Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption

By MARTIN LETTAU AND SYDNEY C. LUDVIGSON*

Both textbook economics and common sense teach us that the value of household wealth should be related to consumer spending. Early academic work by Franco Modigliani (1971) suggested that a dollar increase in wealth (holding fixed labor income) leads to an increase in consumer spending of about five cents. Since then, the so-called "wealth effect" on consumption has increasingly crept into both mainstream and policy discussions of the macroeconomy.[1] Today, it is commonly presumed that significant movements in wealth will be associated with movements in consumer spending, either contemporaneously or subsequently. Quantitative estimates of roughly the magnitude reported by Modigliani are routinely cited in leading macroeconomic textbooks,[2] and are important features of many contemporary macroeconomic models, including those still widely studied by both academic economists and practitioners.[3]

In this paper, we reevaluate the empirical foundation for such estimates of the consumption-wealth link. Contrary to conventional wisdom, we find that a surprisingly small fraction of the variation in household net worth is related to variation in aggregate consumer spending. A variance-decomposition shows that the vast majority of quarterly fluctuations in asset values are attributable to transitory innovations that display virtually no association with consumption, contemporaneously, or at any future date. Such innovations have long-lasting effects on wealth, however, exhibiting a half-life of a little over two years.

We begin by noting that a general household budget constraint, in which asset returns are allowed to fluctuate in an arbitrary manner, implies that log consumption, c_t, log asset wealth (net worth), a_t, and log labor income, y_t, share a common trend (they are cointegrated). We find evidence to support this hypothesis in U.S. data.

Our results can be understood intuitively by observing that, since consumption, wealth, and labor income are cointegrated, their annualized

* Lettau: Department of Finance, Stern School of Business, New York University, 44 West Fourth Street, New York, NY 10012 (e-mail: mlettau@stern.nyu.edu); Ludvigson: Department of Economics, New York University, 269 Mercer Street, 7th Floor, New York, NY 10003 (e-mail: sydney.ludvigson@nyu.edu). Lettau acknowledges financial support from the National Science Foundation; Ludvigson acknowledges financial support from the Alfred P. Sloan Foundation and the National Science Foundation. This material is based upon work supported by the National Science Foundation under Grant No. 0224944. We are especially grateful to Serena Ng and Mark Watson for feedback on this project. We also thank John Y. Campbell, Simon Potter, Richard Rogerson, James Stock, and seminar participants at Boston College, Duke University, Harvard University, the John F. Kennedy School at Harvard University, MIT, New York University, the Federal Reserve Bank of New York, Northwestern University, the University of California-Berkeley, the University of California-Los Angeles, the University of Chicago, the University of Minnesota, and the Wharton School for many helpful comments. Nathan Barczi and Adam Kolasinski provided excellent research assistance. Any errors or omissions are the responsibility of the authors, and do not necessarily reflect the views of the National Science Foundation.

[1] Recent examples in which the link between wealth and consumption was highlighted include the Federal Reserve's conference on "New Challenges for Monetary Policy," Jackson Hole Wyoming, August 26–28, 1999, and public remarks by Alan Greenspan, as in his February 17, 2000 testimony before the House Committee on Banking and Financial Services.

[2] For example, Andrew B. Abel and Benjamin S. Bernanke (2001, Ch. 4) cite a four cent on the dollar figure, while Robert J. Gordon (1993, Ch. 17) says that the marginal propensity to consume out of wealth is between three and six cents on the dollar. James M. Poterba (2000) provides a nontechnical discussion of the wealth effect, and simple calculations of the amount by which a household might increase its consumption if it received a favorable wealth shock. His benchmark calculations for a 30-year planning horizon imply a change in consumption of between 3.8 and 7.5 cents for each $1 increase in wealth.

[3] Examples of such models include the Data Resources, Incorporated (DRI) model, the Washington University Macroeconomic Model (WUMM) model, and the new Federal Reserve (FRB/US) model.

growth rates must be tied together in the very long run, and therefore so must their volatilities. Measured over short horizons, however, wealth growth is far more volatile than both consumption and labor income growth. The short-run and long-run properties of these variables can only be reconciled if either, (i) the volatility of consumption and/or labor income growth increases with the horizon over which they are measured, or (ii) the volatility of wealth growth decreases with the horizon over which it is measured. The second possibility implies that wealth is not a random walk, but instead displays mean-reversion and adjusts over long horizons to match the smoothness of consumption and labor income. Our evidence suggests that the second possibility better describes U.S. data than the first, signaling the existence of a significant transitory component in wealth that is unrelated to consumer spending and labor income.

This intuition can be formalized by employing the methodologies developed in Robert King et al. (1991), Jesus Gonzalo and Clive Granger (1995), and Gonzalo and Serena Ng (2001). These studies show how innovations distinguished by their degree of persistence can be empirically identified in cointegrated systems. We use the insights from these papers to empirically identify the permanent and transitory elements—the "trend" and "cycle"—of household net worth, and investigate how these elements are related to consumer spending. (The terms "permanent" and "transitory" are defined precisely below.) Restrictive assumptions about preferences or the behavior of asset returns are unnecessary, and the findings generated are applicable to a wide variety of theoretical structures. The approach allows us to identify a large fraction of variation in wealth that is ultimately unsustained and unrelated to aggregate consumption, without having to explicitly identify the independent causal effect of wealth on consumption.

This paper reports two main results. First, we find that up to 88 percent of the postwar variation in household net worth is generated by transitory innovations, primarily associated with fluctuations in the stock market component of wealth. Second, although transitory shocks dominate postwar variation in wealth, variation in aggregate consumption is dominated by permanent shocks. It follows that the majority of innovations in household net worth are found to be unrelated to aggregate consumer spending, both contemporaneously and at any future horizon. This does not mean that wealth has no effect on consumer spending, but rather that only permanent changes in wealth are associated with movements in consumption.

Yet the transitory changes in wealth we identify not only account for most of the variation in asset values, they are also quantitatively large, and in some episodes staggeringly so. For example, during the stock market boom of the 1990's, our estimates imply that wealth exceeded its long-run trend by as much as $17,000 per person in 1996 dollars, a magnitude equal to 50 percent of real per capita GDP in 2002.

An implication of these findings is that conventional estimates of the wealth effect greatly overstate the response of consumption to a change in wealth. This is because conventional estimates, like those cited above, are commonly based on parameters of the shared *trend* in consumption, labor income, and wealth. If most changes in wealth are not trend movements but are instead transitory movements unrelated to consumption, as we find here, such estimates will significantly exaggerate the true correlation between consumption and wealth.

Perhaps the more relevant finding of this paper is not the observation that conventional estimates overstate the wealth effect, but rather that consumption responds differently to temporary changes in wealth than to permanent changes. It follows that no single number, or "marginal propensity" (as it is often referred to in textbooks and popular commentary), can accurately summarize the response of consumption to wealth. Although a permanent $1 change in wealth may be associated with a four to five cent change in consumption, the results presented here imply that most movements in wealth are transitory and unrelated to consumer spending.

The transitory component of wealth we uncover is not a feature of daily, weekly, or even monthly volatility, but is instead quite persistent, displaying long-term "bull markets" in the late 1960's and 1990's, and a long-term "bear market" in the 1970's. In each of these episodes, a significant movement in stock prices caused wealth to deviate from its long-run trend with

consumption and labor income, and was ultimately restored to that trend—not by a subsequent movement in consumption or labor earnings—but by a subsequent reversal in wealth. The extraordinary stock market boom of the 1990's differs from other episodes only in that it was the most dramatic example of this historical pattern. We note that most macroeconomic models, including the large-scale models mentioned above, make no allowance for important transitory fluctuations in wealth that are orthogonal to consumption, and are therefore likely to misstate both the timing and magnitude of the consumption-wealth correlation.

The results presented here have implications for financial economics as well as macroeconomics. A large and growing body of empirical literature in financial economics finds that aggregate stock market returns are forecastable over long horizons, implying the existence of transitory variation in wealth of the sort that we uncover here. Such predictability evidence does not reveal how quantitatively important that transitory component is, however.[4] A contribution of this paper is to quantify the relative importance of transitory shocks in the variation of household net worth, and document how closely related are such shocks to aggregate consumer spending.

A number of other papers explore issues related to those considered here. King et al. (1991), Cochrane (1994), Jordi Galí (1999), and Neville Francis and Valery A. Ramey (2001) invoke restrictions implied by cointegration to identify specific innovations in a range of structural models. Cochrane (1994) studies two bivariate, cointegrated systems to characterize the permanent and transitory components in GNP and stock prices. Ludvigson and Charles Steindel (1999) estimate the cointegrating relation between consumption, asset wealth, and labor earnings, but do not investigate further implications of this system. Lettau and Ludvigson (2001) show that the cointegrating residual for this system is informative about the future path returns on aggregate stock market indexes in excess of a Treasury bill rate, but not about future consumption growth.[5] However, Lettau and Ludvigson (2001) do not formally identify the permanent and transitory elements of asset wealth, document their relative quantitative importance, or highlight the consumption implications of those findings—the focus of this paper.

In Section I we motivate our use of the trivariate cointegrating system for consumption, asset wealth, and labor income by deriving cointegration from a linearized budget constraint. Section II presents the econometric framework and Section III presents our main empirical results. Section IV addresses the issue of stability in the cointegrating relation. Section V concludes.

I. The Common Trend in Consumption, Wealth, and Labor Income

It is perhaps obvious that consumption, labor income, and household wealth should move together over the long term. This section motivates this intuition more formally by considering the cointegration implications of a standard budget constraint. We build off of work in Campbell and N. Gregory Mankiw (1989); see Lettau and Ludvigson (2001).

We consider a representative agent economy in which all wealth, including human capital, is tradable. Let W_t be beginning of period aggregate wealth (defined as the sum of human capital, H_t, and nonhuman, or asset wealth, A_t) in period t; $R_{w,t+1}$ is the net return on aggregate wealth. For expositional convenience, we consider a simple accumulation equation for aggregate wealth, written

$$(1) \qquad W_{t+1} = (1 + R_{w,t+1})(W_t - C_t).$$

Labor income Y_t does not appear explicitly in

[4] See John H. Cochrane (2001, Ch. 20) for a comprehensive review of the predictability evidence, and for discussion on the connection between predictability and transitory variation in asset values. Recent theoretical research in financial economics has shown that it is possible to construct models with rational, utility-maximizing investors in which the equilibrium return on risky assets varies over time in a way consistent with the presence of a transitory component in asset values (e.g., George M. Constantinides and Darrell Duffie, 1996; John Y. Campbell and Cochrane, 1999). An open question concerns the extent to which such models are capable of capturing the magnitude of the transitory component documented here.

[5] More recently, similar findings have been reported for U.K. data (Emilio Fernandez-Corugedo et al., 2002).

this equation because of the assumption that the market value of tradable human capital is included in aggregate wealth.[6] Throughout this paper we use lower case letters to denote log variables, e.g., $c_t \equiv \ln(C_t)$.

