the excess sensitivity of consumption growth to forecastable income growth may be explained by nonlinearities in the marginal utility function.

The Review of Economics and Statistics, August 1999, 81(3): 434–447

FIGURE 1.—RATIO, CONSUMER CREDIT TO PERSONAL INCOME

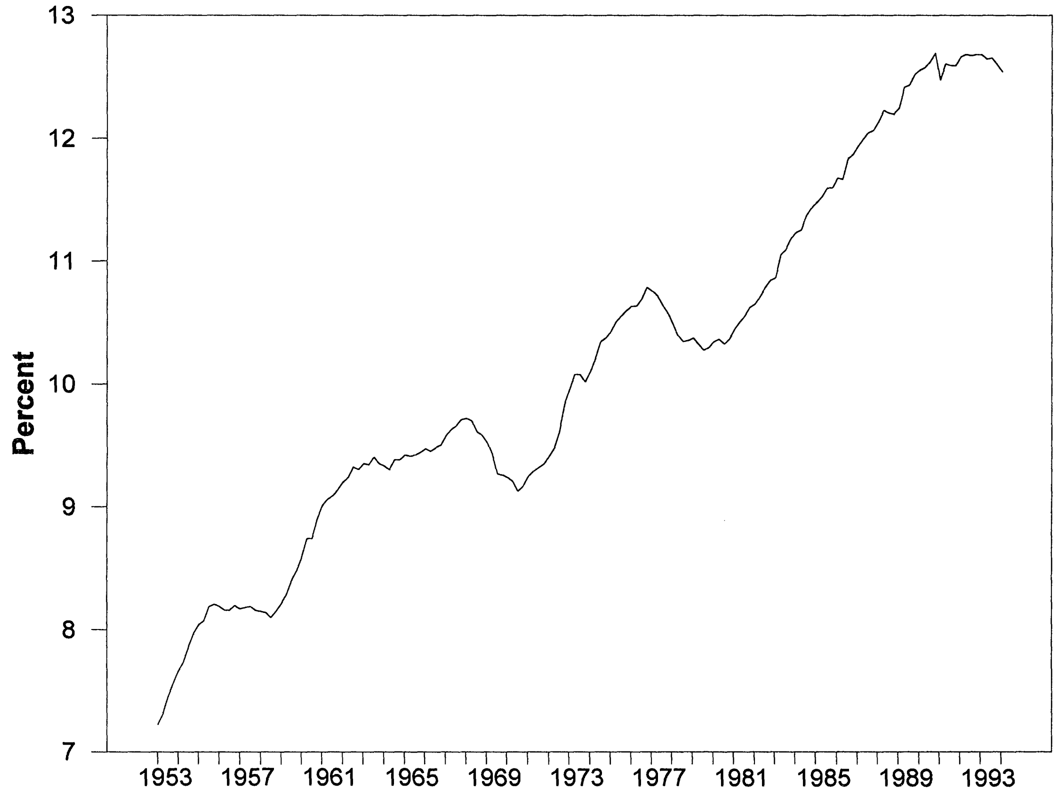

excess sensitivity of consumption growth to forecastable income growth, it does not deliver a correlation between consumption growth and predictable credit growth. Rule-of-thumb consumers are not directly influenced by credit availability because they are assumed not to borrow.

Other researchers have analyzed the behavior of optimizing consumers who face restrictions on their ability to borrow. Deaton (1991) shows that such a model will deliver a correlation between consumption growth and predictable income growth if individuals face a fixed borrowing constraint. Yet fixed borrowing constraints cannot explain the correlation between consumption growth and predictable credit growth. To see why, consider an extreme example in which individuals are so impatient that they are constrained in every period. If a solution to this model is obtainable, consumption growth will be very highly correlated with predictable income growth, but not with predictable movements in credit, because actual borrowing (equal to the upper limit) is fixed.[3]

Thus, a second objective of this paper is to provide a theoretical framework for thinking about the association of consumption growth with predictable movements in credit. My analysis follows in the tradition of studies like Deaton (1991) and Carroll (1997), who employ "buffer stock" versions of the modern-day PIH to explain well-known features of the micro and macro consumption data. My analysis departs from those studies in its treatment of credit availability: whereas those models analyzed the effects of fixed borrowing limits (either exogenously or endogenously imposed), this study focuses on the behavior of consumers whose access to credit varies stochastically with income. A key finding of this paper is that there must be sufficient random variation in the credit ceiling to explain the correlation between consumption and credit documented here.

Allowing the borrowing limit to vary with current income is consistent with the lending practices of banks.[4] Moreover, the notion that individuals' access to external finance should vary with their current financial position is consistent with a large and growing literature on investment behavior. This literature documents the importance of the firm's current cash-flow in determining how much it can borrow and invest (Fazzari et al., 1988; Bernanke & Gertler, 1989; Bernanke et al., 1996).[5]

[3] Even when individuals are less impatient, the next section shows that, when income is nonstationary, there is not a stationary solution to the intertemporal optimization problem when the debt limit is fixed but nonzero; such a solution requires that the debt limit-income ratio be stationary.

[4] For example, Jappelli (1990) studies the *Survey of Consumer Finances* and finds that a major reason households report being denied access to credit is that they had an insufficient level of income.

[5] If lenders cannot costlessly acquire information about the ability of borrowers to repay, then one would expect funds provided by the lender to be associated with some observable measure of the borrower's financial health. Although permanent income is a more-relevant indicator of a consumer's ability to repay longer-term loans, borrowing restrictions could instead vary with current income because creditors do not observe

The model considered here correctly predicts two distinct features of the aggregate data: the statistically significant correlation of consumption growth with predictable credit growth documented in this paper, and the well-known correlation of consumption growth with predictable income growth that has been documented extensively elsewhere.

The rest of this paper is organized as follows. Section II presents the model of individual behavior when the upper borrowing limit varies stochastically with current income. Section III turns to the U.S. data, presents the results of estimating the time-series relationship among aggregate consumption growth, forecastable credit growth, and forecastable income growth, and summarizes the empirical facts. Section IV shows how the model of individual behavior presented in Section II is aggregated, and compares its time-series properties with those of the U.S. aggregate data. Section V concludes.

II. A Model of Time-Varying Liquidity Constraints

Consider a model of intertemporal choice in which individuals maximize the infinite sum of instantaneous utility functions over a nondurable consumption good C_t,

$$E_t\left(\sum_{i=0}^{\infty}\beta^i\left(\frac{C_{t+i}^{1-\gamma}}{1-\gamma}\right)\right) \qquad \beta < 1, \tag{1}$$

with $\beta = (1+\delta)^{-1}$, where δ is the strictly positive rate of time preference. Unlike the certainty-equivalent PIH, this specification yields convex marginal utility, implying that there is a precautionary motive for saving. The consumer can take on debt, which evolves according to

$$D_{t+1} = (1+r)(D_t + C_t - Y_t), \tag{2}$$

where r is the real interest rate at which consumers can borrow, treated as fixed and known.[6]

Individual labor income is uncertain, and I assume it follows an exogenous stochastic process. This process is specified so that (log) income is the sum of a white-noise transitory component, and a random walk with drift "permanent" component. In levels, $Y_t = Y_{1,t}Y_{2,t}$, or in logs, $y_t = y_{1,t} + y_{2,t}$, where $y_{1,t}$ is a random walk with positive drift g, and $y_{2,t}$ is white noise. (Throughout the paper, lowercase letters refer to variables in logs.) This representation is equivalent to the log first difference of income being a moving average,

$$\Delta y_t = g + \eta_t - \psi\eta_{t-1}. \tag{3}$$

Section IV discusses how this income process is decomposed into idiosyncratic and common components.

To complete the description of the consumer's decision problem, one must specify how the borrowing constraint varies over time. I assume individuals are impatient, in the sense that borrowing will be part of the optimal plan.[7] The consumer's upper borrowing limit is stochastic and dependent on current income. However, the upper borrowing limit is not perfectly correlated with income; instead, I assume there is economy-wide variation in the amount individuals can borrow relative to their income, so that the consumer's upper borrowing limit takes the following form:

$$D_{t+1} \leq \overline{D}_{t+1} \equiv \frac{1}{\omega} Y_t \exp(\xi_t), \tag{4}$$

where ω is assumed fixed and known and ξ is a shock to the credit ceiling that is common to all households and independent of individual income. Note also that, because the evolution of debt given in equation (2) specifies that D_t is beginning of period debt, the timing assumption designates that the upper limit on credit at time t is proportional to income at time $t-1$. Deaton's (1991) model arises as a special limiting case when ω approaches infinity.

