Does Buffer-Stock Saving Explain the Smoothness and Excess Sensitivity of Consumption?

By SYDNEY C. LUDVIGSON AND ALEXANDER MICHAELIDES*

One of the lasting contributions to modern-day consumer theory is the idea that consumption is determined by the expected value of lifetime resources, or permanent income. The modern-day specification of this concept is typically presented as a special case of intertemporal-choice theory, where felicity functions are quadratic, labor income is stochastic, there are no restrictions on borrowing, consumers have infinite horizons, and expectations are formed rationally. We refer to this model as the permanent-income hypothesis (PIH).

Although the PIH has considerable intuitive appeal, a series of influential papers published after its inception revealed two notable discrepancies between the model's predictions and aggregate data. First, the model predicts that consumption growth should be more volatile than income growth if aggregate income growth has positive serial correlation (as the quarterly data suggest it does), yet aggregate consumption growth is in fact much smoother than aggregate income growth (Angus Deaton, 1987; John Y. Campbell and Deaton, 1989; Jordi Galí, 1991). Second, the PIH predicts that consumption changes should be orthogonal to predictable, or lagged, income changes, yet the correlation between consumption growth and lagged income growth has been found to be one of the most robust features of aggregate data (for example, Marjorie A. Flavin, 1981; Alan S. Blinder and Deaton, 1985; Campbell and N. Gregory Mankiw, 1989; Orazio P. Attanasio and Guglielmo Weber, 1993). Thus aggregate consumption growth has been described as exhibiting two puzzles: it is both "excessively smooth" relative to current labor-income growth, and "excessively sensitive" to lagged labor-income growth.[1,2]

In response to these and other empirical anomalies, researchers have recently sought out modifications of the PIH framework. Chief

* Ludvigson: Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045; Michaelides: University of Cyprus, P.O. Box 20537, 1678 Nicosia, Cyprus. A previous version of this paper appeared under the title "Can We Explain the Consumption Excesses Yet? Aggregate Consumption Implications of Buffer Stock Saving Behavior." For many helpful comments, we thank Orazio Attanasio, Dave Backus, Ben Bernanke, John Y. Campbell, Christopher Carroll, Todd Clark, Angus Deaton, Pierre-Olivier Gourinchas, Bo Honore, James Kahn, Martin Lettau, Costas Meghir, Christina Paxson, Steve Pischke, Andrew Samwick, Christopher Sims, Mark Watson, Guglielmo Weber, Michael Woodford, Steve Zeldes, and two anonymous referees, conference participants at the National Bureau of Economic Research Summer Workshop on the Aggregate Implications of Microeconomic Consumption Behavior, July 1998, the TMR Workshop on Saving and Pensions-CenTER, Tilburg University, March 1999, and the New York Area Macroeconomics Workshop, October 1998, seminar participants at Columbia University, INSEAD, the New York Federal Reserve, University College London, London Business School, City University Business School, University of Cyprus, Tel Aviv University, and Jeffrey Brown for excellent research assistance. The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

[1] Note that the important insight of Danny Quah (1990), that individuals may be better able than econometricians to distinguish short-run from long-run innovations to their income, cannot resolve the excess-smoothness puzzle laid out in Campbell and Deaton (1989). These authors show that—in the context of the PIH—excess sensitivity and excess smoothness are the same phenomena. In addition, Campbell (1987) demonstrates that superior information of the type emphasized by Quah will show up in households' saving behavior, and using this information allows the econometrician to control for the agent's private information when predicting income. Since excess sensitivity is well established, and controlling for saving does not eliminate it (see Campbell and Deaton, 1989), the observed smoothness of consumption cannot be explained by superior information on the part of permanent-income consumers. See Deaton (1992) for further discussion.

[2] Philippe Bacchetta and Stefan Gerlach (1997) and Ludvigson (1999) document another kind of excess sensitivity: the correlation between consumption growth and predictable changes in consumer credit. Although we do not explore these findings here, Ludvigson finds that a buffer-stock model with time-varying liquidity constraints can replicate such a correlation.

among these is the so-called buffer-stock model, pioneered by the work of Deaton (1991) and Christopher D. Carroll (1992), which modifies the PIH to allow for precautionary saving motives, impatience, and restrictions on borrowing.[3] Several authors have argued that the buffer-stock paradigm provides a good description of the median consumer's behavior (for example, Deaton, 1991; Carroll, 1997a; Carroll and Andrew A. Samwick, 1997, 1998). The buffer-stock model has become a workhorse of modern-day consumer theory.[4]

Can buffer-stock saving explain the smoothness of aggregate consumption and its correlation with lagged income? Several researchers have suggested that buffer stock behavior should be a good candidate for doing so (Deaton, 1991; Carroll, 1992, 1997a). Yet despite the growing popularity of this model, virtually no research has set out to formally investigate whether buffer-stock saving behavior can explain these well-established facts of aggregate data.[5]

In this paper, we address the question posed in the title by investigating two models of buffer-stock behavior. Our first framework follows closely the original models of individual buffer-stock behavior developed in Deaton (1991) and Carroll (1992); we refer to this framework as the *standard buffer-stock model.* We then consider an alternative framework: following Pischke (1995)—who worked entirely within the context of the PIH—we investigate how the predictions of the buffer-stock model vary according to whether households observe economywide variation in their income. Rather than analyzing the theory's implications for individual consumer behavior, as in most of the existing literature, we assume that there are a large number of buffer-stock consumers and aggregate their consumption decisions explicitly.

Our results show that the standard buffer-stock model, which incorporates borrowing restrictions, impatience and precautionary motives, does not generate robust excess sensitivity and aggregate consumption growth that is smoother than aggregate income growth. These implications arise only in the alternate framework which presumes that individuals do not observe the economywide component of their earnings (an assumption we refer to as *incomplete information*). This incomplete-information version of the model goes a long way toward resolving the consumption puzzles discussed above: it produces aggregate consumption growth that is substantially less volatile than the benchmark PIH, and generates robust excess sensitivity to expected income growth. Nevertheless, even this version of the buffer-stock model falls short of matching the data in magnitude: most notably, aggregate buffer-stock consumption remains insufficiently smooth relative to aggregate income. In spite of this lack of smoothness in aggregate consumption, however, there may be considerable smoothing in individual consumption. An important implication of these results is that inferences about aggregate buffer-stock consumption cannot be made by looking at household-level, representative-agent consumption functions.

The excess-sensitivity puzzle has been empirically investigated elsewhere using alternative models of consumer behavior and by emphasizing biases created by aggregation. Attanasio and Martin Browning (1995) argue from

[3] In more recent work, Jeffrey Fuhrer (1999) argues that modifying the PIH to allow for habit formation may improve its quantitative performance by imparting a motive for consumers to smooth the change in consumption. Martin Lettau and Harald Uhlig (1999) show that excess sensitivity may arise as the outcome of boundedly rational learning behavior.

[4] A vast literature has appeared that examines and extends the implications of the standard buffer-stock model of Deaton (1991) and Carroll (1992, 1997a). See, for example, R. Glenn Hubbard et al. (1995); Pierre-Olivier Gourinchas and Jonathan A. Parker (1996); John Heaton and Deborah J. Lucas (1996, 1997); Carroll and Samwick (1997, 1998); David I. Laibson et al. (1998); Michael Haliassos and Michaelides (1999); Ludvigson (1999).

[5] Most work has focused on the model's implications for individual, not aggregate, consumption. A few studies have addressed the aggregate implications of buffer-stock behavior. Carroll (1992) argues that the behavior of a representative buffer-stock agent may explain some features of aggregate data that are not explained by the PIH, but he does not formally address these twin consumption excesses. Deaton (1991) develops an aggregate version of the buffer-stock model in which individual income is subjected to both economywide and idiosyncratic shocks and shows that such a model can produce some smoothing of aggregate consumption, as well as some sensitivity to lagged income. Deaton's results, however, are documented only for a limited set of parameter values and do not reveal how the absence (or ignorance) of contemporaneous information on economywide variables—emphasized in subsequent work by Jörn-Steffen Pischke (1995)—influences the model's predictions.

exploring cohort-mean data (somewhere in between aggregate and individual data) that excess sensitivity may be due to omitted variable bias. The omitted variables can, for instance, be various demographic indicators which are important in models with nonseparable preferences over consumption and leisure. Alternatively, Attanasio and Weber (1993, 1995) show that aggregation problems with Euler equations for consumption can lead to rejections of overidentifying restrictions in those equations, and thus, in principle, to excess sensitivity.

