DISCUSSION PAPER SERIES

No. 3104

UNDERSTANDING TREND AND CYCLE IN ASSET VALUES: BULLS, BEARS AND THE WEALTH EFFECT ON CONSUMPTION

Martin Lettau and Sydney Ludvigson

FINANCIAL ECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP3104.asp**

ISSN 0265-8003

UNDERSTANDING TREND AND CYCLE IN ASSET VALUES: BULLS, BEARS AND THE WEALTH EFFECT ON CONSUMPTION

Martin Lettau, Federal Reserve Bank of New York and Stern School of Business, New York University and CEPR
Sydney Ludvigson, New York University

Discussion Paper No. 3104
December 2001

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 3104

December 2001

ABSTRACT

Understanding Trend and Cycle in Asset Values: Bulls, Bears and the Wealth Effect on Consumption*

This Paper uses restrictions implied by cointegration to identify the permanent and transitory elements (the 'trend' and 'cycle') of household asset wealth. Our empirical analysis yields answers to the following questions:

1. Is there a large transitory component in household net worth or is wealth close to a random walk? Our point estimates imply that a striking 85% of the post-war variation in the growth of household net worth is transitory, attributable to fluctuations in the stock market component of wealth. Transitory wealth shocks are quite persistent, affecting asset values for a number of years. This transitory element picks out the 'bull' markets of the late 1960s and 1990s, and the 'bear' markets of the 1970s. If markets are efficient, these transitory fluctuations must be attributable to time-variation in the required rate of return on assets (discount rates).

2. How is transitory variation in household net worth related to consumer spending? Does consumption adapt with a lag to permanent movements in wealth? Despite their quantitative importance, transitory fluctuations in asset values are found to be unrelated to aggregate consumer spending. Instead, aggregate consumption can be well described as a function of the trend components in wealth and income. We find no evidence that consumption adapts with a long lag to fluctuations in wealth.

3. What kinds of shocks govern the dynamic behaviour of consumption, asset wealth and labour income? We characterize three: a permanent income shock that affects consumption, asset wealth and labour earnings without distorting their long-run equilibrium relation; an income redistributive shock that shifts the composition of income between labour and capital; and a discount rate shock that generates transitory variation in asset values.

JEL Classification: E21, E44 and G10
Keywords: asset values, cointegration, consumption and wealth effect

Martin Lettau
Stern School of Business
New York University
44 West Fourth Street, Suite 9-190
New York, NY 10012-1126
USA
Tel: (1 212) 998 0378
Fax: (1 212) 995 4233
Email: mlettau@stern.nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=128074

Sydney Ludvigson
Department of Economics
New York University
269 Mercer Street, 7th floor
New York, NY 10003
USA
Tel: (1 212) 998 8927
Fax: (1 212) 995 4186
Email: sydney.ludvigson@nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=147057

*We thank Serena Ng, Simon Potter, 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 Wharton School, the University of California Los Angeles and the University of Chicago for many helpful comments. Nathan Barczi and Adam Kolasinski provided 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. Updated versions of this Paper can be found at http://www.stern.nyu.edu/~mlettau and http://www.econ.nyu.edu/user/ludvigsons/.

Submitted 20 November 2001

1 Introduction

Both textbook economics and common sense teach us that the value of household wealth should be related to consumer spending. Yet movements in asset values often seem disassociated with movements in aggregate consumption. For example, in the third quarter of 1998, as a financial crisis associated with the collapse of Long Term Capital Management ensued, the stock market fell precipitously, but aggregate consumption grew briskly over the quarters during and following the crisis. Why? One possibility is that households perceived this stock market decline to be transitory (a perception that was ultimately confirmed) and, as a consequence, it was ignored in spending decisions based on the long-run value of wealth. Other commentators have suggested that consumption adjusts only with a long lag to movements in wealth, implying that sufficiently transitory swings in asset values will have no impact on spending plans.[1]

Although each of these hypotheses are interesting possibilities, no convincing empirical evidence has yet been presented to support either. In particular, there are several elements of both stories for which there is little empirical basis. Each implies that there is a significant transitory component in household net worth (wealth is not a random walk); moreover, the first suggests that consumption is related only to permanent changes in wealth, while the second suggests that consumers base their spending on a cumulation of past changes in wealth.

The question of how wealth is linked to aggregate consumer spending is one of the most age-old in macroeconomics, but it has taken on new importance recently as asset values have surged, personal saving rates have fallen, and stock markets have become increasingly volatile. At the same time, we argue, these issues have become increasingly difficult to address with modern-day models of consumer behavior as evidence mounts that some of their key underlying assumptions are at odds with a large and growing body of empirical work that investigates the time-series behavior of asset *returns*.[2] Although these models have yielded many important insights, they ignore evidence that excess returns on aggregate stock market indexes are not only volatile but *predictable* over long-horizons.[3] Furthermore, extant empirical studies which attempt to estimate the consumption-wealth link present

[1] For example, see the discussion in Abel and Bernanke (2001), chapter 4.

[2] Several models have become standards in the consumption literature, including the specific permanent income model investigated by Flavin (1981), buffer stock saving models explored in Deaton (1991) and Carroll (1997), and, more recently, models with habit formation as in Carroll, Overland and Weil (2000), Carroll (2001), Dynan (2000), and Fuhrer (2000).

[3] Empirical evidence on predictability in asset returns can be found in, among others, Shiller (1984), Campbell and Shiller (1988), Fama and French (1988), Hodrick (1992), Lamont (1998), and Lettau and Ludvigson (2001a). A comprehensive review of this evidence is provided by Cochrane (2001), chapter 20.

their own short-comings since the econometric techniques commonly employed have difficulty identifying how much of the estimated relation is attributable to the effect of wealth on aggregate consumption, rather than the other way around.

In this paper, we approach these issues from a different perspective. We begin by noting that although direct causality in the short-run consumption-wealth linkage is difficult to pin down, innovations to consumption and wealth that are distinguished by their degree of *persistence* can be empirically identified. Such an identification is possible as long as consumption, c_t, asset wealth (net worth), a_t, and labor income, y_t, share a common trend (they are cointegrated), a property that we show can both be verified empirically using standard econometric tools, and derived theoretically from a generalized household budget constraint in which asset returns are allowed to fluctuate in an arbitrary manner. Since the budget constraint is a feature of any consumption model or consumption-based asset pricing model, the framework is quite general, and particular preference specifications appear as special cases. Following the methodology developed in Stock and Watson (1988), King, Plosser, Stock and Watson (1991), and Gonzalo and Granger (1995), we then use the restrictions implied by cointegration 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. An important advantage of this approach is that restrictive assumptions about the behavior of asset returns are unnecessary, and the findings generated are applicable to a wide variety of theoretical structures.

We argue that this approach is of interest not only for its ability to shed light on the aggregate consumption-wealth linkage, but also for what it reveals about the time-series properties of household net worth. For example, the evidence of predictability in broad stock market returns discussed above suggests that conditional expected returns, or discount rates, vary over time, a phenomenon that we show is likely to generate transitory variation in asset values. But none of this evidence reveals how quantitatively large such transitory variation in wealth is. A contribution of this paper is to document the relative importance of transitory shocks in the variation of household net worth, and show to how they are related to aggregate consumer spending.

The empirical approach taken here yields answers to the following questions:

1. *Is there a large transitory component in household net worth or is wealth close to a random walk?* We find that a striking 85 percent of the post-war variation in the growth of household net worth is transitory, generated by fluctuations in the stock market component of wealth. This transitory component is quite persistent, and picks out long-term "bull markets" in the late 1960s and 1990s, and a long-term "bear market" in the 1970s, which includes the sharp decline in stock market wealth that occurred in 1973. We emphasize that

if markets are efficient, this transitory variation must be attributable to time-variation in the required rate of return on assets, or discount rates. Consistent with this hypothesis, we find that the transitory component of wealth is strongly associated with movements in the stock market, but not with transitory movements in earnings or dividends.

2. *How is transitory variation in household net worth related to consumer spending? Does consumption adapt with a lag to permanent movements in wealth?* Although transitory shocks dominate post-war variation in wealth, we find that they are unrelated to aggregate consumer spending at any future horizon. Instead, consumption can be well characterized as a linear combination of the trend components in wealth and labor income. These findings do not mean that wealth has no influence on consumption, but that only permanent changes in wealth affect consumer spending. Transitory fluctuations in wealth also have no impact on labor income, implying that these movements in wealth are not generated by saving out of current income.

We also find no evidence in quarterly data that consumption adapts sluggishly to permanent changes in wealth. We demonstrate that such sluggishness would imply that aggregate consumption contain a significant transitory component (correlated with the permanent components of wealth or labor income), and that deviations from the common trend in consumption, wealth and labor income should forecast consumption growth at some future horizon. Instead, we find that consumption contains virtually no transitory component and that deviations from the common trend appear independent of consumption growth at all future horizons. Thus, of the two hypotheses posed at the beginning of this paper, this evidence is consistent with the former, but not the latter.[4]

3. *What kind of shocks govern the dynamic behavior of consumption, asset wealth and labor income?* We characterize three: a "permanent income" shock that affects consumption, asset wealth and labor earnings without distorting their long-run equilibrium relation; an "income redistributive" shock that persistently shifts the composition of income between labor and capital; and a "discount rate" shock that creates transitory variation in asset values but is not associated with future movements in dividend or earnings growth.

One specific model which captures many of these features of the data is the nonlinear habit

[4]Sluggish adjustment of consumption to permanent shocks is implied by several recently popularized habit-formation models cited in footnote 2. We emphasize that our evidence does not rule out the presence of habit formation *per se*. The implication that habit formation generate sluggish adjustment is quite special to the particular models studied in the consumption literature. If, in contrast to those models, expected returns on some assets are allowed to vary over time, and if the accumulation equation for the habit stock is allowed to be nonlinear, habit formation can imply that consumption adjusts immediately and fully to permanent movements in wealth or income (for example, Campbell and Cochrane (1999)). Our results are broadly consistent with these more general models of habit formation.

framework explored in Campbell and Cochrane (1999). Campbell and Cochrane study an equilibrium economy with efficient markets in which a large transitory component of wealth is generated from time-variation in the equity risk premium, a feature that is in turn driven by time-variation in risk aversion. The transitory component in asset wealth is unrelated to consumer spending, which is based only on permanent shocks to endowment income. Note that this explanation does not imply that households can distinguish permanent from transitory movements in exogenously given asset values. Transitory fluctuations in wealth are not exogenous, but are instead endogenously driven by fluctuations in agents' risk aversion.

The so-called "wealth effect" on consumer spending is a classic research problem at the intersection of finance and macroeconomics. It was explored in early work by Modigliani (1971) who estimated that a dollar increase in wealth holding fixed labor income (the marginal propensity to consume out of wealth) leads to an increase in consumer spending of about five cents, a figure routinely cited in leading macroeconomic text books. The magnitude of this estimate is not trivial quantitatively, and may explain why it is commonly presumed that sharp swings in asset values are likely to generate important movements in aggregate consumption. We show here that these estimates apply only to permanent movements in wealth and, as a consequence, provide a valid summary measure of the wealth effect in an efficient markets setting only if discount rates are constant. Because we find that most changes in wealth are transitory and have no impact on consumption, common estimates of the marginal propensity to consume tend to overstate the wealth effect on consumption. Thus, the results imply that any summary measure of the wealth effect on consumption must take into account the relative importance of permanent and transitory shocks to wealth, and their differential impact on consumer spending. We do so here using the results presented below and find that the marginal propensity to consume out of a "average"movement in wealth is likely to be a far smaller 1.4 to two cents per dollar of wealth gained.

Several other papers explore issues related to those considered here. Cochrane (1994) studies two bivariate, cointegrated systems, one for consumption and GNP, and one for stock prices and dividends. He uses cointegration techniques similar to those employed in this paper to characterize the permanent and transitory components in GNP and stock prices. Ludvigson and 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 (2001a) show that the cointegrating residual for the trivariate system we study embodies rational forecasts of either asset returns, consumption growth or both, and they find strong empirical evidence that this residual is informative about the future path returns on aggregate stock market indexes in excess of a Treasury Bill rate. However, Lettau and Ludvigson (2001a) did not formally identify the permanent and transitory elements of asset wealth, document their relative quantitative importance, or develop the consumption implications of those findings, the focus of this paper. Finally, the early literature on stock market volatility argued that stock returns were too volatile to be accounted for by variation in future dividend growth (LeRoy and Porter (1981); Shiller (1981)). Analogously, thinking of consumption as the dividend payed from aggregate (human plus nonhuman) wealth, the findings we present here imply that aggregate wealth is too volatile to be accounted for by variation in future consumption growth.

The remainder of this paper is organized as follows. In the next section, we motivate our use of the trivariate cointegrating system for consumption, asset wealth and labor income by deriving cointegration from a linearized budget constraint in which asset returns are allowed to vary in an arbitrary manner over time. We consider several examples of specific preference specifications which appear as special cases. Section 3 presents the econometric framework we use to distinguish the permanent from transitory components of consumption, wealth and income. We demonstrate how the assumption of efficient markets, and of particular models of consumer behavior, are related to the permanent-transitory decomposition, and we illustrate the interplay between preferences and time-varying discount rates in determining the wealth effect on consumption. Section 4 presents our main empirical results. Section 5 concludes.

2 The Common Trend in Consumption, Wealth and Labor Income

This section motivates our analysis of the specific cointegrating system we investigate. We build off of work in Campbell and Mankiw (1989) and Lettau and Ludvigson (2001a).

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

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

Labor income Y_t does not appear explicitly in this equation because of the assumption that the market value of tradable human capital is included in aggregate wealth.[5] Throughout this paper we use lower case letters to denote log variables, e.g., $c_t \equiv \ln(C_t)$.

[5]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, $H_t = E_t \sum_{j=0}^{\infty} \prod_{i=0}^{j} (1 + R_{a,t+i})^{-i} Y_{t+j}$.

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, and imposing a transversality condition.[6] The resulting expression is: [7]

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

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 (2001a) 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 y_t. 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, $H_t = E_t \sum_{j=0}^{\infty} \prod_{i=0}^{j} (1 + R_{h,t+i})^{-i} Y_{t+j}$. A log-linear approximation of H_t yields $h_t = \kappa + y_t + v_t$, where κ is a constant and v_t is a mean-zero, stationary random variable given by $v_t = E_t \sum_{j=0}^{\infty} \rho_h^j (\Delta y_{t+1+j} - r_{h,t+1+j})$ 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 equation for the log consumption-aggregate wealth ratio using only observable variables on the left hand side:

$$c_t - \alpha_a a_t - \alpha_y y_t \approx E_t \sum_{i=1}^{\infty} \rho_w^i \Big((1-\nu) r_{at+i} - \Delta c_{t+i} + \nu \Delta y_{t+1+i} \Big). \tag{3}$$

Several points about equation (3) deserve emphasis. First, if labor income follows a random walk and the returns to human capital are constant, $c_t - \alpha_a a_t - \alpha_y y_t$ is proportional to the log consumption-wealth ratio, $c_t - w_t$. Thus we often 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, a simple budget constraint relation implies that the logs of consumption, asset wealth and labor earnings are cointegrated and $c_t - \alpha_a a_t - \alpha_y y_t$ is the cointegrating residual. The cointegrating parameters α_a and α_y should in principle equal the shares $(1-\nu)$ and ν, respectively. As Lettau and Ludvigson (2001a) discuss, however, in practice, these numbers may sum to a number less than one because only a fraction of

[6]This transversality condition rules out rational bubbles.

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

total consumption based on nondurables and services expenditure is observable.[8] Third, note that from the definition of returns to aggregate wealth, (1), $r_{at} \approx \Delta a_t$. This implies that if the cointegrating residual on the left-hand-side of 3 is not constant, it must forecast either changes in asset wealth (returns), changes in labor income, changes in consumption growth, or some combination of the three. Lettau and Ludvigson (2001a) find that it is a strong predictor of excess returns on aggregate stock market indexes, but not consumption growth. In this paper we document that it is also not a predictor of labor income growth.

By combining the approximate budget constraint relation (3) with preferences, we may obtain approximate consumption *functions*. For example, equation (3) provides a loglinear generalization of Campbell's (1987) "rainy-day" equation for the specific formulation of the permanent income hypothesis (hereafter referred to as the PIH) explored in Flavin (1981), Campbell and Deaton (1989), and Galí (1990).[9] In that model, both expected asset returns and expected consumption growth are constant so the left-hand-side of (3) summarizes expectations of future labor income changes (Campbell (1987)).

Another popular preference specification is the time-separable, isoelastic utility specification, written $U_t = \frac{C_t^{1-1/\sigma}}{1-1/\sigma}$. This utility function, and its associated first order condition for optimal consumption choice, yields an expression for expected consumption growth equal to $E_t \Delta c_{t+1} = \mu_{mt} + \sigma E_t r_{w,t+1}$, where μ_{mt} is an intercept term related to the second moments of consumption and the return to aggregate wealth. Campbell (1996) uses this preference specification along with the assumptions that μ_{mt} is a constant and the conditional expected return on asset wealth, r_a, equals the conditional expected return on human wealth, r_h. Taken together, they imply the following cointegrated log-linear model for consumption,

[8] The use of these expenditure categories is justified on the grounds that consumer theory applies to the *flow* of consumption; expenditures on durable goods are not part of this flow since they represent replacements and additions to a stock, rather than a service flow from the existing stock which we do not observe. Thus, nondurables and services expenditure is only a part of total consumption which is unobservable. One way to address this unobservability is to assume that the level of nondurables and services expenditure is a constant fraction of total consumption. However, Blinder and Deaton (1985) report that the share of nondurables and services spending in total expenditure has displayed a secular decline over the post-war period. By contrast, the *log* of total expenditures is a much more stable multiple of the *log* of nondurables and services expenditures. Consequently, we assume that the log of total consumption, c_t, is a constant scale factor, $\lambda \geqslant 1$, times nondurables and services expenditure, c_{nt}, implying that the estimated cointegrating vector for $c_{n,t}$, a_t, and y_t will be given by $[1, -\frac{1}{\lambda}(1-\nu), -\frac{1}{\lambda}v]$, where $1-\nu$ is the average value of A/W. Thus the estimated cointegrating parameters will be given by $\alpha_a = \frac{1}{\lambda}\omega$, and $\alpha_y = \frac{1}{\lambda}(1-\omega)$. Note that $\alpha_a + \alpha_y \leq 1$ identifies $1/\lambda$.