ξ_t is assumed to vary according to an AR(1) process,

$$\xi_{t+1} = \phi\xi_t + u_{t+1}, \tag{5}$$

where u_t is a white-noise random shock with mean zero and constant variance, and $0 < \phi < 1$. Equations (4) and (5) together imply that, on average, the individual's credit limit is a fraction, $1/\omega$, of current income, but is also subject to aggregate shocks that are independent of income and die out gradually.

The random variable ξ introduces economy-wide variation in the credit ceiling that is independent of individual income. Ludvigson (1998) finds evidence that economy-wide variation in the supply of consumer credit exists, and influences aggregate consumer spending. In that paper, movements in the supply of consumer credit are identified

permanent income. In practice, many banks require a minimum of two pay stubs in order to obtain an increased borrowing limit on debt like credit cards, even though an explicit job contract confirming higher income has been agreed upon in advance of the first pay period.

[6] Unless r is specified as the rate at which consumers can borrow and D restricted to be positive, then equation (2) gives the evolution of debt net of assets. In reality, consumer credit cannot be described as the negative of most of the assets held in the U.S. economy, so that assets and debt are distinct stocks. To make equation (2) consistent with the definition of debt as gross debt in the data, I restrict D to be nonnegative. This constraint never binds however, because agents are sufficiently impatient that they always choose to borrow a nonnegative amount, though they will not necessarily borrow up to their credit limit in every period. A challenging direction for future research would be to model the behavior of individuals who choose to hold both consumer credit and positive assets.

[7] Specifically, impatience is defined as follows:

$$\gamma^{-1}(r-\delta) + \gamma\sigma_y^2/2 < g,$$

where Δy is $N(g, \sigma_y^2)$. See Deaton (1991).

by shifts in the ratio (or "Mix") of consumer credit issued by banks, to that issued by non-banks.[8] Variation in this ratio is largely independent of income variables and has an effect on aggregate consumption that is not explained by variation in conventional demand indicators such as income and interest rates.[9] The Mix is also highly correlated with measures of consumer credit availability from the *Senior Loan Officer Opinion Survey,* conducted by the Federal Reserve.[10] Variation in the Mix is used below to calibrate the ξ process, with the persistence parameter and innovation variance chosen to match the persistence and innovation variance of the first difference in the Mix. Calibration assumptions for all of the model's parameters are discussed below.

As there is no known analytical solution to the model presented above, I seek a numerical solution. Deaton (1991) outlines one way to set up the problem when borrowing restrictions are fixed, and I follow a similar procedure here to solve the problem with time-varying constraints.

To illustrate the numerical problem, begin by denoting the instantaneous marginal utility of consumption as $v(c_t)$, i.e.,

$$v(C_t) \equiv u'(C_t), \tag{6}$$

where $u'(\cdot) = C_t^{-\gamma}$. Define the consumer's upper spending limit at time t as X_t (referred to from here on as "current resources"). From equation (2) with $D_{t+1} = \bar{D}_{t+1}$, we have

$$X_t \equiv \frac{\bar{D}_{t+1} - D_t(1 + r)}{(1 + r)} + Y_t. \tag{7}$$

Equation (7) says the most that can be spent on consumption in period t is all of that period's income, and up to the credit limit less the interest and principal owed on outstanding debt.[11] The first-order condition that imposes the constraint (4) can then be written as the following Euler equation

$$C_t^{-\gamma} = \max\,[X_t^{-\gamma}, \beta^* E_t C_{t+1}^{-\gamma}], \tag{8}$$

where $\beta^* = (1 + r)/(1 + \delta)$. Since X_t is the most that can be consumed, marginal utility can never be less than $X_t^{-\gamma}$. Without the borrowing limit, equation (8) collapses to the usual Euler equation where the first element in brackets is eliminated.

Nonstationary income will result in nonstationary processes for consumption and household debt. In order to solve for a stationary policy function over these variables, they must first be transformed into stationary ratios of variables. Define z_{t+1} as

$$z_{t+1} \equiv Y_{t+1}/Y_t = \exp\,(g + \eta_{t+1} - \psi\eta_t).^{12} \tag{9}$$

Using this notation, the Euler equation (8) can be rewritten as

$$\theta_t^{-\gamma} = \max\,[w_t^{-\gamma}, \beta^* E_t z_{t+1}^{-\gamma} \theta_{t+1}^{-\gamma}], \tag{10}$$

where, $\theta_t \equiv C_t/Y_t$, $w_t \equiv X_t/Y_t$.

The resource-to-income ratio, w, evolves according to

$$w_{t+1} = (1 + r)(w_t - \theta_t)z_{t+1}^{-1} + \frac{\bar{D}_{t+2} - \bar{D}_{t+1}(1 + r)}{(1 + r)Y_{t+1}} + 1. \tag{11}$$

The time-varying borrowing limit, $\bar{D}_t$, directly influences the evolution of resources through the second-to-last term in equation (11). This term is absent when borrowing is prohibited ($\bar{D} = 0$). The consumer's upper spending limit increases when the debt ceiling today is higher than the gross interest rate times the debt limit last period.

Equation (11) can be used to show why income must be cointegrated with the borrowing limit to obtain a stationary solution to the optimal consumption policy. Dividing the numerator and denominator of the second-to-last term in equation (11) by $\bar{D}_{t+1}$, and using the definition of $\bar{D}_{t+1}$, one obtains the following expression for w_{t+1}:

$$w_{t+1} = (1 + r)(w_t - \theta_t)z_{t+1}^{-1} + \frac{\exp\,(\Delta\bar{d}_{t+2}) - (1 + r)}{(1 + r)\exp\,(y_{t+1} - \bar{d}_{t+1})} + 1. \tag{12}$$

As equation (12) illustrates, $y_{t+1} - \bar{d}_{t+1}$ must be stationary in order for w_t to be stationary, requiring $\bar{d}_{t+1}$ and y_{t+1} to be cointegrated. Note this implies that, if y_{t+1} is nonstationary, it is not possible to solve for a stationary policy function when the borrowing limit is fixed but nonzero.

[8] It is necessary to use loans for a particular consumer product in constructing the Mix; the most widely available data on consumer credit is for the purchase of automobiles. It is assumed that shifts in the aggregate supply of credit for automobiles are a good proxy for shifts in the supply of consumer credit more generally.

[9] Variance decompositions show that variation in the Mix is almost entirely determined by its own innovations.

[10] The measure of credit availability in this survey is the net percentage of bank loan officers who report a greater willingness to make consumer installment loans relative to the past three months. The first difference in the Mix variable is found to Granger-cause the credit availability measure at better than the 5% level of significance.

[11] Note that equation (7) will never be nonpositive. This is because, with isoelastic preferences, there is an infinite penalty for zero consumption. The consumer knows the distribution of income and so never chooses the endogenous variable D_t to be so large that X_t is nonpositive.

[12] I assume the support of z is finite, and that there is a strictly positive floor on income, forcing z to be bounded away from zero, $z > z_o = 0$. This circumvents the voluntary no-borrowing result in models like those of Carroll (1997), although households in this model will never borrow so much that there is a possibility (given the known distribution of income) for tomorrow's resources to fall to zero after payments on outstanding debt are made.

Given the evolution of the stationary resource ratio, w, it is now possible to look for a stationary solution to the recursive problem in equation (10), in which the consumption-to-income ratio, θ, is a function of the state variables in the model. One such state variable is w_t itself, since its evolution can be used to predict the expected marginal utility of consumption tomorrow. In addition, since income growth (3) is serially correlated, the expected marginal utility of tomorrow's consumption-income ratio will also depend on η_t, the component of income growth known at time t. Similarly, ξ_t is a state variable since autocorrelation in the credit shock process implies that ξ_t forecasts ξ_{t+1}. The optimal consumption decision is function of all three state variables: $\theta = f(w_t, \eta_t, \xi_t)$. If we define $p(w_t, \eta_t, \xi_t) = v[f(w_t, \eta_t, \xi_t)] = f(w_t, \eta_t, \xi_t)^{-\gamma}$, and substitute these definitions along with equation (12) into equation (10), we obtain the following set of recursive equations that defines the individual's decision problem:

$$
\begin{aligned}
p(w_t, \eta_t, \xi_t) \ \max\Bigg[& v(w_t)\beta^* E_t z_{t+1}^{-\gamma} p \\
& \times \Bigg\{ 1 + z_{t+1}^{-1}(1+r)(w_t - v^{-1}p(w_t, \eta_t, \xi_t)) \\
& + \frac{\exp(\eta_{t+1} - \psi\eta_t + \xi_{t+1} - \xi_t + g) - (1+r)}{(1+r)\exp(\eta_{t+1} - \psi\eta_t + g + \ln\omega - \xi_t)}, \\
& \eta_{t+1}, \xi_{t+1} \Bigg\} \Bigg],
\end{aligned}
\tag{13}
$$

where z_{t+1} is given in equation (9), and the conditional expectation is taken with respect to the variables η_{t+1} and ξ_{t+1}.[13] This equation says that consumers take into account both their current income innovation, and the independent shock to their credit ceiling when forming expectations about their credit limit next period.