While these alternate explanations for excess sensitivity are clearly instructive, we think there are several reasons it is useful to explore alternatives. First, the literature cited above has focused on explaining excess sensitivity without addressing the smoothness of consumption; yet other authors have argued that excess sensitivity and excess smoothness are intimately related (Campbell and Deaton, 1989). Accordingly, we think it makes sense to study these phenomena jointly. Second, the results for household-level cohort data using demographic variables, while suggestive, are impossible to replicate in aggregate data, and Attanasio and Weber (1993) find that excess sensitivity to expected income growth remains even after controlling for demographic variables on data averaged over the whole *British Family Expenditure Survey*. This latter finding suggests that factors other than omitted demographic variables and nonseparabilities in the utility function may play a role in generating excess sensitivity.[6] Third, the empirical studies in this literature, typically carried out using linearized Euler equations, are unable to control for the presence of a precautionary saving term (an important feature of buffer-stock consumption), and other authors have argued that omitting this term may lead researchers to overestimate the share of consumption-income tracking that is attributable to labor supply and demographics (Gourinchas and Parker, 1996). Similarly, Carroll (1997b) and Ludvigson and Christina H. Paxson (1999) find that higher-order moments of consumption growth may also be important omitted variables in linearized Euler equations. Finally, our approach is useful because we model aggregate consumption behavior by explicitly summing the consumption outcomes of individual households, thereby allowing us to directly control for the aggregation problems studied in Attanasio and Weber (1993, 1995).

The rest of this paper is organized as follows. Section I presents the buffer-stock model of individual behavior. To conserve space, we focus our analysis here on a few specifications which we argue are most plausible; other specifications along with the properties of U.S. aggregate data are explored in more detail in Ludvigson and Michaelides (1999). Section II discusses the aggregate properties of buffer-stock-saving behavior, comparing the implications of the complete-information case with those when individuals have incomplete information about the economywide component of their earnings. We conclude this analysis with a short discussion of the differences between individual and aggregate consumption and how the aggregate results change if we allow unconstrained PIH consumers to comprise some fraction of the population. Finally, we briefly consider two extensions of these models in Section II, subsection A: allowing for geometric means, and time-averaged data when the decision period of the household is shorter than the sampling interval. Section III concludes.

I. Microeconomic Model of Consumption

This section presents the model of individual consumption behavior. Given that we will compare the model's implications with the properties of aggregate data, we face a preliminary choice on whether to calibrate the household-decision period as a quarter or a year. Two factors favor the consideration of an annual, rather than quarterly, model of behavior. First, there are measurement problems with aggregate quarterly consumption and income data, some of which have been documented by David W. Wilcox (1992);

[6] Consistent with this finding, several papers using aggregate data find no evidence of any important nonseparability in preferences between consumption and leisure, and controlling for these labor-supply indicators has not been found to eliminate the excess sensitivity of consumption growth to expected income growth (see Martin Eichenbaum et al. 1988; Campbell and Mankiw, 1990; Paul Beaudry and Eric van Wincoop, 1996).

Ludvigson and Michaelides (1999) provide more discussion. Second, there are several microeconometric studies that estimate a loglinear labor-income growth process at annual frequency, thereby providing direct evidence for our calibration of individual income.

Other factors argue for the consideration of a quarterly specification. The presence of significant positive autocorrelation in aggregate labor-income growth at quarterly frequency produces interesting differences from the results of the annual model where aggregate labor-income growth is far less autocorrelated.[7] Moreover, almost all of the previous macroeconomic analysis of these issues has been performed on quarterly data and then compared to the predictions of a quarterly framework. Thus we feel that some consideration of both models is warranted.

For models of both frequencies, we assume that the household-decision interval matches the data-sampling interval. (Later, we briefly pursue an extension in which we relax this assumption.) Time is discrete and agents have an infinite horizon. We assume there is one nondurable good and one financial asset (a riskless bond) which yields a constant after-tax real return, r. At time t, agents enter the period with assets held over from last period (A_{it}), and receive Y_{it} units of the nondurable good from inelastically supplying one unit of labor. Households choose the level of nondurable good expenditures (C_{it}) to maximize the present-discounted value of expected utility, subject to an accumulation constraint for assets. Individual income is given exogenously as $Y_{it} = P_{it}U_{it}$ where $P_{it} = G_t P_{it-1} N_{it}$. Following Deaton (1991), assets are exogenously restricted to be nonnegative. We assume that preferences are of the constant relative risk-aversion family; specifically, $U(C_t) = \frac{C_t^{1-\rho}}{1-\rho}$ when $\rho > 0$; if $\rho = 1$, $U(C_t) = \ln C_t$. Unlike the quadratic utility specification in the PIH, this specification has a positive coefficient of relative prudence, equal to $1 + \rho$, leading to a precautionary motive for saving.[8] "Cash-on-hand," X_t, is defined as the sum of current income and assets, $Y_t + A_t$.

Given the restriction on assets, $A_t \geq 0 \; \forall \; t$, the first-order condition for optimal consumption choice may be written:[9]

$$(1) \quad U'(C_t) = \text{MAX}[U'(X_t),\ \beta(1+r)E_t U'(C_{t+1})],$$

where $\beta = \frac{1}{1+\delta}$ is the constant discount factor. If the agent is constrained at time t, the maximum that can be spent on consumption is cash on hand (X_t), implying that marginal utility can never be less than $U'(X_t)$.

The process for individual income specified above was first used in a nearly identical form by Carroll (1992) and is decomposed into a "permanent" component, P_{it}, and a transitory component, U_{it}. We assume that $\ln U_{it}$ and $\ln N_{it}$ are independently and identically distributed (i.i.d.) with mean zero and variances σ_u^2 and σ_n^2, respectively. We assume that $\ln G_t$ is common to all individuals and has an unconditional mean μ_g and constant variance σ_g^2, while the innovation to $\ln G_t$ is uncorrelated with $\ln U_{it}$ and $\ln N_{it}$. As a result, the logarithm of the permanent component of labor income, $\ln P_{it}$, evolves as a random walk with a stochastic drift.

Given these assumptions, the growth in individual labor income follows

$$(2) \quad \Delta \ln Y_{it} = \ln G_t + \ln N_{it} + \ln U_{it} - \ln U_{it-1},$$

where the unconditional mean growth for individual earnings is μ_g, and the unconditional variance $\sigma_g^2 + \sigma_n^2 + 2\sigma_u^2$. Because the last

[7] See Ludvigson and Michaelides (1999) for a recent documentation of the time-series properties of aggregate labor-income growth.

[8] The coefficient of relative prudence is defined by Miles S. Kimball (1990) as $\frac{-U'''(C_t)C_t}{U''(C_t)}$. Note that borrowing restrictions may induce a form of prudence (convexity in the marginal utility of consumption) even if preferences exhibit no prudence (e.g., quadratic preferences). See Deaton (1992).

[9] This form for the Euler equation was first derived in Deaton (1991).

three terms in (2) are idiosyncratic, per capita aggregate labor-income growth inherits the data generating properties of $\ln G_t$.[10]

We make the following additional assumption:

$$\beta(1 + r)E_t[(G_{t+1}N_{t+1})^{-\rho}] < 1. \quad (3)$$

As Deaton and Guy Laroque (1992) have verified, this assumption is a sufficient condition for (1) to have a unique solution. Equation (3) gives the "impatience" condition common to buffer-stock models which insures that borrowing is part of the unconstrained plan. If we let $\beta = \frac{1}{1+\delta}$, use the fact that G and N are independent, and take logs of (3), this condition can be written

$$\frac{r - \delta}{\rho} + \frac{\rho(\sigma_g^2 + \sigma_n^2)}{2} < \mu_g, \quad (4)$$

where the approximation $\ln(1 + x) \approx x$, for x small, has been used.