Defining $r \equiv \log(1 + R)$, Campbell and Mankiw (1989) derive an expression for the log consumption-aggregate wealth ratio by taking a first-order Taylor expansion of (1), solving the resulting difference equation for log wealth forward, imposing a transversality condition and taking expectations. The resulting expression is[7]:

$$(2) \quad c_t - w_t = E_t \sum_{i=1}^{\infty} \rho_w^i (r_{w,t+i} - \Delta c_{t+i}),$$

where $\rho_w \equiv 1 - \exp(\overline{c - w})$. The consumption-wealth ratio embodies rational forecasts of returns and consumption growth.

Although this expression is intuitively appealing, it is of little use in empirical work because aggregate wealth includes human capital, which is not observable. Lettau and Ludvigson (2001) address this problem by reformulating the bivariate cointegrating relation between c_t and w_t as a trivariate cointegrating relation involving three observable variables, namely c_t, a_t, and labor income y_t. To understand this reformulation, denote the net return to nonhuman capital $R_{a,t}$ and the net return to human capital $R_{h,t}$, and assume that human capital takes the form, $H_t = E_t \sum_{j=0}^{\infty} \Pi_{i=0}^{j} (1 + R_{h,t+i})^{-i} Y_{t+j}$. A loglinear approximation of H_t yields $h_t = \kappa + y_t + v_t$, where κ is a constant, v_t is a mean-zero, stationary random variable given by $v_t = E_t \sum_{j=1}^{\infty} \rho_h^j (\Delta y_{t+j} - r_{h,t+j})$ and $\rho_h \equiv 1/(1 + \exp(\overline{y - h}))$. Assume that $\rho_h = \rho_w$. (The equations below can easily be extended to relax this assumption but nothing substantive is gained by doing so.) Then the expression $h_t = \kappa + y_t + v_t$, along with an approximation for log aggregate wealth as a function of its component elements, $w_t \approx (1 - \nu)a_t + \nu h_t$ [where $(1 - \nu)$ is the steady state share A/W] furnish an approximate expression using only observable variables on the left-hand side:

$$(3) \quad c_t - \alpha_a a_t - \alpha_y y_t \approx$$

$$E_t \sum_{i=1}^{\infty} \rho_w^i ((1 - \nu) r_{at+i} - \Delta c_{t+i} + \nu \Delta y_{t+1+i}).$$

Several points about equation (3) deserve emphasis. First, if labor income follows a random walk and the expected return to human capital is either constant or proportional to the expected return to nonhuman wealth, $c_t - \alpha_a a_t - \alpha_y y_t$ is proportional to the log consumption-wealth ratio, $c_t - w_t$. For this reason we loosely refer to $c_t - \alpha_a a_t - \alpha_y y_t$ as a proxy for the log consumption-wealth ratio. Second, under the maintained hypothesis that r_{wt}, Δc_t, and Δy_t are stationary, (3) implies that c_t, a_t, and y_t are cointegrated and $c_t - \alpha_a a_t - \alpha_y y_t$ is a cointegrating residual. The parameters α_a and α_y should in principle equal the shares $(1 - \nu)$ and ν, respectively, but in practice may sum to a number less than one because only a fraction of total consumption based on nondurables and services expenditure is observable (Lettau and Ludvigson, 2001). Third, if the cointegrating residual on the left-hand side of (3) is not constant, it must forecast either changes in log asset wealth (returns), changes in log labor income, changes in log consumption growth, or some combination of the three. Lettau and Ludvigson (2001) find that it is a strong predictor of excess returns on aggregate stock market indexes, a component of r_{at}, but not of consumption growth. In this paper we document that it is also not a predictor of labor income growth.

Notice that the framework above implies an additional restriction, namely that the log consumption-wealth ratio, (2), and therefore the cointegrating residual $c_t - \alpha_a a_t - \alpha_y y_t$ in (3), should be covariance stationary, and not merely stationary around a deterministic trend. If this were not the case, the deterministic trend would imply that either consumption or aggregate wealth must eventually become an infinitesimal

[6] None of the derivations below are dependent on this assumption. In particular, equation (3), below, can be derived from the analogous budget constraint in which human capital is nontradeable: $A_{t+1} = (1 + R_{a,t+1})(A_t + Y_t - C_t)$, where $H_t = E_t \sum_{j=0}^{\infty} \Pi_{i=0}^{j} (1 + R_{a,t+i})^{-i} Y_{t+j}$.

[7] We omit unimportant linearization constants in the equations throughout the paper.

fraction of the other, violating the budget constraint identity. It follows that the cointegrating parameters which eliminate the common trends in c_t, a_t, and y_t should eliminate both the stochastic and deterministic trends, corresponding to the concept of deterministic cointegration emphasized by Masao Ogaki and Joon Y. Park (1997).

By combining the approximate budget constraint relation (3) with models of consumer preferences, one may obtain approximate consumption functions. For example, equation (3) provides a loglinear generalization of Campbell's (1987) "rainy-day" equation when combined with the specific formulation of preferences explored in Marjorie Flavin (1981), Campbell and Angus Deaton (1989), and Galí (1990). In that model, both expected asset returns and expected consumption growth are constant, so the right-hand side of (3) is a function only of expected future labor income changes.

Importantly, however, the empirical approach described next does not require imposing such additional structure. The empirical results we obtain exploit only cointegration, a phenomenon that can be motivated by the logic of a simple budget constraint identity, applicable to a wide variety of theoretical structures.

II. Econometric Framework

The contribution of this paper is the use of cointegration to identify permanent and transitory components of consumption, wealth, and labor earnings. To explain our approach, this section describes how we isolate the permanent and transitory shocks of a generic cointegrated vector of variables, $\mathbf{x}_t$, that has n elements. In our application, $\mathbf{x}_t = (c_t, a_t, y_t)'$. In the discussion below, we refer the reader to the papers cited for a detailed description of the permanent-transitory decomposition, and only briefly summarize the methodology here.

A. *Data and Preliminary Analysis*

Appendix A contains a detailed description of the data used in this study. The log of asset wealth, a_t, is a measure of real, per capita household net worth, which includes all financial wealth, housing wealth, and consumer durables. Durable goods are properly accounted for as part of nonhuman wealth, A_t, a component of aggregate wealth, W_t, and so should not be accounted for as part of consumption or treated purely as an expenditure.[8] The budget constraint (1) therefore applies to the *flow* of consumption, C_t; durables expenditures are excluded in this definition because they represent replacements and additions to a capital stock (investment), rather than a service flow from the existing stock. The total flow of consumption is unobservable because we lack observations on the service flow from much of the durables stock. We therefore follow Alan S. Blinder and Deaton (1985) and Campbell (1987) and use the log of real, per capita, expenditures on nondurables and services (excluding shoes and clothing), as a measure of c_t.[9] An internally consistent cointegrating relation may then be obtained if we assume that the log of (unobservable) real total flow consumption is cointegrated with the log of real nondurables and services expenditures. The log of after-tax labor income, y_t, is also measured in real, per capita terms. Our data are quarterly and span the fourth quarter of 1951 to the first quarter of 2003.

Wealth is measured at the end of the period. A timing convention for wealth is needed because the level of consumption is a flow during the quarter rather than a point-in-time estimate as is wealth (consumption data are time-averaged). If we think of a given quarter's consumption data as measuring spending at the beginning of the quarter, then wealth for the quarter should be measured at the beginning of the period. If we think of the consumption data as measuring spending at the end of the quarter, then wealth for the quarter should be measured at the end of the period. None of our main findings discussed below (estimates of the cointegrating parameters, error-correction specification, or permanent-transitory decomposition) are sensitive to this timing convention.

[8] Treating durables purchases purely as an expenditure (by, e.g., removing them from A_t and including them in C_t) is improper because it ignores the evolution of the asset over time, which must be accounted for by multiplying the stock by a gross return. (In the case of many durable goods this gross return would be less than one and consist primarily of depreciation.)

[9] This measure of consumption typically comprises over 85 percent of total personal consumption expenditures.

Given our finding that most of the variation in wealth is not associated with consumption, this timing convention is conservative in that the use of end-of-period wealth produces a higher contemporaneous correlation between consumption growth and wealth growth. The contemporaneous correlation between quarterly growth in household net worth, Δa_t, and the real return on the Center for Research on Security Prices (CRSP) value-weighted stock market index is 0.87 in our sample. This shows that quarterly fluctuations in household net worth are dominated by movements in stock market returns.

How can we identify the permanent and transitory components of the three variable system, $\mathbf{x}_t$? Identification is possible because cointegration places restrictions on the long-run multipliers of the shocks in a structural model where innovations are distinguished by their degree of persistence, (King et al., 1991; Gonzalo and Granger, 1995). The procedure has several steps. The first is to estimate a vector-error-correction model (VECM) for the cointegrated system. The estimated VECM parameters may then be used to back out the long-run restrictions.

To obtain a correctly specified error-correction model, we begin by testing for both the presence and number of cointegrating relations in $\mathbf{x}_t$. These results are contained in Appendix B. We assume all of the variables contained in $\mathbf{x}_t$ are first-order integrated, or I(1), an assumption confirmed by unit root test results, available upon request. In addition, the results presented in the Appendix B suggest the presence of a single cointegrating vector; we impose this in our VECM specification from now on. The cointegrating coefficient on consumption is normalized to one, and we denote the single cointegrating vector for $\mathbf{x}_t = [c_t, a_t, y_t]'$ as $\boldsymbol{\alpha} = (1, -\alpha_a, -\alpha_y)'$.

The cointegrating parameters α_a and α_y must be estimated. We use a dynamic least-squares procedure, which generates "superconsistent" estimates of α_a and α_y (James H. Stock and Mark W. Watson, 1993).[10] We estimate $\hat{\boldsymbol{\alpha}} = (1, -0.30, -0.60)'$. The Newey-West corrected t-statistics for these estimates are 12 and 27, respectively. As discussed in Lettau and Ludvigson (2001), these coefficients are not expected to sum to unity because total consumption is unobservable and nondurables and services expenditure (a subset of the total) is used as a proxy. Yet the estimates themselves are intuitively appealing: they imply that the share of capital income in total output is about 0.34, very close to values used in the real-business-cycle literature.

We are now in a position to estimate the VECM representation of $\mathbf{x}_t$ which takes the form

$$\Delta \mathbf{x}_t = \boldsymbol{v} + \boldsymbol{\gamma}\hat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1} + \boldsymbol{\Gamma}(L)\Delta\mathbf{x}_{t-1} + \mathbf{e}_t, \quad (4)$$

where $\Delta\mathbf{x}_t$ is the vector of log first differences, $(\Delta c_t, \Delta a_t, \Delta y_t)'$, $\boldsymbol{v}$, and $\boldsymbol{\gamma} \equiv (\gamma_c, \gamma_a, \gamma_y)'$ are (3×1) vectors, $\boldsymbol{\Gamma}(L)$ is a finite-order distributed lag operator, and $\hat{\boldsymbol{\alpha}} \equiv (1, -\hat{\alpha}_a, -\hat{\alpha}_y)'$ is the (3×1) vector of previously estimated cointegrating coefficients.[11] Throughout this paper, we use "hats" to denote the estimated values of parameters.