The optimal consumption policy of an individual household, $f(w_t, \eta_t, \xi_t)$, is obtained by solving the set of recursive equations in (13). If the underlying distributions to η and ξ are continuous, the expectational integral in set equations (13) can in principle be solved using Simpson's rule or a similar approximation. In practice, this turns out to be computationally intractable since the policy function has to be evaluated over three state variables. To reduce computational complexity, η_t and ξ_t are specified as discrete Markov processes by making an m point discrete approximation of an underlying normal distribution for η_t and ξ_t, each with mean zero and constant variances, σ^2 and σ_ξ^2, respectively. Setting $m = 5$, the discrete approximation produces $m * m = 25$ income-credit states over which set equations (13) is evaluated. Details about the numerical solution procedure are provided in appendix A.

[13] Deaton and Laroque (1992) have already demonstrated that, if the impatience condition discussed above holds, a unique optimum policy will exist for the type of recursive equation in (13) in the case of a fixed borrowing constraint that prohibits borrowing.

A. *Characteristics of the Optimal Consumption Function*

This section discusses properties of the optimal consumption function of a single household. Before this solution can be obtained, values must first be assigned to the underlying parameters. I discuss calibration assumptions next and then describe the household's optimal consumption policy.

Calibration: As a baseline, I begin by using estimates from MaCurdy (1982) and Pischke (1995) who study individual income using data from the *Panel Study of Income Dynamics.* Both of these studies find that ψ, the moving average parameter in equation (3), is approximately 0.44. Although MaCurdy estimated the standard deviation of income growth to be about 0.25, I follow others (e.g., Deaton, 1991; Attanasio et al., 1995) in considering much lower values for this parameter in the simulations.[14] The simulations below use values for σ equal to 0.05 and 0.10. I also set $r = 0.03$, $g = 0.02$, $\gamma = 2$, and $\delta = .15$.[15] The values of these parameters are varied in several simulations reported below.

To calibrate parameters of the aggregate credit shock, ξ_t, a first-order autoregressive process for the change in the Mix variable, discussed above, is estimated. The persistence of ξ is set equal to 0.6 to match the persistence of the first difference in the Mix, while the standard deviation of the innovation, u_t, is set equal to 0.01 to match the innovation variance of the first difference in the Mix.

To obtain a reasonable value for the parameter ω, summary statistics from three waves of the *Survey of Consumer Finances* (SCF) were computed. Appendix B contains details on the variables used from this survey. This survey was chosen because it is the only one that provides direct information about the maximum quantity of credit that consumers can borrow. In particular, the 1989, 1992, and 1995 waves contain a variable that asks households what their total credit limit was on all of their credit-card accounts. For each wave of the survey, this variable was divided by the household's wage and salary income for that year to obtain a value for the household's credit limit-income ratio. In rough analogy to the notation used in equation (5), I will denote the summary value, across households, of the inverse of this ratio as $\hat{\omega}$, even though, in the model, this ratio also varies over time for any given household. By looking across waves of the SCF, one can get an idea of the time-average value of $\hat{\omega}$.

In using these data from the SCF, several measurement issues should be taken into account. First, in order to obtain a less-noisy measure of the net worth characteristics of households, the SCF heavily oversamples wealthy and high-income households. To compensate for this, the survey provides a weighting scheme that I use in calculating $\hat{\omega}$.

[14] As Deaton (1991) and Pischke (1995) point out, estimates of the standard error of income growth are likely inflated as a result of substantial measurement error in recorded income.

[15] Note in the steady state defined by $\eta_t = 0 \forall t$, and where the debt to income ratio is constant, the presumption that $r > g$ rules out Ponzi games.

FIGURE 2.—STOCHASTIC DEBT LIMIT—INCOME RATIO $\omega = 15$, $\sigma = 0.05$, $\sigma_u = 0.01$

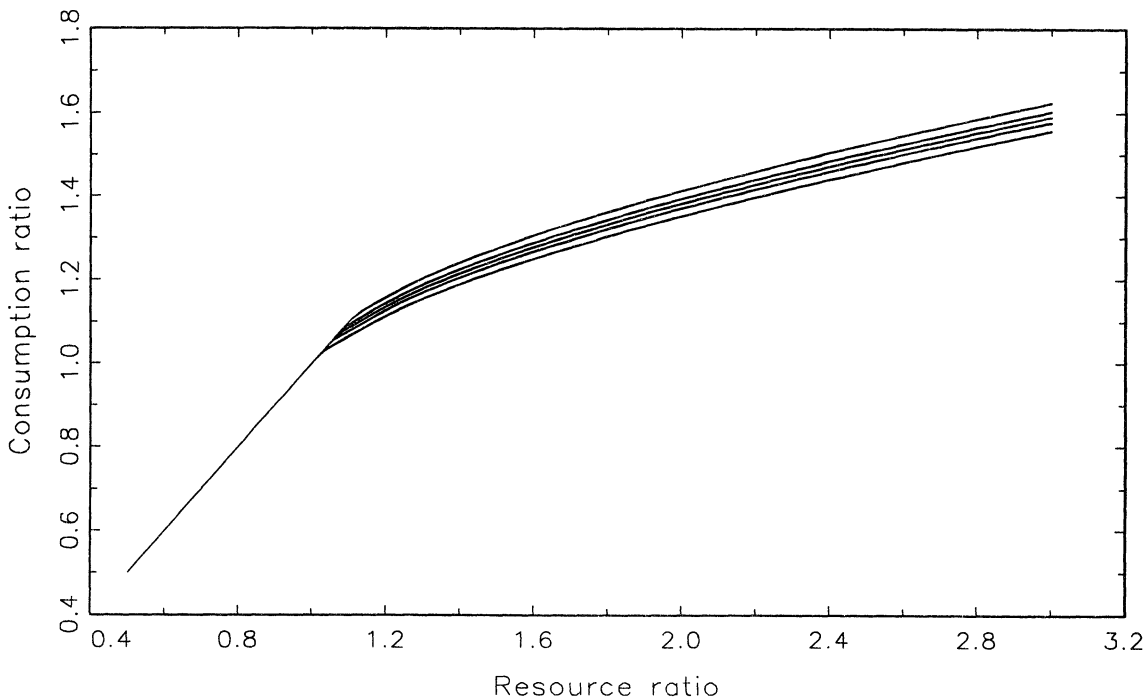

Second, the buffer stock model is not appropriate for describing the behavior of individuals who accumulate large amounts of liquid assets. Moreover, if buffer stock savers tend to hold few assets, they might have higher values of ω (lower debt limit-to-income ratios) than the median household in the SCF. Third, the SCF provides information on whether households were actively denied credit, suggesting one way to split the sample according to groups that were actually constrained. With these considerations in mind, we looked at the median value and the first quantile (25%) of the ratio described above, for the whole sample, and for various subgroups of households, split according to liquid asset holdings, or according to whether they reported being turned down for credit or discouraged from borrowing.

The typical value of $\hat{\omega}$ in the SCF varies from year to year—consistent with the hypothesis that there is income-independent variation in the credit ceiling. In 1989, the median value for the full sample was approximately 8, while the value for the first quantile was about 14. In 1992 and 1995, these figures were 7 and 14, and 5, and 10, respectively. For individuals who held less than 2% of their annual wage and salary income in liquid assets,[16] the ratios were generally higher, equal to 10 and 17 in 1989, 10 and 18 in 1992, and 6 and 13 in 1995. Selecting out households who reported being either turned down for credit, or approved for less credit than they requested, the 50th and 25th percentiles of $\hat{\omega}$ were 8 and 17 in 1989 and 1992, and 6 and 13 in 1995. Finally, for households who reported being discouraged from borrowing because they anticipated being denied credit, the analogous values for $\hat{\omega}$, are 8 and 15 in 1989, 9 and 20 in 1992, and 8 and 18 in 1995. Summarizing these results, it seems reasonable to consider a range of values for ω in the time-series simulations reported below. Section IV considers values for ω equal to 8, 10 and 15.