[9] The specific formulation of this model 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.

asset wealth and labor income:

$$c_t - \alpha_a a_t - \alpha_y y_t \approx E_t \sum_{i=1}^{\infty} \rho_w^i \Big((1 - \sigma - \nu) r_{at+i} + \nu \Delta y_{t+i} \Big), \tag{4}$$

where σ is the intertemporal elasticity of substitution in consumption (IES).

Equation (4) says that an increase in expected returns raises or lowers the cointegrating error $c_t - \alpha_a a_t - \alpha_y y_t$ depending on whether $\sigma + \nu$ is less than one. If the IES is small (i.e., consumption is close to a random walk), expected consumption growth is close to constant and the cointegrating residual will be positively related to expected future returns (income effects dominate substitution effects), consistent with empirical evidence in Lettau and Ludvigson (2001a).[10] Recall that, in the PIH, both expected returns and expected consumption growth are constant, so (4) provides a loglinear generalization of the PIH to allow for time-variation in expected returns. Thus, it is clear from both (3) and (4) that Campbell's rainy day equation is quite special; only under the assumption (made in the PIH and in most other consumption models) that expected asset returns are constant, will the cointegrating error in these models depend solely on expected income growth. In more general settings, expected returns to future asset wealth will also be part of the cointegrating error.

A third preference specification that has been popularized more recently assumes that utility is formed, not over consumption itself, but over consumption relative to some reference level or *habit* stock.[11] The habit specifications explored in the consumption literature assume that the habit stock is a linear function of past consumption, and that the rate of return on all assets is constant. As an example, consider the habit model studied in Dynan (2000), in which the utility function takes the form, $U_t = \frac{(C_t - \phi C_{t-1})^{1-1/\sigma}}{1-1/\sigma}$.[12] The first-order condition for optimal consumption choice implies that $E_t \Delta c_{t+1} = \gamma_0 + \phi \Delta c_t$, where γ_0 and ϕ are constants and $\phi \geq 0$. Combining this expression with (3), we may derive the following cointegrated model

$$c_t - \alpha_a a_t - \alpha_y y_t = E_t \sum_{i=1}^{\infty} \rho_w^i \Big(\gamma_0 + \nu \Delta y_{t+i} - \phi \Delta c_{t+i} \Big). \tag{5}$$

An important implication of this formulation is that consumption adjusts sluggishly and predictably to permanent changes in wealth and income. In this model, a *permanent*

[10] The IES relates consumption growth to the conditional expectation of next period's interest rate. Many researchers have found little connection between the rate of aggregate consumption growth and *ex-ante* real interest rates, suggesting that the IES is close to zero; see Campbell and Mankiw (1989), Hall (1988), Attanasio and Weber (1993), Ludvigson (1999), and the international evidence in Campbell (1999).

[11] See, for example, Muellbauer (1988), Carroll et al. (2000), Dynan (2000), Fuhrer (2000), and Carroll (2001).

[12] This function is a simplified version of the preference specification considered by Dynan (2000) which also allowed for "taste-shifters."

shock to a or y will not be immediately accompanied by a full adjustment in c. Instead, consumption adjusts slowly over time, so that the deviation from the common trend is, at least in part, ultimately eliminated by subsequent movement in consumption. Note that this is precisely the prediction given by the appearance of expected consumption growth on the right-hand-side of (5). Accordingly, deviations from the common trend in c_t, a_t, and y_t (the cointegrating residual) should forecast consumption growth over some future horizon, in addition to possibly forecasting labor income growth.

Not all habit models imply that consumption growth should be forecastable by the cointegrating residual. A notable exception is the model explored by Campbell and Cochrane (1999), which is specified so that expected consumption growth is constant, $E_t \Delta c_{t+1} = g$, (c_t is a random walk), a specification they argue is roughly consistent with post-war U.S. data. In their model, it is not consumption growth that is predictable by the cointegrating error, but asset returns. Campbell and Cochrane point out that constant expected consumption growth can be obtained in habit models by specifying the habit as a nonlinear function of past consumption. We refer to the habit model in (5) as the *linear habit model*, to distinguish it from the nonlinear Campbell-Cochrane model. Although the Campbell-Cochrane model does not have labor income, it is straightforward to modify it to allow explicitly for labor income. In this case, it too can be expressed as a special case of (3):

$$c_t - \alpha_a a_t - \alpha_y y_t \approx E_t \sum_{i=1}^{\infty} \rho_w^i \Big((1-\nu) r_{at+i} + \nu \Delta y_{t+i} \Big). \tag{6}$$

Both expected consumption growth and the risk-free rate are constant in the Campbell Cochrane model because precautionary savings terms completely offset the intertemporal substitution terms in the loglinear Euler equation. Thus, the Campbell-Cochrane model is another example of a general permanent income framework that allows expected returns, or discount rates, to be time-varying.[13]

The point of these examples is to illustrate the role of preferences in determining the functional form of the cointegrating residual, $c_t - \alpha_a a_t - \alpha_y y_t$. If preferences are such that consumption adapts sluggishly to permanent changes in wealth and income and discount

[13] We use the term "permanent income framework" loosely to refer to models in which the log of consumption is approximately a random walk. The Campbell-Cochrane model is presented as an asset pricing model in which consumption is specified exogenously and the behavior of asset returns is derived given consumption. They point out that their model can instead be closed as a production economy with a linear technology and their framework reinterpreted as a theory of consumption. This follows because the behavior of returns is derived from the consumer's first-order-condition for optimal consumption choice. Thus, if asset returns are specified exogenously to match the behavior implied by the Campbell-Cochrane model, a random walk for log consumption is its testable implication, given returns.

rates are constant (as in (5)), movements in this residual should be related to future movements in consumption. If instead, consumption is close to a random walk but discount rates vary (as in (6)), then the residual should be related to future movements in asset returns. In addition, each of these examples demonstrate that it is straightforward to generalize conventional models of consumer behavior to allow for arbitrary variation in expected returns. Of course, most of these examples make no attempt to model the determination of the equilibrium return. (The Campbell-Cochrane model is an exception because the behavior of returns is derived given an empirically plausible process for consumption.) Instead, asset returns may be viewed as equilibrium outcomes, and the examples illustrate the implications of such outcomes for consumption behavior.

3 Econometric Framework

This section describes our approach to isolating the permanent and transitory shocks of a cointegrated vector, $\mathbf{x_t}$, that has n elements. In our application, $n = 3$, and $\mathbf{x_t} = (c_t, a_t, y_t)'$. 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.

3.1 Data and Preliminary Analysis

Appendix A contains a detailed description of the data used in this study. The log of aggregate consumption, c, is measured as real, per capita, nondurables and services expenditure, excluding shoes and clothing. Theories of consumer expenditure apply to the *flow* of consumption; durables expenditures are excluded in this definition because they represent replacements and additions to a capital stock, rather than a service flow from the existing stock. The log of asset wealth, a, is total household net worth measured at the beginning of the period, in real, per capita terms.[14] The log of labor income, y, is also in real, per

[14] A timing convention is needed because the level of consumption is a flow during the quarter rather than a point-in-time estimate (consumption data are time-averaged). To investigate the structural relation between consumption and wealth (as opposed to questions of pure forecasting) wealth must be in the information set of households at the time consumption decisions are made. If we think of consumption for a given quarter as measuring spending at the beginning of the quarter, then the appropriate measure of wealth is beginning-of-period wealth. If we think of consumption for a given quarter as measuring spending at the end of the quarter, then the appropriate measure of wealth is end-of-period wealth. We argue that the former is more reasonable since in this scenario households can "stock their refrigerator" at the beginning of the period and consume over the period by running down that stock during the period. The end-of-period only allows consumption to occur in one instant on the last day of the period after the markets close. Nevertheless, we note that none of our conclusions hinge on this timing convention: as a robustness check, we performed our

capita terms. Our data are quarterly and span the first quarter of 1953 to the first quarter of 2001. Table 1 reports summary statistics of the data. We note one important statistic: the correlation between the quarterly growth in household net worth, Δa_t, and the return on the Center for Research on Security Prices (CRSP) value weighted stock market index is 0.87 in our sample. Thus, it is clear that quarterly fluctuations in household net worth are dominated by movements in stock market returns.

Our empirical approach is to use the restrictions implied by cointegration to 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 (Gonzalo and Granger (1995); King et al. (1991)). The procedure has several steps, the first of which 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 strongly suggest the presence of a single cointegrating vector for these three variables; we impose this in our VECM specification from now on.[15] 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 super-consistent estimates of α_a and α_y (Stock and Watson (1993)).[16] We estimate $\widehat{\boldsymbol{\alpha}} = (1, -0.30, -0.60)'$. The Newey-West corrected t-statistics

empirical tests under both timing assumptions and find that the conclusions we present here are not altered by whether wealth is measured at the beginning or end of the period.

[15]It is perhaps somewhat surprising that we find no evidence of bivariate cointegration between asset wealth and labor income, as would be implied in simple general equilibrium models with a Cobb-Douglas production function. Of course, finding no evidence of cointegration is not the same as finding evidence against cointegration; therefore the result could merely reflect the test's failure to distinguish a stationary process from a nonstationary one in a finite sample. The data do suggest that deviations from any common trend for a_t and y_t, should one exist, must be very persistent. For this reason, allowing the statistical tests to determine the number of cointegrating relations is likely to be superior to imposing a theoretical cointegrating relation for which there is no empirical evidence. Doing so would essentially *force* shocks that cause deviations between a_t and y_t to be eliminated, despite evidence to the contrary in our sample. As Campbell and Perron (1991) emphasize, this is often a reasonable approach, since unit root/cointegration tests may provide a primer for judging which asymptotic distribution gives a better approximation to the true finite sample distribution, even if the test results are not what would be obtained in an infinite sample.

[16]We use eight leads and lags of the first differences of Δy_t and Δa_t in the dynamic least squares regression.

for these estimates are 6 and 14, respectively. As discussed in Lettau and Ludvigson (2001a), these coefficients are not expected to sum to one because total consumption is unobservable and nondurables and services expenditure (a subset of the total) is used as a proxy.

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

$$\Delta\mathbf{x}_t = \boldsymbol{\upsilon} + \boldsymbol{\gamma}\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1} + \boldsymbol{\Gamma}(L)\boldsymbol{\Delta}\mathbf{x}_{t-1} + \mathbf{e}_t, \tag{7}$$

where $\Delta\mathbf{x}_t$ is the vector of log first differences, $(\Delta c_t, \Delta a_t, \Delta y_t)'$, $\boldsymbol{\upsilon}$, and $\boldsymbol{\gamma} \equiv (\gamma_c, \gamma_a, \gamma_y)'$ are $(3{\times}1)$ vectors, $\boldsymbol{\Gamma}(L)$ is a finite order distributed lag operator, and $\widehat{\boldsymbol{\alpha}} \equiv (1, -\widehat{\alpha}_a, -\widehat{\alpha}_y)'$ is the $(3{\times}1)$ vector of previously estimated cointegrating coefficients.[17] Throughout this paper, we use "hats" to denote the estimated values of parameters.

The term $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$ gives last period's equilibrium error, or cointegrating residual; $\boldsymbol{\gamma}$ is the vector of "adjustment" coefficients that tells us which variables react to last periods cointegrating error; that is which variables 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 (7). 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 (7).

The results of estimating a second-order specification of (7) are presented in Table 2.[18] The regressions verify that quarterly consumption growth is slightly forecastable, but lags of asset and labor income growth have little predictive capacity. The small amount of predictability that is present is attributable to the modest first-order serial correlation present in quarterly spending growth. The estimates of the adjustment parameters in $\boldsymbol{\gamma}$ are given in the last row of Table 2; the 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 asset wealth: γ_a is estimated to be about 0.4, with a t-statistic equal to 3.3. These results imply that when log consumption deviates from its habitual ratio with log labor income and log assets, it is asset wealth, not consumption

Monte Carlo simulation evidence in both Ng and Perron (1997) and our own, calibrated to match the data generating process of the variables studied in this paper, suggested that, in samples of the size we encounter here, the DLS procedure can be made more precise with larger lag lengths. In our Monte Carlo simulation, we obtain extremely precise estimates of the cointegrating parameters and a lag length of eight was sufficient to remove all of the bias in the DLS estimator. By contrast, a lag length of two left some bias.

[17]Standard errors do not need to be adjusted to account for the use of the generated regressor, $\boldsymbol{\alpha}'\mathbf{x}_t$ in (7) because estimates of the cointegrating parameters converge to their true values at rate T, rather than at the usual rate $\sqrt{T}$ (Stock (1987)) .

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

or labor income, that is forecast to adjust until the equilibrating relationship is restored.[19] The next section shows that the parameters in $\boldsymbol{\gamma}$ play a key role in identifying permanent and transitory components of the system $\mathbf{x}_t$.

3.2 Permanent and Transitory Identification

We use cointegration to decompose of $\mathbf{x}_t$ 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 (Stock and Watson (1988)). We discuss one interpretation of these shocks below. Following King et al. (1991), 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. Second, the transitory innovation is assumed to be orthogonal to the two permanent innovations. This assumption implies that identification of the space spanned by the two permanent shocks follows only from restrictions implied by cointegration, and is not dependent on any particular ordering of the variables in $\mathbf{x}_t$.

To understand the results below, it is useful briefly review this methodology and explain how it is related to our application. This methodology requires that the number of permanent and transitory shocks equal the number of variables in the system. Thus, we seek to create three transformed, or "structural-form," innovations that are distinguished by whether they have permanent or transitory effects. We denote these transformed innovations $\boldsymbol{\eta}_t \equiv (\eta_{1t}, \eta_{2t}, \eta_{3t})'$ where two are permanent and one is transitory. In what follows, without loss of generality, these shocks are ordered so that the first two of them, η_{1t}, and η_{2t}, have permanent effects; the third, η_{3t}, has transitory effects. Following Gonzalo and Granger (1995), we 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}, is transitory if $\lim_{h\to\infty} \partial E_t(\mathbf{x}_{t+h})/\partial\eta_{Tt} = 0$.

Applying the methodology in Gonzalo and Granger (1995), the permanent and transitory innovations may be identified using the estimated parameters $\widehat{\boldsymbol{\gamma}}$ and $\widehat{\boldsymbol{\alpha}}$ from the VECM in the following way: for a vector, $\mathbf{x}_t$, of $I(1)$ processes that have r cointegrating relations, write the reduced-form Wold representation of this system, $\Delta\mathbf{x}_t = \boldsymbol{\delta} + \mathbf{C}(L)\mathbf{e}_t$, where $\mathbf{C}(L)$ is a distributed lag operator and the VECM parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$, both of rank r, satisfy

[19] We also find that the *four*-quarter lagged value of the cointegrating error strongly predicts asset growth, implying that the forecasting power of the cointegrating error for asset growth is not attributable to interpolation procedures which compute the quarterly housing service flow component of services expenditure from annual surveys.

$\boldsymbol{\alpha}'\mathbf{C}(1) = 0$ and $\mathbf{C}(1)\boldsymbol{\gamma} = 0$. Let

$$\mathbf{G} = \begin{bmatrix} \boldsymbol{\gamma}'_{\perp} \\ \boldsymbol{\alpha}' \end{bmatrix}, \tag{8}$$

where $\boldsymbol{\gamma}'_{\perp}\boldsymbol{\gamma} = 0$. Define a new distributed lag operator $\mathbf{D}(L)$ equal to $\mathbf{C}(L)\mathbf{G}^{-1}$. Then the transformed residuals in $\boldsymbol{\eta}_t$ are equal to $\mathbf{G}\mathbf{e}_t$, and their relation to $\mathbf{x}_t$ is given by the Wold representation

$$\Delta\mathbf{x}_t = \boldsymbol{\delta} + \mathbf{D}(L)\boldsymbol{\eta}_t, \tag{9}$$

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

What's the intuition for this decomposition? Equation (8) says that, if the γ_j (the jth element of the adjustment vector $\boldsymbol{\gamma}$ in (7)) is small in absolute value (i.e., the jth variable does not participate in the error correction), the element of $\boldsymbol{\gamma}'_{\perp}$ that multiplies e_{jt} in $\mathbf{G}\mathbf{e}_t$ will be large in absolute value, giving the x_j a large weight in the permanent innovations. By contrast, if γ_j is large, the element of $\boldsymbol{\gamma}'_{\perp}$ that multiplies e_{jt} in $\mathbf{G}\mathbf{e}_t$ will be small in absolute value, giving the x_j a small weight in the permanent innovations and a large weight in the transitory innovations. Thus, variables have a large transitory component when they do much of the adjusting needed to restore a transitory "gap," or cointegrating error, subsequent to an equilibrium distorting shock. This makes sense because variables which participate in such adjustment must, by definition, deviate from trend, and hence contain a transitory component. In our system, the elements of the adjustment vector $\boldsymbol{\gamma}$ corresponding to c_t and y_t are close to zero (Table 2), implying that these variables will 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 (Table 2), 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 component, where k is equal to the number of common trends.[20] The random walk component of the $I(1)$ common factors from the Gonzalo-Granger procedure is the trend concept provided by the multivariate Beveridge–Nelson (Beveridge and Nelson (1981)) decomposition investigated in Stock and Watson (1988). We now show how this permanent-transitory decomposition is related to efficient markets theory and is affected by particular preference specifications. In addition

[20] The k common factors in the Granger-Gonzalo decomposition 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.

we show how the presence of time variation in discount rates affects common measures of the wealth effect on consumption.