The term $\hat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$ gives last period's equilibrium error, or cointegrating residual; $\boldsymbol{\gamma}$ is the vector of "adjustment" coefficients that tells us which variables subsequently adjust to restore the common trend when a deviation occurs. The Granger Representation Theorem states that, if a vector $\mathbf{x}_t$ is cointegrated, at least one of the adjustment parameters, γ_c, γ_a, or γ_y must be nonzero in the error-correction representation (4). Thus if x_j does at least some of the adjusting needed to restore the long-run equilibrium subsequent to a shock that distorts this equilibrium, γ_j should be different from zero in the equation for Δx_j of the error-correction representation (4).

The results of estimating a first-order specification of (4) are presented in Table 1.[12] The estimates of the adjustment parameters in $\boldsymbol{\gamma}$ are given in the next-to-last row of Table 1 and are the coefficients on the lagged cointegrating residual $\hat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$. Consumption and labor income are predictable by lagged consumption and wealth growth, but not by the cointegrating

[10] We use eight leads and lags of the first differences of Δy_t and Δa_t in the dynamic least-squares (DLS) regression. Monte Carlo simulation evidence in both Ng and Pierre Perron (1997) and our own suggested that the DLS procedure can be made more precise with larger lag lengths.

[11] Standard errors do not need to be adjusted to account for the use of the generated regressor, $\boldsymbol{\alpha}'\mathbf{x}_t$ in (4) because estimates of the cointegrating parameters converge to their true values at rate T, rather than at the usual rate $\sqrt{T}$ (Stock, 1987).

[12] This first-order lag length was chosen in accordance with the Akaike and Schwarz criteria.

TABLE 1—ESTIMATES FROM A COINTEGRATED VAR

	Equation		
Dependent variable	Δc_t	Δa_t	Δy_t
Δc_{t-1}	**0.203**	0.063	**0.499**
	(2.659)	(0.184)	(3.197)
Δa_{t-1}	**0.043**	0.099	**0.090**
	(2.696)	(1.375)	(2.783)
Δy_{t-1}	0.067	−0.041	−0.123
	(1.674)	(−0.226)	(−1.500)
$\hat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$	−0.036	**0.387**	0.008
	(−1.380)	(3.236)	(0.159)
$\bar{R}^2$	0.157	0.042	0.086

Notes: The table reports the estimated coefficients from cointegrated vector autoregressions of the column variable on the row variable; *t*-statistics are in parentheses. Estimated coefficients that are significant at the 5-percent level are highlighted in bold face. The term $c_t - \hat{\alpha}_a a_t - \hat{\alpha}_y y_t = \hat{\boldsymbol{\alpha}}'\mathbf{x}_t$ is the estimated cointegrating residual. The sample spans the fourth quarter of 1951 to the first quarter of 2003.

residual $\hat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$. Thus, estimates of γ_c and γ_y are economically small and insignificantly different from zero. By contrast, the cointegrating error is an economically large and statistically significant determinant of next quarter's wealth growth: γ_a is estimated to be about 0.36, with a *t*-statistic above 3. Although there is some short-run predictability in the growth of consumption and labor income (as exhibited by the dependence of these variables on lagged growth rates), it is wealth growth that exhibits error-correction behavior and therefore predictability over long horizons. Wealth is mean reverting and adapts over long horizons to match the smoothness in consumption and labor income.[13] These findings were highlighted in Lettau and Ludvigson (2001); we now extend these results to quantify the permanent and transitory components in wealth and consumption.

B. *Permanent and Transitory Shocks*

Cointegration may be used to decompose $\mathbf{x}_t$ into innovations that are distinguished by their degree of persistence. Because $\mathbf{x}_t$ has three elements and a single cointegrating vector, there are two permanent shocks, or common trends, and one transitory shock (Stock and Watson, 1988). Following King et al. (1991), Gonzalo and Granger (1995), and Gonzalo and Ng (2001), identification is achieved in two steps. First, cointegration restricts the matrix of long-run multipliers of shocks in the system, which identifies the permanent components. The transitory component is identified as the residual. Second, in order to study the dynamic impact of the transitory innovation, it is assumed to be orthogonal to the two permanent innovations. This second assumption is not necessary to obtain our main results, and is made only to interpret impulse response functions. Our main conclusions can be obtained from a variance-covariance decomposition, which requires no such orthogonalization assumption.

A summary of this methodology is as follows. Under the maintained hypothesis that the growth rates of the variables in $\mathbf{x}_{t-1}$ are covariance stationary, a multivariate Wold representation exists taking the form

$$\Delta\mathbf{x}_t = \boldsymbol{\delta} + \mathbf{C}(L)\mathbf{e}_t,$$

where $\mathbf{e}_t$ is an $n \times 1$ vector of innovations, and where $\mathbf{C}(L)$ is a distributed lag operator. We seek to identify $n = 3$ transformed, or structural-form, innovations distinguished by whether they have permanent or transitory effects. Denote these transformed innovations $\boldsymbol{\eta}_t \equiv (\eta_{1t}, \eta_{2t}, \eta_{3t})'$, where two are permanent and one is transitory. Without loss of generality, shocks are ordered so that the first two have permanent effects, the third transitory effects. Following Gonzalo and Granger (1995), define a shock, η_{Pt}, as permanent if $\lim_{h\to\infty}\partial E_t(\mathbf{x}_{t+h})/\partial\eta_{Pt} \neq 0$, and a shock, η_{Tt}, as transitory if $\lim_{h\to\infty}\partial E_t(\mathbf{x}_{t+h})/\partial\eta_{Tt} = 0$.

The permanent and transitory innovations may be identified using the estimated parameters $\hat{\boldsymbol{\gamma}}$ and $\hat{\boldsymbol{\alpha}}$ from the error-correction representation. Write the reduced-form Wold representation of this system as above, $\Delta\mathbf{x}_t = \boldsymbol{\delta} + \mathbf{C}(L)\mathbf{e}_t$. By the Granger Representation Theorem, the parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$, both of rank *r*, satisfy $\boldsymbol{\alpha}'\mathbf{C}(1) = 0$ and $\mathbf{C}(1)\boldsymbol{\gamma} = 0$ (Robert F. Engle and Clive W. J. Granger, 1987).

[13] We also find that the *four*-quarter lagged value of the cointegrating error strongly predicts asset growth. This shows that the forecasting power of the cointegrating residual for future asset growth cannot be attributable to interpolation procedures used to convert annual survey data to a quarterly housing service flow estimate, part of the services component of c_t.

Let

$$\mathbf{G} = \begin{bmatrix} \boldsymbol{\gamma}_\perp' \\ \boldsymbol{\alpha}' \end{bmatrix}, \tag{5}$$

where $\boldsymbol{\gamma}_\perp'\boldsymbol{\gamma} = 0$. Define a new distributed lag operator $\mathbf{D}(L)$ equal to $\mathbf{C}(L)\mathbf{G}^{-1}$. Gonzalo and Ng (2001) show that the transformed (permanent and transitory) residuals are given by

$$\boldsymbol{\eta}_t = \mathbf{G}\mathbf{e}_t,$$

and their relation to $\mathbf{x}_t$ is given by the Wold representation

$$\begin{aligned} \Delta\mathbf{x}_t &= \boldsymbol{\delta} + \mathbf{C}(L)\mathbf{G}^{-1}\mathbf{G}\mathbf{e}_t \\ &= \boldsymbol{\delta} + \mathbf{D}(L)\boldsymbol{\eta}_t, \end{aligned} \tag{6}$$

where $\boldsymbol{\delta}$ is a constant vector. Each element of $\Delta\mathbf{x}_t$ has been decomposed into a function of two permanent shocks and a single transitory shock.

This decomposition can be understood intuitively by noting that it gives the jth variable a large weight in the permanent innovations and a small weight in the transitory innovations when γ_j is small, implying that the variable participates little in the error-correction required to eventually restore the variables to their common trend. Conversely, it gives the jth variable a small weight in the permanent innovations and a large weight in the transitory innovations when γ_j is large, implying that the variable plays an important role in the error correction required to restore the variables to their common trend. In the application studied here, the elements of the adjustment vector $\boldsymbol{\gamma}$ corresponding to c_t and y_t are statistically indistinguishable from zero (Table 1), implying that these variables have a large weight in the permanent innovations and a small weight in the transitory innovations. By contrast, the element of the adjustment vector $\boldsymbol{\gamma}$ corresponding to a_t is large in absolute value and strongly statistically significant (Table 1), implying that a_t will have a large weight in the transitory innovations and a small weight in the permanent innovations.

With this decomposition, the level of $\mathbf{x}_t$ can be written as the sum of k I(1) common factors (permanent component), and $n - k$ I(0) transitory components, where k is equal to the number of common trends, in our case two.[14]

III. Empirical Results

A. *Permanent and Transitory Components of Consumption, Wealth, and Income*

Using the permanent-transitory decomposition discussed above, it is straightforward to investigate how each of the variables in our system is related to permanent and transitory shocks. Table 2 displays the fraction of the total variance in the forecast error of Δc_t, Δa_t, and Δy_t that is attributable to the two permanent shocks combined (we do not attempt to identify the impact of the two permanent shocks separately), and to the single transitory shock. The transitory shock is constructed to be orthogonal to the two permanent shocks; we use this orthogonal shock in the impulse responses presented below. Later we present variance-covariance decompositions which do not assume that the transitory shock is orthogonal to the two permanent shocks. To quantify the sampling uncertainty of the variance decompositions, we compute cumulative distribution functions for each variance decomposition using a bootstrapping procedure. These results are presented in Appendix C and Table C1.

Gonzalo and Ng (2001) recommend restricting the values of the parameters in $\boldsymbol{\gamma}$ to zero where they are statistically insignificant at the 5-percent level. Their simulation evidence suggests that doing so as a matter of practice generates much more stable estimates of the permanent-transitory decomposition. This is also done elsewhere in other applications of this methodology (Cochrane, 1994, and Gonzalo and Granger, 1995). In the computations that follow, we set γ_c and γ_y to zero in order to match the evidence from Table 1 that these

[14] The k common factors in the Granger-Gonzalo/Gonzalo-Ng decomposition described above are determined by $\boldsymbol{\gamma}_\perp'\mathbf{x}_t$, where $\boldsymbol{\gamma}_\perp'\boldsymbol{\gamma} = 0$. This permanent component may contain serial correlation around the random walk component given by the multivariate Beveridge-Nelson decomposition. The random walk component of the I(1) common factors from the Gonzalo-Granger/Gonzalo-Ng procedure is the trend concept provided by the multivariate Beveridge-Nelson (Stephen Beveridge and Charles R. Nelson, 1981) decomposition investigated in Stock and Watson (1988).