Policy Functions: Figure 2 shows the set of optimal policy functions corresponding to baseline parameters and with $\sigma = 0.05$ and $\omega = 15$. θ_t, the consumption-income ratio, is plotted against ω_t, the resource-income ratio. There are actually 25 different functions relating these two variables, one for each value of η and ξ. However, the functions that vary according to ξ are sufficiently similar that they are not visibly discernable on the graph, so they appear to lie along a single line corresponding to each value of the income shock, η. This is because ξ has a much smaller variance than η. The functions form five clumps of five branches, where each clump corresponds to a particular value of η, and each branch in each clump corresponds to a particular value of ξ.

The figure illustrates several important properties of optimal consumption subject to a time-varying borrowing constraint. First, when resources are sufficiently low, their marginal value today is higher than the marginal value of consumption tomorrow, and the policy function is a 45 deg. line. Although this feature is also present in the simpler case when borrowing is prohibited, in this case, the individual spends all of current income and borrows up to the current credit limit net of the interest and principle owed on accrued debt along the 45 deg. line. Consequently, spending over this region is directly determined by the time-varying credit

[16] The definition of this variable, along with the other variables constructed from the SCF, is given in appendix B.

ceiling. It is this feature of the model that is crucial for replicating the empirical correlations discussed in the next section.

Second, the policy functions show that, when resources are sufficiently high, the relationship between consumption and the individual's current financial position flattens out and consumers are unconstrained. Along this portion of the policy function, transitory changes in income are smoothed out by accumulating and decumulating debt: a high innovation in η_t implies low income growth next period, so that the lowest value for η corresponds to the highest "clump," and the highest value for η corresponds to the lowest clump. But consumers don't simply smooth out transitory changes in income when unconstrained, they also respond optimally to transitory changes in their current credit limit. Notice from equation (13) that the resource ratio, w_t, depends on the log difference in the credit market process ξ_t, which contains a negatively autocorrelated component as long as ξ_t itself is stationary and positively autocorrelated. Consequently, when growth in ξ_t is high today, it is expected to be low tomorrow, so that, for any given value of η, the highest outcome for ξ corresponds to the lowest branch and vice versa.

There are two channels through which variation in the credit limit affects consumption. The variable on the horizontal axis, ω_t, depends directly on the time-varying credit limit, so variation in the credit ceiling affects consumption through its influence on current resources. The credit limit also affects consumption independently of current resources: for any given value of ω_t, individuals alter their consumption in response to innovations in ξ_t. Even when consumers are currently unconstrained, today's value of ξ affects current consumption because it affects the likelihood of being constrained tomorrow.

How do the optimal consumption functions change when the average debt limit-income ratio, ω, is varied? Perhaps ironically, increasing the average ability to borrow by decreasing ω induces individuals to behave more prudently. Although consumers can borrow a larger amount on average when ω is lower, the credit limit is now more sensitive to any given change in η and ξ; small shocks to income or the credit limit translate into large swings in the quantity of credit available, introducing greater volatility into the consumer's resources. Thus, when ω is lowered from its baseline value, households smooth out a larger portion of any given transitory shock, so that policy functions (not shown) display branches that are more widely spaced than in the baseline case, even though the variance of η and ξ is held fixed at the baseline value.

III. Empirical Results on Aggregate Consumption and Aggregate Credit

This section is devoted to establishing time-series properties of aggregate consumption and consumer credit in U.S. aggregate data. Section IV will ask whether the empirical findings documented in this section arise as implications of the model.

The empirical strategy follows much previous literature by estimating consumption growth equations in which the null hypothesis is set up as a loglinear version of the PIH. One case that has received a great deal of prior consideration assumes that asset returns and consumption are jointly log normal, implying the following relationship between ex ante consumption growth and the expected real interest rate:

$$E_{t-1}\Delta c_t = \mu_t^* + \sigma E_{t-1} r_t, \tag{14}$$

where r_t denotes the net real interest rate, $\sigma \equiv 1/\gamma$, and μ_t^* is typically assumed constant $= \mu^*$.[17] Many researchers have augmented equation (14) by adding either lagged or predictable income growth as an additional explanatory variable (Flavin, 1981, Campbell & Mankiw 1989, 1991; Deaton, 1992). For example, Campbell and Mankiw (1989) add expected income growth to the equation:

$$\Delta c_t = \mu + \lambda E_{t-1}\Delta y_t + \pi E_{t-1} r_t + \epsilon_t. \tag{15}$$

Previous empirical studies have assumed that μ is constant at the aggregate level, implying that equation (15) reduces to a "loglinear" PIH when $\lambda = 0$, and the error term, ϵ_t, is orthogonal to all variables known at time $t - 1$ or before. All of these studies find that they can strongly reject the null hypothesis of $\lambda = 0$.[18]

This paper extends the empirical analysis above to include predictable credit growth as an additional explanatory variable. I estimate a debt-augmented version of equation (15): [19]

$$\Delta c_t = \mu + \lambda E_{t-1}\Delta y_t + \pi E_{t-1} r_t + \alpha E_{t-1}\Delta d_{t+1} + \epsilon_t.^{19} \tag{16}$$

Equation (16) is estimated using U.S. quarterly times-series data. As is customary in the macro literature, I use nondurables and services expenditure as a measure of consumption in these tests.[20] Two alternate measures of income are used for y_t: disposable personal income, and "labor income," composed of wages and salaries plus transfers minus personal contributions for social insurance.

[17] In fact, μ^* is not constant because it includes the conditional variance of consumption growth. Assuming it is constant is equivalent to assuming asset returns and consumption growth are homoskedastic. However, Carroll (1997) and Carroll and Kimball (1996) have recently shown that, with isoelastic utility, precautionary motives imply that the variance of consumption growth is endogenous and depends on wealth. It is not possible, however, to measure this variance term in aggregate data, because the appropriate variable is the variance of consumption growth at the micro level, a variable that cannot be obtained from macro data since the aggregate of the variances is not equal to the variance of the aggregate.

[18] Campbell and Mankiw interpret their results as supporting rule-of-thumb consumption and λ as approximating the fraction of income accruing to agents who consume their current income in every period.

[19] Debt can be measured at the beginning of the period or the end of the period. I adopt the timing convention that D_t is beginning-of-period debt. That is, D_{t+1} accrues when C_t is greater than Y_t. The rest of the paper uses this timing convention in denoting predictable debt growth as $E_{t-1}\Delta d_{t+1}$.

[20] All variables are real, in per capita terms. Consumer installment credit is obtained from the Federal Reserve Board's G19 statistical release.

For comparison with the previous literature, the interest rate, r_t, used to obtain the results reported below is the quarterly average of the three-month Treasury bill rate, using the consumer price index to convert to real rates. As a robustness check, two lending rates were substituted for the Treasury bill rate: the average prime lending rate, and the average credit card rate at commercial banks.[21] Results (not reported) indicated no appreciable difference in the outcome based on the choice of interest rate. In particular, neither of the ex ante borrowing rates exhibited a significant effect on consumption growth.

D_t is consumer installment credit collected by the Board of Governors of the Federal Reserve. Installment credit is defined as short- and intermediate-term loans, not secured by real estate, and scheduled to be repaid in two or more installments. Results reported below use both total consumer installment credit, and the revolving credit component (credit-card debt) of this series.[22]

In using the revolving credit component, an issue arises about when to start the sample. Unlike the broader measure of installment credit, the widespread use of credit cards is thought to be closely related to a 1978 Supreme Court ruling that relaxed and removed interest-rate ceilings on bank cards, and allowed issuers to charge a variety of fees for them. This deregulatory shift not only increased the amount of credit available to those who already had access to it, but it greatly widened the availability of revolving credit to individuals who are more likely to face binding constraints, including young consumers, and those with lower incomes (Staten & Johnson, 1995).[23] Consequently, it seems appropriate to use the revolving credit component in the regressions above since 1978:1. Using the total installment credit series, the sample period is 1953:1 to 1993:1.

Before equation (16) can be estimated, two specification issues need to be addressed. First, instrumental variables estimation must be used because the conditional expected values of the explanatory variables are not observable. Although the actual values might be good proxies, OLS is inappropriate since the error term will be correlated with the explanatory variables if the PIH is true. I include lagged values of consumption growth, income growth, interest rates, credit growth, and an "error correction" term, the lagged log consumption to income ratio as instruments in the estimation.[24]

A second specification issue concerns the use of time-averaged data. It is well known that data at quarterly frequency is time averaged, leading the error term in equation (21) to follow an MA(1) process (Christiano et al., 1991). This first-order serial correlation in the error term may lead to inconsistent estimates if one-period lagged instruments are used. One possible estimation strategy is to proceed with linear instrumental-variables (IV) estimation, but use only instruments that have been lagged at least two periods.[25] This strategy is taken below.