We begin by solving the model under a set of "baseline" parameter assumptions. We set the rate of time preference, δ, equal to 0.05, and the constant real interest rate, r, equal to 0.02, both at annual rates. Carroll (1992) estimates the variances of the idiosyncratic income shocks using data from the *Panel Study of Income Dynamics*, an annual household data set, and our baseline simulations use values very close to those: 0.10 percent per year for σ_u and 0.08 percent per year for σ_n.[11] These estimates are consistent with those from other household studies (see Thomas E. MaCurdy, 1982; John Abowd and David Card, 1989). Aggregate labor-income growth, $\ln G_t$, is well described by an i.i.d. process in annual data, and a positively autocorrelated AR(1) process (with an autoregressive coefficient equal to 0.23) in quarterly data.[12] We use these distributions in the simulations discussed below.

To our knowledge, there is no microeconometric study that estimates a loglinear specification of household earnings at quarterly frequency; therefore some assumption about the quarterly process for the idiosyncratic components of income growth must be made in order to model decision periods as quarters. A natural starting point might be to think of annual income growth as the log difference in the sum of income over four quarters. Given the annual income process specified in (2), however, this strategy turns out to imply an unreasonably small permanent component in the quarterly income process. To resolve this dilemma, we instead model annual income growth as the year-over-year change in log income. We may then derive the quarterly process for individual income growth that gives the annual process as the year-over-year change in log income, a strategy that reduces the role of the permanent component but does not eliminate it.[13] Since the year-over-year log difference in income is just the sum of the quarterly log differences during the year, the idiosyncratic components of the quarterly process for individual income growth will consist of an innovation to a random walk, "permanent" component of income, $\ln N_{is}$, plus the first difference of white noise, transitory income, $\ln U_{is} - \ln U_{is-1}$, where s denotes the time interval in terms of quarters. Because quarterly permanent shocks to income growth accumulate over a year, we divide the standard deviation of the idiosyncratic permanent shock estimated from annual data by two (the variance

[10] It can be shown that $\ln\left(\frac{1}{N}\sum_{i=1}^{N} Y_{it}\right) - \ln\left(\frac{1}{N} \times \sum_{i=1}^{N} Y_{it-1}\right) = \ln G_t + 0.5\sigma_n^2$.

[11] Carroll uses $\sigma_n = \sigma_u = 0.1$. We deviate slightly from these values by using $\sigma_n = 0.08$, and $\sigma_u = 0.1$ to be consistent with the earnings-process specification in the incomplete-information case (see discussion below). Using $\sigma_n = 0.1$ instead of $\sigma_n = 0.08$ did not change the conclusions reported later in the paper.

[12] Ludvigson and Michaelides (1999) report several different selection criteria to evaluate the best ARMA process for aggregate labor-income growth.

[13] Because time aggregation of permanent shocks induces positive serial correlation, taking the log difference in the sum of income requires the virtual elimination of the permanent component in quarterly income in order to match the observed negative autocorrelation in annual income growth. Such a specification did not seem sensible given that the little evidence we have on quarterly household income suggests a substantial permanent component (Pischke, 1995). Of course, if households receive exactly one-fourth of their annual income every quarter, the year-over-year log change in income will be identical to the log change in the sum of income over four quarters.

by four) to get quarterly permanent innovations, so that $\sigma_n = 0.04$ for the baseline quarterly model.[14]

Empirical results in Ludvigson and Michaelides (1999) on aggregate data show that the standard deviation of aggregate labor-income growth, denoted, σ_g, is about 0.025 at an annual rate in annual data (0.0115 at a quarterly rate in quarterly data). In addition, the mean of the aggregate income-growth process, denoted μ_g, is found to be 0.02 at an annual rate (0.005 at a quarterly rate). We use these parameter values below.

There is no known analytical solution to the problem above; thus we seek a numerical solution described in the Appendix.

A. *Information Assumptions*

In a recent paper, Pischke (1995) argues that economywide shocks account for a very small fraction of the variance in individual earnings growth. As such, households have little incentive to distinguish aggregate from idiosyncratic shocks to their income. Although Pischke does not address the impact of this form of incomplete information on buffer-stock-saving behavior, he shows that informational assumptions can have important effects on aggregate consumption when individual households behave according to the PIH.

We apply Pischke's insights to the buffer-stock model by considering two frameworks. In the first, which we define as the standard buffer-stock model, we assume that individuals observe each component of their earnings separately (*complete information*); in the second, we assume that households have no way of separating the individual and aggregate components of their earnings-growth process (*incomplete information*), but rather can only observe how much their income changed in a given period.

The growth in income is assumed to be stationary, so that the aggregate and individual components can be described by their respective stationary Wold representations,

$$(5)\qquad \Delta \ln Y_{it} = A(L)\varepsilon_t + B(L) \ln N_{it} + C(L) \ln U_{it},$$

where $A(L)$, $B(L)$, and $C(L)$ are distributed lag operators and the subscript i denotes a household-specific variable. For the standard buffer-stock model, we assume that individuals can distinguish each of the separate Wold representations for income growth given in (5). Note that this definition of complete information implies not only that individuals understand the source of each shock to their income, but also that they comprehend perfectly the implications of the aggregate shock for their future income growth.

Given stationarity, individual income changes also have a single Wold representation

$$(6)\qquad \Delta \ln Y_{it} = D(L)\eta_{it}.$$

If individuals can observe only the log difference in their income each period and do not distinguish any of the separate components given in (5), the process for individual income growth appears just like the single Wold representation in (6).

When the aggregate component of the earnings growth process is i.i.d., individual income growth follows

$$(7)\qquad \Delta \ln Y_{it} = \mu_g + \varepsilon_t + \ln N_{it} + \ln U_{it} - \ln U_{it-1},$$

where $\ln G_t = \mu_g + \varepsilon_t$, and $A(L) = 1$, $B(L) = 1$, and $C(L) = 1 - L$. Given incomplete information, individual earnings growth looks just like an MA(1), a specification that has been estimated in many empirical studies using annual data on household income:[15]

$$(8)\qquad \Delta \ln Y_{it} = \mu_g + \eta_{it} - \psi\eta_{it-1},$$

where $\psi > 0$.

We consider models of both annual and quarterly decision periods. In the annual incomplete-

[14] The idiosyncratic transitory shocks, once accumulated, completely cancel each other out since the transitory component is white noise. Accordingly, we keep the standard deviation of this component the same at both quarterly and annual frequency. This procedure follows Carroll and Samwick (1997).

[15] For example, MaCurdy (1982), Abowd and Card (1989).

information version of the model, aggregate income growth appears well described by an i.i.d. process, and we solve for the consumption policy rule assuming that individuals only observe the log first difference in income [in this case the MA(1) process in (8)] even though shocks to this process are actually generated by innovations to each of the separate components given in (7). The parameters of (8) are calibrated by setting $\psi = 0.44$, and $\sigma_\eta = 0.15$, in accordance with evidence from annual earnings in household data.[16] These values, along with an estimate of σ_g taken from aggregate annual data, generate values for σ_u, and σ_n by matching variances and covariances between (7) and (8). Note that the estimate of ψ pins down the relative variance of the permanent to transitory component, implying $\sigma_n = 0.08$ and $\sigma_u = 0.10$, values that are consistent with Carroll's (1992) estimates obtained using annual PSID household data.

In quarterly data, aggregate income growth can be well described by a stationary AR(1) process, thus we specify a quarterly process for individual income as

$$(9) \qquad \Delta \ln Y_{it} = \mu_g + \frac{\varepsilon_t}{1 - \phi L} + \ln N_{it} + \ln U_{it} - \ln U_{it-1}.$$

As with the annual incomplete-information model, we assume that individuals in the quarterly model with incomplete information only observe the log difference of their income, even though shocks to this process are actually generated by innovations to each of the separate components in (9).