TABLE 2—VARIANCE DECOMPOSITION (ORTHOGONALIZED)

	$\Delta c_{t+h} - E_t\Delta c_{t+h}$		$\Delta y_{t+h} - E_t\Delta y_{t+h}$		$\Delta a_{t+h} - E_t\Delta a_{t+h}$	
Horizon h	P	T	P	T	P	T
			$\gamma_c = \gamma_y = 0$			
1	1.000	0.000	1.000	0.000	0.097	0.903
2	0.967	0.033	0.968	0.032	0.109	0.891
3	0.964	0.036	0.968	0.032	0.111	0.889
4	0.964	0.036	0.968	0.032	0.113	0.887
∞	0.961	0.039	0.966	0.034	0.119	0.881
	(0.92, 0.99)	(0.01, 0.08)	(0.93, 0.99)	(0.01, 0.07)	(0.11, 0.20)	(0.80, 0.89)
			γ_c and γ_y estimated			
1	0.894	0.106	0.999	0.001	0.418	0.582
2	0.895	0.105	0.995	0.005	0.424	0.576
3	0.893	0.107	0.995	0.005	0.421	0.579
4	0.893	0.107	0.995	0.005	0.418	0.582
∞	0.892	0.108	0.994	0.006	0.412	0.588
	(0.63, 0.96)	(0.05, 0.37)	(0.81, 0.99)	(0.01, 0.19)	(0.25, 0.87)	(0.44, 0.76)

Notes: The table reports the fraction of the variance in the h step-ahead forecast error of the variable listed at the head of each column that is attributable to innovations in the permanent shocks, P, and the transitory shock, T. Horizons are in quarters, and the underlying VECM is of order one. The last row reports the bootstrapped, described in Appendix C, 95-percent confidence intervals for the $h = \infty$ case. The sample spans the fourth quarter of 1951 to the third quarter of 2003.

variables are small and statistically indistinguishable from zero, but we also present results using the actual point estimates of $\boldsymbol{\gamma}$.

Table 2 displays the fraction of h-step ahead forecast error in consumption, labor income, and wealth growth that is attributable to the two permanent shocks combined, and to the single transitory shock. For $h \rightarrow \infty$, we compute the portion of the total variance of each variable attributable to each disturbance. The top panel of Table 2 shows the results when γ_c and γ_y are set to zero; the bottom panel shows the results when all elements of $\boldsymbol{\gamma}$ are set to their point estimates.

First consider the results in the top panel. For consumption growth, the two permanent shocks in this system explain over 95 percent of the variance in the forecast error at all horizons. The bootstrapped 95-percent confidence interval is (0.92, 0.99). Thus, consumption is a function of only the permanent components in a_t and y_t: less than 5 percent of the variation in consumption growth is attributable to the transitory shock. Similarly, almost all of the variance in income growth is attributable to the two permanent shocks. Together they account for over 96 percent of the variation in the long-run forecast error of Δy_t, with the bootstrapped 95-percent confidence interval equal to (0.93, 0.99).

The findings are quite different for asset wealth. Notice that the orthogonalization of transitory and permanent shocks assumed above orders the transitory shock last, thereby giving it the smallest possible role in the transitory component of a_t. Despite this ordering, we find that transitory shocks dominate changes in wealth: the estimates reported in the top panel imply that 88 percent of the variation in the growth of asset wealth is attributable to this shock; only 12 percent is attributable to permanent shocks. It follows that the single transitory shock in this system is essentially a wealth shock. This finding implies that the vast majority of variability in consumption, driven by permanent shocks, is disassociated with the vast majority of variability in wealth, driven by transitory shocks. This does not mean that wealth has no impact on consumption, but rather that only permanent changes in wealth are related to consumer spending.

Turning to the bottom panel of Table 2, it is evident that when all elements of $\boldsymbol{\gamma}$ are set to their point estimates, the transitory component of wealth is smaller, but still comprises the

TABLE 3—VARIANCE DECOMPOSITION (UNORTHOGONALIZED)

	$\Delta c_{t+h} - E_t\Delta c_{t+h}$			$\Delta y_{t+h} - E_t\Delta y_{t+h}$			$\Delta a_{t+h} - E_t\Delta a_{t+h}$		
Horizon h	$\tilde{P}$	$\tilde{T}$	$\tilde{P}, \tilde{T}$	$\tilde{P}$	$\tilde{T}$	$\tilde{P}, \tilde{T}$	$\tilde{P}$	$\tilde{T}$	$\tilde{P}, \tilde{T}$
				$\gamma_c = \gamma_y = 0$					
1	1.000	0.000	0.000	1.000	0.000	0.000	0.932	1.663	−1.595
2	0.990	0.013	−0.003	1.026	0.045	−0.071	0.932	1.582	−1.514
3	0.991	0.014	−0.005	1.025	0.046	−0.071	0.888	1.571	−1.459
4	0.990	0.014	−0.005	1.025	0.046	−0.071	0.851	1.560	−1.411
∞	0.990	0.016	−0.005	1.022	0.049	−0.071	0.732	1.531	−1.262
				γ_c and γ_y estimated					
1	1.149	0.153	−0.303	0.949	0.002	0.049	0.608	0.843	−0.451
2	1.134	0.153	−0.287	0.951	0.006	0.043	0.608	0.811	−0.419
3	1.132	0.154	−0.285	0.951	0.006	0.043	0.581	0.830	−0.411
4	1.131	0.154	−0.285	0.951	0.006	0.043	0.560	0.844	−0.404
∞	1.131	0.154	−0.285	0.950	0.007	0.043	0.512	0.880	−0.391

Notes: The table reports the fraction of the variance in the h step-ahead forecast error of the variable listed at the head of each column that is attributable to innovations in the permanent shocks, $\tilde{P}$, the transitory shock, $\tilde{T}$, and two times the covariance between $\tilde{P}$ and $\tilde{T}$. Horizons are in quarters, and the underlying VECM is of order 1. The sample spans the fourth quarter of 1951 to the first quarter of 2003.

majority of its fluctuations: in this case, about 60 percent of the variation in wealth growth is estimated to be transitory in this case. For consumption and labor income, permanent shocks are still the dominant source of variation, although the confidence intervals are wider.

So far we have restricted the single transitory shock to be orthogonal to the two permanent shocks. By doing so, we cannot rule out the possibility that consumption (or labor income) contains a transitory component that is correlated with the permanent components. To address this possibility, Table 3 provides the results of a variance-covariance decomposition, which requires no orthogonalization. The table displays the fraction of the h-step ahead forecast error that is attributable to the variance of the permanent shocks, the transitory shocks, and to two times the covariance between the permanent shocks and the transitory shock. The permanent components of consumption and labor income are virtually uncorrelated with the transitory component; therefore it is not the case that consumption contains a large transitory component that is correlated with the permanent shocks. The transitory component in wealth is correlated with the permanent shocks, but this does not alter the conclusion that the majority of variation in a_t is attributable to transitory shocks and therefore unrelated to the majority of variation in quarterly consumption growth.

An implication of these last findings is that they provide no support for the hypothesis that consumption takes many periods to adapt to permanent innovations in wealth or labor income. Suppose consumption did adapt sluggishly to trend movements in wealth and income. Then it would have to contain a transitory component, correlated with the permanent shocks in wealth or income. Permanent movements in wealth or income would not be immediately accompanied by a full adjustment in consumption to its trend level, generating a transitory component in consumption and a temporary deviation in the cointegrating residual from its mean.

This point can be made another way. If consumption adjusts slowly to shocks, such temporary deviations in the cointegrating residual from its mean should be eliminated by a subsequent movement in consumption, as it sluggishly adapts to the permanent innovation in wealth or income. Deviations from the common trend in c, a, and y should then tell us something about the future path of consumption growth; i.e., the cointegrating residual should have long-horizon forecasting power for consumption growth. Table 4 reports the results of regres-

TABLE 4—LONG-HORIZON REGRESSIONS

Horizon H	Δc_t	Δy_t	Δa_t	$\hat{\boldsymbol{\alpha}}'\mathbf{x}_t$	$\bar{R}^2$
Panel A: $\sum_{h=1}^{H} \Delta c_{t+h}$ regressed on					
1.00	**0.20**	0.07	**0.04**	−0.04	0.17
	(2.82)	(1.58)	(2.96)	(−1.21)	
4.00	**0.61**	0.04	**0.13**	−0.06	0.14
	(3.63)	(0.37)	(3.74)	(−0.42)	
8.00	0.58	0.07	0.09	−0.04	0.04
	(1.64)	(0.40)	(1.74)	(−0.18)	
12.00	0.61	0.14	0.00	−0.17	0.04
	(1.47)	(0.70)	(0.02)	(−0.61)	
16.00	0.56	0.15	0.02	−0.34	0.04
	(1.06)	(0.66)	(0.24)	(−0.94)	
Panel B: $\sum_{h=1}^{H} \Delta y_{t+h}$ regressed on					
Horizon H	Δc_t	Δy_t	Δa_t	$\hat{\boldsymbol{\alpha}}'\mathbf{x}_t$	$\bar{R}^2$
1.00	**0.50**	−0.12	**0.09**	0.01	0.10
	(3.81)	(−1.01)	(2.47)	(0.20)	
4.00	**1.36**	−0.20	**0.15**	0.05	0.12
	(3.95)	(−1.21)	(2.45)	(0.28)	
8.00	1.37	−0.32	0.02	−0.10	0.04
	(2.97)	(−1.27)	(0.26)	(−0.35)	
12.00	1.79	−0.20	−0.03	−0.26	0.06
	(3.17)	(−0.61)	(−0.24)	(−0.73)	
16.00	1.46	0.02	−0.07	−0.45	0.04
	(1.89)	(0.05)	(−0.46)	(−0.84)	
Panel C: $\sum_{h=1}^{H} \Delta a_{t+h}$ regressed on					
Horizon H	Δc_t	Δy_t	Δa_t	$\hat{\boldsymbol{\alpha}}'\mathbf{x}_t$	$\bar{R}^2$
1.00	0.06	−0.04	0.10	**0.39**	0.06
	(0.20)	(−0.23)	(1.20)	(2.61)	
4.00	0.45	0.11	**0.26**	**1.46**	0.17
	(0.70)	(0.36)	(2.01)	(2.67)	
8.00	−0.63	0.76	0.42	**2.84**	0.30
	(−0.61)	(1.77)	(1.62)	(3.86)	
12.00	0.03	**1.04**	0.18	**3.81**	0.37
	(0.02)	(2.19)	(0.59)	(3.96)	
16.00	−0.08	0.96	0.15	**3.72**	0.29
	(−0.06)	**(2.00)**	(0.53)	(2.96)	

Notes: The table reports output from long-horizon regressions of consumption, labor income, and asset wealth on lags of these variables and the cointegrating residual $\hat{\boldsymbol{\alpha}}'\mathbf{x}_t$. The dependent variables in the h-period regressions are $\Delta x_{t+1} + \cdots + \Delta x_{t+h}$, where $x \in \{c, y, a\}$. For each regression, the table reports ordinary least-squares (OLS) estimates of the regressors, Newey-West corrected t-statistics (in parentheses) and adjusted R^2 statistics. Significant coefficients at the 5-percent level are highlighted in bold face. The sample spans the fourth quarter of 1951 to the first quarter of 2003.

sions of long-horizon consumption growth, Δc_{t+h} (defined as $c_{t+h} - c_t$), long-horizon labor income growth, Δy_{t+h}, and long-horizon asset wealth growth, Δa_{t+h}, on the estimated cointegrating residual $\hat{\boldsymbol{\alpha}}'\mathbf{x}_t$, controlling for Δc_t, Δy_t, and Δa_t, over horizons, h, ranging from one to 16 quarters.