Table 1 presents first- and second-stage results using total installment credit, and analogous results using revolving credit.

As table 1 shows, when consumption growth is estimated as a function of expected credit growth, the point estimates of α are approximately 0.10, and statistically significant at better than the 0.2% level. Moreover, predictable credit growth remains a strongly statistically significant determinant of consumption growth regardless of the particular instrument combination used, or whether expected income growth and the ex ante real interest rate are included as additional explanatory variables. Consumption is "excessively sensitive" to ex ante credit growth.

Table 1 also shows that predictable credit growth and predictable income growth influence consumption independently. The last three columns of table 1 indicate that the coefficients on both of these variables are individually statistically significant, with the coefficient on predictable income growth estimated to be approximately 0.3. An exception occurs in the fifth column of results where income is measured by personal disposable income; in this case, the explanatory power of predictable income growth is compromised by the inclusion of expected credit growth.[26] The reason for this appears evident from the first-stage regression results: disposable income growth is poorly forecast by the instrument set used in that column. To better forecast the

[21] The prime lending rate is a relevant borrowing rate for consumers since approximately 70% of bank card debt is tied to an adjustable index, typically the prime rate (Staten & Johnson, 1995).

[22] A discrepancy exists between the model and the data because some of the credit in this series is used to purchase durable goods while the model posits a single nondurable good. Implicit in the use of the installment-credit series is the assumption that some of the debt is used to purchase nondurables and services. The fact that the installment-credit series contains debt for durables may introduce some noise between credit and consumption of nondurables and services, arguably making it less likely to find a correlation between the two in the data. Unfortunately, it is not possible to dissagregate credit into components that are used solely for the purchase of nondurables and services. Even the revolving-credit component, which excludes auto debt, is used to make both durable and nondurable expenditures. Credit for durable goods may still influence consumption of nondurables and services however; if consumers receive separable utility from both durables and nondurables, consumer credit is "fungible" in the sense that credit for one good simply frees up income for another. Moreover, although the fungibility of consumer credit may not be quite applicable in borrowing for durable goods that can be collateralized (such as automobiles), this is not an issue for the results reported below using revolving credit because that component consists almost entirely of unsecured loans.

[23] In 1968, revolving credit was less than 2% of total installment credit; in 1977, it was still only 7% of total, but, by 1978, it jumped to 15% of total. In 1996, revolving credit was the largest component of total installment credit at 42% of total.

[24] A constant term is always included as both an instrument and a regressor, although it is not reported in the tables.

[25] Alternatively, linear IV can be abandoned for nonlinear instrumental variables estimation that explicitly models the error term as obeying an MA(1) process, $\epsilon_t = v_t + pv_{t-1}$. This strategy is due to the insight of Carroll et al. (1994). By modeling the error term as an MA(1), one can enforce the restriction that variables dated at $t - 1$ should be orthogonal to v_t, because pv_{t-1} is explicitly controlled for. The advantage of the second strategy is that one does not need to eliminate the most up-to-date predictor from the instrument list. Results using nonlinear instrumental variables estimation were quite similar to those using linear IV; thus, those results are not reported.

[26] A previous draft of this paper that reported only results using disposable income emphasized this result.

TABLE 1.—RESULTS FROM U.S. DATA, DEPENDENT VARIABLE IS ΔC_t

Row		d_t Is Total Installment Credit							d_t is Revolving Credit		
		Y_t = DI	Y_t = DI	Y_t = DI	Y_t = DI	Y_t = DI	Y_t = LI	Y_t = LI	Y_t = DI	Y_t = LI	Y_t = LI
		Second Stage									
1	$E_{t-1}\Delta y_t$	0.360*	—	0.385*	0.237	0.313*	0.333*	0.307*	0.265**	0.294**	0.366*
		(3.2)		(3.5)	(1.5)	(3.2)	(2.21)	(2.41)	(1.6)	(1.91)	(2.59)
2	$E_{t-1}r_t$	—	—	0.017	0.008	0.005	−0.01	−0.03	0.012	0.036	0.024
				(0.90)	(0.48)	(0.29)	(−1.01)	(−1.23)	(0.52)	(0.68)	(0.469)
3	$E_{t-1}\Delta d_t$	—	0.102*	—	0.085*	0.085*	0.089*	0.082*	0.110*	0.078*	0.066*
			(3.2)		(2.6)	(2.6)	(3.0)	(3.17)	(2.34)	(3.3)	(3.08)
		First Stage									
4	$\overline{R}^2$ OLS (Δy_t)	0.056*	—	0.050**	−0.003	0.080*	0.110*	0.180*	−0.02	0.02	0.026
5	$\overline{R}^2$ OLS (Δd_t)	—	0.430*	—	0.497**	0.572*	0.483*	0.658*	0.465*	0.265*	0.386*
6	Test of Restrictions	7.02	5.82	12.30	12.34	18.21	6.076	9.055	13.24	7.496	9.633
7	Instruments	set 1	set 1	set 2	set 3	set 4	set 3	set 5	set 5	set 3	set 5

Notes: *Significant at the 5% or better level.
** Significant at the 10% or better level.
IV estimation d_t is either log installment or revolving credit as indicated; DI is disposable personal income; LI is labor income, measured as wages and salaries plus transfers minus personal contributions for social insurance. The sample period in the regressions using installment credit and disposable income is 1953:1 to 1993:1, using installment credit and labor income is 1959:4 to 1993:1, and using revolving credit the sample is 1978:1 to 1993:1. Instruments: *set 1:* lags 2 through 4 of income growth, consumption growth, and lags 2 of the error-correction term (EC, the lag two log-difference of consumption and income); *set 2:* adds lags two through four of the real interest rate to set 1; *set 3:* lags 2 through 4 of credit growth, income growth, the real interest rate, and the EC term; *set 4:* lags 2 through 5 of consumption growth, income growth, credit growth, the real interest rate, and the EC term; *set 5:* lags 2 through 4 of consumption growth, income growth, credit growth, the real interest rate, and the EC term; Rows 1, 2, and 3 report the instrumental variables estimates of λ, θ, α; in parenthesis are the *t*-statistics for the null hypothesis that the coefficient is zero. Rows 4 and 5 report the adjusted R^2 statistic from an OLS regression of income and credit growth on the instruments; asterisks indicate the joint marginal significance value for the null hypothesis that all the coefficients except the constant are zero. Row 6 reports $T^* \overline{R^2}$ from an OLS regression of the residual from the IV regression, on the instruments. Hypothesis tests were conducted using a heteroskedasticity and serial correlation robust covariance matrix.

growth in disposable income, it is necessary to either use more lags of disposable income in the instrument set—as in column 6—or to use labor income instead of disposable income, as in columns 7 and 8. In all three cases, the instruments are jointly significant in predicting income growth at better than the 5% level. Credit growth is well forecast by all of the instrument sets considered.[27]

As table 1 reveals, there is no evidence that the ex ante real interest rate is associated with the growth rate of consumption, a finding that has been well documented elsewhere. Moreover, the row labeled "Test of Restrictions" displays no evidence against the over identifying restrictions.[28]

The last three columns of table 1 show the results of estimating equation (16) using revolving credit instead of total installment credit. The results are similar to those obtained using total installment credit. In particular, both expected credit growth and expected income growth are significant explanatory variables in each regression on consumption growth.[29]

One possible reaction to these findings is that preferences are more complicated than those specified in the model discussed above. For example, non-time-separability between durables and nondurables implies that nondurable consumption growth should be related to the growth in durables consumption. If so, expected installment credit could simply be proxying for expected durables consumption in equation (16). Further calculations (not reported) showed that, unlike expected credit growth, growth in durables expenditure is not significantly correlated with the growth in nondurables and services expenditure once predictable income growth and the expected real interest rate are controlled for. This result corroborates the findings in other studies (Campbell & Mankiw, 1990; Beaudry & van Wincoop, 1996) which find no evidence for this sort of nonseparability.

In summary, the instrumental variables tests in this section show that consumption growth is correlated with not just predictable income growth (as has been extensively documented), but also with predictable consumer credit growth.[30] The result is striking because it does not represent a simple contemporaneous correlation between credit and spending. The next section asks whether time-varying borrowing constraints can rationalize these correlations.