3.2.1 How Do Efficient Markets relate to the Permanent-Transitory Decomposition?

Efficient markets theory implies that asset values are equal to the present discounted value of expected future dividends from nonhuman wealth. Campbell and Shiller (1988) show that a loglinear version of this theory can be expressed approximately in terms of the log dividend-asset value ratio:

$$a_t - d_t \approx \kappa + E_t \sum_{j=1}^{\infty} \rho_a^j \Delta d_{t+j} - E_t \sum_{j=1}^{\infty} \rho_a^j r_{a,t+j}, \tag{10}$$

where d_t is log dividends, r_a is the log return to a_t, and $\rho_a \equiv 1/(1+\exp(\overline{d-a}))$ and κ are constants with the former less than one.[21] If a_t and d_t follow loglinear unit root processes, (10) implies that a_t and d_t are cointegrated with cointegrating vector $(1,-1)'$. Note that $\Delta a_t \approx r_{at}$.[22]

A special case of efficient markets theory holds when expected asset returns are constant. Equation (10) shows that, in this case, the cointegrating error $a_t - d_t$ is a function only of expected future dividend growth. It follows that all of the error-correction subsequent to a equilibrium-distorting shock is done by dividends and the adjustment parameter for asset values is exactly zero. Thus, a_t has no transitory component. Instead, the transitory component in the $(a_t, d_t)'$ system will be entirely associated with dividends, and a_t will define the stochastic trend in d_t. On the other hand, if expected returns vary, some of the error-correction subsequent to a equilibrium-distorting shock may be done by a_t. Recalling the discussion above, this implies that wealth must have a transitory component. (A temporarily high value of $a_t - d_t$ may foretell lower future returns, and therefore lower a_t.) It follows that, if markets are efficient, asset values, a_t, can have a transitory component if and only if expected returns vary.

[21] The formulation of (10) is a renormalization of the expression in Campbell and Shiller (1988). Campbell and Shiller measure returns on a per-share basis, in which case the price per share would appear in place of total asset value. In (10), the return is instead measured on a total value basis, so that the total asset value appears on the left-hand-side.

[22] See Campbell, Lo and MacKinlay (1997), chapter 7 for an explanation of this approximation. For our normalization, this approximation assumes that the shares of assets held remain relatively constant from quarter to quarter. This approximation works well in practice: the stock market component of $\Delta a_t \equiv \Delta s_t$ has a correlation coefficient of 0.95 with the return on the universe of stocks given by the CRSP value weighted composite index.

3.2.2 How Do Preferences Relate to the Permanent-Transitory Decomposition?

Equations (4)-(6) show how the cointegrating residual $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t = c_t - \widehat{\alpha}_a a_t - \widehat{\alpha}_y y_t$ is related to future expected values of Δc_t, Δy_t, and Δa_t (via returns) in various consumption models. It is straightforward to see how each of the examples above will affect the permanent-transitory decomposition. For example, in the loglinear PIH, expected returns and expected consumption growth are presumed constant, thus the cointegrating error in (3) is a function only of expected future labor income growth. This implies that the adjustment parameters for a_t and y_t, γ_a and γ_c, are zero, so that neither a_t or c_t can have an important transitory component. Instead, only y_t can have a transitory component. When this permanent income framework has time-varying expected returns, as in (4) or (6), $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$ is related to future asset returns implying that γ_a nonzero; thus a_t must contain a transitory component, in addition to a possible transitory component in y_t, but there is still no transitory component in consumption. (If y_t is close to a random walk, a_t may contain the only transitory component in the system.) By contrast, in the linear habit model (5), movements in the cointegrating residual $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$ are by assumption related to future consumption growth since spending adjusts only sluggishly to permanent movements in wealth and income. As a consequence the adjustment parameter for consumption, γ_c, will be negative, implying that c_t must have a transitory component. If spending adjusts sluggishly to permanent movements a_t or y_t, there must be a transitory component in c_t, possibly correlated with the permanent components. An equivalent implication of this sluggish adjustment is that the cointegrating residual, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$, must *forecast* consumption growth over some future horizon. We test both of these implications below.

3.2.3 How Do Time-Varying Discount rates Relate to the Wealth Effect on Consumption?

Estimates of how wealth affects aggregate consumer spending have traditionally been produced from time-series regressions of consumption on wealth, controlling for labor income. These estimates are then converted into a marginal propensity to consume (MPC) out of wealth, giving the effect of a dollar increase in wealth on the level of consumption. In a classic paper, Modigliani (1971) performed such a calculation and found that a dollar increase in wealth leads to an increase in consumer spending of about 5 cents. This similar estimation strategies using more recent data find an MPC between three and five cents on the dollar (Ludvigson and Steindel (1999); Poterba (2000)).

These MPC estimates are commonly interpreted as applying to *every* movement in wealth. Yet this interpretation is generally valid only if discount rates are constant, and all movements in wealth are permanent. To see why, consider the relation between log consumption and log wealth when expected returns are not constant. After rearranging, the approximate consumption function in (4) may be written as a regression equation taking the form

$$\begin{aligned} c_t &= (1-\nu)a_t + \nu y_t + u_t \\ u_t &= (1-\nu-\sigma)E_t \sum_{i=1}^{\infty} \rho_w^i r_{wt+i} + \nu E_t \sum_{i=1}^{\infty} \rho_a^i \Delta y_{t+i}, \end{aligned} \tag{11}$$

where we have used the equalities, $\alpha_a = (1-\nu)$ and $\alpha_y = \nu$. Estimates of the MPC out of wealth are based on estimates of α_a in a regression of c_t on a_t and y_t; the part of the consumption function associated with expected returns and expected labor income growth appear in the regression error, u_t. Thus, the parameter α_a gives the marginal impact of wealth on consumption, holding fixed labor income. Notice, however, that, under the maintained hypothesis that returns and labor income growth are stationary, the coefficients α_a and α_y are cointegrating coefficients (they measure a common stochastic trend). Thus they give the marginal impact, on consumption, of *permanent* changes in wealth and income, but are silent about the possible influence of transitory changes. Consequently, they are only valid as a description of the effect of a typical change in wealth if all movements in wealth are permanent. If markets are efficient, this will occur in models with constant discount rates.

If discount rates vary, not all movements in wealth will be permanent. Movements in a_t generated by movements in expected returns will be transitory, and their impact on consumption is determined by preferences. For example, if consumption is a random walk ($\sigma \to 0$ in (11)) shocks to expected returns will affect asset values but have *no* effect on consumption. To see this, note that, for a given dividend or earnings stream, movements in expected returns must be *negatively* correlated with movements in a_t.[23] Thus, when $\sigma \to 0$ the coefficients on $E_t \sum_{i=1}^{\infty} \rho_w^i r_{wt+i}$ and a_t in (11) are precisely the same, and a negative shock to expected returns is exactly offset by an increase in asset values today. By contrast, if consumption is not a random walk because the IES, σ, is greater than zero, a decrease in expected returns will not be entirely offset by the temporary rise in asset values, and shocks to expected returns will affect consumption. But whatever the value of σ, the impact of a transitory movement in wealth on consumption will not be reflected in estimated parameter α_a, upon which standard estimates of the MPC are based. Even though transitory variation in wealth implies that the first term in u_t will generally be negatively correlated with a_t, estimates of α_a will not be affected because cointegrating coefficients converge to

[23] This is evident from (10). Fixing dividends, a decrease in expected returns can only be generated by future asset price depreciation from higher current asset values.

their true values at a rate proportional to the sample size T, implying that they are robust to regressor endogeneity in sufficiently large samples.[24] It follows that estimates of the marginal propensity to consume based on the cointegrating parameter α_a reveal little about the relation between transitory changes in wealth and consumption or the relative importance of permanent and transitory movements, a task we take up now using the empirical approach described above.

4 Empirical Results

4.1 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 are related to permanent and transitory shocks. Table 3 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, and to the single transitory shock η_{3t}. We begin with this analysis in order to emphasize the point that cointegration alone (along with the assumption that the transitory shock is orthogonal to the two permanent shocks), allows us to identify variance decompositions with respect to the space spanned by the two permanent shocks and the transitory shock, as well as impulse responses to the transitory shock, η_{3t}. In the table, the former are denoted P, and the latter is denoted T. To quantify the sampling uncertainty of these variance decompositions, we compute cumulative distribution functions for each variance decomposition using a bootstrapping procedure. These results are presented in Appendix C. In the next subsection, we analyze variance decompositions and impulse responses to the two permanent shocks, η_{1t}, and η_{2t} individually. In the computations that follow, we set γ_c and γ_y to zero in order to match the evidence from Table 2 that these variables are small and statistically indistinguishable from zero.[25]

Table 3 shows that the two permanent shocks in this system explain over 99 percent of

[24] As mentioned, Monte Carlo analysis calibrated to match the data generating process of our data suggests that samples of the size currently encountered are sufficiently large to obtain extremely accurate estimates of the cointegrating parameters.

[25] Gonzalo and Ng (2001) recommend restricting the values of these parameters to zero where they are statistically insignificant at the five percent level. The same approach is followed in similar applications by Cochrane (1994) and Gonzalo and Granger (1995). Results, available upon request, show that nothing essential in our findings hinges on this restriction. In addition, as we show later in Table 5, the cointegrating error has no forecasting power for consumption growth or labor income growth at any horizon in the future, reinforcing the conclusion that γ_c and γ_y are indeed close to zero.

the variance in the forecast error for consumption growth at all horizons. Consumption is a function of only the permanent components in a_t and y_t. Less than one 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: again together they account for more than 99 percent of the variation in the long-run forecast error of Δy_t.

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 estimate suggests that 85 percent of the variation in the growth of asset wealth is attributable to this shock; only 15 percent is attributable to permanent shocks. This finding implies that the vast majority of variability in consumption, driven by permanent shocks, is completely 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 that only permanent changes in wealth are related to consumer spending.

These findings are broadly consistent with a loglinear permanent income framework generalized to allow for time-varying expected returns. The Campbell and Cochrane (1999) model, for example, generates a random walk for consumption in the face of substantial time-variation in expected returns (transitory variation in a_t), consistent with what we find. Note that the Campbell-Cochrane model does not imply that agents can distinguish permanent from transitory shocks in a_t that are taken as *exogenous*; instead, the transitory component in a_t is generated endogenously by shifts in risk aversion.

These results show that consumption has virtually no transitory component that is orthogonal to the permanent shocks in the system. We argued above that consumption must contain a transitory component if it adapted sluggishly to permanent changes in a_t or y_t.The results we have presented so far, however, restrict η_{3t} to be orthogonal to the permanent shocks, and so do not rule out the possibility that consumption contains a transitory component that is correlated with the permanent components. To address this possibility, we provide, in Table 4, a variance-covariance decomposition, which requires no orthogonalization. The table shows that 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 somewhat correlated with the second permanent shock, but this does not alter the conclusion that the vast majority of variation in a_t is attributable to transitory shocks. In summary, the evidence from quarterly data suggests that consumption displays virtually no transitory variation, implying that it adapts within the span of about one quarter to permanent changes in wealth or income.

One intuitively appealing way to investigate whether consumption adjusts sluggishly to permanent shocks in wealth or income is to consider the *long-horizon forecasting* power of the cointegrating residual for consumption growth. If consumption adapts with a lag, permanent shocks to a_t or y_t will create a temporary cointegrating disequilibrium which must eventually be eliminated by an adjustment in c_t. This is precisely the prediction given by the presence of expected consumption growth on the right-hand-side of (5): the cointegrating residual must forecast consumption growth at some future horizon. Table 5 reports the results of regressions of Δc_{t+h} (defined as $c_{t+h} - c_t$), Δy_{t+h}, and Δa_{t+h} on $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$, controlling for Δc_t, Δy_t, and Δa_t, for horizons, h, ranging from one to 16 quarters. (Results for longer horizons are available on request and do not change the conclusions drawn from the findings presented in Table 5.)

The first panel of Table 5 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 4 quarters out, reflecting the modest degree of serial correlation in aggregate expenditure growth. These findings provide no support, however, for the proposition that consumption adjusts sluggishly to permanent income shocks. The coefficients on $\widehat{\alpha}'\mathbf{x}_t$ are always statistically insignificant and they explain a negligible fraction of the variation in future consumption growth at all horizons. The only variable for which $\widehat{\alpha}'\mathbf{x}_t$ has any forecasting power at any horizon is asset wealth; it has marginal predictive power even at a horizon of 16 quarters and beyond, with the adjusted R^2 statistic peaking at about 32 percent, at a 12 quarter horizon. Moreover, the adjusted R^2 statistic is unaffected by removing the other regressors Δc_t, Δy_t, and Δa_t from the forecasting regression. It follows that all of the long-horizon forecasting power for the growth in asset wealth is attributable to $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$.

In summary, these results indicate that when there is a transitory deviation from the common trend in c_t, a_t and y_t, it is not consumption or labor income that is forecast to adjust until the equilibrium is restored, but wealth. Since consumption has virtually no transitory component and does not adjust to restore a cointegrating disequilibrium, these results provide little evidence that consumption adapts sluggishly to permanent innovations in wealth or labor income, as would be the case in the linear habit model discussed above.

But if consumption is dominated by permanent shocks, how can we explain the serial correlation of consumption growth (Table 2) and the hump-shaped response of consumption to income shocks that are so well captured by the linear habit model (5)? One possibility, a possibility raised by the results presented here, is that these properties of measured expenditure growth may be primarily attributable to data construction methodologies. As Dynan

(2000) emphasizes, aggregate spending data are time averaged and much of the quarterly service flow from durable goods is computed by interpolation from annual surveys. These data construction procedures can create serial correlation in measured spending growth without generating a relation between the cointegrating residual, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$, and future consumption growth. Time averaging of aggregate data, for example, influences the first difference, not the level, of measured consumption. Accordingly, if the true log level of consumption were a random walk, permanent innovations would be reflected in consumption immediately, and the cointegrating error would not forecast time-averaged consumption growth, consistent with what we find. Similarly, the time-averaging of a random walk for log consumption can produce a hump-shaped response of measured consumption to wealth or income shocks, a feature of aggregate spending data documented by Fuhrer (2000). We illustrate these points in Appendix C, by performing a simple Monte-Carlo exercise for the case where log consumption is a random walk and therefore by construction responds only to permanent shocks and does so immediately and fully within the period. This evidence is also consistent with the results of Dynan (2000), who finds no evidence of serial correlation in household level spending.

Despite the serial correlation in measured spending growth, it still displays a correlation of 92 percent with its random walk component, given by the multivariate Beveridge-Nelson decomposition for this system (Table 6).[26] Similarly labor income growth displays a 91 percent correlation with its random walk component. By contrast, asset wealth is far from a random walk, displaying a correlation of just 12 percent with its random walk component.

How persistent is the transitory wealth shock, η_{3t}? Figure 1 shows the cumulative responses of Δc_t, Δa_t, and Δy_t, to a one-standard deviation shock in η_{3t}. Standard errors for these responses are presented in Appendix C. The responses of the estimated cointegrating error, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t = c_t - \widehat{\alpha}_a a_t - \widehat{\alpha}_y y_t$, are also plotted. The long-run response of $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$ to all three shocks is zero. Since consumption, asset wealth and labor income are cointegrated, deviations from the common trend in these variables must eventually be eliminated.

Figure 1 depicts graphically what the variance decompositions depicted numerically: an increase in the transitory wealth leads to a sharp increase in asset wealth, but has virtually no impact on consumption and labor earnings at any 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 at short to intermedi-

[26] 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.

ate horizons, but is eventually eliminated, as it must be, because the shock is transitory. More significantly, the figure shows that this transitory shock is quite 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. Yet despite their persistent effect on asset values, they bear virtually no relation to consumption at any future horizon.

So far, we have treated nonhuman wealth, a_t, as an aggregate. We do this because the budget constraint does not typically rationalize distinct roles for particular components of a_t. Nevertheless, it is reasonable to ask whether this transitory variation is driven by the stock market. Several pieces of evidence that indicate this transitory variation we find is attributable to the stock market component of household net worth. First, if we split wealth into its stock and nonstock components, we can again compute variance decompositions with respect to the identified permanent and transitory shocks for the four variable system $(c_t, s_t, n_t, y_t)'$, where s_t is the log of stock market wealth, S_t, and n_t is the log of all other wealth, N_t (Appendix A provides a description of these categories). Appendix C presents a detailed analysis of that four-variable variance decomposition, which shows clearly that the transitory variation in quarterly net worth is driven by volatility in the stock market, not by volatility in other forms of household net worth. In contrast to stock market wealth, nonstock wealth is dominated by permanent shocks. Second, Table 7 shows that the cointegrating residual for c_t, a_t, and y_t is a strong univariate predictor of Δs_t at horizons ranging from 1 to 24 quarters, with R^2 statistics that peak at 40 percent for a 12 quarter horizon. By contrast, the cointegrating residual has virtually no predictive power for Δn_t; the R^2 statistics are negligible and the coefficient estimates are not statistically different from zero. Since nonstock forms of wealth do not adjust to close a transitory gap in the cointegrating relation, they cannot have an important transitory component. Third, the evidence in Table 7 is consistent with evidence presented in Lettau and Ludvigson (2001a) which shows that the cointegrating residual for this system has strong forecasting power for returns on aggregate stock market indexes in excess of a short term interest rate but not for interest rates themselves.

How should we interpret the transitory component in wealth? From (10), we know that if asset markets are efficient, transitory variation in wealth must be attributable to time-variation in the required rate of return on assets, or discount rates. The contribution of this paper is empirical, and we do not attempt to explain why the required return might vary over time. Nevertheless, our results do make it possible to test for some forms of market inefficiency. For example, transitory variation in wealth could be related to transitory variation in dividend or earnings growth. We explore this possibility in Table 8 by asking whether the transitory component of asset values we identify is related to future movements in dividend or earnings growth. (The transitory component of wealth, which we describe further below, is defined as the residual of asset values over the trend component in a_t, where the trend component is given by the Beveridge-Nelson decomposition for our system.) Table 8 shows that this transitory component is a strong marginal predictor of real returns on the Standard and Poor Composite Index of 500 stocks (S&P 500). For example, at a four year horizon the R^2 statistic is 20 percent and the t-statistics exceed two at every horizon. By contrast, the R^2 statistics and t-statistics for predicting future earnings or dividend growth on the S&P 500 Index are unimpressive: The R^2 statistic never exceeds 10 percent and the Hodrick-corrected t-statistics (described in the Table) never exceed 2. These results imply that transitory net worth shocks are primarily shocks to asset prices, fixing future dividends and earnings.[27] Transitory shocks to a_t therefore have a natural interpretation not as dividend or earnings shocks, but as discount rate shocks. We now consider what these results mean for commonly used metrics of the wealth effect on consumption, and how one might interpret the two permanent shocks in our system.