The first panel of Table 4 displays the long-horizon forecastability of consumption growth. The coefficients on time t consumption growth are small but statistically significant predictors of consumption growth up to four quarters out, reflecting serial correlation in aggregate expenditure growth. The R^2 statistics imply that the growth of consumption is not highly predictable over horizons in excess of one year, however. More to the point, the cointegrating residual is always statistically indistinguishable from zero in these regressions, and explains a negligible fraction of the variation in future consumption growth at all horizons. The results for labor income growth are similar. These regressions

provide no evidence that consumption adapts sluggishly to permanent innovations in wealth or income. The findings underscore the point that the modest short-horizon predictability in consumption growth evident in the data (and primarily attributable to serial correlation) does not imply that consumption slowly adjusts to permanent innovations in wealth or income, and is therefore forecastable over long horizons by the cointegrating residual.

The only variable for which $\hat{\boldsymbol{\alpha}}'\mathbf{x}_t$ has any forecasting power at any horizon is wealth growth, consistent with the finding reported above that it has a significant transitory component. The residual has marginal predictive power even at a horizon of 16 quarters and beyond, with the adjusted R^2 statistic peaking at about 31 percent at a 12-quarter horizon. These adjusted R^2 statistics are unaffected by removing the other regressors Δc_t, Δy_t, and Δa_t from the forecasting regression, thus all of the long-horizon forecasting power for the growth in asset wealth is attributable to $\hat{\boldsymbol{\alpha}}'\mathbf{x}_t$. Note also that wealth growth is substantially more forecastable than consumption or labor income growth as the horizon over which these variables are measured increases. Wealth is mean reverting and adapts over long horizons to match the smoothness in consumption and labor income.

The permanent-transitory decomposition employed here allows us to identify the random walk component of each variable, given by the multivariate Beveridge-Nelson decomposition for this system. This allows us to assess how highly correlated is each variable with its random walk component. Not surprisingly, consumption and labor income are highly correlated with their random walk components, while wealth is not. Even with the serial correlation in measured spending growth, consumption still displays a correlation of 91 percent with its random walk component (Table 5).[15] Similarly, labor income growth displays an 87-percent correlation with its random walk component. By contrast, asset wealth is far from a random walk, displaying a correlation of just 18 percent with its random walk component.

TABLE 5—CORRELATION OF GROWTH RATES WITH RANDOM WALK COMPONENTS

Variable	Correlation
Δc_t	0.911
Δy_t	0.871
Δa_t	0.182

Note: The sample period is the fourth quarter of 1951 to the first quarter of 2003.

To characterize the dynamic impact of the transitory wealth shock, Figure 1 shows the cumulative responses of Δc_t, Δa_t, and Δy_t, to a one-standard deviation innovation in the transitory shock. Standard errors for these responses are presented in Appendix C. The responses of the estimated cointegrating error, $\hat{\boldsymbol{\alpha}}'\mathbf{x}_t = c_t - \hat{\alpha}_a a_t - \hat{\alpha}_y y_t$, are also plotted. Figure 1 depicts graphically what the variance decompositions depict numerically: an increase in the transitory shock leads to a sharp increase in asset wealth, but has virtually no impact on consumption and labor earnings at any future horizon. The consumption and labor income responses are statistically insignificant (Appendix C), and economically negligible. The effect of a transitory wealth shock on a_t is strongly significant over periods from a quarter to several years, but is eventually eliminated, as it must be, since the shock is transitory. The long-run response of $\hat{\boldsymbol{\alpha}}'\mathbf{x}_t$ to this shock is zero, since deviations from the common trend in c_t, a_t, and y_t must eventually be eliminated. Note that these transitory movements in wealth cause the cointegrating residual to deviate from its mean, forecasting a subsequent adjustment in wealth, which restores the cointegrating relation.

Figure 1 also shows that the transitory shock to wealth is very persistent; a typical shock continues to affect asset values for a little over four years. Thus, transitory variation in wealth is not characterized by mere day-to-day or even quarter-to-quarter volatility. Instead, temporary shocks can lead wealth to deviate for a number of years from its long-run trend. Despite their persistent effect on asset values, such shocks bear virtually no relation to consumption at any future horizon.

It is reasonable to ask whether this transitory

[15] This figure does not contradict the result in Table 2, that 99 percent of the variation in consumption growth is attributable to permanent shocks. The reason is that the permanent shocks in Table 2 are defined as those that have permanent effects, and therefore allow for serial correlation around the random walk innovation.

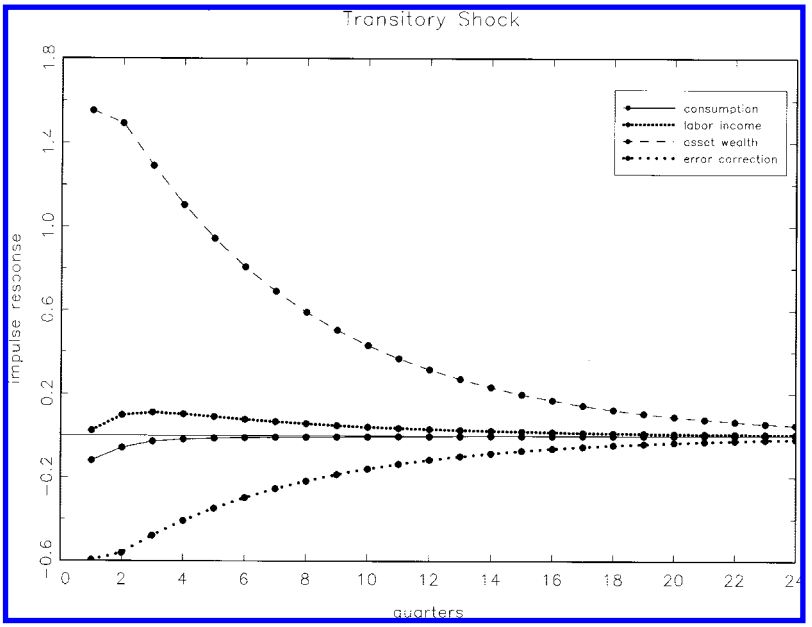

FIGURE 1. TRANSITORY SHOCK

Note: Impulse responses to a one-standard-deviation transitory shock.

variation is driven by the stock market component of wealth. At least two pieces of evidence indicate that it is. First, if we split wealth into its stock and nonstock components, a four-variable variance decomposition shows clearly that, in contrast to stock market wealth, nonstock wealth is dominated by permanent shocks (results available on request). Second, the cointegrating residual for c_t, a_t, and y_t is a strong univariate predictor of stock market wealth growth, but has virtually no predictive power for nonstock wealth growth (Table 6). Since nonstock forms of wealth do not adjust to close a transitory gap in the cointegrating relation, they cannot have an important transitory component.

We close this section by considering what these findings imply for conventional estimates of the wealth effect on consumption. Estimates of this effect, often referred to as the marginal propensity to consume out of wealth, are typically formed from estimates of the cointegrating coefficients for consumption, wealth, and income. This makes some sense because such cointegrating coefficients arise from a regression of consumption on wealth, controlling for the influence of labor income. Moreover, because these estimates are based on cointegrating coefficients, they are superconsistent, and therefore robust to the presence of regressor endogeneity. Modigliani (1971) used this approach to calculate that a dollar increase in wealth leads to an increase in consumer spending of about five cents. Using our own estimates of the cointegrating coefficients for log variables (i.e., $\hat{\alpha}_a = 0.30$), we find that—once translated as a relation between levels of variables—the marginal impact of a dollar increase in wealth on consumption is about 4.6 cents, in line with customary estimates.[16] Similar figures are reported in Ludvigson and Steindel (1999). Such estimates are valid descriptions of the trend relation among consumption, wealth, and income. The magnitude of these estimates is not trivial quantitatively and may explain why it is commonly presumed that sharp swings in asset values will generate important changes in consumer spending.

[16] This number is obtained by multiplying $\hat{\alpha}_a = 0.30$ by the most recent value of C_t/A_t.

TABLE 6—LONG-HORIZON REGRESSIONS: STOCK MARKET WEALTH AND NON-STOCK-MARKET WEALTH

Forecast Horizon H						
1	2	4	8	12	16	24
Panel A: Stock Market Wealth $\sum_{h=1}^{H} \Delta s_{t+h}$						
2.01	3.86	6.93	12.38	16.28	17.85	23.57
(4.28)	(4.00)	(4.14)	(6.10)	(7.79)	(6.84)	(5.71)
[0.08]	[0.14]	[0.23]	[0.42]	[0.52]	[0.46]	[0.43]
Panel B: Non-Stock-Market Wealth $\sum_{h=1}^{H} \Delta n_{t+h}$						
−0.12	−0.21	−0.39	−0.61	−0.83	−1.24	−1.79
(−1.93)	(−1.65)	(−1.55)	(−1.27)	(−1.22)	(−1.40)	(−1.56)
[0.04]	[0.04]	[0.05]	[0.04]	[0.04]	[0.05]	[0.06]

Notes: The table reports results from long-horizon regressions of the log first difference of the stock market component of asset wealth, Δs_t, in Panel A and the log first difference of the non-stock-market component, Δn_t, in Panel B on the lagged cointegrating residual $\hat{\boldsymbol{\alpha}}'\mathbf{x}_t$. In each column, the first number is the OLS coefficient; the second number, in parentheses, is the Newey-West corrected t-statistic; the third number, in square brackets is the adjusted R^2 statistic for the regression. The sample spans the fourth quarter of 1951 to the first quarter of 2003.

But are these estimates appropriate as a summary measure of the consumption-wealth link? It is commonly said that every dollar increase in wealth leads to a four- to five-cent increase in consumption. Because these estimates are based on parameters of a common trend, however, they are informative only about the correlation between consumption and *permanent* movements in wealth, not every movement in wealth. They reveal nothing about the relation between consumption and transitory changes in wealth. If most movements in wealth are transitory, and if transitory movements have a very different impact on consumption than do permanent movements, such an estimate will poorly describe the true correlation between wealth and consumption.

A better summary statistic would account for the finding that much of the variation in wealth is transitory and uncorrelated with consumer spending. If 88 percent of the variance in wealth is transitory, as Table 2 reports, only 12 percent of the total variation in wealth will be associated with a four to five cent on the dollar change in consumption; the remaining 88 percent of the variation in wealth will be associated with no change in consumption. Although one could combine these estimates into a single summary statistic of the average response of consumption to a dollar change in wealth, it may not be desirable to do so. After all, such an exercise must necessarily obfuscate the finding that consumption responds differently to permanent changes in wealth than it does to transitory changes in wealth. Instead, we emphasize that conventional estimates of the wealth effect are likely to be quite misleading, because they are found to apply only to a small fraction of the variation in wealth.

B. *Time-Series Analysis of the Trend in Asset Values*

The empirical procedure employed here can be used to decompose any of the variables in our system into a "trend" and "cyclical" component. A natural definition of trend in each variable is the long-run forecast of the variable, furnished by the trend component from the multivariate Beveridge-Nelson decomposition for the cointegrated system $(c_t, a_t, y_t)'$.