[27] Recently, some have suggested that a better measure of instrument relevance in multivariate linear models is not the simple adjusted R^2 statistic form regressing the explanatory variable on the instruments, but an alternative measure sometimes called the "partial" R^2 (e.g., Shea, 1993). I computed this measure for a variety of instrument sets considered here and found that the degree of relevance for each explanatory variable was very similar to that indicated by the conventional measure. For example, when the instruments are lags two through four of consumption growth, credit growth, income growth, the real interest rate, and lag two of the log consumption to income ratio, the partial R^2 measure for income growth, debt growth, and the real interest rate is 0.09, 0.42, and 0.44, respectively.

[28] To test the overidentifying restrictions, an LM test statistic is formed by regressing the residual from the IV regression on the instruments. Taking T times R^2 from this regression (where T is the sample size), produces a test statistic that is distributed χ^2 with $K - N$ degrees of freedom, where K is the number of instruments and N is the number of independent variables. The corresponding number in the table is the $T * R^2$ statistic for the OLS regression of the IV residual on the instruments.

[29] One difference from table 1 is that the instrument sets for expected income growth are very weak predictors over the smaller sample period used in the revolving credit regressions, regardless of what combination of instruments is used.

[30] Two dummy variables were included in the above regressions, with little impact on the results: 1975:2 (social security and income-tax rebate) and 1980:2 (credit controls).

IV. Time-Series Predictions from the Model

To compare the model's predictions with the U.S. aggregate data, this section summarizes the results of time-series simulations computed by using the optimal consumption functions discussed in section II.[31] To do so, some method of aggregating over many consumers must be specified. I discuss this next.

A. Aggregation

A fundamental problem with using microeconomic models of consumption behavior to explain macroeconomic phenomena concerns whether they are capable of "scaling up" so that their implications are preserved at a macroeconomic level. While it's possible, in principle, to think about a representative agent whose behavior is guided by microeconomic decision rules, in this model, policy functions for individual consumption are highly nonlinear, making it difficult to infer how the time-series behavior arising from many aggregated decision rules should be related to that arising from a representative agent.

An additional difficulty with using a representative agent framework concerns the differences between aggregate and individual income processes. The income process specified above—in which income is the sum of a random walk and white noise—is a plausible approximation of individual income with its large transitory and permanent components; thus individual income growth is found to be negatively autocorrelated. By contrast, this process yields a very poor description of aggregate income growth which is positively autocorrelated.

One possibility for avoiding the problems associated with nonlinear decision rules—and for reconciling the income processes at the two levels of aggregation—is to undertake explicit aggregation of individual consumption outcomes, an approach taken in Deaton (1991). In this section, I follow his suggestion to consider H individuals, with household-specific variables denoted by h. The income growth of each individual can be represented by a decomposition of equation (3) into three separate components:

$$\Delta y_t^h - g = \epsilon_{1,t} + \chi\epsilon_{1,t-1} + \epsilon_{2,t}^h + \epsilon_{3,t}^h - \epsilon_{3,t-1}^h. \tag{17}$$

The first two terms on the right-hand side, along with the growth rate g, make up a growth component common to all consumers and capture the positive autocorrelation in aggregate income growth with positive parameter χ. The third term, ϵ_2^h, is the innovation in an idiosyncratic random walk, and the last two terms are the first difference of idiosyncratic transitory income. Total income for each consumer is the sum of a common $ARIMA(0, 1, 1)$, an idiosyncratic random walk, and transitory white noise. Once income is aggregated over many households, the idiosyncratic components cancel out, and aggregate income growth, equal to $g + \epsilon_{1,t} + \chi\epsilon_{1,t-1}$, is positively autocorrelated.[32]

From the aggregate data, σ_1 is set equal to 0.01, allowing the two idiosyncratic shocks to account for most of the income variation at the individual level. χ is set equal to 0.5, implying an autocorrelation coefficient of 0.4 for aggregate income growth. Given the equivalence between equations (17) and (3), and given σ, the standard deviation of η_t, the standard deviation of each idiosyncratic innovation in equation (17) can be obtained by matching variances and covariances between equations (17) and (3).[33] σ is varied in several cases below.

As discussed above, income-independent shocks to individual credit limits, ξ_t, are assumed to be a common across households.

The aggregation process then proceeds as follows. Income and credit histories are simulated for each individual. Using these histories, optimal consumption decisions are computed for each individual, and aggregated explicitly over many households. This produces a theoretical time series on consumption, income, and credit that can be compared with the aggregate data.

B. Theoretical Time Series Properties of Aggregate Consumption, Credit, and Income

This section summarizes the results of IV regressions using simulated, aggregate data from the model presented in section II. I assume that there are $H = 1{,}000$ households, and use the simulated data to perform IV regressions of consumption growth on expected income and credit growth, using one lag of consumption growth, income growth, and credit growth as instruments.

Table 2 summarizes the time-series relationship among consumption growth, expected credit growth, and expected income growth implied by the model. For the baseline parameter values, and for selected changes from the baseline in particular parameters, the table presents the mean, median, and standard deviation for the regression coefficient and its t-statistic over 100 simulations of 160 periods (the size of the large U.S. data set).

The top panel of table 2 shows the results for the baseline case with $\omega = 15$. The model produces a statistically significant correlation between consumption growth and both predictable income growth and predictable credit growth. The mean coefficients on predictable income growth

[31] This section focuses on comparing the model's predictions about the high-frequency correlations between consumption, income, and credit with those documented in section III. I do not address the apparent secular change (or at least the very low-frequency movements) in the debt-to-income ratio displayed in figure 1, and difficult to explain in models with constant, balanced growth in the steady state.

[32] To think of what follows as an aggregate model, assume the existence of firms, whose behavior I do not model, with access to a linear-investment technology, allowing them to earn a constant rate of return by either investing directly, or, alternatively, lending to consumers at this constant rate.

[33] I also follow Deaton by assuming that individuals do not observe the three components separately, but do observe their sum which is equivalent to the representation in equation (3).

TABLE 2.—IV RESULTS USING SIMULATED AGGREGATE DATA FROM THE MODEL: $\overline{D_t^i} = (1/\omega)\, Y_t^i \exp(\xi_t)$ EMPIRICAL MODEL: $c_t = \mu + \lambda E_{t-1}\Delta y_t + \alpha E_{t-1}\Delta d_{t+1} + \epsilon_t$

	Baseline Parameter Values			
	$\hat{\lambda}$	$\hat{\alpha}$	$t-\lambda$	$t-\alpha$
Mean	0.59	0.31	4.46	2.50
Median	0.58	0.33	4.20	2.49
Standard	0.55	0.44	3.21	0.94
	$\delta = 0.05$			
Mean	0.53	0.36	5.07	2.49
Median	0.56	0.34	3.75	2.51
StDev	0.24	0.22	4.25	1.20
	$\gamma = 4$			
Mean	0.49	0.38	4.36	2.65
Median	0.54	0.35	3.71	2.66
StDev	0.25	0.22	3.45	0.91
	$\sigma = 0.10$			
Mean	0.44	0.43	3.20	2.09
Median	0.48	0.38	2.46	2.13
StDev	0.29	0.29	2.90	1.06
	$\omega = 10$			
Mean	0.23	0.60	1.70	1.62
Median	0.27	0.57	0.88	1.65
StDev	0.89	0.91	2.40	1.32
	$\omega = 8$			
Mean	0.33	0.42	1.08	1.03
Median	0.23	0.63	0.61	1.33
StDev	1.81	1.72	1.92	1.39

Notes: $\xi_t \sim N(0, \sigma\xi)$. The first panel shows results obtained using baseline parameter values: $g = 0.02$, $\chi = 0.5$, $\psi = 0.44$, $\sigma = 0.05$, $\sigma_1 = 0.01$, $\sigma_u = 0.01$, $\phi = 0.6$, $\omega = 15$, $r = 0.03$, $\rho = 2$, $\delta = 0.15$. Cases that represent a change from one of the baseline parameter values are listed above the other panels. Data are simulated for 1,000 households; i indicates a household specific parameter. Entries in each cell give the mean, median, and standard deviation of each coefficient across 100 simulations of 160 periods each. The column labeled "t-" give the t-statistic for that coefficient. Instruments: lag 1 of income growth, credit growth, and consumption growth; a constant is included as a regressor and an instrument.

and predictable credit growth are 0.59 and 0.31, respectively, with both estimates statistically significant at better than the 5% level. Although these results are qualitatively consistent with the data, they tend to overpredict the magnitude of the correlations. (Estimates from the U.S. data for $\hat{\lambda}$ ranged from 0.3 to 0.4 and for $\hat{\alpha}$ from 0.06 to 0.14.)