4.1.1 The Marginal Propensity to Consume Out of Wealth

The wealth effect on consumption is often summarized by computing a marginal propensity to consume out of wealth from the cointegrating coefficients for consumption, wealth and income, in the manner described above. Using our own estimates of these cointegrating coefficients (i.e., $\widehat{\alpha}_a = 0.32$) we may follow the standard procedure and calculate the marginal impact of a *permanent* dollar increase in wealth on consumption, a figure that comes to about 4.6 cents, in line with values that represent a consensus for the MPC out of wealth.[28]

It is common practice to apply the MPC estimated in this manner to an every-day change in wealth. The results presented here, however, suggest that this practice is misleading because it fails to distinguish wealth movements that are permanent from those that are transitory. Not only are most movements in wealth are transitory, but these transitory changes–no matter how persistent–bear virtually no relation to consumer spending. Thus, any metric which seeks to summarize the wealth effect on consumption must take into account the relative importance of permanent and transitory fluctuations in asset values, and their differential impact on consumption.

One way to do this is to use our estimates from the variance decompositions presented

[27] These findings are consistent with those in Campbell (1991) and Cochrane (1991) who use data on stock prices and dividends, rather than household net worth and consumption, as we do here.

[28] This number is obtained by multiplying $\widehat{\alpha}_a = 0.32$ by the most recent value of A_t/C_t.

in Table 3. Those results suggest that the vast majority of variation in asset wealth–indeed 85 percent if one takes the point estimates at face value–is transitory and has no impact on consumption. On the other hand, a permanent one dollar changes in wealth is associated with a 4.6 cent movement in consumption. Thus the MPC out of an "average" movement in wealth will be a weighted average of zero and 4.6 cents: $MPC = \pi \cdot 0 + (1 - \pi) \cdot 0.046$. If 85 percent of the variance in wealth is transitory, $\pi = 0.71$, implying $MPC = 1.4$ cents per dollar of wealth gained.[29] If we instead use the variance decomposition estimates at the lower end of the 95 percent confidence interval in Table 3, one obtains $MPC = 2$ cents per dollar of wealth gained. It follows that common estimates of the MPC out of wealth tend to overstate the wealth effect on consumption because it applies only to permanent changes in asset values, movements which constitute a very small fraction of the historical variation in asset values.

4.1.2 Interpretation of the Two Permanent Shocks

We now turn to an analysis of the two permanent shocks. To interpret impulse response functions and variance decompositions with respect to the two permanent shocks individually, we must transform the permanent innovations in (9) so that they are mutually orthogonal. We do this mechanically using the procedure developed in Gonzalo and Ng (2001) who provide the following practical rule: if H is the Cholesky decomposition of $cov(\boldsymbol{\eta}_t)$, then $\widetilde{\boldsymbol{\eta}} \equiv H^{-1}\boldsymbol{\eta}_t$ yields a vector of mutually uncorrelated innovations, under the assumption that permanent shocks are uncorrelated with transitory shocks.[30]

How can we interpret a particular Cholesky decomposition which would separately identify our two permanent innovations? Recall that a random walk for c_t it is a good approximation of measured spending (Table 6). This evidence is consistent with the generalized permanent income model in (4) (when $\sigma \to 0$) and in the Campbell-Cochrane model (6). Thus we label one of the two orthogonalized permanent shocks, $\widetilde{\eta}_1$, as a "consumption" shock and interpret it as an innovation to permanent income (including both capital and labor income). Such a permanent income shock could be driven, for example, by a shock to productivity growth. This labelling is in the spirit of Cochrane (1994) who investigated a

[29] Since a variance decomposition gives the fraction of variability in squared changes, we take the square root of these figures to get the impact on actual changes, and rescale so the weights sum to one. Taking the estimates in Table 3 at face value, they imply that 85 percent of the variance in Δa_t is attributable to transitory shocks; 15 percent is attributable to the permanent shocks. Thus we compute, $\pi = \frac{\sqrt{0.85}}{\sqrt{0.85}+\sqrt{0.15}} = 0.71$.

[30] Note that the orthogonalized transitory shock, $\widetilde{\eta}_3$, is by construction identical to the transitory wealth shock, η_3 analyzed above. This follows because the space spanned by the transitory shocks was already assumed to be orthogonal to that of the permanent shocks.

bivariate VECM for consumption and GNP. Cochrane characterized a transitory movement in GNP as one with no contemporaneous movement in consumption, an interpretation motivated by the PIH. Similarly, in our application the permanent income interpretation suggests that innovations in wealth or income with consumption held contemporaneously fixed can have no effect on permanent income, so a shock to consumption is a shock to permanent income.[31]

Given that we interpret $\widetilde{\eta}_1$ as a permanent income shock, the second mutually orthogonal permanent shock in our system, $\widetilde{\eta}_2$, must be one that has no contemporaneous affect on permanent income. It follow that the second permanent shock must shift the composition of income between labor and capital in such a way that total permanent income is left unaffected. We label this second shock an "an income redistributive" shock. These considerations imply the following ordering of the variables in the vector, $\mathbf{x}_t$, for the Cholesky decomposition of $\boldsymbol{\eta}_t$: $\mathbf{x}_t = (c_t, y_{t,} a_t)'$.

We do not designate names for these orthogonalized shocks that convey more than "permanent income" shock, and "income redistributive" shock. The purpose of this paper is not to identify the underlying economic shocks for this system, a task which would require imposing a structural model with exactly three innovations, two of which that have permanent effects and one of which that has transitory effects. We also do not present an analysis of the other possible orthogonalizations. The permanent income interpretation suggests the orthogonalization discussed above; alternative orthogonalizations are difficult to interpret because they produce impulse responses that are linear combinations of responses to the permanent income shock and to the income-neutral shock.

Table 9 shows that the first permanent shock, $\widetilde{\eta}_1$ (labelled P1 in the table), explains about 97 percent of the variance in the long-run forecast error for consumption growth; if log consumption were exactly a random walk, this number would be 100 percent. The three percent attributable to the second permanent shock, $\widetilde{\eta}_2$ (P2 in the table), and to the transitory shock arises because the terms for lagged asset and income growth are not precisely zero in the consumption growth equation of the underlying VECM (Table 2). This evidence is consistent with the permanent income interpretation given above. The permanent income shock explains about 30 percent of the long-run forecast error variance of labor income growth, while the income-redistributive, $\widetilde{\eta}_2$ shock explains about 69 percent.

Figure 2 shows the cumulative responses of Δc_t, Δa_t, and Δy_t, to a one-standard deviation

[31] Note that Cochrane's bivariate consumption-GNP model does not in general permit an investigation of the relation between transitory variation in wealth and consumption. This follows because unspent capital gains (losses) are not counted as part of GDP. Yet if consumption is close to a random walk and there is a large transitory component in wealth, these unspent gains will comprise the bulk of variation in asset values.

shocks in $\widetilde{\eta}_{1t}$, and $\widetilde{\eta}_{2t}$. Standard errors for these responses are presented in Appendix C. As before, the responses of the estimated cointegrating error, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$, are also plotted. Notice that consumption responds to a permanent income shock by jumping immediately to a level that is close to its long-run position, indicating that log consumption is close to a random walk. The slight sluggishness in this adjustment is a feature of the modest degree of serial correlation in measured spending growth. Labor income and asset wealth also respond positively to $\widetilde{\eta}_{1t}$; thus our identified permanent income shock leads both to higher labor earnings and higher asset wealth.

The second panel of Figure 2 plots the response of each variable to a one standard deviation increase in the income redistributive shock, $\widetilde{\eta}_{2t}$. This shock has virtually no impact on consumption at any horizon and the consumption response is never statistically different from zero at any horizon (consistent with the permanent income interpretation), but instead causes labor income to increase and asset wealth to decrease. The initial impact of the shock is to drive up labor earnings and, within two quarters, drive down asset values. Thus the cointegrating error deviates from its historical norm in response to an $\widetilde{\eta}_{2t}$ shock because asset wealth adjusts sluggishly to its new long-run level; the shock forecasts a decline in asset values. This evidence suggests the presence of a persistent component that shifts the composition of income between labor and capital, reminiscent of the low frequency movements in labor's share observed in post-war data.

Why should such an income-redistributive shock forecast asset values? One possible answer is given in a recent paper by Santos and Veronesi (2000). They propose a simple general equilibrium model in which consumption is funded by dividends and labor income, which they assume follow a joint stochastic process. An implication of their model is that movements in the labor income-consumption ratio–a ratio they specify as stationary but extremely persistent–should predict asset wealth because such movements change the degree to which asset returns are correlated with consumption. In their model, higher labor income-consumption ratios predict a decline in asset wealth at long horizons. This implication is borne out nicely in Figure 2: an innovation to the second permanent shock, $\widetilde{\eta}_{2t}$, drives up labor income but has no impact on consumption; thus it immediately increases the labor income-consumption ratio. The impulse response shows that asset wealth responds slowly to such a shock, implying that the labor income-consumption ratio forecasts asset wealth at long horizons, consistent with the Santos and Veronesi model. A wrinkle in this story is that Santos and Veronesi presume that the labor income-consumption ratio is persistent but nonetheless stationary, whereas the data cannot distinguish such persistence from nonstationarity. For this reason, our model suggests that an $\widetilde{\eta}_{2t}$ shock leads to a permanent shift in the labor income-consumption ratio, which is associated with a permanent shift downward in asset wealth.

Finally, notice that Figures 1 and 2 both show that, no matter what the shock, fluctuations in the cointegrating residual always tell us something about the future path of asset values, not about the future path of consumption or labor income. Both the transitory shock and the redistributive shock cause $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$ to deviate from its historical norm, and this deviation leads a subsequent movement in asset values.

4.2 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. We define the trend in each variable as the long-run forecast of that variable, given by the trend component from the multivariate Beveridge-Nelson decomposition for the cointegrated system $(c_t, a_t, y_t)'$.[32]

The three panels of Figure 3 plot the resulting trend components of consumption, asset wealth and labor earnings from our three-variable system, along with the actual series for each variable. The plot spans the period beginning in the fourth quarter of 1952 to the fourth quarter of 2001:1. Consumption and labor earnings display little deviation from their respective trend components–the series are visually indistinguishable from their trends in the figure. This is not surprising since we already know that they are highly correlated with each other (Table 6). Of greater interest is the panel showing the trend for asset wealth. Here we see that there have been many periods in the post-war period during which asset wealth has diverged substantially from the estimated trend for this variable.

A clearer picture of the extent to which this is true is given in Figure 4, which shows the difference between the trend and actual value of asset wealth (the transitory component of wealth), in percent of the trend component. The figure plots two series: the transitory component of household net worth and the transitory component of the stock market component

[32] The trend component of each variable is obtained by using the structural model $\Delta\mathbf{x}_t = \boldsymbol{\delta} + \mathbf{D}(\mathbf{L})\boldsymbol{\eta}_t$, where $\boldsymbol{\eta}_t$ are the permanent and transitory shocks and $\boldsymbol{\delta} \equiv E(\Delta\mathbf{x}_t)$. The trend component is a random walk, $\mathbf{x}_t^{rw}$, given by $\mathbf{x}_t^{rw} = \mathbf{x}_0^{rw} + t\boldsymbol{\delta} + \mathbf{D}(1)\sum_{s=1}^{t}\boldsymbol{\eta}_s$, where $\mathbf{x}_0^{rw}$ is the initial value of $\mathbf{x}_t^{rw}$ and the parameters in $\boldsymbol{\delta}$ and $\mathbf{D}(1)$ have been estimated. This definition of trend, as the long-run forecast of each variable, implies that the unconditional mean of each variable will equal the unconditional mean of its trend component (Beveridge and Nelson (1981)). This identifies the initial values as $\mathbf{x}_0^{rw} = \overline{\mathbf{x}}_t - \boldsymbol{\delta}(T+1)/2$, where $\overline{\mathbf{x}}_{\mathbf{t}}$ is the sample mean of $\mathbf{x}_t$. The drift components in $\boldsymbol{\delta}$ corresponding to Δc_t and Δy_t are estimated as the sample means of these series. Because Δa_t is quite volatile, its drift is much more difficult to estimate using the sample mean. Thus, to insure that our estimates of the trend component in a_t will not be highly sensitive to the sample we use, we choose a value for δ_a that satisfies $\widehat{\boldsymbol{\alpha}}'\boldsymbol{\delta} = 0$, given $\widehat{\boldsymbol{\alpha}}$ and the estimated drifts in Δc_t and Δy_t. The restriction $\widehat{\boldsymbol{\alpha}}'\boldsymbol{\delta} = 0$ is part of the cointegration assumption and is required to rule out deterministic trends in $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$.

of net worth. The first is computed directly from the BN decomposition for our trivariate system. From this decomposition, an estimate of the transitory component in stock market wealth is obtained by dividing the transitory component of A_t by the average share of stock wealth, S_t, in total net worth, A_t, in our sample.[33] This is reasonable because the evidence presented above suggests that the transitory component in net worth is entirely associated with the stock market component of wealth. Both series in the figure have been normalized so that when they are above zero, wealth is estimated to be above its long-term level; when they are below zero, wealth is estimated to be below its long-term level. It is important to make clear that this transitory component is measured as a *percent of trend*; since the trend rises on average, if wealth is X percent above trend, a decline of less than X percent is typically required to restore wealth to its long run value.

Figure 3 shows that asset values were above their long-term trend during the mid 1950s and late 1960s, and below their long-term trend in many quarters from the mid 1970s through 1997. Consistent with popular impression, the estimate picks out the "bull markets" of the late 1960s and 1990s, and the "bear markets" of the 1970s. Note the sharp decline in wealth in 1973, a year in which stock market wealth relative to GDP fell by a factor of two. Some researchers, (for example, Hobijn and Jovanovic (2000)) have attributed this decline to a resistance on the part of incumbent firms to embrace new information technologies. While this hypothesis an interesting possibility, the evidence presented here suggests that movements in the transitory component of stock market wealth are typically generated by movements in stock prices for a *given* dividend or earnings stream (Table 8). This evidence is hard to square with the information technology hypothesis since that hypothesis suggests the temporary decline in stock values in 1973 was attributable to a temporary decline in

[33] In principal, one could compute the transitory component in stock market wealth directly from a BN decomposition for the four-variable system $(c_t, s_t, n_t, y_t)'$. In practice, this approach can suffer from important problems. To see why, note that, this decomposition requires an additional approximation taking the form $\ln(S_t + N_t) \approx \kappa + \alpha_s \ln S_t + \alpha_n \ln N_t$ (where α_s and α_n are asset shares and κ is a constant). Such an approximation will be reliable only when the S_t/N_t is close to its mean, so that the resulting approximation error is small. Assuming such approximation error is mean zero, covariance stationary, it will not effect calculations like variance decompositions, because those statistics represent averages over the whole sample. The difficulty occurs when the statistic one is interested in takes a value at every point in time. The transitory component itself is one such statistic. The statistic will give a poor estimate of the transitory component during episodes for which the approximation holds poorly. For example, the ratio S_t/N_t became particularly digressed from its sample mean in the last part of the 1990s, owing to the extraordinary surge in stock values during that period. As a consequence, the approximation $\kappa + \alpha_s \ln SW_t + \alpha_n \ln NSW_t$ considerably understates $\ln(SW_t + NSW_t)$ in recent data, both in absolute terms and relative to consumption and income. Episodes like these are likely to grossly understate the transitory value in stock market wealth when it is computed from the four variable BN decomposition directly.

earnings by listed firms.

In general periods of above-trend asset wealth were typically followed by episodes of negative excess returns on the S&P 500; below-trend asset values were 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 1960s, and by the bear markets of the 1970s. Comparably, the decline in asset wealth relative to trend in 1994 was followed by the bull market of the late 1990s. This estimate suggests that, by the end of 1997, the surge had driven stocks to above trend levels, sending a bearish signal for stock values going forward.

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 1990s. Interestingly, the estimate suggests that as late as 1997–well into the bull market of the last half of the 1990s and well after the 1996 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 ratios, 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.

Why was wealth so far below trend in 1995? The answer, mechanically, lies with what was happening to aggregate consumption relative to aggregate wealth over the previous 5 years. Despite modest gains in the stock market over this period, total household net worth actually fell slightly from 1990 to 1995, a result of the decline in nonstock wealth (primarily home equity) over this period. In the face of this decline in net worth however, consumption continued to grow at a respectable average annual rate of just over three percent from 1990 to 1995. The strength in consumption in the face of flat to declining asset values suggested that equity values were due for an upward correction, consistent with what occurred. The figure underscores the importance of the information in consumption and labor income, in addition to conventional indicators such as dividends and earnings, for determining the long-run value of asset wealth.

5 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. In this paper, we present assumptions under which the trend and cyclical components in household net worth may be identified, and empirically investigate how they are related to aggregate consumer spending.

Several aspects of these results stand out.

Transitory variation in asset wealth is both quantitatively large and 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 2 years. Yet despite their quantitative importance, transitory fluctuations in asset values–no matter how persistent–are found to be unrelated to aggregate consumer spending. Instead, aggregate consumption is close to a random walk, and can be well described as a function of the trend components in wealth and income. The conclusion that consumption may be well described by a random walk has been reached elsewhere (e.g., Harvey and Stock (1988); Cochrane (1994)), but unlike previous studies, we show that it holds even once the substantial volatility household net worth is accounted for.

These findings have important implications for understanding the consumption-wealth link: permanent changes in wealth *do* affect consumer spending, but most changes in wealth are transitory and have no impact on consumption. Moreover, we find no evidence that spending adjusts with a lag to movements in wealth: instead, consumer spending adapts to permanent shocks within the span of about one quarter. The results indicate that commonly used metrics for summarizing the wealth effect on consumption tend to overstate the consumption-wealth link, and need to be modified to account for the differential impact of permanent and transitory changes in asset values. 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 are left 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 Gertler (1999)). The results in this paper underscore the importance of heeding this recommendation: most changes in asset values are transitory and have no impact on consumer spending, the largest component of aggregate demand.

Finally, the results presented here also have implications for the analysis of asset markets. One is that when stock prices are high relative to fundamentals, asset wealth is also typically high relative to its common trend with consumption and labor income. The coincidence of these two phenomena creates a certain irony: high prices relative to fundamentals are often justified as the rational response of forward-looking markets to the anticipation of permanently higher earnings growth. But if earnings are expected to be permanently higher, forward looking investors should raise consumption immediately, and there should be no deviation of asset wealth from its shared trend with consumption and labor income. A second implication is that, when the cointegrating residual, $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$ is dominated by error-correction in wealth (as we find here), $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$ forms a proxy for expected returns and can be used to define the "state" in a conditional version of consumption-based capital asset pricing model (Lettau and Ludvigson (2001b)). Such a conditional framework holds out hope that conditional consumption-based models can explain the cross-section of average stock returns where unconditional models have failed.