The three panels of Figure 2 plot the resulting trend components of consumption, asset wealth, and labor earnings, along with the actual series for each variable. The plot spans the period beginning in the second quarter of 1952 to the first quarter of 2003. Consumption and labor earnings are visually indistinguishable from their trends. This is not surprising since we already know that they are highly correlated (Table 5). For asset wealth, by contrast, there are many times in the postwar period when asset wealth has diverged substantially from its estimated trend.

A clearer picture of the extent to which this is true is given in Figure 3, which shows the

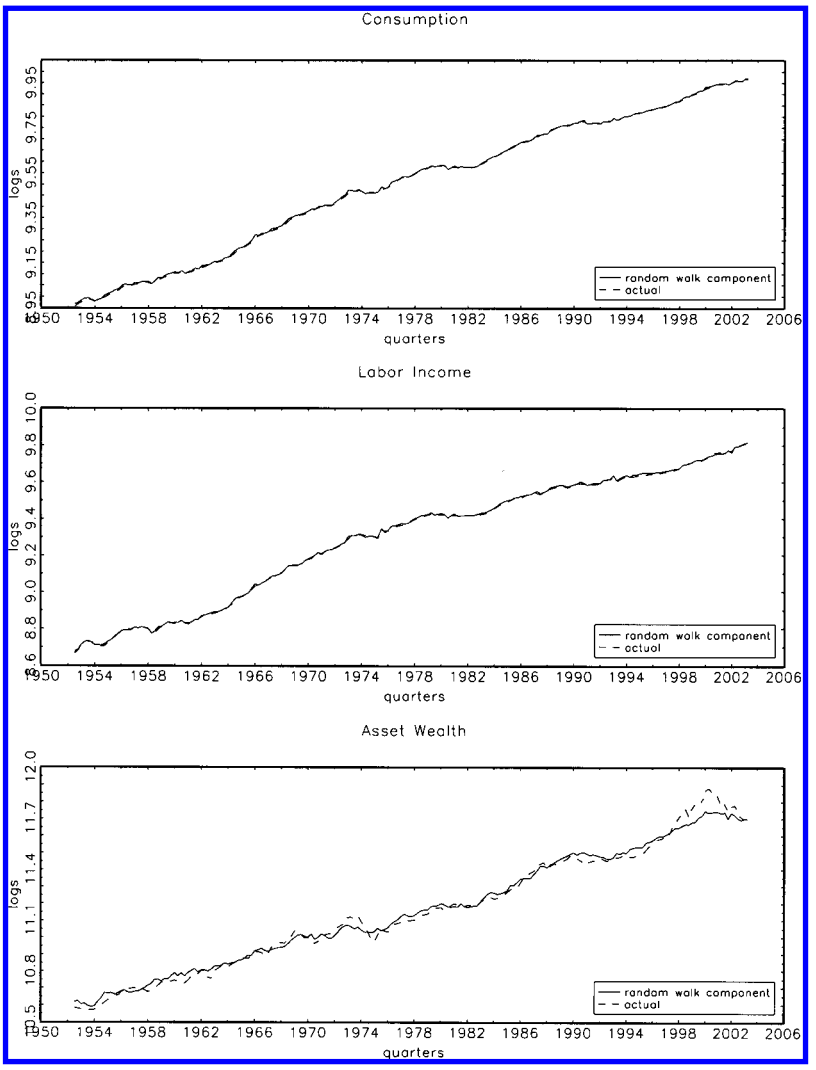

FIGURE 2. ESTIMATES OF THE TRENDS IN *c*, *y* AND *a*

Note: The trend is defined as the long-run forecast of each variable given the multivariate Beveridge-Nelson decomposition for the trivariate system c_t, y_t, and a_t. The variables are measured in log real per capita units.

difference between the trend and actual value of asset wealth (the transitory component of wealth), in percent of the trend component. The series displayed in the figure has been normalized so that when it is above zero, wealth is estimated to be above its long-term trend; when it is below zero, wealth is estimated to be below its long-term trend.

Figure 3 shows that asset values were above their long-term trend during the mid-1950's and late 1960's, and below their long-term trend in many quarters from the mid-1970's through 1997. Consistent with popular impression, the estimate picks out the "bull markets" of the late 1960's and 1990's, and the "bear markets" of the 1970's. Note the sharp decline in wealth in 1973, a year in which stock market wealth relative to GDP fell by a factor of two.

The figure also shows clearly that transitory swings in wealth are both quantitatively large

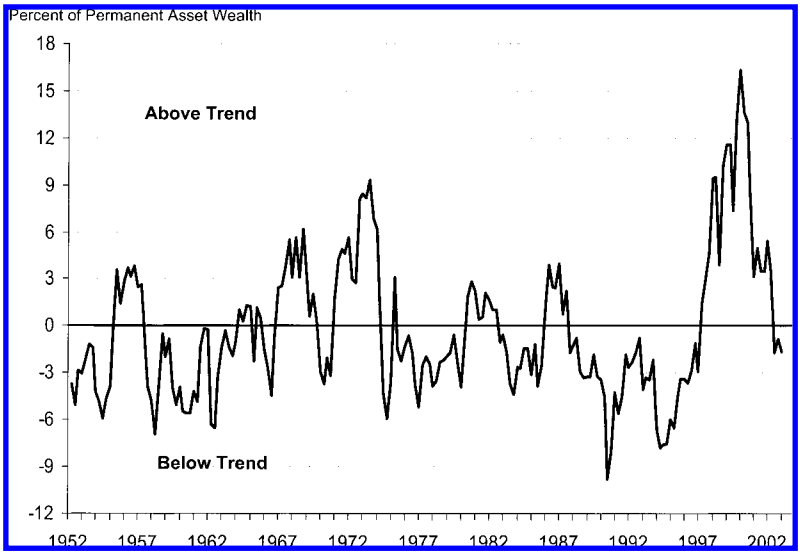

FIGURE 3. THE TRANSITORY COMPONENT OF ASSET WEALTH

Note: This plot shows the difference between the trend and actual value of asset wealth (the transitory component of wealth), in percent of the trend component.

and persistent. In the 1990's, the transitory component reached as high as 16 percent of the permanent component of wealth. Translated into dollar amounts, this implies that wealth exceeded its long-run trend by as much as $16,781 per capita in 1996 dollars.

In general, periods of above-trend asset wealth are typically followed by episodes of negative excess returns on the aggregate stock market, while below-trend asset values are typically followed by episodes of positive excess stock returns. For example, the spikes upward in asset wealth relative to trend in 1956 and 1973, were followed, respectively, by a sequence of negative excess returns in the 1960's, and by the bear markets of the 1970's. The decline in asset wealth relative to trend in 1994 was followed by the bull market of the late 1990's. This estimate suggests that, by the end of 1997, the stock market boom had driven wealth to above-trend levels, anticipating the broad decline in the stock market wealth since 2000.

Of particular interest in Figure 3 is the level of transitory wealth in 1995, a period that marked the beginning of the extraordinary surge in equity values that occurred in the latter half of the 1990's. As late as 1997—well into the bull market of the 1990's and well after the onset of worries shared by many analysts that the market was "irrationally exuberant"—wealth was still below its long-term trend. In contrast to other indicators such as the dividend-yield and price-earnings ratio, this estimate suggests that wealth was not due for a correction until much later, when the stock market had reached the lofty levels it obtained by 1998 and 1999. The figure underscores the importance of using information contained in consumption and labor income, in addition to conventional indicators such as dividends and earnings, for determining the long-run value of asset wealth.

IV. Stability

In this section, we address an important objection to our empirical approach: the possibility of instability in the cointegrating relation. The meaningful question is not whether the econometric model exhibits any instability at all (few models aimed at fitting data for the entire postwar period would share such a property), but whether the instability present is sufficiently large to alter our main econometric conclusions. We address this issue from a number of different angles.

One possible angle is simply to employ formal tests of parameter instability for cointegrated systems. These tests are not without their limitations. The issue of identifying structural breaks in common trends with finite samples is a subtle and tricky one. Long data spans are often required to obtain consistent estimates of cointegrating coefficients, yet instability tests require those parameters to be estimated by splitting an already finite sample into even smaller subsamples. This requirement has led to a well-known criticism of the entire structural break approach, namely that the data-driven specification searches inherent in the methodology can bias inferences dramatically toward finding breaks where none exist (see Edward Leamer, 1978; Andrew Lo and A. Craig MacKinlay, 1990).

With these caveats in mind, we present the results of three tests of parameter stability in cointegrated regression models. The procedures are from Bruce E. Hansen (1992). For all three tests, the null hypothesis is that the cointegrating vector is constant over the sample, although the tests differ in their treatment of the alternative hypothesis. The first test, called the *SupF* test, models the cointegrating vector as obeying a single structural break with unknown break date. The second two tests, called the *MeanF* and the L_c tests, model the cointegrating vector as a martingale process. For the trivariate cointegrated system we investigate, the following values for the three statistics, with *p*-values in parentheses, were obtained:

$$SupF = 9.67 \quad (>0.20)$$

$$MeanF = 5.85 \quad (0.064)$$

$$L_c = 0.477 \quad (0.112).$$

The results present a somewhat mixed picture concerning the presence of instability. No test presents overwhelming evidence of instability, and the *SupF* test in particular does not suggest any instability—this test statistic is not significant at the 20-percent level. On the other hand, the *MeanF* test statistic is significant at the 10-percent level, though not at the 5-percent level; the L_c test statistic is at the 20-percent level but not the 10-percent level. These last two test results may hint at some instability in the cointegrating parameters, but for all tests the data are not sufficiently informative to reject the null of stability at the 5-percent level. A possibility is that some instability is present, but is sufficiently small that a stable cointegrating relation is a reasonable approximation of the data. This possibility is consistent with the finding in Lettau and Ludvigson (2001) that the cointegrating residual is sufficiently stable so as to display statistically significant out-of-sample forecasting power for asset returns.

TABLE 7—COINTEGRATION VECTOR—SUBSAMPLES

Sample	α_a	α_y
1951Q4-2003Q1	0.269	0.621
1951Q4-1995Q4	0.332	0.571
1977Q3-2003Q1	0.077	0.978

Notes: This table presents DLS estimates of the cointegration vector α_y and α_a for various samples. The variables are consumption c_t, labor income y_t, and asset wealth a_t. The sample spans the fourth quarter of 1951 to the first quarter of 2003.

A little digging indicates that the best case for instability in this cointegrated system can be made when data since 1995 are included in the sample. For example, if the sample is ended at 1995:Q4, all of the test statistics reported above have *p*-values in excess of 20 percent. Thus, formal tests suggest no instability whatsoever in the 44 years of data up to 1995. Another way to see the influence of recent data is to estimate the cointegrating vector over subsamples. Table 7 shows that only in subsamples that include data from the late 1990's do the cointegrating coefficients look drastically different from the other subsamples. For example, estimates for the period 1951:Q4-1995:Q4 look almost the same as estimates from 1951:Q4-1977:Q2, but are quite different from 1977:Q3-2003:Q1, which includes post-1995 data and exhibits a much smaller coefficient on wealth and larger coefficient on income. Of course, a caveat with these estimates is that they are not really proper, since it is never appropriate to throw away information when estimating the parameters of a common trend. Nevertheless they give a rough idea of where in the sample instability might lie.