Other panels in table 2 present the results with variations on the baseline parameter values. Panel 2 decreases the rate of time preference, δ to 0.05 from its baseline value of 0.15. The decreased degree of impatience relative to imprudence has the effect of making the point estimates on the income variable more strongly significant. The regression can more precisely pick up the independent effects of income and credit when individuals are less patient because predictable income growth and predictable credit growth are less collinear in that case.

Panels 4, 5, and 6 of table 2 show the results when the standard deviation of η_t, σ, is raised to 0.10 from 0.05, and when ω is lowered to 10 and to 8, respectively. The point estimates do not appear very sensitive to σ, though they are estimated less precisely when σ is higher. The point estimates and t-statistics are sensitive to the value of ω, however. In particular, when ω shrinks to 10, the point estimate on income growth shrinks to 0.23, while the point estimate on credit growth increases to 0.60. Note, however, that both coefficients are estimated far less precisely than in the baseline case, and, when ω is as low as 8, neither point estimate is statistically significantly different from zero at the 5% level.

Several properties of the theoretically generated data bear noting. First, there must be sufficient variation in the credit ceiling to rationalize a correlation between consumption growth and expected credit growth. Results (not shown) indicated that, with other parameters set at their baseline values, the coefficient on expected credit growth approaches zero as ω approaches infinity. At the same time, the coefficient on expected income growth approaches one, while the t-statistic becomes very large.[34]

Second, some variation in the borrowing limit is necessary to produce a correlation between predictable credit growth and consumption growth. Nevertheless, the last panel of table 2 shows that too much variation in the borrowing limit also makes it difficult to rationalize the data. When ω is lowered, the credit ceiling is higher on average, but it is also more volatile, and precautionary motives may make individuals more reluctant to borrow up to their limit. This has the effect of reducing the sensitivity of consumption growth to both predictable changes in income, and predictable changes in credit.

It is sometimes argued that increased access to credit should allow consumption to be less dependent on predictable changes in income, as formerly binding constraints are relaxed. At issue, however, is whether an increase in credit availability would simply lead to a very transitory relaxation of constraints as impatient households rapidly run up debt to its new, higher limit. In this case, a relaxation of formerly binding constraints may cause no appreciable decline in the sensitivity of consumption growth to expected income because the extra credit would be consumed immediately. Thus, a remaining question is whether impatient households who experience increased access to credit should be expected to display consumption that is subsequently less sensitive to predictable components of income and credit.

Table 2 helps shed light on this question. The last panel suggests that higher credit ceilings may indeed lead to less sensitivity of consumption growth, to both predictable components of income and predictable components of consumer credit.[35] Note that this result occurs in a sample of impatient households. Consumers who can borrow a larger fraction of their income on average (those who have smaller values of ω), display consumption that is less strongly correlated with predictable components of income and credit. Consumption is not significantly correlated with either expected income growth, or expected credit growth

[34] For example, when $\omega = 1{,}000$, the mean coefficient on expected credit growth is equal to 0.001, while the coefficient on income growth is equal to 0.991.

[35] The exercise of comparing excess sensitivity across households that have different values for ω is one of comparative statics and can be interpreted only as suggestive of how a forward-looking consumer might react to an expected change in ω.

when ω is roughly half its baseline value. This occurs because, in this model, impatient consumers cannot increase the average amount they borrow without increasing their exposure to credit market risk at the same time. The increased exposure to risk leads to more-prudent behavior and, in aggregate, less excess sensitivity. By contrast, a perfectly riskless increase in the credit limit would be consumed immediately, and the dynamic behavior of consumption would quickly return to its previous state, with consumers simply carrying around higher debt balances. In that case, consumers with higher credit limits would not be expected to display consumption that is less sensitive to predictable income or credit.[36]

The prediction that consumption should display less excess sensitivity when households can borrow a larger fraction of their income is consistent with recent empirical evidence. Bacchetta and Gerlach (forthcoming) find that the sensitivity of consumption growth to predictable components of income and credit in the U.S.—a country that has experienced a large increase in the availability of consumer credit over the last twenty years—has been declining gradually over time since the late 1970s.

V. Summary and Concluding Remarks

This paper finds evidence that consumption growth is correlated with predictable consumer credit growth in U.S. aggregate data. The finding presents a puzzle for students of consumption behavior because it is not readily explained by existing models of agents, constrained or unconstrained, forming rational expectations over the extended future.

The theoretical results in this paper show that a correlation between consumption and predictable consumer credit is consistent with a buffer-stock model in the presence of time-varying borrowing constraints. In addition to rationalizing the new evidence presented here on consumption and consumer credit, this framework is also consistent with the previously well-documented finding that consumption growth is independently correlated with predictable income growth. This so-called excess sensitivity to predictable income—a violation of the certainty-equivalent PIH—has prompted an on-going debate over the most likely explanation for this result.[37] The evidence presented here reinforces the hypothesis that borrowing constraints are an important part of the story.

Although the model in this paper correctly predicts the independent influence of forecastable income growth and forecastable credit growth on consumption growth, and yields predictions that are in the correct direction for a range of parameter values, some aspects of the data are less well captured. Quantitatively, the model tends to overpredict the correlations discussed above. Nevertheless, stochastic variation in credit limits coupled with buffer stock saving behavior seems to account for a substantial portion of the time-series properties reported here on aggregate consumption, income, and consumer credit.

The findings have several economically meaningful implications. First, they suggest that adverse shocks to the economy are likely to be amplified by the type of procyclical borrowing restrictions considered here, creating a "financial accelerator" in the consumption sector, analogous to that documented for the investment sector by others (e.g., Bernanke et al., forthcoming). Second, independent variation in credit availability that is common across households will produce an additional source of volatility in aggregate consumption that is not associated with shocks to conventional indicators such as income and interest rates. Even consumers who are not currently constrained may adjust their consumption in response to changes in credit-market conditions, because it affects the probability that they will be constrained in the future. Finally, the results suggest that increased access to consumer credit may lead to more-prudent behavior and less excess sensitivity if consumers' new-found ability to borrow also increases their exposure to credit-market risk.

[36] An earlier version of this paper demonstrated this result by analyzing the effects of a riskless increase in the debt limit and looking at transitions between steady states.

[37] For example, researchers have variously attributed the correlation of consumption growth with predictable income growth to liquidity constraints (Zeldes, 1989a; Deaton, 1991), aggregation bias (Attanasio & Weber, 1993; Pischke, 1995), myopia (Hall & Mishkin, 1982; Campbell & Mankiw, 1989), habit formation (Deaton, 1992; Boldrin et al., 1995), and precautionary motives (Zeldes, 1989b; Carroll, 1997).

REFERENCES

Attanasio, O. P., and G. Weber, "Consumption Growth, the Interest Rate, and Aggregation," *Review of Economic Studies* 60 (1993), 631–649.

Attanasio, O. P., J. Banks, C. Meghir, G. Weber, "Humps and Bumps in Life Time Consumption," NBER working paper no. 5350 (November, 1995).

Bacchetta, Phillippe, and Stefan Gerlach, "Consumption and Credit Constraints: International Evidence," *Journal of Monetary Economics* (forthcoming).

Beaudry, Paul, and Eric van Wincoop, "The Intertemporal Elasticity of Substitution: An Exploration Using a U.S. Panel of State Data," *Economica* 63, No. 251 (1996).

Bernanke, Ben, and Mark Gertler, "Agency Costs, Net Worth, and Business Fluctuations," *American Economic Review* LXXIX (1989), 14–31.

Bernanke, Ben, Mark Gertler, and Simon Gilchrist, "The Financial Accelerator and the Flight to Quality," *Review of Economics and Statistics,* 78(1) (1996), 1–15.

Boldrin, M., L. J. Christiano, and J. Fisher, "Asset Pricing Lessons for Modeling Business Cycles," NBER working paper no. 5262 (1995).

Campbell, John Y., and Greg Mankiw, "*Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence,*" in O. Blanchard and S. Fischer (eds.), NBER Macroeconomics Annual 1989 (Cambridge, MA: MIT press, 1989).

—— "Permanent Income, Current Income and Consumption," *Journal of Business and Economic Statistics* 8 (1990), 269–279.

—— "The Response of Consumption to Income, A Cross-Country Investigation," *European Economic Review* 35 (1991), 723–767.