II. Aggregate Consumption Properties of Buffer-Stock Saving Behavior

To study the properties of aggregate buffer-stock consumption, we employ an aggregation procedure that accumulates the consumption decisions of many buffer-stock consumers. The procedure takes the following steps. First, for a large number of households, we use Monte Carlo simulations to generate the idiosyncratic shocks to their income for 200 periods (roughly the size of our quarterly data set). In each period, every household also receives a common aggregate shock to their income growth governed by the process for $\ln G_t$. The number of households is determined by increasing the population until the individual income draws, once aggregated, match the appropriate aggregate income process. This procedure showed that 2,000 households was sufficient; using more households did not change the results.

Second, given an initial value for normalized cash-on-hand, x_{it}, we use the policy functions computed from the microeconomic model of consumption behavior to determine the value of individual consumption, c_{it}, for each household.[17] An aggregate time series on consumption and labor income is constructed by taking the cross-sectional average of the simulated data at each point in time. Aggregate consumption is simply the sum of many individual consumption outcomes.

Throughout the paper we define the *smoothness ratio* as the ratio of the standard deviation of consumption growth to that of labor-income growth. The *sensitivity coefficient* is defined as the ordinary least squares (OLS) coefficient from a regression of aggregate consumption growth on lagged aggregate labor-income growth.[18] These properties from the model can be compared with the analogous statistic computed from U.S. aggregate data, shown in Table 1.

[16] See MaCurdy (1982), Abowd and Card (1989), and Pischke (1995), which all find that ψ is estimated around 0.44. The choice of 0.15 for σ_η follows Deaton (1991). Although this value is lower than what MaCurdy, Abowd and Card, and Pischke estimated, Deaton argues that this variance should be reduced from its estimated value which is likely inflated by substantial measurement error in recorded income.

[17] We used several different procedures to find the initial cash-on-hand to income ratio. In practice, the results are not sensitive to the initial asset-income ratio because x_{it} converges to an invariant distribution after about 12 periods. This characteristic of buffer-stock saving behavior has been documented elsewhere. See Deaton (1991, 1992); Carroll (1997a).

[18] Many of the early tests for excess sensitivity consisted of regressing the change in consumption on the lagged *level* of income. As Mankiw and Matthew D. Shapiro (1985) point out, however, if income is not trend stationary, these tests will be biased in favor of finding excess sensitivity even if none exists. Since we find strong evidence of a unit root in aggregate income, we use lagged values of the *growth* in income, thereby avoiding inference problems due to the nonstationarity of lagged income.

TABLE 1—RELATIVE SMOOTHNESS AND EXCESS SENSITIVITY: U.S. AGGREGATE DATA

	Relative smoothness	Excess sensitivity
A. Annual Data		
$\Delta C_t/C_{t-1}$	0.48	0.17
	(0.04)	(0.07)
$\Delta C_t^{ND}/C_{t-1}^{ND}$	0.61	0.18
	(0.06)	(0.08)
$\Delta C_t^{S}/C_{t-1}^{S}$	0.43	0.14
	(0.04)	(0.06)
B. Quarterly Data		
$\Delta C_t/C_{t-1}$	0.47	0.16
	(0.04)	(0.03)
$\Delta C_t^{ND}/C_{t-1}^{ND}$	0.68	0.16
	(0.05)	(0.05)
$\Delta C_t^{S}/C_{t-1}^{S}$	0.46	0.15
	(0.03)	(0.03)

Notes: The sample period begins with the first quarter of 1947 and ends with the fourth quarter of 1997. The column labeled "Relative smoothness" reports the ratio of the standard deviation of the aggregate consumption growth measure in the row, to the standard deviation of aggregate labor-income growth. In parentheses are the standard errors for this ratio, computed by GMM. The column labeled "Excess sensitivity" reports the OLS coefficient of consumption growth on lagged labor-income growth. OLS standard errors are in parentheses. The consumption growth measure $\Delta C_t/C_{t-1}$ is growth in real, per capita nondurables and services expenditure less shoes and clothing. $\Delta C_t^{ND}/C_{t-1}^{ND}$ is nondurables expenditure growth less shoes and clothing and $\Delta C_t^{S}/C_{t-1}^{S}$ is growth in services expenditure. Labor income is compiled from the NIPA components as wages and salaries plus other labor income, minus personal contributions for social insurance minus taxes. Taxes are defined as the fraction of wage and salary income in total income, times personal tax and nontax payments. This measure is also per capita and is deflated by the PCE chain-type price deflator.

The table gives the smoothness ratio and sensitivity coefficient for three measures of real, per capita consumption expenditure: nondurables and services less shoes and clothing (C), nondurables less shoes and clothing (C^{ND}), and services (C^S);[19] standard errors for the smoothness ratio and for the excess-sensitivity coefficient are given in parentheses.[20]

Table 1 shows that consumption growth is about half as volatile as the growth in aggregate labor income in both quarterly and annual data, with a smoothness ratio equal to 0.48 in annual data and 0.47 in quarterly data. Much of the smoothness is in services expenditure: the smoothness ratio for this component of spending is 0.46. By comparison, the smoothness ratio for nondurables is 0.68. For all of these categories of consumer expenditure, the smoothness ratio is statistically significantly smaller than one. Consumption growth is also positively correlated with lagged income growth in both annual and quarterly data. The sensitivity coefficients are statistically significant at better than the 5-percent level in data of both frequencies, and the point estimates are virtually identical regardless of which measure of consumption is used, equal to 0.16 and 0.17 for quarterly and annual data respectively.

To investigate how closely the aggregate buffer-stock models come to matching these features of aggregate data, we perform 100 simulations of the models described above and report the mean smoothness ratio and sensitivity coefficients in Tables 2 and 3. The standard deviation across simulations for each statistic is given in parentheses. Table 2 presents results from the standard, complete-information framework for both the annual model [where aggregate labor-income growth is assumed to be an AR(1) process]. Table 3 presents the analogous results for the alternate model with incomplete information. The simulated statistics reported in the first two columns are given for two values of the risk-aversion parameter, $\rho = 1$, and $\rho = 2$, and for a low- and high-variance case. The high-variance case utilizes the baseline values for the standard deviation of each component of income; for the annual model these are, $\sigma_g =$

[19] These measures of consumption are part of the personal consumption expenditure series, seasonally adjusted at annual rates, in billions of chain-weighted 1992 dollars. Labor income is defined as wages and salaries plus other labor income minus personal contributions for social insurance minus taxes. Taxes are defined as the fraction of wage and salary income in total income, times personal tax and nontax payments. This measure of labor income is in per capita terms and is deflated by the PCE chain-type price deflator. The source for both the consumption and income components is the Bureau of Economic Analysis (BEA). The construction of all of these variables is detailed in Ludvigson and Michaelides (1999).

[20] The standard errors for the smoothness ratio are computed by estimating population moments of consumption and income growth using Generalized Method of Moments (GMM) and then applying the delta method.