What's going on? The late 1990's were an extraordinary episode in which the price-dividend ratio on aggregate stock market indexes increased

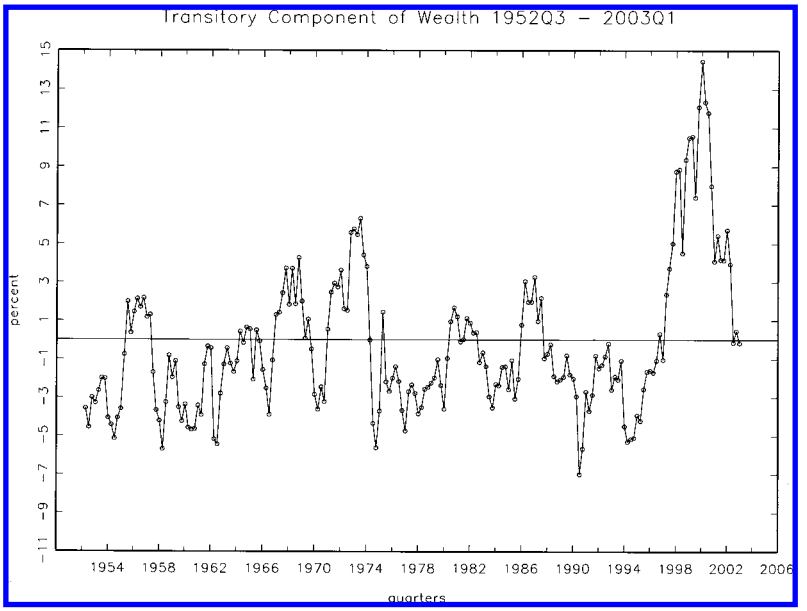

FIGURE 4. COINTEGRATION RESIDUAL

Note: The plot shows the cointegrating residual $c_t - \alpha_y y_t - \alpha_a a_t$ from 1951:Q4 to 2003:Q1. The cointegration vector $[\alpha_y, \alpha_a]$ is estimated using only data up to 1995:Q4.

more than threefold between 1995 and 1999. During this period, wealth remained far above its previously estimated long-run trend with consumption and labor income, and persistently so. Figure 4 demonstrates graphically by plotting the cointegrating residual when the cointegrating vector is estimated on data through 1995:Q4. The residual takes on large and sustained negative values during the late 1990's, as wealth moved well above its long-run trend with consumption and labor income.

The irony of this episode is that, while it weakens some formal statistical evidence in favor of a stable cointegrating vector, it serves to reinforce the main conclusions of this paper: as Figure 4 shows, the large and sustained negative values in the cointegrating residual are entirely eliminated by 2003, and the residual restored to its postwar mean. The restoration of the cointegrating residual to its mean clearly coincides with a large error correction in wealth, a direct result of the broad stock market declines since 2000. By using only data up to 1995 to estimate the cointegrating vector, we can be sure that this restoration is not a result of estimating a new mean and an entirely new cointegrating vector on the unusual data post-1995. (Nevertheless, the figure based on full-sample estimates looks quite similar.)

Including data from the last half of the 1990's creates instability, not because the period was fundamentally different from historical experience, but because, on the contrary, the episode was a more extreme version of the historical record, generating a transitory movement in wealth that was larger and more persistent than previously observed in our sample. Yet the essence of the episode bears out the historical experience to a tee: wealth surged well above its long-run trend with consumption and labor income only to be ultimately restored to that trend because of a subsequent decline in asset values. The important question going forward is whether this episode represents a break toward a fundamentally different long-run relation between consumption, wealth, and income, or whether it was simply an outlier, merely the largest of many such episodes in a familiar pattern of events. It is too early to know for sure, but Figure 4 suggests that the latter possibility may be an important part of the story.

We close this section by noting a number of

additional robustness checks that were performed. The permanent/transitory decomposition is robust to changes in lag length in the VECM, to the measure of consumption used (total personal consumption expenditures rather than nondurables and services), to the timing convention for wealth, and to the procedure used to estimate the VECM (two-stage DLS/OLS versus the Søren Johansen, 1991, one-step maximum likelihood estimator). These results are available upon request.

V. Conclusion

The empirical linkage between wealth and consumption is a classic research problem at the intersection of finance and macroeconomics. We argue here that this linkage cannot be understood without distinguishing trend from cycle in asset values. Transitory variation in asset wealth is both quantitatively large and highly persistent. Indeed, our estimates imply that transitory shocks constitute the vast majority of fluctuations in quarterly net worth and have a half-life of about two years. Yet despite their quantitative importance, transitory shocks in wealth are found to be unrelated to aggregate consumer spending, contemporaneously and at any future date. Aggregate consumption is well described as a function of the trend components in wealth and income and is dominated by permanent shocks. In summary, permanent changes in wealth *do* affect consumer spending, but most changes in wealth are transitory and are uncorrelated with consumption. A contribution of this paper is to document the sheer quantity of variation in asset values that is ultimately unsustained, as well as the extent to which macroeconomic aggregates such as consumption and labor income appear unaffected by this variation.

These findings have at least one important implication for monetary policy. Recent research has suggested that central banks pursuing inflation targets should ignore movements in asset values that do not influence aggregate demand (Bernanke and Mark Gertler, 2001). The results in this paper underscore the relevance of this recommendation, since they suggest that most changes in asset values are transitory and unrelated to consumer spending, the largest component of aggregate demand.

Appendix A: Data Description

CONSUMPTION

Consumption is measured as either total personal consumption expenditure or expenditure on nondurables and services, excluding shoes and clothing. The quarterly data are seasonally adjusted at annual rates, in billions of chain-weighted 1996 dollars. The components are chain-weighted together, and this series is scaled up so that the sample mean matches the sample mean of total personal consumption expenditures. Our source is the U.S. Department of Commerce, Bureau of Economic Analysis.

AFTER-TAX LABOR INCOME

Labor income is defined as wages and salaries + transfer payments + other labor income − personal contributions for social insurance − taxes. Taxes are defined as [wages and salaries/(wages and salaries + proprietors' income with IVA and Ccadj + rental income + personal dividends + personal interest income)] times personal tax and nontax payments, where IVA is inventory valuation and Ccadj is capital consumption adjustments. The quarterly data are in current dollars. Our source is the Bureau of Economic Analysis.

POPULATION

A measure of population is created by dividing real total disposable income by real per capita disposable income. Our source is the Bureau of Economic Analysis.

WEALTH

Total wealth is household net worth in billions of current dollars, measured at the end of the period. Stock market wealth includes direct household holdings, mutual fund holdings, holdings of private and public pension plans, personal trusts, and insurance companies. Nonstock wealth includes tangible/real estate wealth, nonstock financial assets (all deposits, open market paper, U.S. Treasuries and Agency securities, municipal securities, corporate and foreign bonds, and mortgages), and also includes ownership of privately traded companies in noncorporate equity, and other. Subtracted off are liabilities, including mortgage loans and loans made under home equity lines of credit and secured by junior liens, installment consumer debt, and other. Our source is the Board of Governors of the Federal Reserve System. A complete description of these data may be found

TABLE B1—COINTEGRATION TESTS

	L-max test			Trace test			*t*-test		
	$r =$ 0	1	2	$r =$ 0	1	2	H_0: no cointegration	AIC	SIC
10-percent c.v.	18.70	12.10	2.82	26.79	13.34	2.82	−3.52		
5-percent c.v.	20.78	14.04	3.96	29.51	15.12	3.96	−3.80		
1-percent c.v.	25.52	17.94	6.94	35.34	19.31	6.94	−4.36		
1951Q4-2003Q1 sample									
1 lag	**30.33**	5.98	0.95	**37.26**	6.93	0.95	**−4.36**	−19.84	−19.69
2 lags	17.78	4.09	1.05	22.93	5.15	1.05	**−3.86**	−19.92	−19.63
1951Q4-1995Q4 sample									
1 lag	**31.86**	4.80	0.63	**37.30**	5.44	0.63	**−4.77**	−20.13	−19.97
2 lags	**21.98**	3.79	0.94	26.72	4.74	0.94	**−4.44**	−20.18	−19.86

Notes: The first two columns report the L-max and Trace test statistics described in Johansen (1988; 1991). The former tests the null hypothesis that there are r cointegrating relations against the alternative of $r + 1$; the latter tests the null of r cointegrating relations against the alternative of p, where p is the number of variables in the cointegrated system. The last column reports the Phillips-Ouliaris (1990) cointegration test. The critical values (c.v.) for the Phillips-Ouliaris tests allow for trends in the data. The null hypothesis is no cointegration; significant statistics at the 10-percent level are highlighted in bold face. The number of lags in the Johansen tests refers to the VAR specification. AIC is the Akaike information criterion and SIC is the Schwartz criterion. The sample spans the fourth quarter of 1951 to the first quarter of 2003.

at http://www.federalreserve.gov/releases/Z1/Current/.

PRICE DEFLATOR

The nominal after-tax labor income and wealth data are deflated by the personal consumption expenditure chain-type deflator (1996 = 100), seasonally adjusted. In principle, one would like a measure of the price deflator for total flow consumption here. Since this variable is unobservable, we use the total expenditure deflator as a proxy. Our source is the Bureau of Economic Analysis.

APPENDIX B: TESTS FOR COINTEGRATION

This Appendix presents the results of cointegration tests. Tests for the presence of a unit root in c, a, and y (not reported) are consistent with the hypothesis of a unit root in those series and are available upon request.

Table B1 reports test statistics corresponding to two cointegration tests. Reported in the far right column are residual-based cointegration test statistics (Peter C. B. Phillips and S. Ouliaris, 1990). The table shows both the Dickey-Fuller t-statistic and the relevant 5- and 10-percent critical values. The test is carried out without a deterministic trend in the static regression, allowing for trends in the raw data. We applied the data dependent procedure suggested in Campbell and Perron (1991) for choosing the appropriate lag length in an augmented Dickey-Fuller t-test. This procedure suggested that the appropriate lag length was one for the $(c, a, y)'$ system. The tests reject the null of no cointegration at the 5-percent level.

Table B1 also reports the outcome of testing procedures suggested by Johansen (1988, 1991) that allow the researcher to estimate the number of cointegrating relationships. This procedure presumes a p-dimensional vector autoregressive model with k lags, where p corresponds to the number of stochastic variables among which the investigator wishes to test for cointegration. For our application, $p = 3$. The Johansen procedure provides two tests for cointegration: under the null hypothesis, H_0, that there are exactly r cointegrating relations, the "Trace" statistic supplies a likelihood ratio test of H_0 against the alternative, H_A, that there are p cointegrating relations, where p is the total number of variables in the model. A second approach uses the "L-max" statistic to test the null hypothesis of r cointegrating relations against the alternative of $r + 1$ cointegrating relations.