Carroll, Christopher, "Buffer Stock Saving and the Permanent Income Hypothesis," *Quarterly Journal of Economics* 112(1) (1997), 1–56.

Carroll, Christopher, and Miles Kimball, "On the Concavity of the Consumption Function," *Econometrica* LXIV (1996), 981–922.

Carroll, Christopher, Jeffrey Fuhrer, and David Wilcox, "Does Consumer Sentiment Forecast Household Spending? If So, Why?" *American Economic Review,* 84(5) (1994), 1397–1408.

Christiano, Lawrence J., M. Eichenbaum, and M. David, "The Permanent Income Hypothesis Revisited," *Econometrica* 59(2) (1991), 397–423.

Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica* 59, No. 5 (1991), 1221–1248.

——— *Understanding Consumption,* (Oxford University Press, 1992).

Deaton, Angus, and Guy Laroque, "On the Behavior of Commodity Prices," *Review of Economic Studies* 59 (1992), 1–23.

Fazzari, Steven M., Glenn Hubbard, and Bruce Petersen, "Financing Constraints and Corporate Investment," *Brookings Papers on Economic Activity* 1:1988 (1988), 141–206.

Flavin, Marjorie, "The Adjustment of Consumption to Changing Expectations About Future Income," *Journal of Political Economy* 89 (1981), 974–1009.

Hall, Robert E., "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy* 86 (1978), 971–987.

Hall, Robert E., and Frederic S. Mishkin, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica* 50(2) (1982), 461–81.

Hansen, Lars P., and Kenneth J. Singleton, "Stochastic Consumption, Risk Aversion, and the Temporal behavior of Asset Returns," *Journal of Political Economy* 91 (1983), 249–265.

Jappelli, Tullio, "Who is Credit Constrained in the U.S. Economy?" *Quarterly Journal of Economics* 105 (February, 1990), 219–234.

Ludvigson, Sydney C. "The Mechanism of Monetary Transmission to Demand: Evidence from the Market for Automobile Credit," *Journal of Money, Credit, and Banking* 30, no. 3 (August, 1998, Part I).

MaCurdy, T. E., "The Use of Time Series Processes to Model the Error Structure of Earnings in Longitudinal Data Analysis," *Journal of Econometrics* 18 (1982), 83–114.

Mankiw, N. Gregory, "The Permanent Income Hypothesis and the Real Interest Rate," *Economic Letters* 7 (1981), 307–311.

Pischke, Jörn-Steffen, "Individual Income, Incomplete Information and Aggregate Consumption," *Econometrica* 63(4) (1995), 805–840.

Shea, John, "Instrument Relevance in Linear Models: A Simple Measure," Unpublished manuscript, University of Wisconsin-Madison (1993).

Staten, Michael E., and Robert W. Johnson, *Household Credit Data Book* (West Lafayette, Indiana: Credit Research Center Purdue University, 1995).

Zeldes, Steve, "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy* 97 (1989a), 305–346.

——— "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics* CIV (1989b), 275–298.

APPENDIX A

This appendix briefly describes the recursive technique used to solve equation (13).

In making a discrete approximation to the underlying distributions of η_t and ξ_t, transition probabilities must be computed for η and ξ next period, conditional on their values this period. Since the innovation to income growth, η_t, is i.i.d., transition probabilities for this variable are equal to $\pi = 1/m$ regardless of the value of η this period. In contrast, ξ_t is not i.i.d., but positively autocorrelated.

Consequently, the transition probabilities, $\pi_{lj} = prob[\xi_{t+1} = \xi_l | \xi_t = \xi_j]$ for all l and j, are set equal to the transition probabilities from one discrete interval to the next of the underlying autoregressive process for ξ_t. Specifically, the distributions are discrete approximations of normal distributions. The procedure can be summarized as follows. Suppose that η is normally distributed, $N(0, \sigma^2)$. First, m points are chosen so that successive areas under the standard normal between each of these points are equal to $1/m$. Then, the m conditional means within each of these intervals are chosen as the m equiprobable values of a discrete process that approximates $N(0, 1)$. Each of these values are then multiplied by σ so that Δy_t in equation (3) is replaced by a discrete, first-order Markov process consisting of innovations η that approximate $N(0, \sigma^2)$. A similar approach is taken in modeling ξ_t as a five-point discrete Markov process, taking into account the autocorrelation in ξ_t. The joint density of ξ_{t+1} and ξ_t is specified as a m by m discrete approximation to a bivariate normal, so that, conditional on any value of ξ_t, there are m possible outcomes for ξ_{t+1}.

With η_t and ξ_t set up as discrete Markov processes, the recursive problem can be rewritten as follows, with $i = 1, \ldots, m$ income states η_i, and $j = 1, \ldots m$ credit states ξ_j, defining $m * m$ functions, $p(w, i, j)$,

$$p(w, i, j) = \max \Bigg[v(w), \beta^* \sum_l \pi_{lj} \sum_k \pi \rho_{ik}^{\gamma} p \times \{[1 + \rho_{ik}(1 + r)(w - v^{-1}p(w, i, j) + \zeta_{ijlk}], k, l\} \Bigg],$$

where

$$\rho_{ik} \equiv \exp(-\eta_k + \psi\eta_i - g)$$

$$\zeta_{ijlk} \equiv \frac{\exp(\eta_k - \psi\eta_i + g + \xi_l - \xi_j) - (1 + r)}{(1 + r)\exp(\eta_k - \psi\eta_i + g + \ln\omega - \xi_j)}.$$

The policy functions $p(w, i, j)$ are found numerically by backward recursion.

Now consider n functions, $p_o, p_1, \ldots, p_n$ where $p = \nu[f(\cdot)]$. Instead of the infinite horizon, for the moment suppose that $n = 0$ corresponds to a final period, so that period 1 is one period before the final period and period n is n periods before the final period, or the first period. The set of equations can be rewritten as an updating rule,

$$p_n(\cdot) = \max \Bigg[v(w), \beta^* \sum_l \pi_{lj} \sum_k \pi \rho_{ij}^{\gamma} p_{n-1} \times \{[1 + \rho_{ik}(1 + r)(w - v^{-1}p_n(\cdot) + \zeta_{ijlk}], k, l\} \Bigg], \tag{A.1}$$

where p_{ij} and ζ_{ijlk} are as given in the text. Suppose in the final period, all resources are spent, so that $\theta = (f(\omega, \eta) = \omega$, and $p_o \equiv \nu(\omega)$. In principle, one can work backwards. Given the first iteration p_n on the right-hand side, the quantity in brackets can be explicitly computed, yielding a new p_n. Comparing the two values one then determines whether a convergence criterion has been satisfied, in which case we have a solution, p, to the infinite-horizon problem. If the convergence criterion has not been satisfied, iteration continues so that the second value of p_n is substituted on the right-hand side of (A.1), and so on. Under appropriate conditions (already established by Deaton and Laroque (1992) for this type of problem), the mapping is a contraction, and the backward solution will converge in a finite number of iterations to a unique solution. In practice, the policy is solved over a grid of values for ω, and it is easier to replace p_n on the right-hand side with p_{n-1}, setting $p_0 = \nu(\omega)$. Note that this implicitly imposes a no-Ponzi condition. If the convergence criterion can be satisfied with this replacement, we have found the optimal policy, since equation (A.1) is a contraction and guarantees a unique solution for p.

APPENDIX B

This appendix contains information about the variables from the SCF used to compute the summary statistics for $\hat{\omega}$.

Credit-Card Limits

Question 414 asks what the maximum amount the household could borrow on all of its credit-card accounts.

Wage and Salary Income

Question 5702 asks the household how much income from wages and salaries it received during the year of the survey, before deductions from taxes and anything else.

Questions about Recent Credit Applications

Questions 407 and 408 ask the household whether a particular lender or creditor has turned down any request made for credit, or did not give it as much credit as it applied for; and whether the household was discouraged from borrowing because it might be turned down for credit.

Liquid Assets

Liquid assets consist of the following variables, a subset of total assets: Checking accounts 3506, 3510, 3514, 3518, 3521, 3525, 3529; Market value of money market accounts 3706, 3711, 3716, 3718; Market value for CDS 3721 savings accounts 3804, 3807, 3810, 3813, 3816, 3818; Market value for mutual funds 3822, 3824, 3828, 3830, 6704; Face value of savings bonds 3902 (no market value available); Market value of bonds 7635, 7636, 7637, 7638, 7639, 6706; Market value of stocks 3915; Cash or call money accounts 3930; Market value of trusts, annuities, and managed investment accounts 3942.