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 wealth in billions of current dollars, measured at the end of the period. We lag this series one period to produce a measure of beginning of period wealth. Stock market wealth includes direct household holdings, mutual fund holdings, holdings of private and public pension plans, personal trusts, and insurance companies. Nonstock wealth is the residual of total wealth minus stock market wealth, and includes ownership of privately traded companies in noncorporate equity. Our source is the Board of Governors of the Federal Reserve System.

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. Our source is the Bureau of Economic Analysis.

Appendix B: Tests for Cointegration

This appendix describes procedures we use to test for cointegration among consumption, labor income and household wealth.The output from these tests is contained in Tables *B.1* and *B.2*.

We report the results of two types of cointegration tests: residual based tests designed to distinguish a system without cointegration from a system with at least one cointegrating relationship, and tests for cointegrating rank designed to estimate the number of cointegrating relationships. The former requires that each individual variable pass a unit root test and are conditional on this pretesting procedure. Dickey-Fuller tests for the presence of a unit root in c, y, and a (not reported) are consistent with the hypothesis of a unit root in those series.

Table *B.1* reports test statistics corresponding to the Phillips and Ouliaris (1990) residual based cointegration tests. The approach applies the augmented Dickey-Fuller unit root test to the residuals from a regression of consumption on labor income and household wealth, assuming trending series. The table shows both the Dickey-Fuller t-statistic and the relevant five and 10 percent critical values. For the trivariate system with c_t, a_t, and y_t, the hypothesis of no cointegration is rejected at the five percent level by the augmented Dickey-Fuller test with one, or two lags, and at the 10 percent level for all lags considered (one through four). We applied the data dependent procedure suggested in Campbell and Perron (1991) for choosing the appropriate lag length in an augmented Dickey-Fuller test. This procedure suggested that the appropriate lag length was one, implying that test results favoring cointegration should be accepted.

The Phillips–Ouliaris tests provide no support for the hypothesis that any of the pairs $\{c, y\}$; $\{y, a\}$; or $\{c, a\}$ are cointegrated. The Dickey-Fuller t-statistics for these tests are not close to significant even at the 10 percent level.

Next we employ, for the trivariate system, testing procedures suggested by Johansen (1988, 1991) that allow the researcher to estimate the number of cointegrating relationships (the bivariate results using this procedure yield the same conclusions as the Phillips–Ouliaris procedure). This technique 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.

The test procedure depends on the number of lags assumed in the vector autoregressive structure. The table presents the test results obtained under a number of lag assumptions. The same effective sample (1954:1 to 1999:4) was used in estimating the model under each lag assumption.

The critical values obtained using the Johansen approach also depend on the trend characteristics of the data. We present results allowing for linear trends in data, and assuming that the cointegrating relation has a constant (Table *B.2*) (see Johansen (1988) and Johansen (1991) for a more detailed discussion of these trend assumptions). In choosing the appropriate trend model for our data, we are guided by both theoretical considerations and statistical criteria. Theoretical considerations imply that the long-run equilibrium relationship between consumption, labor income and wealth does not have deterministic trends, although each individual data series may have deterministic trends. Moreover, statistical criteria suggests that modelling a trend in the cointegrating relation is not appropriate: the normalized cointegrating equation under this assumption suggests that the parameters of the cointegrating vector are negative, at odds with any sensible model of consumer behavior. The Table also reports the 90 percent critical values for these statistics.

The Johansen L-max test results establish strong evidence of a single cointegrating relation among log consumption, log labor income, and the log of household wealth. Table *B.2* shows that, for every lag specification we consider, we may reject the null of no cointegration against the alternative of one cointegrating vector. In addition, we cannot reject the null hypothesis of one cointegrating relationship against the alternative of two or three. While the evidence in favor of cointegration is somewhat weaker according to the Trace statistic (we cannot reject the null of no cointegration against the alternative of three cointegrating relations), this evidence is contradicted by the unit root tests which suggest that each variable contains a unit root. Moreover, according to the Trace statistic, we may not reject the null of one (or two) cointegrating relations against the alternative of three.

Appendix C: Alternative Specifications and Extensions

I. Standard Errors for Impulse Response Functions

This appendix presents 95% confidence intervals for the impulse response functions in Figures 1 and 2 (see Table C.1). 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, $\widehat{e}_t$, are obtained and a new sample of data is constructed using the initial VECM parameter estimates by random sampling of $\widehat{e}_t$ with replacement. Given this new sample of data, all the parameters are reestimated, holding fixed the number of cointegrating vectors, and the impulse responses stored. This is repeated 1,000 times. The empirical 95% confidence intervals are evaluated from these 1,000 samples of the bootstrapped impulse response functions.

II. Cumulative Distribution Function of the P/T Variance Decompositions

To get a sense of the sampling uncertainty of the P/T variance decompositions, we compute cumulative distribution functions from a bootstrapping procedure. For each simulation, we draw randomly from the residuals from the estimated VECM and construct new data series. Then we reestimate the VECM and compute the P/T variance decomposition. We repeat this procedure 5,000 and plot the cumulative distribution function of the variance decomposition in Figure C.1. We set the coefficient γ_j to zero if the estimated γ_j in a particular simulation is insignificant at the 95% level.

III. Variance Decompositions for Net Worth Split Into Stock and Nonstock Wealth

Appendix A provides a description of the assets that fall into each category (stock wealth and nonstock wealth) as we have defined them. The purpose of investigating this four variable system is to investigate the component source of the transitory variation in quarterly asset values documented above. Tests for the number of cointegrating vectors in this four-variable system again suggested the presence of a single cointegrating vector, with little evidence of pair-wise cointegration. Consequently, the cointegrated system has three permanent shocks and one transitory shock. The cointegrating vector, $\boldsymbol{\alpha}$, is once more normalized so that the coefficient on consumption is unity; the estimated cointegrating coefficient on labor income is somewhat higher than it is in the three-variable system, equal to 0.65 instead of 0.58; the estimated coefficient on stock wealth is 0.05, which is about one-fourth of that for nonstock wealth, estimated to be 0.20.

Variance decompositions for this four-variable system are provided in Table C.2. The estimated portion of variation in consumption and labor income growth that is attributable to permanent shocks is little changed from that given by the three-variable system.[34] What is new is the finding that nonstock wealth also displays little transitory variation: only about one percent of the total variation in the growth of this variable is attributable to the transitory shock. By contrast, 82 percent of the variation in stock market wealth is attributable to the transitory shock.

IV. The hump-shape Response of Consumption

This section shows that time-averaged consumption data can generate the type of hump-shape response of consumption in standard recursive VARs, even if actual consumption is a random walk. This hump-shape response is produced by nonstructural, recursive VARs where a Choleski decomposition is employed to orthogonalize of the VAR residuals. A simple Monte-Carlo exercise demonstrates this for the case where log consumption is a random walk, and therefore by construction responds only to permanent shocks and does so immediately and fully within the period. In addition, if we assume that log consumption shares a single cointegrating relation with log labor income and log wealth (consistent with our empirical evidence), an impulse response function from a recursive VAR may be generated with simulated data from such a system. To illustrate this point, we generate data for a simple system calibrated to match the data generating process and empirical properties of our system. Notice that consumption adjusts somewhat sluggishly to its own innovations, but this sluggishness is slight, and transitory shocks have virtually no impact on consumption. The basic empirical pattern can be well captured by the following structural VECM representation

$$\begin{aligned} \Delta c_t &= \phi_{cc}\Delta c_{t-1} + u_{ct} \\ \Delta y_t &= \phi_{cy}\Delta c_{t-1} + u_{yt} + \rho_{cy}u_{ct} \\ \Delta a_t &= \phi_{ca}\Delta c_{t-1} + u_{at} + \rho_{ya}u_{yt} + \rho_{ca}u_{ct} + \gamma(c_{t-1} - \alpha_a a_{t-1} - \alpha_y y_{t-1}), \end{aligned} \tag{12}$$

where ϕ_{ij} and ρ_{ij}, $i, j = c, y, a$, are constants, γ corresponds to the third element of $\boldsymbol{\gamma}$ in (7) and u_{ct}, u_{yt}, u_{at} correspond to the three innovations $\widetilde{\eta}_{1t}, \widetilde{\eta}_{2t}$, $\widetilde{\eta}_{3t}$, respectively. With ϕ_{cc} set to about 0.2 (consistent with the data), Monte Carlo simulations show that the impulse responses and variance decompositions of this system to permanent and transitory shocks do a good job of replicating those presented above using actual data.in (12). As discussed however, the serial correlation in spending growth inherent in finding that $\phi_{cc} > 0$, may be attributable to data construction methodologies, such as time-averaging. To show that

[34] Since the purpose of analyzing this four-variable system is to determine the component source of the transitory variation in asset wealth, we are not concerned with labelling the individual permanent shocks, but only with the space spanned by the permanent shocks. The ordering of the variables for this exercise is therefore irrelevant.

such construction methodologies can generate a hump-shape response of consumption in standard recursive VARs (even if actual consumption is a random walk) we set $\phi_{cc} = 0$, and generate 100,000 Monte Carlo draws, averaging over every 20 consumption and income outcomes to produce an artificial time-series of time-averaged data which would approximate continuous decision making (consumption and income data are time-averaged, but wealth data are not).[35] We do this 100 times. Notice that, in this artificial system, consumption and labor income are both specified as random walks, so that the single transitory shock in the system only affects a_t.

Figure C.2 provides the impulse responses of each variable to a wealth shock from a simple recursive VAR where the variables are ordered: y, a, and c respectively. Two points about this figure bear noting. First, the response of measured (time-averaged) spending to a wealth shock is hump-shaped. This occurs despite the fact that consumption is a random walk and exhibits no fundamental stickiness or serial correlation. The left-side of the hump is produced by time-averaging, which induces serial correlation in the growth of consumption. The right-side of the hump is generated because the response of any variable to a shock in any other variable in this recursive VAR is a mixture of the responses to permanent and transitory shocks (a simple recursive VAR cannot separately identify the permanent and transitory components in this system). Thus, measured spending can respond slowly and predictably to a wealth shock even when true consumption responds fully within the period to all permanent shocks. Second, notice that the asset wealth response to its own innovation decays slowly over time. This response misleadingly suggests that consumption responds to a wealth shock that is largely transitory even though actual consumption is a random walk and by construction responds only to permanent shocks. Again, this occurs because the data are time-averaged and because each response generated by this recursive VAR is mixture of responses to permanent and transitory shocks. Thus, the wealth shocks from this recursive system do not identify a transitory innovation in wealth (part of that innovation is permanent), despite the visible decay in the response of a_t to its own innovation. These findings underscore the danger of drawing conclusions about the dynamic impact of permanent and transitory innovations from simple recursive VAR impulse response functions.

Finally, we note that, in these artificial data, the time-aggregation generates first order autocorrelation in consumption growth of about 0.24, yet the cointegrating residual $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$ has no forecasting power at any horizon for time-averaged consumption growth, consistent with the empirical results presented in this paper.

[35]The values of the other parameters in (12) (ϕ_{ca}, ϕ_{cy}, ρ_{cy}, ρ_{ca}, and ρ_{ya}) are set to match the impulse response and variance decompositions of c, a, and y found in the data to the two permanent and one transitory shock. These values are, $\phi_{ca} = 0.2$; $\phi_{cy} = 0$; $\rho_{cy} = 0.5$; $\rho_{ca} = 0.1$; $\rho_{ya} = 0.0$.

References

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

Attanasio, Orazio P. and Guglielmo Weber, "Consumption Growth, the Interest Rate and Aggregation," *Review of Economic Studies*, 1993, *60* (3), 631–649.

Bernanke, Ben S. and Mark Gertler, "Monetary Policy and Asset Price Volatility," *Federal Reserve Bank of Kansas City Economic Review, 4th Quarter*, 1999, pp. 17–53.

Beveridge, Stephen and Charles R. Nelson, "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*, 1981, *7*, 151–174.

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

Campbell, John Y., "Does Saving Anticipate Declining Future Labor Income? An Alternate Test of the Permanent Income Hypothesis," *Econometrica*, 1987, *55*, 1249–73.

——, "A Variance Decomposition for Stock Returns," *Economic Journal*, 1991, *101*, 157–179.

——, "Understanding Risk and Return," *Journal of Political Economy*, 1996, *104*, 298–345.

——, "Asset Prices, Consumption and the Business Cycle," in John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, Amsterdam: North-Holland, 1999, pp. 1231–1303.

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

—— **and Gregory Mankiw**, "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," in Oliver Blanchard and Stanley Fischer, eds., *NBER Macroeconomics Annual: 1989*, Cambridge, MA: MIT Press, 1989, pp. 185–216.

—— **and John H. Cochrane**, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 1999, *107*, 205–251.

——— **and Pierre Perron**, "Pitfalls and Opportunities: What Macroeconomists Should Know About Unit Roots," in Olivier J. Blanchard and Stanley Fischer, eds., *NBER Macroeconomics Annual: 1991*, Cambridge, MA: MIT Press, 1991.

——— **and Robert J. Shiller**, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1988, *1*, 195–227.

——— **, Andrew W. Lo, and Craig MacKinlay**, *The Econometrics of Financial Markets*, Princeton, NJ: Princeton University Press, 1997.

Carroll, Christopher D., "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 1997, *112* (1), 1–55.

——— , "Risky Habits and the Marginal Propensity to Consume out of Permanent Shocks," *International Economic Journal,* forthcoming, 2001.

——— **, Jody Overland, and David N. Weil**, "Savings and Growth with Habit Formation," *American Economic Review*, 2000, *90*, 341–355.

Cochrane, John H., "Explaining the Variance of Price-Dividend Ratios," *Review of Financial Studies*, june 1991, *5* (2), 243–280.

——— , "Permanent and Transitory Components of GDP and Stock Prices," *Quarterly Journal of Economics*, 1994, *109*, 241–265.

Cochrane, John H, *Asset Pricing*, Princeton, NJ: Princeton University Press, 2001.

Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica*, 1991, *59* (5), 1221–48.

Dynan, Karen E., "Habit Formation in Consumer Preferences: Evidence from Panel Data," *American Economic Review*, 2000, *90*, 391–406.

Fama, Eugene F. and Kenneth French, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 1988, *22*, 3–27.

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

Fuhrer, Jeffrey, "Habit Formation in Consumption and Its Implications for Monetary-Policy Needs," *American Economic Review*, 2000, *90*, 367–390.

Galí, Jordi, "Finite Horizons, Life-Cycle Savings, and Time-Series Evidence on Consumption," *Journal of Monetary Economics*, 1990, *26*, 433–452.

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

—— **and Serena Ng**, "A Systematic Framework for Analyzing the Dynamic Effects of Permanent and Transitory Shocks," *Journal of Economic Dynamics and Control,* forthcoming, 2001.

Hall, Robert, "Intertemporal Substitution in Consumption," *Journal of Political Economy*, 1988, *96*, 221–273.

Harvey, David and James H. Stock, "Continuous Time Autoregressive Models with Common Stochastic Trends," *Journal of Economic Dynamics and Control*, 1988, *12*, 365–384.

Hobijn, Bart and Boyan Jovanovic, "The Information Technology Revolution and the Stock Market: Evidence," *American Economic Review,* forthcoming, 2000.

Hodrick, Robert, "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement," *Review of Financial Studies*, 1992, *5*, 357–386.

Johansen, Soren, "Statistical Analysis of Cointegrating Vectors," *Journal of Economic Dynamics and Control*, 1988, *12*, 231–254.

—— , "Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models," *Econometrica*, 1991, *56*, 1551–1580.

King, Robert, Charles Plosser, James H. Stock, and Mark W. Watson, "Stochastic Trends and Economic Fluctuations," *American Economic Review*, 1991, *81*, 819–840.

Lamont, Owen, "Earnings and Expected Returns," *Journal of Finance*, 1998, *53*, 1563–87.

LeRoy, Stephen and Richard Porter, "The Present Value Relation: Tests Based on Variance Bounds," *Econometrica*, 1981, *49*, 555–557.

Lettau, Martin and Sydney Ludvigson, "Consumption, Aggregate Wealth and Expected Stock Returns," *Journal of Finance*, June 2001, *56* (3), 815–849.

—— **and** —— , "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia are Time-Varying," *Journal of Political Economy,* forthcoming, 2001.

Ludvigson, Sydney, "Consumption and Credit: A Model of Time-Varying Liquidity Constraints," *Review of Economics and Statistics*, 1999, *81*, 434–447.

——— **and Charles Steindel**, "How Important Is the Stock Market Effect on Consumption?," *Economic Policy Review, Federal Reserve Bank of New York*, 1999, *5*, 29–51.

Modigliani, Franco, "Consumer Spending and Monetary Policy: The Linkages," *Federal Reserve Bank of Boston Conference Series*, 1971, *1971* (5).

Muellbauer, John, "Habits, Rationality and Myopia in the Life Cycle Consumption Function," *Annales des Economie et de Statistique*, 1988, *9*, 47–70.

Newey, Whitney K. and Kenneth D. West, "A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 1987, *55*, 703–708.

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

Phillips, Peter and Sam Ouliaris, "Asymptotic Properties of Residual Based Tests for Cointegration," *Econometrica*, 1990, *58*, 165–193.

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

Santos, Jesus and Pietro Veronesi, "Labor Income and Predictable Stock Returns," 2000. Manuscript, University of Chicago.

Shiller, Robert J., "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," *American Economic Review*, 1981, *71*, 421–436.

———, "Stock Prices and Social Dynamics," *Brookings Papers on Economic Activity*, 1984, *2*, 457–498.

Stock, James H., "Asymptotic Properties of Least Squares Estimators of Cointegrating Vectors," *Econometrica*, 1987, *55*, 113–144.

Stock, James H and Mark Watson, "Testing for Common Trends," *Journal of the American Statistical Association*, 1988, *83*, 1097–1107.

Stock, James H. and Mark Watson, "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems," *Econometrica*, 1993, *61*, 783–820.