TABLE 2—SIMULATED RELATIVE SMOOTHNESS AND EXCESS SENSITIVITY: COMPLETE INFORMATION

Risk preference	Low variance		High variance		High relative variance	
A. Annual Model, I.I.D. Income Growth						
$\rho = 1$	0.99	0.004	0.99	0.00	0.98	0.00
	(0.004)	(0.070)	(0.009)	(0.071)	(0.008)	(0.067)
$\rho = 2$	0.98	0.006	0.98	0.009	0.97	0.012
	(0.004)	(0.067)	(0.010)	(0.070)	(0.009)	(0.063)
PIH	1.00	0.00	1.00	0.00	1.00	0.00
B. Quarterly Model, AR(1) Income Growth						
$\rho = 1$	1.06	0.063	1.06	0.059	1.05	0.051
	(0.017)	(0.083)	(0.019)	(0.072)	(0.025)	(0.081)
$\rho = 2$	1.09	0.055	1.08	0.041	1.07	0.035
	(0.018)	(0.080)	(0.022)	(0.078)	(0.024)	(0.079)
PIH	1.26	0.00	1.26	0.00	1.26	0.00

Notes: The first number in each cell is the mean smoothness ratio (the ratio of the standard deviation of consumption growth to the standard deviation of income growth) over 100 simulations; the second number is the mean excess-sensitivity coefficient (the OLS coefficient estimate from a regression of consumption growth on lagged income growth). The standard deviation across simulations for each parameter is given in parentheses. The column labeled "Risk preference" indicates whether these statistics are reported for a buffer-stock model with coefficient of relative risk aversion, ρ, or for a representative-agent PIH model, where the representative agent receives the aggregate income process. The column labeled "Low variance" reports these statistics for cases where $\sigma_u = 0.07$, $\sigma_n = 0.05$; the column labeled "High variance" reports these statistics for cases where $\sigma_u = 0.1$, $\sigma_n = 0.08$. The column labeled "High relative variance" reports these statistics for cases where the ratio of the standard deviation of the transitory to permanent shock, σ_u/σ_n, is double that of the high-variance case. The AR parameter, ϕ, is set equal to 0.23.

0.025, $\sigma_n = 0.08$; $\sigma_u = 0.10$, implying the standard deviation of the innovation to overall income growth, σ_η equals 0.15; the low-variance case sets σ_η one-third lower than this value.

By varying ρ and σ_η in the manner discussed above, we may investigate the properties of buffer-stock saving over a wide range of permissible parameter values that adhere to the impatience condition (4). Note that, for example, in the high-variance case and with δ and r set at their baseline values, calibrating the value of the risk-aversion parameter, ρ, higher than two would lead to a violation of (4). Although we could raise risk aversion as long as we lowered r or raised δ by a sufficient amount, such permutations of the parameter space do not produce results that are substantively different from our current baseline case. The reason is that each of these parameter combinations yield very similar degrees of impatience relative to prudence. The degree by which the constraint (4) binds (that is, how close the inequality is to being strict) is about the same in each case, and it is this relationship among the parameters, rather than the value of any one parameter, that has the most important influence on the results. Similarly, given this range for ρ and σ_η, varying r, δ, and g from their baseline values by amounts that do not breach (4) has little effect on the outcomes reported below.

In the third column of Tables 2 and 3, we report an additional case in which the standard deviation of the transitory component relative to the standard deviation of the permanent component is twice as high as that in the high-variance, baseline case. The results in this column, along with the motivation for considering such a case, are discussed later. We begin by discussing the first two columns of Tables 2 and 3.

The first two columns of Panel A, Table 2, show the results from the annual model with complete information. Over a range of parameter values, simulated aggregate consumption is just as volatile as aggregate income—the smoothness ratio is very close to one in each case, and, in all but one case, is not statistically different from one at the 5-percent level of significance. Consumption and income have roughly the same volatility in this model because the variation in aggregate consumption

TABLE 3—SIMULATED RELATIVE SMOOTHNESS AND EXCESS SENSITIVITY: INCOMPLETE INFORMATION

Risk preference	Low variance		High variance		High relative variance	
A. Annual Model, I.I.D. Income Growth						
$\rho = 1$	0.96	0.047	0.92	0.114	0.84	0.180
	(0.005)	(0.064)	(0.011)	(0.066)	(0.016)	(0.061)
$\rho = 2$	0.92	0.104	0.85	0.168	0.77	0.212
	(0.008)	(0.062)	(0.016)	(0.066)	(0.022)	(0.058)
PIH	1.00	0.00	1.00	0.00	1.00	0.00
B. Quarterly Model, AR(1) Income Growth						
$\rho = 1$	0.93	0.365	0.91	0.381	0.88	0.400
	(0.014)	(0.056)	(0.019)	(0.056)	(0.017)	(0.049)
$\rho = 2$	0.91	0.433	0.84	0.409	0.77	0.400
	(0.016)	(0.054)	(0.019)	(0.049)	(0.017)	(0.046)
PIH	1.26	0.00	1.26	0.00	1.26	0.00

Notes: The first number in each cell is the mean smoothness ratio (the ratio of the standard deviation of consumption growth to the standard deviation of income growth) over 100 simulations; the second number is the mean excess-sensitivity coefficient (the OLS coefficient estimate from a regression of consumption growth on lagged income growth). The standard deviation across simulations for each parameter is given in parentheses. The column labeled "Risk preference" indicates whether these statistics are reported for a buffer-stock model with coefficient of relative risk aversion, ρ, or for a representative-agent PIH model, where the representative agent receives the aggregate income process. The column labeled "Low variance" reports these statistics for cases where $\sigma_u = 0.07$, $\sigma_n = 0.05$; the column labeled "High variance" reports these statistics for cases where $\sigma_u = 0.1$, $\sigma_n = 0.08$. The column labeled "High relative variance" reports these statistics for cases where the ratio of the standard deviation of the transitory to permanent shock, σ_u/σ_n, is double that of the high-variance case. The AR parameter, ϕ, is set equal to 0.23.

reflects those saving and dissaving decisions that do not cancel out in the aggregate, namely the responses of individual consumption to economywide shocks. Since (for the results reported in Panel A) consumers correctly perceive that shocks to the aggregate component of their income are permanent, it is undesirable for households to undertake any smoothing out of aggregate shocks. Similarly, there is no sensitivity of aggregate consumption growth to lagged income growth—the OLS point estimate on lagged aggregate income growth is not statistically different from zero.

To form a benchmark of comparison, we contrast these results with the smoothness ratio and sensitivity coefficient that would be implied by a representative-agent version of the PIH, where the representative agent receives the aggregate income process.[21] These figures are presented in the third row of Table 2, Panels A and B. The benchmark PIH statistics depend only on the value of the real interest rate and on the persistence of the aggregate income-growth process, and therefore do not vary over all the parameter permutations given in Table 2. Note that the excess-sensitivity coefficient is always zero in the PIH, and when income growth is i.i.d., the smoothness ratio is unity. Thus, the complete-information, annual model of aggregate buffer-stock behavior produces a smoothness ratio and sensitivity coefficient that is virtually indistinguishable from that which would be implied by the representative-agent PIH model.

How do the results in the standard model differ at quarterly frequency when the process for aggregate income growth is positively autocorrelated? Panel B of Table 2 presents the smoothness and sensitivity coefficients when aggregate labor-income growth follows the

[21] The PIH is formulated in levels, rather than in log levels. It relates changes in the level of consumption to changes in the level of labor income given by $\Delta C_t = A\left(\frac{1}{1+r}\right)\varepsilon_t$, where income has the Wold representation $\Delta Y_t = A(L)\varepsilon_t$. For the purposes of providing a rough benchmark, we assume the process for the first difference in aggregate income can be described by the process for log changes used above, so that the smoothness ratio in this case is the standard deviation of aggregate consumption changes to that of aggregate income changes.

AR(1) discussed above. Positively autocorrelated aggregate income growth produces aggregate consumption that is more volatile than in the annual model. The smoothness ratios range from 1.05 to 1.09, so that consumption is always more variable than income. The standard errors for this statistic show that, in all cases, we cannot reject the hypothesis that the ratio is greater than one at conventional significance levels. The intuition for why consumption is so volatile in this case is simple: since shocks to the aggregate component of income growth are expected to persist, individuals respond to positive (negative) innovations by reducing (increasing) their savings when they are not constrained, thereby causing consumption to move by more than one-for-one in response to a given change in income.

Again, these values can be compared with the smoothness ratios that would be implied by the representative-agent PIH given in the last row of Panel B. In this case, the representative agent receives the aggregate, AR(1) process for the first difference in income. Unlike the annual, complete-information model, the quarterly model is noticeably different from the PIH. With $r = 0.02$, the PIH would generate a smoothness ratio of about 1.26, compared to the range of 1.05 to 1.09 generated by the aggregate buffer-stock model. Thus, although the standard buffer-stock model fails to generate consumption smoothing that is close to empirical values, in the face of positively autocorrelated shocks to aggregate income growth, it does generate consumption that is considerably smoother than a representative-agent, PIH benchmark. Moreover, the standard deviation across simulations suggests that the differences in smoothness are statistically significant. If we compare these figures to the smoothness ratio for expenditures on nondurables (less shoes and clothing—Table 1), we find that the standard buffer-stock model produces a smoothness ratio that is about a third of the way between that of the representative agent, PIH, and the U.S. data.