Because these particular test results are more sensitive to the sampling interval, we present results for both the full sample, 1951:Q4-2003:

TABLE B2—CCR COINTEGRATION TESTS

Sample	χ^2	p-value
1951Q4-2003Q1	0.175	0.675
1951Q4-1995Q4	0.869	0.351

Notes: The table reports results from Park's (1992) canonical cointegration regression (CCR) test. The null hypothesis is the presence of cointegration. The test statistic is has a χ^2 distribution with one degree of freedom. The variables are consumption c_t, labor income y_t, and asset wealth a_t. The sample spans the fourth quarter of 1951 to the first quarter of 2003.

Q1, as well as the sample 1951:Q4-1995:Q4, which exclude the unusual late 1990's data. In the latter sample, both the L-max and Trace test results establish evidence of cointegration among log consumption, log labor income, and the log of household wealth. The L-max test implies that we may reject the null of no cointegration against the alternative of one cointegrating vector. In addition, the Trace test shows that we cannot reject the null hypothesis of one cointegrating relationship against the alternative of two or three. In the full sample, the results are more sensitive to the lag length in the underlying levels vectorautoregression. When the lag length is set to unity, both tests establish evidence of a single cointegrating relation, but not when the lag length is set to two.

When the economic model implies cointegration, it has been argued that it is more appropriate to test the null hypothesis of cointegration than the null of no cointegration (Park, 1990, 1992; Ogaki and Park, 1997). Tests of the no cointegration null are known to have low power against some alternatives and can fail to reject the null with high probability even though the variables are actually cointegrated (Ogaki and Park, 1997). As a consequence, a number of researchers have relied on tests of the null of cointegration to establish evidence of cointegration, even in applications where Dickey-Fuller and Johansen tests are unable to reject the null of no cointegration (for example, Hsiang-Ling Han and Ogaki, 1997; Ogaki and Park, 1997; Ogaki and Carmen M. Reinhart, 1998). A test based on the canonical cointegrating regression, developed in the work of Park (1990; 1992), is available to test the null of cointegration. Park's $H(0, 1)$ test has an asymptotic chi-square distribution with one degree of freedom, and is a test of the null hypothesis of deterministic cointegration (i.e., that the cointegrating vector which eliminates the stochastic trends also eliminates the deterministic trends). The results of this test for c, a, and y are reported in Table B2. The $H(0, 1)$ test fails to reject the null of cointegration at a very high significance level.

TABLE C1—IMPULSE RESPONSE FUNCTION

Horizon	c_t	y_t	a_t
1	0.000	0.000	1.835
	(0.000, 0.000)	(0.000, 0.000)	(1.628, 1.963)
4	0.113	0.178	1.437
	(0.060, 0.159)	(0.101, 0.249)	(1.085, 1.747)
8	0.073	0.113	0.876
	(0.0333, 0.117)	(0.054, 0.179)	(0.468, 1.392)
12	0.044	0.069	0.535
	(0.015, 0.091)	(0.026, 0.134)	(0.203, 1.163)
16	0.027	0.042	0.326
	(0.007, 0.071)	(0.011, 0.107)	(0.087, 0.970)
∞	0.000	0.000	0.000
	(0.000, 0.001)	(0.000, 0.002)	(0.000, 0.018)

Notes: The table reports the impulse response function of *c*, *y*, and *a* to the transitory shock. The 95-percent confidence intervals are reported in parentheses below the point estimates. Standard errors are computed using the bootstrap procedure described in Gonzalo and Ng (2001) using 10,000 replications. Horizons *h* are in quarters, and the underlying VECM is of order 1. The sample spans the fourth quarter of 1951 to the first quarter of 2003.

APPENDIX C: STANDARD ERRORS FOR IMPULSE RESPONSE FUNCTIONS AND VARIANCE DECOMPOSITIONS

This Appendix presents 95-percent confidence intervals for the impulse response functions and variance decompositions. The confidence intervals are generated from a bootstrap as described in Gonzalo and Ng (2001). The procedure is as follows. First, the cointegrating vector is estimated, and conditional on this estimate, the remaining parameters of the VECM are estimated. The fitted residuals from this VECM, $\hat{e}_t$, are obtained and a new sample of data is constructed using the initial VECM parameter estimates by random sampling of $\hat{e}_t$ with replacement. Given this new sample of data, all the parameters are reestimated, holding fixed the number of cointegrating vectors, and the impulse responses and variance decompositions stored. This is repeated 5,000 times. The empirical 95-percent confidence intervals are evaluated from these 5,000 samples of the bootstrapped impulse response functions are presented in Table C1. The empirical 95-percent confidence intervals are evaluated from these 5,000 samples of the bootstrapped variance decompositions as presented in Table 2.

REFERENCES

Abel, Andrew B. and Bernanke, Ben S. *Macroeconomics*, 4th Ed. New York: Addison Wesley Longman, Inc., 2001.

Bernanke, Ben S. and Gertler, Mark. "Should Central Banks Respond to Movements in Asset Prices?" *American Economic Review*, May 2001 (*Papers and Proceedings*), *91*(2), pp. 253–57.

Beveridge, Stephen and Nelson, Charles R. "A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the 'Business Cycle'." *Journal of Monetary Economics*, March 1981, *7*(2), pp. 151–74.

Blinder, Alan S. and Deaton, Angus. "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 Alternate Test of the Permanent Income Hypothesis." *Econometrica*, November 1987, *55*(6), pp. 1249–73.

Campbell, John Y. and Cochrane, John H. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy*, April 1999, *107*(2), pp. 205–51.

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 J. Blanchard and Stanley Fischer, eds., *NBER macroeconomics annual: 1989*. Cambridge, MA: MIT Press, 1989, pp. 185–216.

Campbell, John Y. and Perron, Pierre. "Pitfalls and Opportunities: What Macroeconomists Should Know about Units Roots," in Olivier Jean Blanchard and Stanley Fischer, eds., *NBER macroeconomics annual: 1991*. Cambridge, MA: MIT Press, 1991, pp. 141–201.

Cochrane, John H. "Permanent and Transitory Components of GDP and Stock Prices." *Quarterly Journal of Economics*, February 1994, *109*(1), pp. 241–65.

_____. *Asset pricing*. Princeton, NJ: Princeton University Press, 2001.

Constantinides, George M. and Duffie, Darrell. "Asset Pricing with Heterogeneous Consumers." *Journal of Political Economy*, April 1996, *104*(2), pp. 219–40.

Engle, Robert F. and Granger, Clive W. J. "Cointegration and Error Correction: Representation, Estimation, and Testing." *Econometrica*, March 1987, *55*(2), pp. 251–76.

Fernandez-Corugedo, Emilio; Price, Simon and Blake, Andrew. "The Dynamics of Consumer Expenditures: The UK Consumption ECM Redux." Working paper, Bank of England, 2002.

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

Francis, Neville and Ramey, Valerie A. "Is the Technology Driven Real Business Cycle Dead?" Unpublished manuscript, University of California-San Diego, 2001.

Galí, Jordi. "Finite Horizons, Life-Cycle Savings, and Time-Series Evidence on Consumption." *Journal of Monetary Economics*, December 1990, *26*(3), pp. 433–52.

_____. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, March 1999, *89*(1), pp. 249–71.

Gonzalo, Jesus and Granger, Clive. "Estimation of Common Long-Memory Components in Cointegrated Systems." *Journal of Business and Economic Statistics*, January 1995, *13*(1), pp. 27–35.

Gonzalo, Jesus and Ng, Serena. "A Systematic Framework for Analyzing the Dynamic Effects of Permanent and Transitory Shocks." *Journal of Economic Dynamics and Control*, October 2001, *25*(10), pp. 1527–46.

Gordon, Robert J. *Macroeconomics*, 6th Ed. New York: HarperCollins, 1993.

Han, Hsiang-Ling and Ogaki, Masao. "Consumption, Income and Cointegration." *International Review of Economics and Finance*, 1997, *6*(2), pp. 107–17.

Hansen, Bruce E. "Tests for Parameter Instability in Regressions with I(1) Processes." *Journal of Business and Economic Statistics*, July 1992, *10*(3), pp. 321–35.

Johansen, Søren. "Statistical Analysis of Cointegrating Vectors." *Journal of Economic Dynamics and Control*, June/September 1988, *12*(2/3), pp. 231–54.

_____. "Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models." *Econometrica*, November 1991, *56*(6), pp. 1151–80.

King, Robert; Plosser, Charles; Stock, James H. and Watson, Mark W. "Stochastic Trends and Economic Fluctuations." *American Economic Review*, September 1991, *81*(4), pp. 819–40.

Leamer, Edward. *Specification searches*. New York: John Wiley and Sons, 1978.

Lettau, Martin and Ludvigson, Sydney C. "Consumption, Aggregate Wealth and Expected Stock Returns." *Journal of Finance*, June 2001, *56*(3), pp. 815–49.

Lo, Andrew and MacKinlay A. Craig. "Data-Snooping Biases in Tests of Financial Asset Pricing Models." *Review of Financial Studies*, Fall 1990, *3*(3), pp. 431–67.

Ludvigson, Sydney C. and Steindel, Charles. "How Important Is the Stock Market Effect on Consumption?" *FRBNY Economic Policy Review*, July 1999, *5*(2), pp. 29–51.

Modigliani, Franco. "Consumer Spending and Monetary Policy: The Linkages." *Federal Reserve Bank of Boston Conference Series*, June 1971, *5*, pp. 9–84.

Ng, Serena and Perron, Pierre. "Estimation and Inference in Nearly Unbalanced Nearly Cointegrated Systems." *Journal of Econometrics*, July 1997, *79*(1), pp. 53–81.

Ogaki, Masao and Park, Joon Y. "A Cointegration Approach to Estimating Preference Parameters." *Journal of Econometrics*, January 1997, *82*(91), pp. 107–34.

Ogaki, Masao and Reinhart, Carmen M. "Measuring Intertemporal Substitution: The Role of Durable Goods." *Journal of Political Economy*, October 1998, *106*(5), pp. 1078–96.

Park, Joon Y. "Testing for Unit Roots and Cointegration by Variable Addition," in Thomas B. Fomby and George F. Rhodes, eds., *Advances in econometrics*, Vol. 8. London: JAI Press, 1990, pp. 107–33.

_____. "Canonical Cointegration Regressions." *Econometrica*, January 1992, *60*(1), pp. 119–43.

Phillips, Peter C. B. and Ouliaris, S. "Asymptotic Properties of Residual Based Tests for Cointegration." *Econometrica*, January 1990, *58*(1), pp. 165–93.

Poterba, James M. "Stock Market Wealth and Consumption." *Journal of Economic Perspectives*, Spring 2000, *14*(2), pp. 99–118.

Stock, James H. "Asymptotic Properties of Least Square Estimators of Cointegrating Vectors." *Econometrica*, September 1987, *55*(5), pp. 113–44.

Stock, James H. and Watson, Mark W. "Testing for Common Trends." *Journal of the American Statistical Association*, December 1988, *83*(404), pp. 1097–107.

_____. "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems." *Econometrica*, July 1993, *61*(4), pp. 783–820.