Table 1: Summary Statistics

	Δc_t	Δy_t	Δa_t	r_t
Univariate Summary Statistics				
Mean (in %)	0.523	0.555	0.653	2.001
Standard Deviation (in %)	0.472	0.868	1.889	8.005
Autocorrelation	0.334	0.084	0.061	0.042
Correlation Matrix				
Δc_t	1.000	0.536	0.274	0.107
Δy_t		1.000	0.155	0.067
Δa_t			1.000	0.872
r_t				1.000

Notes: This table reports summary statistics for quarterly growth of consumption Δc_t, labor income Δy_t, asset wealth Δa_t, and the quarterly real return of the CRSP value-weighted index, r_t. The sample spans the fourth quarter of 1952 to the first quarter of 2001.

Table 2: Estimates From a Cointegrated VAR

	Equation		
	Δc_t	Δa_t	Δy_t
Dependent variable			
$\Delta c_{t-i,i=1,2}$	**0.302**	0.712	**0.512**
(t-stat)	(2.570)	(1.499)	(2.259)
$\Delta a_{t-i,i=1,2}$	0.018	0.079	0.053
(t-stat)	(0.670)	(0.734)	(1.021)
$\Delta y_{t-i,i=1,2}$	0.060	0.124	-0.036
(t-stat)	(0.867)	(0.442)	(-0.266)
$\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$	-0.020	**0.393**	0.075
(t-stat)	(-0.683)	(3.288)	(1.313)
$\bar{R}^2$	0.117	0.104	0.034

Notes: The table reports the sum of estimated coefficients from cointegrated vector autoregressions (VAR) of the column variable on the row variable; t-statistics for the sum are in parentheses. Estimated coefficients that are significant at the 5% level are highlighted in bold face. The term $c_t - \widehat{\alpha}_a a_t - \widehat{\alpha}_y y_t = \widehat{\boldsymbol{\alpha}}'\mathbf{x}_t$ is the estimated cointegrating residual. Our sample spans the fourth quarter of 1952 to the first quarter of 2001.

Table 3: 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
1	1.000	0.000	1.000	0.000	0.056	0.944
2	1.000	0.000	1.000	0.000	0.138	0.862
3	0.995	0.005	0.996	0.004	0.149	0.851
4	0.994	0.006	0.996	0.004	0.152	0.848
∞	0.994	0.006	0.996	0.004	0.159	0.841
	(0.896, 0.999)	(0.002, 0.038)	(0.853, 0.999)	(0.003, 0.055)	(0.120, 0.296)	(0.734, 0.882)

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 2. The last row reports the bootstrapped 95% confidence intervals for the $h = \infty$ case. Appendix C describes the bootstrap. The sample period is the first quarter of 1953 to the first quarter 2001.

Table 4: Variance-Covariance Decomposition (Unorthogonalized)

Horizon	$\widetilde{\text{P1}}$	$\widetilde{\text{P2}}$	$\widetilde{\text{T}}$	$\widetilde{\text{P1}}$, $\widetilde{\text{P2}}$	$\widetilde{\text{P1}}$, $\widetilde{\text{T}}$	$\widetilde{\text{P2}}$, $\widetilde{\text{T}}$
			Panel A: $\Delta c_{t+h} - E_t \Delta c_{t+h}$			
1	1.000	0.000	0.000	0.000	0.000	0.000
2	0.986	0.006	0.000	0.007	-0.000	0.001
3	0.989	0.007	0.002	0.004	-0.001	-0.001
4	0.989	0.007	0.002	0.004	-0.001	-0.001
∞	0.989	0.007	0.002	0.004	-0.001	-0.001
			Panel B: $\Delta y_{t+h} - E_t \Delta y_{t+h}$			
1	0.000	1.000	0.000	0.000	0.000	0.000
2	0.035	0.974	0.000	-0.009	0.000	0.000
3	0.056	0.960	0.006	-0.017	-0.002	-0.004
4	0.059	0.957	0.007	-0.016	-0.002	-0.004
∞	0.060	0.956	0.007	-0.016	-0.002	-0.004
			Panel C: $\Delta a_{t+h} - E_t \Delta a_{t+h}$			
1	0.696	0.886	1.613	-0.809	-0.211	-1.175
2	0.706	0.713	1.280	-0.599	-0.164	-0.936
3	0.651	0.663	1.229	-0.555	-0.148	-0.840
4	0.628	0.639	1.219	-0.535	-0.142	-0.808
∞	0.548	0.551	1.219	-0.466	-0.131	-0.722

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 panel that is attributable to innovations in the unorthogonalized shocks. The unorthogonalized permanent shocks are denoted $\widetilde{\text{P1}}$ and $\widetilde{\text{P2}}$, and the unorthogonalized transitory shock is $\widetilde{\text{T}}$. The entries in the columns denoted x refer to the share of the total variance that is due to shock x while the entries in the columns denoted x, z refer to the share of the total variance that is due the covariance of shock x and shock z. Horizons are in quarters, and the underlying VECM is of order 2. The sample period is the first quarter of 1953 to the first quarter 2001.

Table 5: Long-Horizon Regressions

Horizon H	Δc_t	Δy_t	Δa_t	$\widehat{\boldsymbol{\alpha}}'\,\mathbf{x}_t$	$\bar{R}^2$
Panel A: $\sum_{h=1}^{H} \Delta c_{t+h}$ regressed on					
1	**0.24**	0.09	0.00	-0.03	0.14
	(2.79)	(1.80)	(0.25)	(-0.73)	
	{2.80}	{2.01}	{0.18}	{-0.89}	
4	**0.71**	0.10	0.05	-0.04	0.12
	(3.74)	(0.97)	(1.23)	(-0.29)	
	{3.62}	{1.14}	{1.30}	{-0.40}	
8	**0.74**	0.07	-0.03	-0.11	0.04
	(1.90)	(0.47)	(-0.49)	(-0.42)	
	{2.67}	{0.17}	{-0.04}	{-0.17}	
12	0.67	0.17	-0.04	-0.17	0.03
	(1.39)	(0.87)	(-0.36)	(-0.60)	
	{2.10}	{1.31}	{-0.62}	{-0.78}	
16	0.59	0.22	-0.07	-0.20	0.03
	(1.05)	(1.01)	(-0.61)	(-0.55)	
	{1.67}	{1.60}	{-0.96}	{-0.81}	
Panel B: $\sum_{h=1}^{H} \Delta y_{t+h}$ regressed on					
Horizon H	Δc_t	Δy_t	Δa_t	$\widehat{\boldsymbol{\alpha}}'\,\mathbf{x}_t$	$\bar{R}^2$
1	**0.43**	-0.03	0.02	0.06	0.06
	(3.12)	(-0.23)	(0.58)	(1.22)	
	{2.85}	{-0.27}	{0.58}	{1.10}	
4	**1.28**	-0.05	0.09	0.23	0.12
	(3.56)	(-0.29)	(1.42)	(1.12)	
	{3.77}	{-0.26}	{0.86}	{1.36}	
8	**1.43**	-0.22	-0.09	0.06	0.04
	(2.74)	(-0.95)	(-1.01)	(0.17)	
	{2.55}	{-0.89}	{-0.53}	{0.16}	
12	**1.75**	-0.13	-0.11	0.16	0.04
	(2.49)	(-0.41)	(-0.81)	(0.37)	
	{2.65}	{-0.54}	{-0.68}	{0.34}	
16	1.51	0.08	-0.22	0.38	0.04
	(1.73)	(0.19)	(-1.42)	(0.64)	
	{2.04}	{0.31}	{-1.34}	{0.67}	
Panel C: $\sum_{h=1}^{H} \Delta a_{t+h}$ regressed on					
Horizon H	Δc_t	Δy_t	Δa_t	$\widehat{\boldsymbol{\alpha}}'\,\mathbf{x}_t$	$\bar{R}^2$
1	0.66	0.35	0.06	**0.42**	0.13
	(2.16)	(2.02)	(0.61)	(3.23)	
	{1.84}	{1.61}	{0.57}	{2.77}	
4	**1.56**	0.28	0.10	**1.23**	0.17
	(2.02)	(0.92)	(0.61)	(2.09)	
	{1.90}	{0.72}	{0.54}	{2.04}	
8	0.43	0.75	0.20	**2.34**	0.22
	(0.45)	(1.79)	(0.86)	(2.81)	
	{0.36}	{1.48}	{0.67}	{2.25}	
12	0.26	1.01	0.35	**3.64**	0.32
	(0.21)	(2.33)	(0.98)	(4.51)	
	{0.19}	{1.77}	{0.91}	{2.80}	
16	0.53	0.70	0.20	**3.89**	0.31
	(0.41)	(1.74)	(0.60)	(3.50)	
	{0.36}	{1.12}	{0.47}	{2.79}	

Notes: See next page.

Notes for Table 5: The table reports output from long-horizon regressions of consumption, labor income and asset wealth on lags of these variables and the cointegrating residual $\widehat{\boldsymbol{\alpha}}'\,\mathbf{x}_t$. The dependent variables in the h-period regressions are $\Delta x_{t+1} + ... + \Delta x_{t+h}$, where $x \in \{c, y, a\}$. For each regression, the table reports OLS estimates of the regressors, Newey and West (1987) corrected t-statistics (in parentheses), Hodrick (1992) corrected t-statistics (in curly brackets) and adjusted R^2 statistics. Significant coefficients at the 5% level are highlighted in bold face. The sample period is the first quarter of 1953 to the first quarter 2001.

Table 6: Correlation of Growth Rates with Random Walk Components

Variable	Correlation
Δc_t	0.922
Δy_t	0.906
Δa_t	0.123

Notes: The sample period is the first quarter of 1953 to the first quarter 2001.

Table 7: 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 $\sum_{h=1}^{H} \Delta s_{t+h}$						
1.98	3.81	6.43	10.96	14.99	16.73	22.74
(4.06)	(3.45)	(2.96)	(3.51)	(5.12)	(5.04)	(4.66)
{3.61}	{3.58}	{3.12}	{2.95}	{3.11}	{3.18}	{3.56}
[0.08]	[0.14]	[0.21]	[0.32]	[0.40]	[0.38]	[0.36]
Panel B: Non-Stock Market Wealth $\sum_{h=1}^{H} \Delta n_{t+h}$						
-0.12	-0.16	-0.25	-0.23	-0.20	-0.41	-1.18
(-1.96)	(-1.21)	(-0.96)	(-0.49)	(-0.33)	(-0.56)	(-1.22)
{-2.68}	{-1.80}	{-1.50}	{-0.82}	{-0.52}	{-0.96}	{-1.28}
[0.04]	[0.02]	[0.02]	[0.01]	[0.00]	[0.01]	[0.03]

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 $\widehat{\boldsymbol{\alpha}}' \mathbf{x}_t$. In each column, the first number is the OLS coefficient; the second number, in parentheses, is the Newey and West (1987) corrected t-statistic; the third number, in curly brackets, is the Hodrick (1992) corrected t-statistic; and the fourth number, in square brackets is the adjusted R^2 statistic for the regression. The sample period is the first quarter of 1953 to the first quarter 2001.

Table 8: Long-Horizon Regressions: Returns, Dividends and Earnings

Forecast Horizon H						
1	2	4	8	12	16	24
Panel A: Stock Returns $\sum_{h=1}^{H} r_{t+h}$						
-0.01	-0.01	-0.02	-0.02	-0.03	-0.04	-0.05
(-2.86)	(-2.70)	(-2.54)	(-3.08)	(-3.27)	(-3.11)	(-3.45)
{-3.35}	{-2.96}	{-2.57}	{-2.20}	{-2.32}	{-2.39}	{-2.83}
[0.07]	[0.09]	[0.13]	[0.16]	[0.21]	[0.20]	[0.23]
Panel B: Dividend Growth $\sum_{h=1}^{H} \Delta d_{t+h}$						
-0.00	-0.01	-0.01	-0.02	-0.03	-0.04	-0.04
(-0.56)	(-1.69)	(-1.17)	(-2.70)	(-2.52)	(-2.29)	(-2.00)
{-0.30}	{-0.68}	{-0.47}	{-0.66}	{-0.89}	{-1.09}	{-1.04}
[-0.00]	[0.01]	[0.02]	[0.07]	[0.10]	[0.10]	[0.07]
Panel C: Earnings Growth $\sum_{h=1}^{H} \Delta e_{t+h}$						
-0.01	-0.01	-0.01	-0.03	-0.09	-0.14	-0.11
(-0.95)	(-0.62)	(-0.25)	(-0.45)	(-2.54)	(-2.92)	(-1.30)
{-0.59}	{-0.38}	{-0.15}	{-0.22}	{-0.55}	{-0.81}	{-0.63}
[-0.00]	[-0.00]	[-0.00]	[-0.00]	[0.03]	[0.07]	[0.03]

Notes: The table reports results from long-horizon regressions of stock returns, r_t, in Panel A, growth rates of dividends, Δd_t, in Panel B and earnings growth rates, Δe_t, in Panel C. The independent variable is the transitory component of wealth (i.e. the difference between wealth and the random walk component of wealth.) Stock returns are from the S&P 500 Index. Dividends are total CRSP dividends and earnings are measured per share from the S&P 500 Index. In each column, the first number is the OLS coefficient; the second number, in parentheses, is the Newey and West (1987) corrected t-statistic; the third number, in curly brackets, is the Hodrick (1992) corrected t-statistic; and the fourth number, in square brackets is the adjusted R^2 statistic for the regression. The sample period is the first quarter of 1953 to the first quarter 2001.

Table 9: Variance Decomposition of the Permanent Shocks

	$\Delta c_{t+h} - E_t \Delta c_{t+h}$		$\Delta y_{t+h} - E_t \Delta y_{t+h}$		$\Delta a_{t+h} - E_t \Delta a_{t+h}$	
Horizon h	P1	P2	P1	P2	P1	P2
1	1.000	0.000	0.265	0.735	0.055	0.000
2	0.979	0.021	0.289	0.711	0.136	0.002
3	0.975	0.020	0.299	0.698	0.134	0.016
4	0.974	0.020	0.301	0.695	0.133	0.019
∞	0.973	0.021	0.302	0.684	0.127	0.032
	(0.889, 0.984)	(0.007, 0.075)	(0.217, 0.396)	(0.580, 0.755)	(0.073, 0.197)	(0.029, 0.099)

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 an innovation in the first (permanent income) shock, P1, and the second permanent (income-neutral) shock, P2. Horizons are in quarters, and the underlying VECM is of order 2. The ordering of the variables is (c, y, a). The last row reports the bootstrapped 95% confidence intervals for the $h = \infty$ case. Appendix C describes the bootstrap. The sample period is the first quarter of 1953 to the first quarter 2001.

Table B.1: Phillips-Ouliaris Test for Cointegration

variables	Dickey-Fuller t-statistic				Critical Values	
	Lag=1	Lag=2	Lag=3	Lag=4	5% Critical Level	10% Critical Level
c, a, y	-4.053	-3.901	-3.560	-3.532	-3.80	-3.52
c, a	-1.529	-1.396	-1.381	-1.315	-3.42	-3.13
c, y	-0.748	-0.384	-0.590	-0.473	-3.42	-3.13
a, y	-0.820	-0.770	-0.899	-0.825	-3.42	-3.13

Notes: The first row presents Dickey-Fuller tests statistic that has been applied to the fitted residuals from the cointegrating regression of a trivariate system including consumption c on labor income y and wealth a. The second to fourth rows report results using pairs of these variables. Critical values assume trending series. "Lags" refers to the number of lags of first differences used in the regression of residuals on the lagged residual, and on lags of first differences of the residual. The sample period is the first quarter of 1953 to the first quarter 2001.

Table B.2: Johansen Cointegration Test

Lag in VAR Model= 1				
L-Max		Trace		$H_o = r$
Test Statistic	90% CV	Test Statistic	90% CV	$r =$
20.61	13.39	25.47	26.70	0
4.44	10.60	4.86	13.31	1
0.42	2.71	0.42	2.71	2
Lag in VAR Model = 2				
L-Max		Trace		$H_o = r$
Test Statistic	90% CV	Test Statistic	90% CV	$r =$
20.96	13.39	25.55	26.70	0
3.97	10.60	4.59	13.31	1
0.62	2.71	0.62	2.71	2
Lag in VAR Model = 3				
L-Max		Trace		$H_o = r$
Test Statistic	90% CV	Test Statistic	90% CV	$r =$
17.05	13.39	21.65	26.70	0
4.17	10.60	4.61	13.31	1
0.43	2.71	0.43	2.71	2
Lag in VAR Model = 4				
L-Max		Trace		$H_o = r$
Test Statistic	90% CV	Test Statistic	90% CV	$r =$
15.85	13.39	20.38	26.70	0
4.46	10.60	4.52	13.31	1
0.06	2.71	0.06	2.71	2

Notes: See Table B.1. I(1) analysis with linear trend in the data and a constant in the cointegrating relation. The columns labeled “Test Statistic” give the value for the test named in the row above; “90% CV” gives the 90 percent confidence level of that statistic. The sample period is the first quarter of 1953 to the first quarter 2001.

Table C.1: Impulse Response Function

	c_t			y_t			a_t		
h	P1	P2	T	P1	P2	T	P1	P2	T
1	0.439	0.000	0.000	0.438	0.729	0.000	0.371	-0.004	1.621
	(0.38,0.48)	(-0.00,0.00)	(0.00,0.00)	(0.30,0.56)	(0.62,0.80)	(0.00,0.00)	(0.16,0.56)	(-0.18,0.17)	(1.37,1.78)
4	0.696	0.078	0.049	0.777	0.710	0.072	0.931	-0.296	1.253
	(0.53,0.81)	(-0.02,0.17)	(0.04,0.13)	(0.51,1.00)	(0.51,0.85)	(0.05,0.20)	(0.53,1.26)	(-0.61,-0.00)	(0.83,1.53)
8	0.730	0.066	0.038	0.821	0.688	0.054	0.889	-0.640	0.697
	(0.53,0.89)	(-0.05,0.18)	(0.02,0.10)	(0.52,1.10)	(0.47,0.85)	(0.03,0.14)	(0.52,1.21)	(-0.95,-0.36)	(0.29,0.95)
12	0.730	0.055	0.021	0.821	0.673	0.031	0.846	-0.835	0.383
	(0.53,0.90)	(-0.06,0.18)	(0.01,0.05)	(0.52,1.12)	(0.45,0.85)	(0.02,0.08)	(0.49,1.14)	(-1.13,-0.55)	(0.06,0.63)
16	0.728	0.049	0.012	0.819	0.665	0.017	0.822	-0.943	0.210
	(0.53,0.90)	(-0.07,0.17)	(0.00,0.03)	(0.52,1.12)	(0.45,0.85)	(0.01,0.04)	(0.46,1.13)	(-1.21,-0.65)	(0.01,0.43)
∞	0.727	0.042	0.000	0.817	0.654	0.000	0.792	-1.072	0.000
	(0.54,0.89)	(-0.10,0.17)	(0.00,0.00)	(0.54,1.08)	(0.43,0.86)	(0.00,0.00)	(0.43,1.12)	(-1.33,-0.78)	(0.00,0.00)

Notes: The table reports the impulse response function of c, y, and a to the three shocks. The first (permanent income) shock is denoted P1, the second permanent (income-neutral) shock is P2, and the transitory shock is T. The 95% confidence intervals are reported in brackets below the point estimates. Standard errors are computed using the bootstrap procedure described in Gonzalo and Ng (2001) using 1000 replications. Horizons h are in quarters, and the underlying VECM is of order 2. The sample period is the first quarter of 1953 to the first quarter 2001.