Table 2 shows that the quarterly model with complete information does not produce a statistically significant correlation between consumption growth and lagged income growth. We conclude that the mere presence of liquidity constraints, coupled with impatience and precautionary motives, is not enough to generate robust excess sensitivity or consumption that is substantially smoother than income.

How does the alternate model with incomplete information fare? Table 3 shows that the introduction of incomplete information into the standard framework causes the smoothness ratio to fall substantially. In the annual model (Panel A), the smoothness ratio is as low as 0.84 in one case, and both economically and statistically lower than the benchmark PIH figure of 1.0. Furthermore, the alternate framework produces aggregate consumption which is far less volatile than the analogous case with complete information. This occurs because individuals mistake part of the macroshock as idiosyncratic and therefore transitory. Smoothing at the individual level gets transferred to aggregate consumption as a result of an (unwitting) smoothing of the aggregate shock.

In the quarterly model (Panel B), the alternate framework also displays more consumption smoothing than the standard model. In this case, the smoothness ratios in the first two columns range between 0.80 and 0.92. In contrast to the analogous complete-information model, consumption growth is always smoother than income growth. In addition, the presence of incomplete information coupled with positive serial correlation in the aggregate income process produces robust excess sensitivity. If a positive aggregate shock to income occurs, individuals mistake the shock for transitory and do not raise consumption by the amount warranted by the persistence of the shock. This misperception creates a sluggish response of consumption to aggregate income shocks: next period's consumption will be raised again when income is higher than expected. The sluggishness of consumption in turn produces a direct correlation between consumption growth and lagged income growth and explains why consumption growth may display excess sensitivity even when income growth follows an i.i.d. process, as it does in the annual buffer-stock model, Panel A. The phenomenon of excess sensitivity coupled with i.i.d. income growth is consistent with what is found in the annual data. Notice, however, that the excess-sensitivity coefficient for the quarterly model with incomplete information ranges from about 0.36 to 0.44 and is strongly significant for every parameter combination we consider. These point estimates are considerably higher than the 0.16 estimated from aggregate data.

In summary, the buffer-stock model with incomplete information goes a long way toward resolving the consumption excesses: it produces aggregate consumption growth that is substantially less volatile than the benchmark PIH, and generates robust excess sensitivity to expected income growth. Despite the improvements over the standard model, however, the incomplete-information buffer-stock model still falls short of matching the data. Consumption growth is both too volatile and too highly correlated with lagged income growth.

Why isn't aggregate buffer-stock consumption smoother? It is tempting to attribute this finding to the fact that buffer-stock consumers hold few assets and face constraints on their liquidity, thereby making it difficult for them to smooth out the effects of transitory shocks to their income. This is not the case, however, because with such transitory noise in individual income, buffer-stock households can achieve a great deal of smoothing with few assets.[22] Inferences about aggregate buffer-stock consumption cannot be made by looking at household-level consumption functions, or vice versa.

Would aggregate consumption be smoother in an economy that also included more patient households who may at times hold large amounts of assets and who do not face constraints on their ability to borrow? It is clear that the answer to this question would be "no" if those households were PIH consumers with complete information about the income processes we have been considering. This follows because we know from the benchmark results in Tables 2 and 3 that those households will have consumption that is at least as volatile as that of an aggregate of buffer-stock consumers. Moreover, results in Ludvigson and Michaelides (1999) show that adding some fraction of PIH consumers to the population who have the same incomplete information about their earnings process as do the buffer-stock consumers typically does not improve the aggregate model's performance. Indeed, when the economy is populated entirely by PIH consumers, the smoothness ratio is farther from its empirical counterpart than for any other mix of these two types of consumers, equal to 0.92 for the baseline annual model, compared to 0.84 when the economy is populated entirely with buffer-stock consumers. We conclude that the lack of smoothness at the aggregate level cannot be explained by an absence of unconstrained households.

It might seem that the insufficient smoothness of aggregate consumption could be resolved by changing the specification of the income process. For example, decreasing the relative variance of the permanent component reduces the perceived persistence of the individual's income process when households cannot observe these components separately, thereby magnifying mistakes households make when responding to permanent aggregate shocks. The exact amount of additional smoothness produced, however, will depend on parameter values. To investigate the importance of the relative variance of transitory to permanent shocks in the models studied above, column 3 of Tables 2 and 3 shows the results from doubling the standard deviation of the transitory component relative to that of the permanent component over the (high-variance) baseline case. The standard deviations of the individual components are adjusted so that the overall variance is the same in columns 2 and 3. Table 2 demonstrates that increasing the importance of the transitory shock does not have a noticeable affect on the smoothness ratio or sensitivity coefficient in the complete-information model. This is not surprising since complete-information households do not misperceive the source of any shock to their income. By contrast, in the incomplete-information model, a larger transitory component leads to smoother consumption for all the parameter values we consider (compare columns 2 and 3, Table 3). The decrease in the volatility of consumption is not large, however. The lowest smoothness ratio obtained from a doubling of σ_u/σ_n over the baseline case is about 0.77, a value that is still substantially above that in the data. Moreover, there are at least two caveats with accepting an arbitrary reduction in the relative variance of the permanent component as a resolution to the smoothness puzzle. First, such a

[22] Ludvigson and Michaelides (1999) show that individual consumption growth in these models is much smoother relative to individual income growth than is aggregate consumption growth relative to the aggregate income growth.

reduction does not appear to be supported by existing evidence from studies using household-income data (MaCurdy, 1982; Abowd and Card, 1989; Carroll, 1992; Pischke, 1995). Second, decreasing the perceived persistence of the household's income process often generates greater excess sensitivity (e.g., rows 1, 2, and 3 of Table 3), already too high in the incomplete-information model.

A. *Extensions*

We have explored whether two variants of buffer-stock-saving models can account for the smoothness and excess sensitivity of aggregate consumption data. In this section we briefly investigate two extensions of these models and ask whether these extended versions allow the model to do a better job of replicating the data.

The first extension follows Attanasio and Weber (1993) who stress the possible role of aggregation bias in generating excess sensitivity in linearized Euler equations. Since the first-order condition for optimal consumption choice is often approximated as a loglinear function, they argue that loglinear Euler equations should ideally be estimated using the average of the logarithms of expenditures, rather than the logarithm of the average. While this estimation is not possible with aggregate data, Attanasio and Weber carry out this exercise using household-level data and find that the overidentifying restrictions of the standard Euler equation may be rejected simply because logs are taken over arithmetic averages rather than taking geometric averages. This finding suggests that some amount of excess sensitivity may be attributable to aggregation bias.

In the simulation exercises performed above, we compiled results by calculating the logarithm of the arithmetic average since this replicates the procedure used to construct actual aggregate data. Although it is impossible to control for this possible source of bias in the data, there is of course no restriction in our simulations and it is straightforward to recompute the statistics using geometric means. We did not find that this change qualitatively influenced the results. For example, using geometric means and our baseline parameter values, the quarterly model with complete information produces an excess-sensitivity coefficient equal to 0.029 and is statistically insignificant, while it produces a smoothness ratio of 1.08. Similarly, the incomplete-information version of this model produces an excess-sensitivity coefficient of 0.41 that is strongly statistically significant, while the smoothness ratio is 0.83. These results are almost identical to those presented in Tables 2 and 3 for the same cases using arithmetic means. It appears that specific features of the buffer-stock model dominate these aggregation effects in determining the degree of excess sensitivity and consumption smoothing relative to that of PIH.[23]

A second extension we consider is to alter the timing of households' decisions so as to create a mismatch between their decision interval and the data-sampling interval. The possibility that these intervals do not coincide has been explored elsewhere as an explanation for excess sensitivity found in aggregate data (for example, Luigi Ermini, 1989, 1993; Lawrence J. Christiano et al., 1991; Heaton, 1993). These authors point out that if the decision period of households is more frequent than the data-sampling interval, measured consumption will be the time-averaged value of multiple consumption decisions, a phenomenon that may introduce spurious serial correlation in consumption growth. For example, if the original consumption process follows a martingale (as in the PIH), the first difference of time-averaged consumption will follow a first-order moving average process (Holbrook Working, 1960). Thus, in empirical tests, this type of temporal aggregation bias could create excess sensitivity even if consumption growth is in fact unforecastable over the interval in which households make decisions.[24]

[23] This may be partly a result of the fact that the typical household in our sample is constrained about 35 percent of the time.