Table C.2: Variance Decomposition: Stock Market Wealth and Non-Stock Market Wealth

	$\Delta c_t - E_{t-1}\Delta c_t$		$\Delta y_t - E_{t-1}\Delta y_t$	
Horizon	P	T	P	T
1	1.000	0.000	1.000	0.000
2	0.999	0.001	1.000	0.000
3	0.999	0.001	0.998	0.002
4	0.998	0.002	0.998	0.002
∞	0.998	0.002	0.997	0.003
	$\Delta n_t - E_{t-1}\Delta n_t$		$\Delta s_t - E_{t-1}\Delta s_t$	
Horizon	P	T	P	T
1	0.998	0.002	0.055	0.945
2	0.994	0.006	0.104	0.896
3	0.988	0.012	0.138	0.862
4	0.989	0.011	0.149	0.851
∞	0.989	0.011	0.178	0.822

Notes: The table reports variance decompositions of forecast errors at various horizons. P(T) denotes permanent (transitory) shocks. n_t denotes the log of non-stock market wealth (defined in Appendix A), s_t denotes the log of stock market wealth. The sample period is the first quarter of 1953 to the first quarter 2001.

Figure 1: Impulse Response Function to Transitory Shock

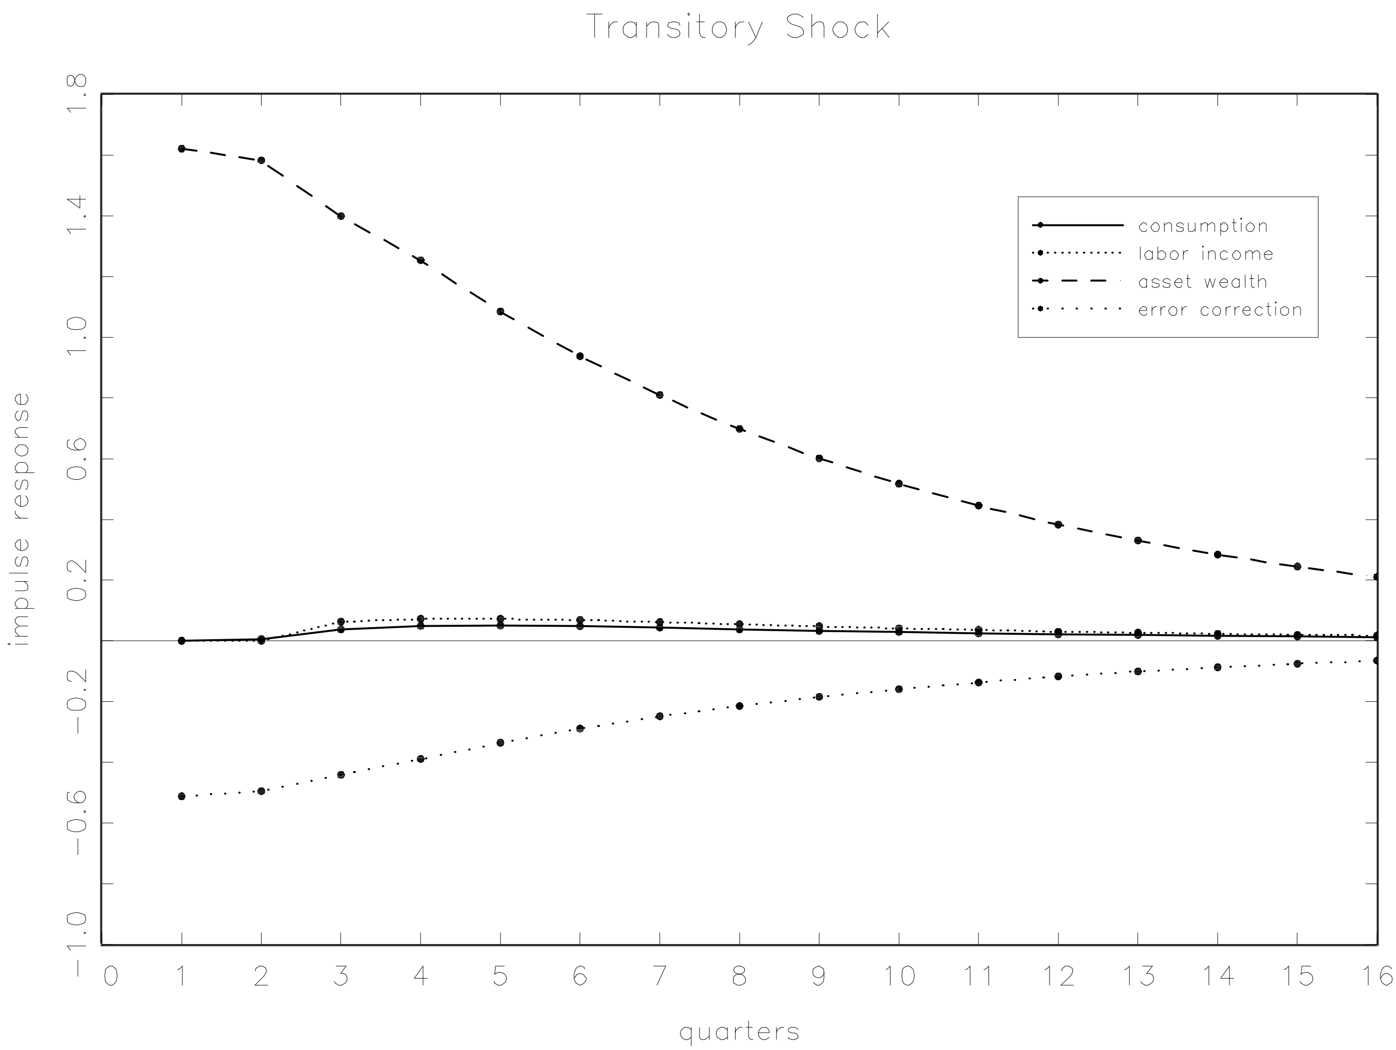

Note: Impulse responses to a one-standard deviation transitory shock.

Figure 2: Impulse Response Functions to Permanent Shocks

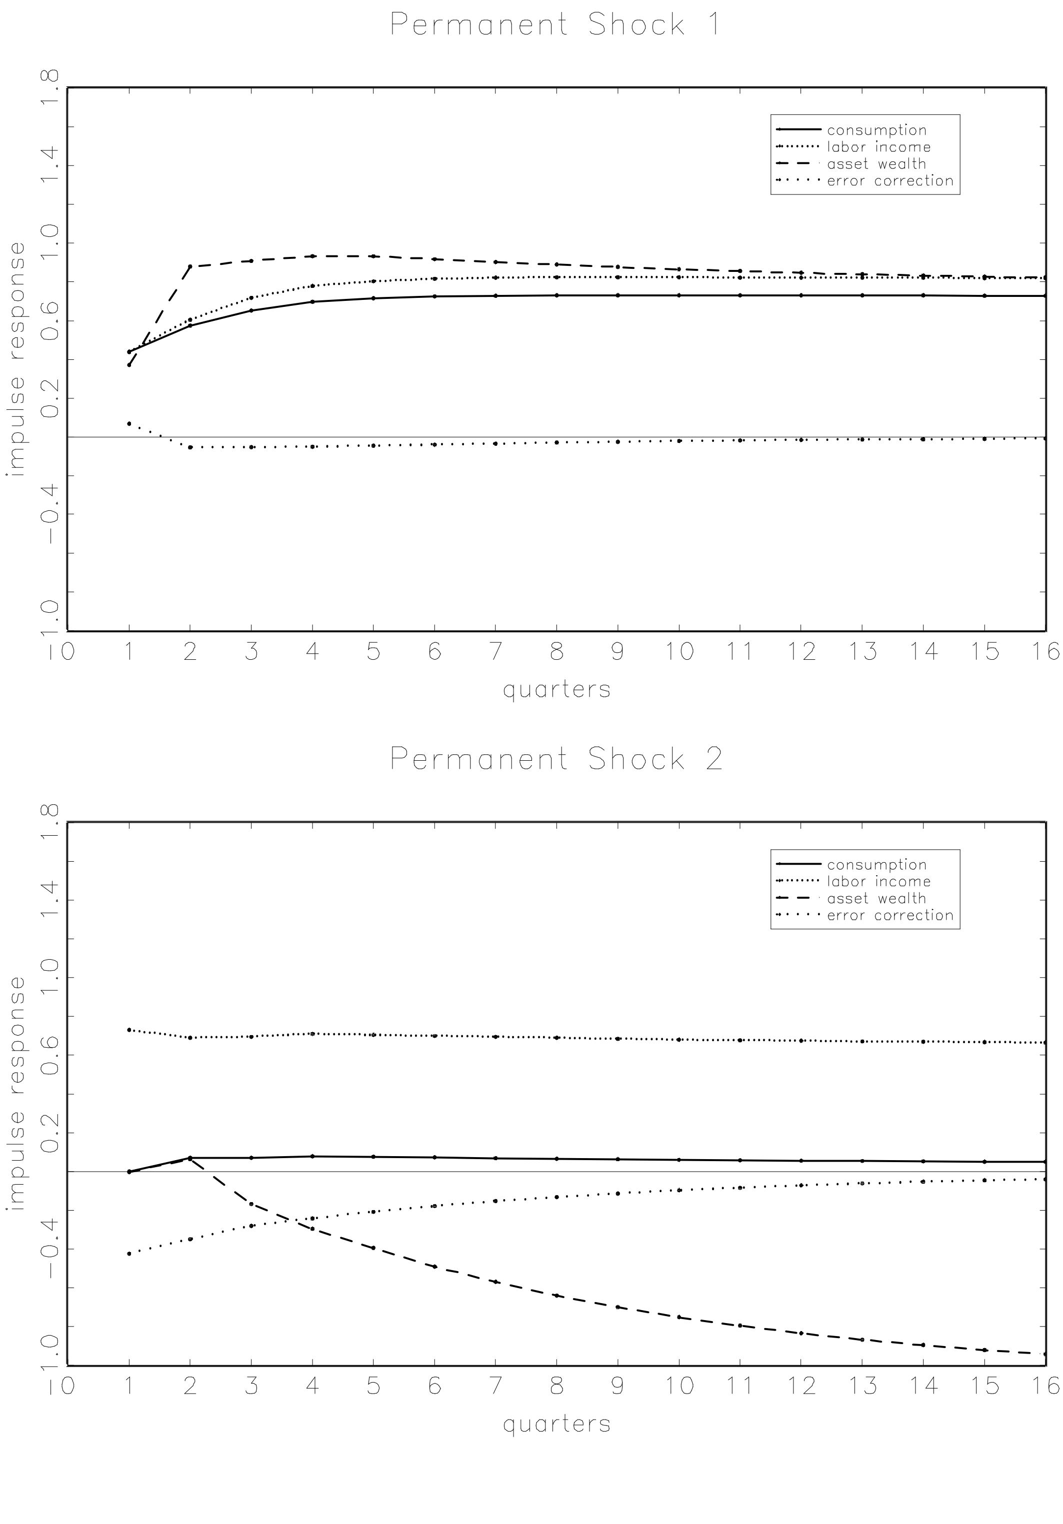

Note: Impulse responses to one-standard deviation permanent shocks.

Figure 3: Estimates of the Trends in c, y and a

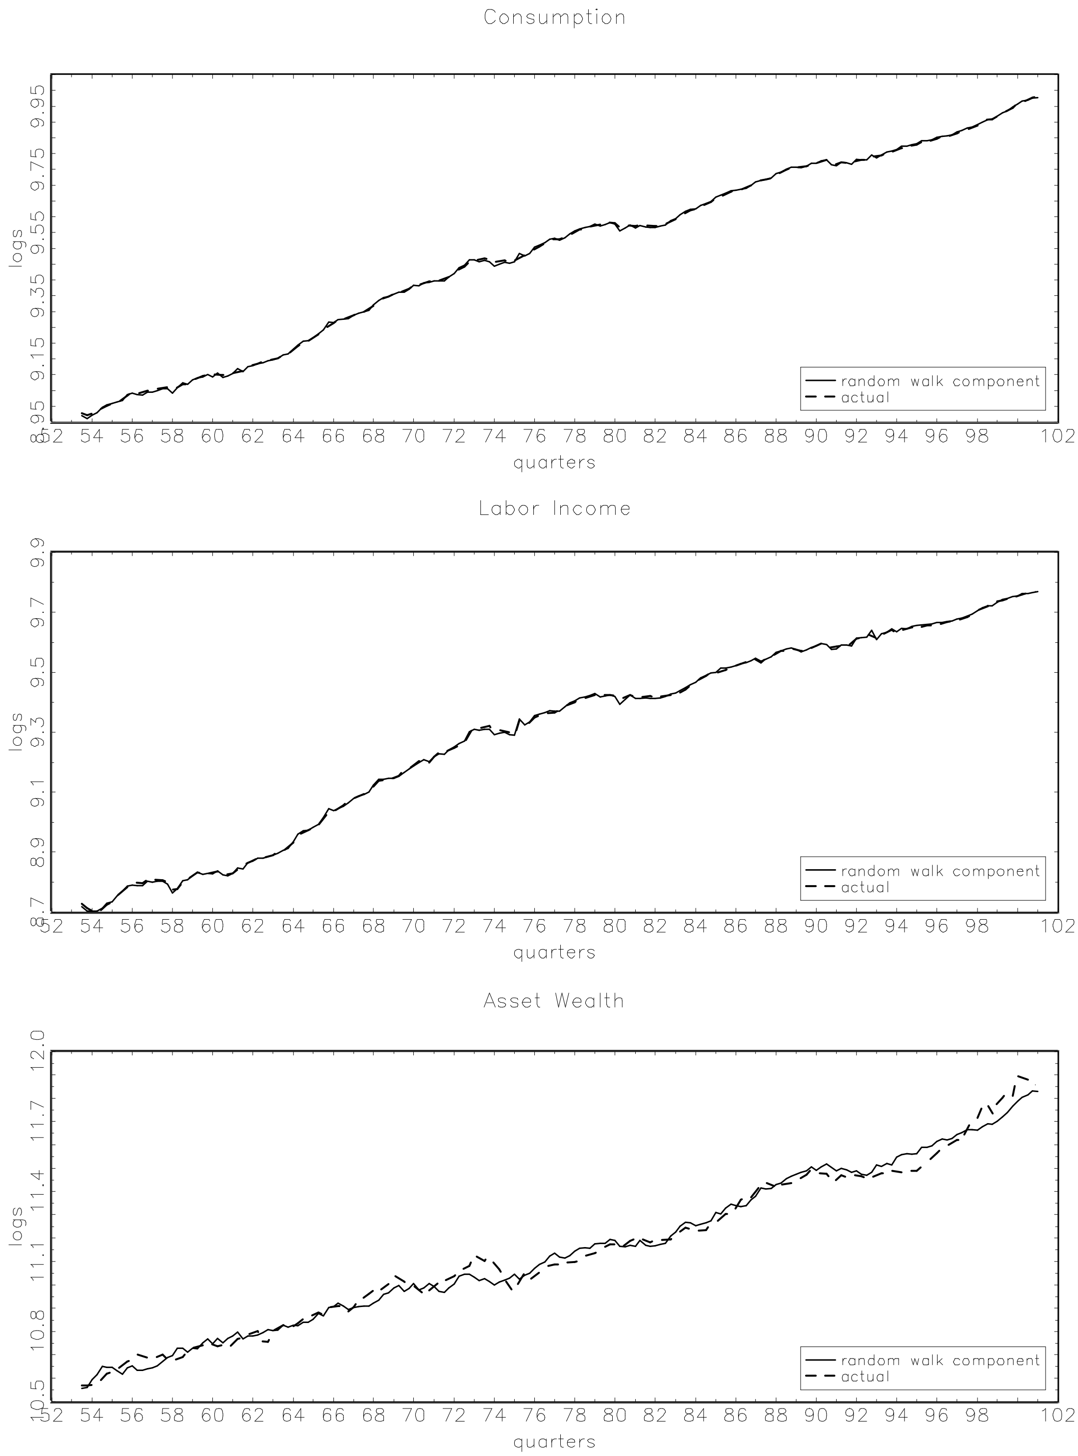

Note: The trend is defined as the long-run forecast of each variable given the multivariate Beveridge-Nelson decomposition for the trivariate system c_t, y_t and a_t. The variables are measured in log real per-capita units.