[24] In practice, time averaging has not been found to explain excess sensitivity in aggregate data. Time-averaging problems induce spurious correlations for adjacent observations of a series that has been first-differenced: there is no overlap between differences two or more periods apart. Researchers have instrumented for income growth using variables lagged at least two periods to avoid a spurious finding of excess sensitivity. Using twice-lagged instruments for income growth, however, has not been found to eliminate excess sensitivity in aggregate data (for example, Campbell and Mankiw, 1989; Deaton, 1992; Ludvigson, 1999).

The smoothness and excess sensitivity of consumption growth documented above for the incomplete-information models cannot be explained by this type of temporal aggregation because we assumed above that households'-decision intervals match the data-sampling interval. Moreover, since there is no martingale presumption in the buffer-stock model, the findings of Working (1960) are not directly applicable to the framework we consider. Nevertheless, it is possible that the standard complete-information model—which displays little excess sensitivity in the results above—might match the data better if we modeled households'-decision intervals so that they occurred more frequently than the sampling interval. To consider this possibility, we assume that households make decisions on a quarterly basis, but that the stylized consumption facts are computed from annual averages of quarterly data. To proceed, we simulate the quarterly model, compute annual consumption and labor income as the sum of quarterly data in levels, and then take log differences of these time-averaged data. We then recompute our excess-sensitivity and smoothness statistics using these annual averages, which should be compared with the results from the annual model above in which the data are not time averaged. Because the difference in results this extension produces are not highly sensitive to variations in the parameter values, we only discuss the findings under our baseline parameter assumptions.

The results of this exercise show that time aggregation of household-consumption decisions does improve the standard model's predictions along the excess-sensitivity dimension, increasing the excess-sensitivity coefficient to 0.20 from about zero. Even with this time aggregation, however, the coefficient is not statistically significant at the 5-percent level. Moreover, the smoothness ratio is hardly changed, remaining about 0.98 and not statistically different from 1. The results are qualitatively similar for the model with incomplete information: using annual averages, the smoothness ratio is found to be 0.98, while the sensitivity coefficient is found to be 0.29.[25] Without time averaging, these figures were 0.85 and 0.17, respectively. In summary, time averaging increases the sensitivity coefficient in both models, but the point estimate remains statistically insignificant in the complete-information version of the model and the smoothness ratios are in the same range as that produced by the non-time-averaged data.

III. Conclusions

One of the primary objectives of modern-day consumer theory has been to understand the degree of smoothness and predictability of aggregate consumption. Buffer-stock models with precautionary motives and borrowing restrictions have recently become popular tools for thinking about the way the typical consumer behaves. Yet existing research gives little indication of how close buffer-stock saving behavior, once aggregated, might come to quantitatively matching these well-known features of aggregate consumption.

Can buffer-stock saving explain the twin "consumption excesses?" Our results show that incomplete information about the aggregate component of individual earnings may be an important factor in explaining the smoothness of aggregate consumption growth and its correlation with lagged labor-income growth. Only the incomplete-information version of the aggregate buffer-stock model is capable of simultaneously producing some smoothing of aggregate consumption in the face of persistent aggregate income shocks (excess smoothness), and a robust correlation between consumption growth and lagged income growth (excess sensitivity).

Yet even the buffer-stock model with incomplete information creates a smoothness puzzle that remains to be explained. Aggregate consumption growth generated by the model is, at best, only about 80 percent as volatile as aggregate income growth, compared to about 50 percent in aggregate data. Further, the incomplete-

[25] Comparing this estimate of excess sensitivity with that from the *quarterly* model with incomplete information, it may appear that the model with annual averages generates less excess sensitivity than the quarterly model without averaging. This is not the case, however, since the standard error of the excess-sensitivity coefficient computed from the averaged data is more than twice as large as that computed from the nonaveraged data. Thus the point estimate of the latter is well within the 95-percent confidence interval of the point estimate of the former.

information model often produces estimates of excess sensitivity that are too large to fall within empirical ranges, while cases which deliver estimates closer to empirical values do not deliver the most favorable degrees of consumption smoothing. Lack of smoothing at the individual level or the absence of unconstrained permanent-income consumers does not appear to explain the insufficient smoothness at the aggregate level.

Appendix: Numerical Procedures

Define the marginal utility of money (price of consumption) $p(x_t)$ by

$$p(x_t) = U'(f(x_t)),$$

or equivalently

$$c_t = f(x_t) = U'^{-1}(p(x_t)).$$

We seek the solution to the functional equation

$$p(x) = \text{MAX}[U'(x), \zeta(x)],$$

where the second argument on the right-hand side is defined by

$$\beta \int_\eta \int_g \int_y (1 + r)(\exp(g\eta))^{-\rho} \times p\{(1 + r) * (\exp(-g\eta)) \times (x - U'^{-1}(p(x))) + y\} \times dY(y)\, dG(g)\, dN(\eta).$$

Ten policy functions (one for each aggregate state) are defined by:

$$p(x, i) = \max[U'(x), \beta \sum_j \sum_k \sum_l \pi_{ij}\pi_k\pi_l(1 + r)(G_j N_l)^{-\rho} \times p\{(1 + r)(G_j N_l)^{-1} \times (x - U'^{-1}(p(x, i))) + U_k, j\}].$$

We assume that each shock can be well described by a ten-point discrete Markov process that approximates the Normal Distribution, thus π_{ij}, π_l, and π_k are transition probabilities associated with G_t, N_t, and U_t respectively. We discretize the state variable x by dividing it into 100 equidistant grid points. Starting from an initial guess for the policy function on the right-hand side, we update the function using the above functional equation until convergence, the convergence criterion being a difference of six decimal points between a guess and its update at all grid points. Whenever the need arises to evaluate a candidate policy function in between grid points, we use an interpolation scheme such as cubic splines.

An additional complexity in the AR(1) case involves the transition probabilities for the common shock. Given the autoregression's Markov structure, knowing the current state contains information about next period's state that is used to compute expectations. The transition probabilities are set identical to the transition probabilities of the true underlying autoregressive process.[26] We numerically solve for the transition probabilities π_{ij} of moving from interval i to interval j set to be identical to the transition probabilities from interval to interval of the underlying normal autoregressive process for aggregate income growth. The aggregate shock is approximated using $\{\mu_g + \vartheta z_i\}_{i=1}^{i=10}$ where $\vartheta = \dfrac{\sigma_\varepsilon}{\sqrt{1 - \phi^2}}$. From the properties of the normal distribution of the error term perturbing the aggregate shock, we have

$$\pi_{ij} = \Pr(\vartheta z_j \geq g_t - \mu_g \geq \vartheta z_{j-1} | \vartheta z_i \geq g_t - \mu_g \geq \vartheta z_{i-1}) = \frac{1}{\vartheta\sqrt{2\pi}} \int_{\vartheta z_{i-1}}^{\vartheta z_i} \exp\left(\frac{-x^2}{2\vartheta^2}\right)\left\{\Phi\left(\frac{\vartheta z_j - \phi x}{\sigma}\right) - \Phi\left(\frac{\vartheta z_{j-1} - \phi x}{\sigma}\right)\right\} dx.$$

[26] See Deaton (1991) and Deaton and Laroque (1995) for a lengthier analysis.

For any given $\{\sigma, \phi\}$, this integral can be calculated directly using GAUSS routines that approximate the cumulative normal (Φ).

The optimal individual policy function under incomplete information is obtained by the following functional equations:

$$p(x, i) = \max[\lambda(x), \beta \sum_j \pi_{ij}\gamma_{ij}^{\rho} \times p\{[1 + \gamma_{ij}(1 + r)(x - \lambda^{-1}(p(x, i)))], j\}],$$

where β is the discount factor, $\gamma_{ij} = \exp(-Y_{it+1}/Y_{it})$ and π_{ij} is the probability from income state i given that income state j has occurred.

REFERENCES

Abowd, John M. and Card, David. "On the Covariance Structure of Earnings and Hours Changes." *Econometrica*, March 1989, *57*(2), pp. 411–45.

Attanasio, Orazio P. and Browning, Martin. "Consumption over the Life Cycle and over the Business Cycle." *American Economic Review*, December 1995, *85*(5), pp. 1118–37.

Attanasio, Orazio P. and Weber, Guglielmo. "Consumption Growth, the Interest Rate and Aggregation." *Review of Economic Studies*, July 1993, *60*(3), pp. 631–49.

______ . "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *Journal of Political Economy*, December 1995, *103*(6), pp. 1121–57.

Bacchetta, Philippe and Gerlach, Stefan. "Consumption and Credit Constraints: International Evidence." *Journal of Monetary Economics*, October 1997, *40*(2), pp. 207–38.

Beaudry, Paul and van Wincoop, Eric. "The Intertemporal Elasticity of Substitution: An Exploration Using a U.S. Panel of State Data." *Economica*, August 1996, *63*(251), pp. 495–512.

Blinder, Alan S. and Deaton, Angus. "The Time Series Consumption Function Revisited." *Brookings Papers on Economic Activity*, 1985, (2), pp. 465–511.

Campbell, John Y. "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis." *Econometrica*, November 1987, *55*(6), pp. 1249–73.

Campbell, John Y. and Deaton, Angus. "Why Is Consumption So Smooth?" *Review of Economic Studies*, July 1989, *56*(3), pp. 357–73.

Campbell, John Y. and Mankiw, N. Gregory. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," in Olivier Jean Blanchard and Stanley Fischer, eds., *NBER macroeconomics annual: 1989*. Cambridge, MA: MIT Press, 1989, pp. 185–216.

______ . "Permanent Income, Current Income, and Consumption." *Journal of Business and Economic Statistics*, July 1990, *8*(3), pp. 265–79.

Carroll, Christopher D. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity*, 1992, (2), pp. 61–135.

______ . "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics*, February 1997a, *112*(1), pp. 1–55.

______ . "Death to the Loglinearized Euler Equation! (And Very Poor Health to the Second-Order Approximation)." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 6298, 1997b.

Carroll, Christopher D. and Samwick, Andrew A. "The Nature of Precautionary Wealth." *Journal of Monetary Economics*, September 1997, *40*(1), pp. 41–71.

______ . "How Important Is Precautionary Saving?" *Review of Economics and Statistics*, August 1998, *80*(3), pp. 410–19.

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

Deaton, Angus. "Life-Cycle Models of Consumption: Is the Evidence Consistent With the Theory?" in Truman F. Bewley, ed., *Advances in economic theory: Fifth World Congress*. Cambridge: Cambridge University Press, 1987, pp. 121–48.

______ . "Saving and Liquidity Constraints." *Econometrica*, September 1991, *59*(5), pp. 1221–48.

______ . *Understanding consumption*. Oxford: Clarendon Press, 1992.

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

______ . "Estimating a Nonlinear Rational Expectations Commodity Price Model with Unobservable State Variables." *Journal of Applied Econometrics*, December 1995, Supp. *10*, pp. S9–40.

Eichenbaum, Martin S.; Hansen, Lars Peter and Singleton, Kenneth J. "A Time Series Analysis of Representative Agent Models of Consumption and Leisure Choice under Uncertainty." *Quarterly Journal of Economics*, February 1988, *103*(1), pp. 51–78.

Ermini, Luigi. "Some New Evidence on the Timing of Consumption Decisions and on Their Generating Process." *Review of Economics and Statistics*, November 1989, *71*(4), pp. 643–50.

______ . "Effects of Transitory Consumption and Temporal Aggregation on the Permanent Income Hypothesis." *Review of Economics and Statistics*, November 1993, *75*(4), pp. 736–40.

Flavin, Marjorie A. "The Adjustment of Consumption to Changing Expectations about Future Income." *Journal of Political Economy*, October 1981, *89*(5), pp. 974–1009.

Fuhrer, Jeffrey. "An Optimizing Model for Monetary Policy Analysis: Can Habit Formation Help?" Working paper, Federal Reserve Bank of Boston, 1999.

Galí, Jordi. "Budget Constraints and Time-Series Evidence on Consumption." *American Economic Review*, December 1991, *81*(5), pp. 1238–53.

Gourinchas, Pierre-Olivier and Parker, Jonathan A. "Consumption Over the Life Cycle." Working paper, Princeton University, 1996.

Haliassos, Michael and Michaelides, Alexander. "Portfolio Choice and Liquidity Constraints." Working Paper No. 99–18, University of Cyprus, 1999.

Heaton, John. "The Interaction between Time-Nonseparable Preferences and Time Aggregation." *Econometrica*, March 1993, *61*(2), pp. 353–85.

Heaton, John and Lucas, Deborah J. "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing." *Journal of Political Economy*, June 1996, *104*(3), pp. 443–87.

______ . "Market Frictions, Savings Behavior, and Portfolio Choice." *Macroeconomic Dynamics*, March 1997, *1*(1), pp. 76–101.

Hubbard, R. Glenn; Skinner, Jonathan and Zeldes, Stephen P. "Precautionary Saving and Social Insurance." *Journal of Political Economy*, April 1995, *103*(2), pp. 360–99.

Kimball, Miles S. "Precautionary Saving in the Small and in the Large." *Econometrica*, January 1990, *58*(1), pp. 53–73.

Laibson, David I.; Repetto, Andrea and Tobacman, Jeremy. "Self-Control and Saving for Retirement." *Brookings Papers on Economic Activity*, 1998, (1), pp. 91–172.

Lettau, Martin and Uhlig, Harald. "Rules of Thumb versus Dynamic Programming." *American Economic Review*, March 1999, *89*(1), pp. 148–74.

Ludvigson, Sydney C. "Consumption and Credit: A Model of Time-Varying Liquidity Constraints." *Review of Economics and Statistics*, August 1999, *81*(3), pp. 434–47.

Ludvigson, Sydney C. and Michaelides, Alexander. "Can Buffer Stock Saving Explain the Consumption Excesses?" Working paper, Federal Reserve Bank of New York, 1999.

Ludvigson, Sydney C. and Paxson, Christina H. "Approximation Bias in Linearized Euler Equations." National Bureau of Economic Research (Cambridge, MA) Working Paper No. T0236, 1999.

MaCurdy, Thomas, E. "The Use of Time Series Processes to Model the Error Structure of Earnings in Longitudinal Data Analysis." *Journal of Econometrics*, January 1982, *18*(1), pp. 83–114.

Mankiw, N. Gregory and Shapiro, Matthew D. "Trends, Random Walks, and Tests of the Permanent Income Hypothesis." *Journal of Monetary Economics*, September 1985, *16*(2), pp. 165–74.

Pischke, Jörn-Steffen. "Individual Income, Incomplete Information and Aggregate Consumption." *Econometrica*, July 1995, *63*(4), pp. 805–40.

Quah, Danny. "Permanent and Transitory Movements in Labor Income: An Explanation For 'Excess Smoothness' in Consumption." *Journal of Political Economy*, June 1990, *98*(3), pp. 449–75.

Wilcox, David W. "The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work." *American Economic Review*, September 1992, *82*(4), pp. 922–41.

Working, Holbrook. "Note on the Correlation of First Differences of Averages in a Random Chain." *Econometrica*, October 1960, *28*(4), pp. 916–18.