Figure 5: The Transitory Components of Net Worth and Stock Market Wealth

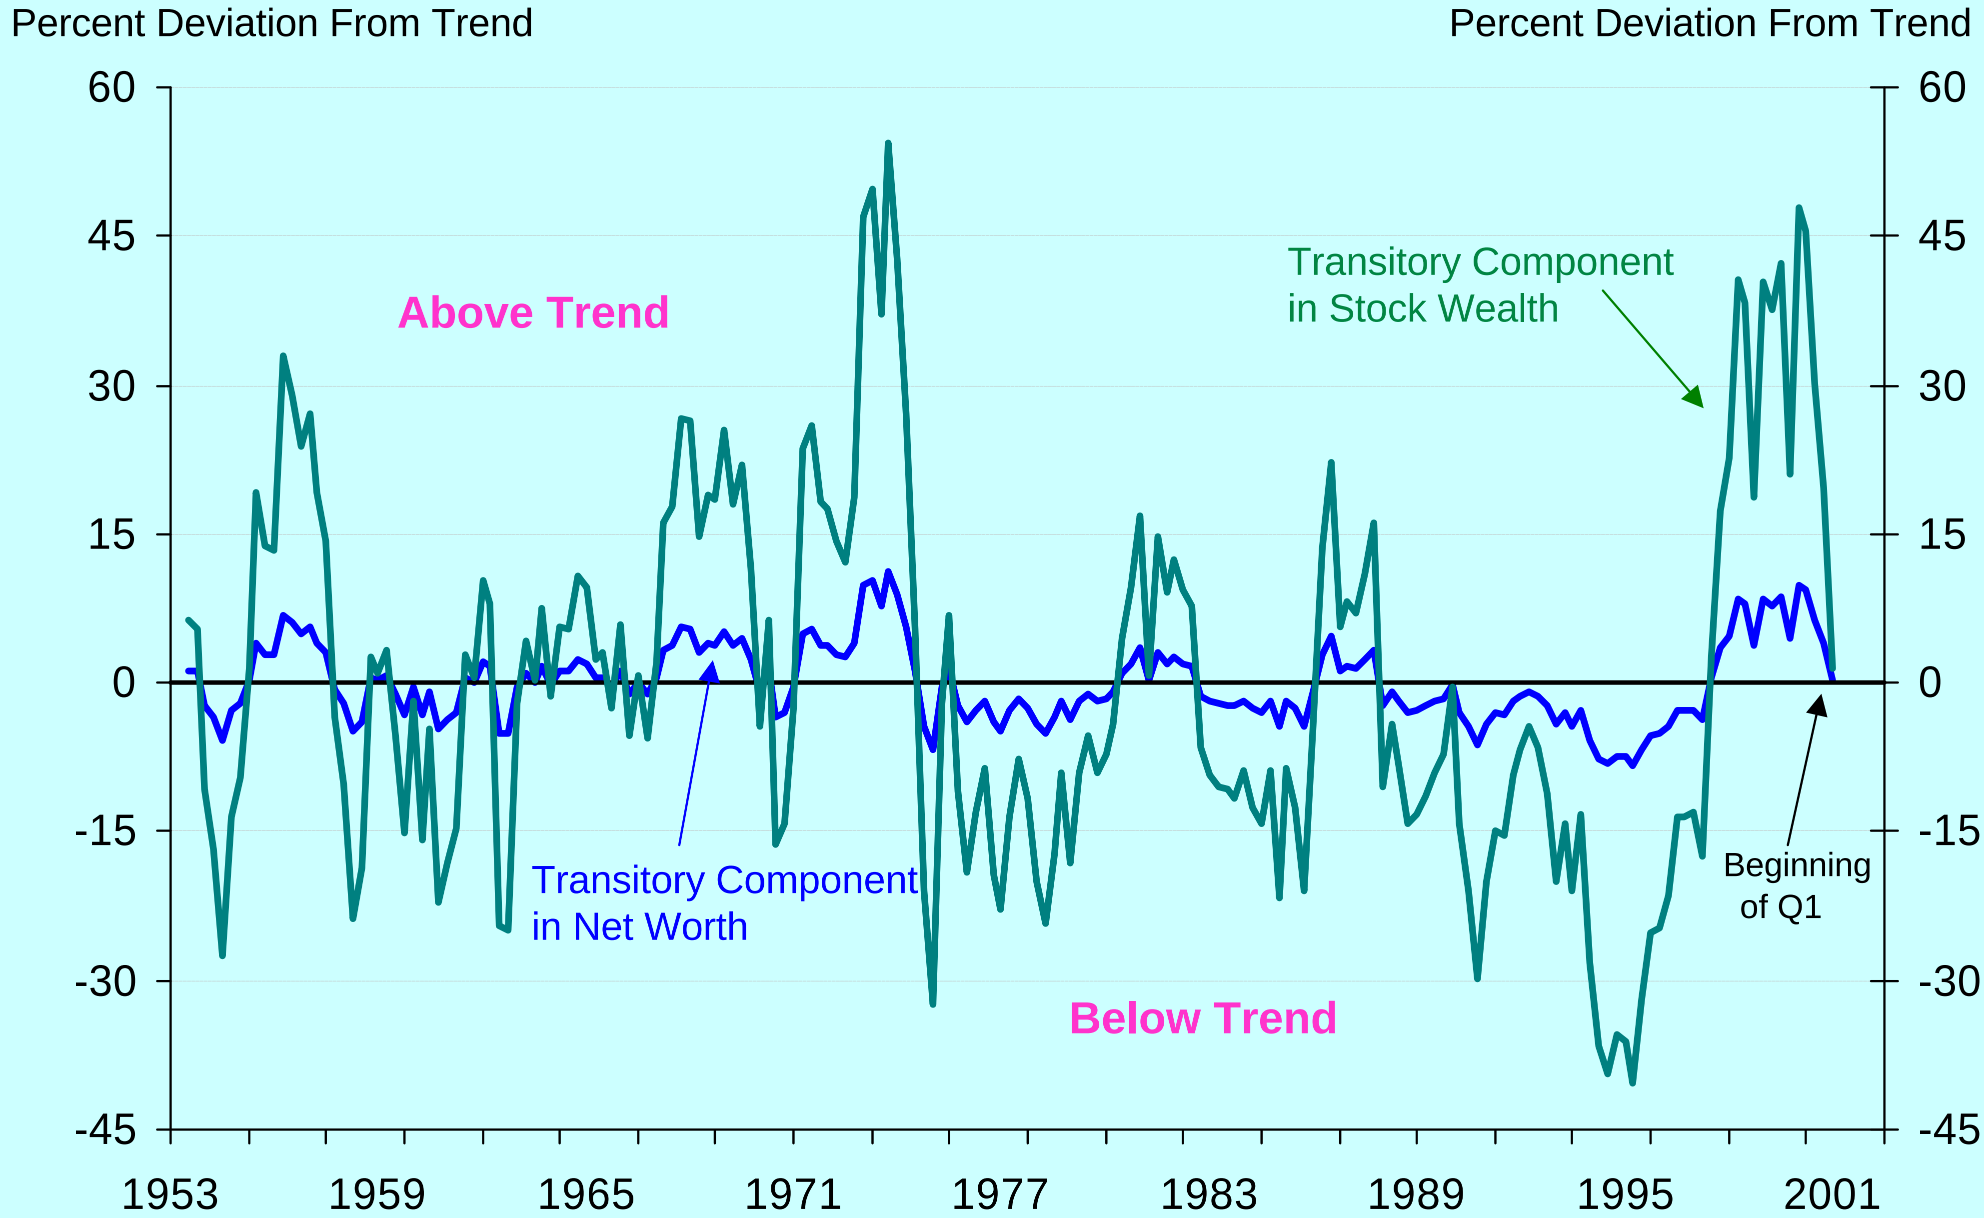

Note: The transitory component of wealth is defined as the difference between the actual value and the random walk component. The transitory component of stock wealth is the transitory component of wealth divided by the share of stock wealth in asset wealth in the sample, 0.208.

Figure C.1: Cumulative Distribution Functions of P/T Variance Decomposition

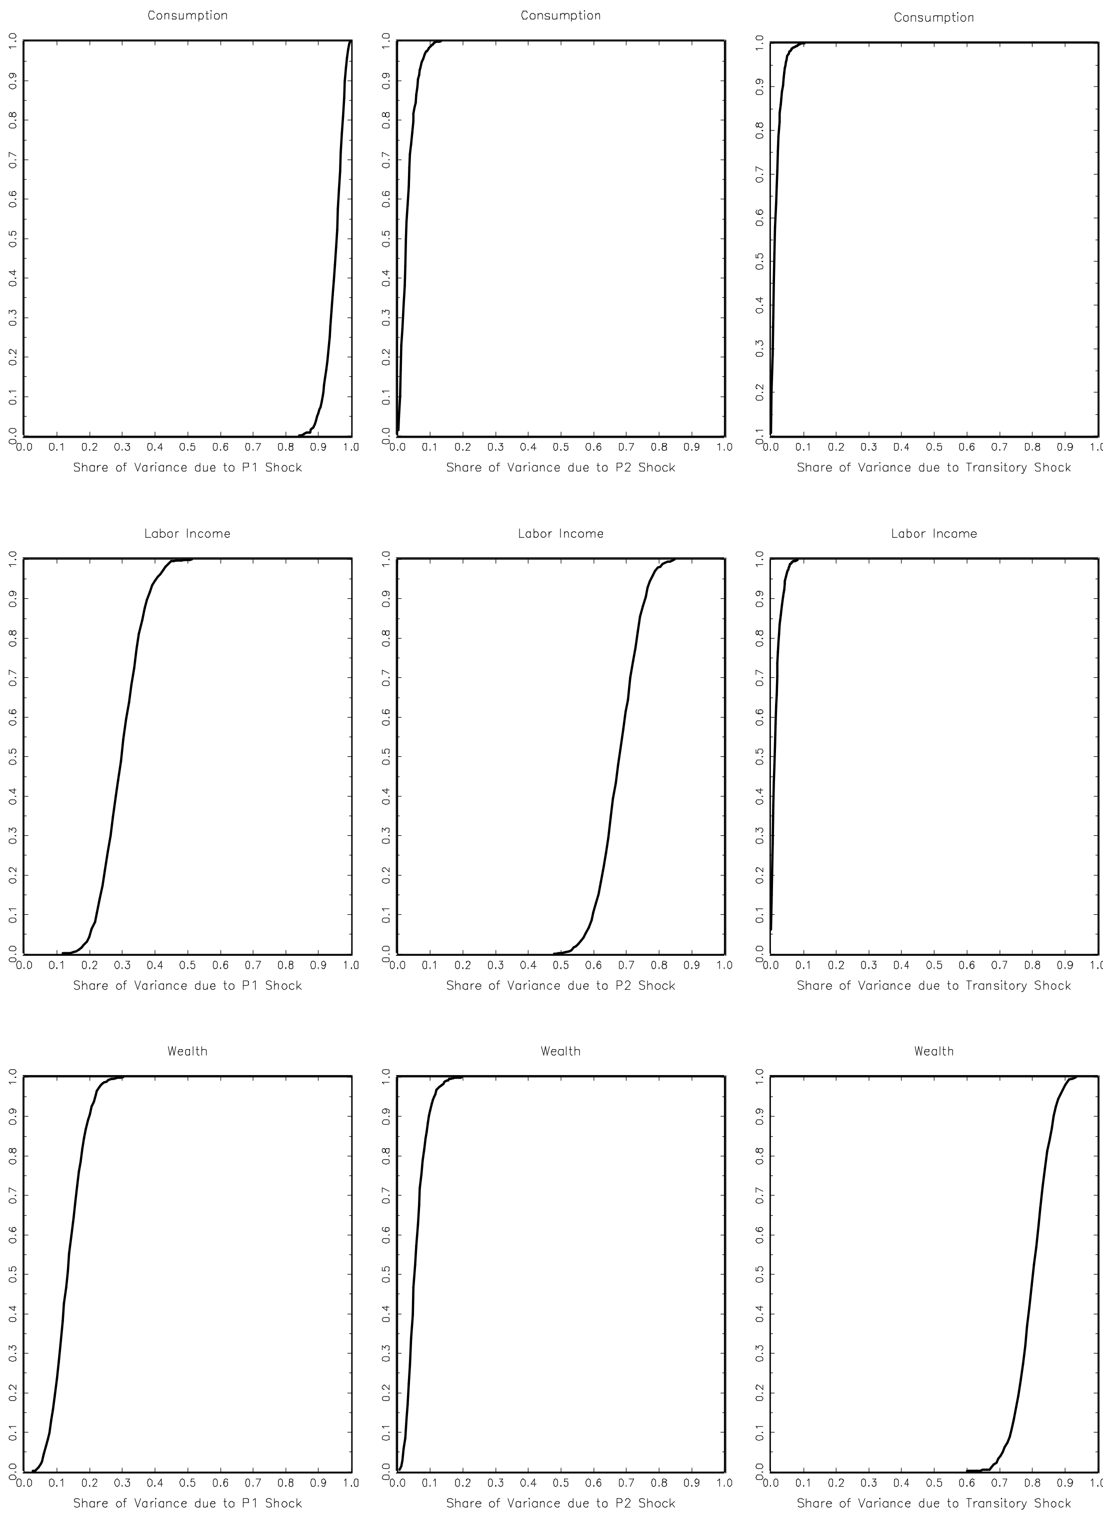

Note: This figure plots the cumulative distribution function of the P/T variance decomposition computed from a bootstrap procedure with 5,000 replications. In each simulation, the γ coefficients in the VECM are set to zero if they are insignificant at the 5% level.

Figure C.2: Impulse Response for recursive VAR

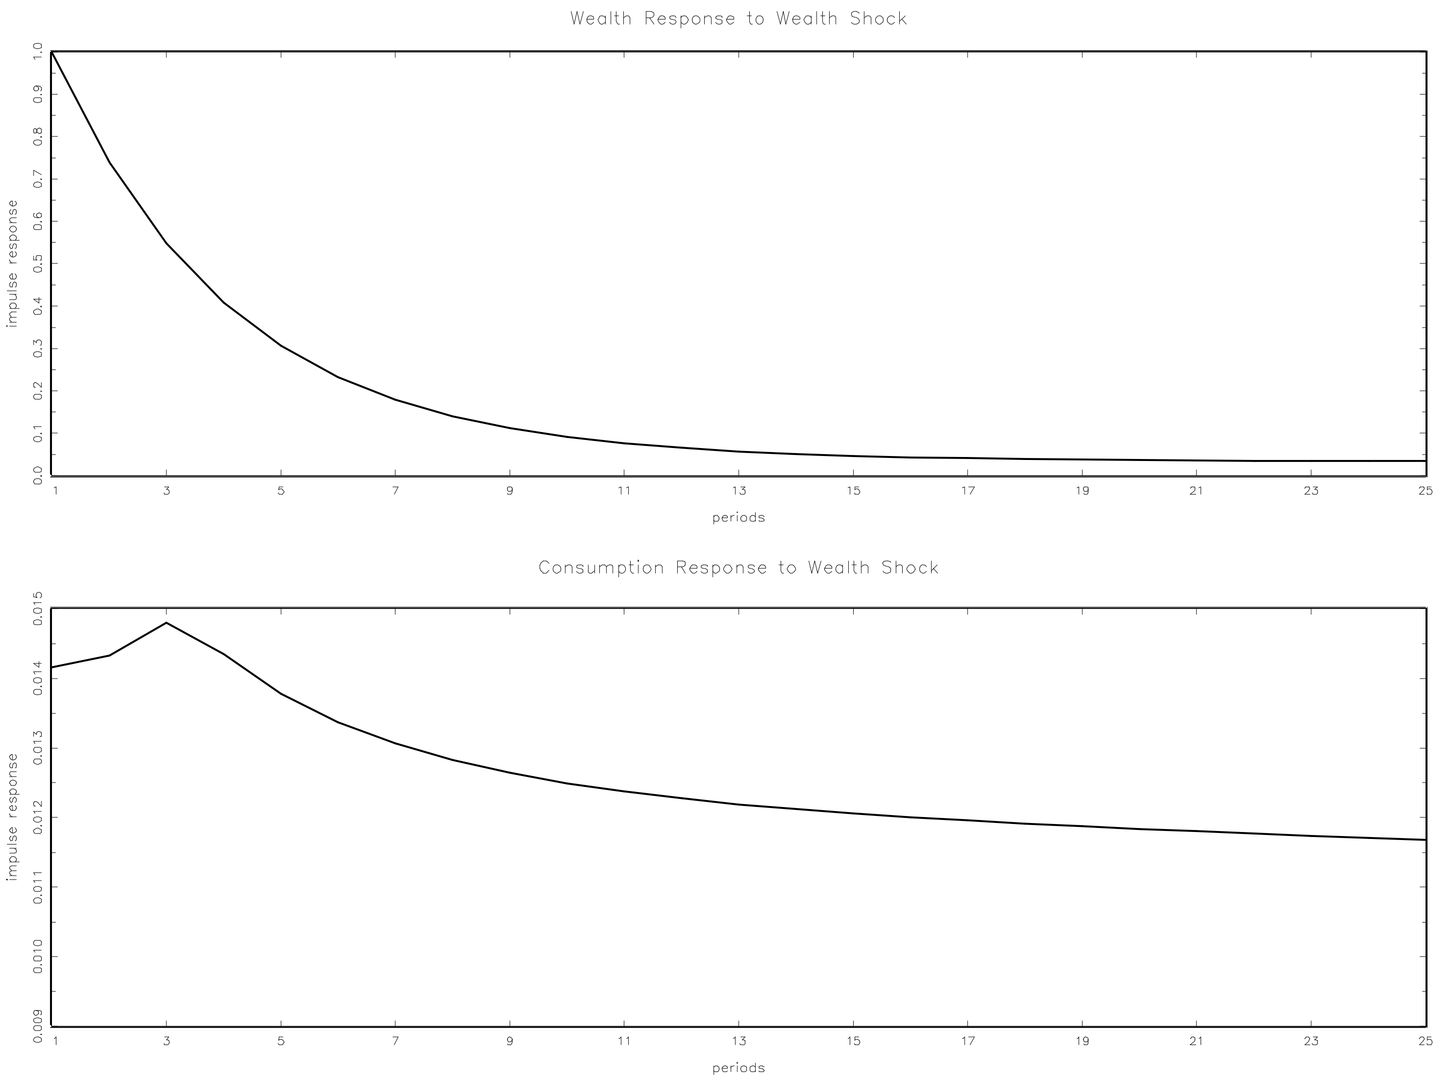

Note: This figure plots the impulse response function to a shock in a_t in a simple recursive VAR with the ordering (y_t, a_t, c_t) using simulated data. The data is generated using the structural VECM (12) with $\phi_{cc} = \phi_{cy} = 0$ and $\phi_{ca} = 0.2$. The simulated data for y and c is time-averaged using 20 observation per period. The sample size of the simulation is 50,000 and the Monte Carlo simulations are run 100 times.