Available online at www.sciencedirect.com
SCIENCE DIRECT®
Review of Economic Dynamics 9 (2006) 34–51
Review of Economic Dynamics
www.elsevier.com/locate/red

Empirical proxies for the consumption–wealth ratio ☆

Jeremy Rudd [a,*], Karl Whelan [b]

[a] *Board of Governors of the Federal Reserve System, 20th and C Streets NW, Washington, DC 20551, USA*
[b] *Central Bank and Financial Services Authority of Ireland, P.O. Box 559, Dame Street, Dublin 2, Ireland*

Received 26 June 2003; revised 12 August 2005

Available online 4 October 2005

Abstract

Using a log-linearized approximation to an aggregate budget constraint, it is possible to show that the ratio of consumption to total (human and non-human) wealth summarizes agents' expectations concerning both future labor income and future asset returns. In a series of recent papers, Lettau and Ludvigson construct an empirical analogue to the consumption–wealth ratio by approximating total wealth with a linear combination of labor income and observable non-human wealth. If valid, this framework suggests that consumption, assets, and labor income will be cointegrated. We demonstrate, however, that standard tests fail to reject the hypothesis of no cointegration once one employs measures of consumption, assets, and labor income that are jointly consistent with an underlying budget constraint. We also show that deviations of consumption, assets, and income from an estimated common trend are unable to predict future excess returns on stocks out of sample once theoretically consistent measures are used.
Crown Copyright © 2005 Published by Elsevier Inc. All rights reserved.

JEL classification: E21

Keywords: Budget constraint; Return forecastability; Cointegration; *cay*

1. Introduction

In a number of important studies, John Campbell has shown how a log-linearized version of the household budget constraint can be used to obtain the prediction that the ratio of consumption to total household wealth (defined as the sum of observable tangible wealth and unobservable

☆ The views expressed are those of the authors and do not necessarily reflect the views of the Board of Governors, the staff of the Federal Reserve System, or the Central Bank of Ireland.

* Corresponding author. Fax: (202) 872 4927.
E-mail addresses: jeremy.b.rudd@frb.gov (J. Rudd), karl.whelan@centralbank.ie (K. Whelan).

1094-2025/$ – see front matter Crown Copyright © 2005 Published by Elsevier Inc. All rights reserved.
doi:10.1016/j.red.2005.08.003

human wealth) is determined by the path of expected future returns on total wealth.[1] Thus, consumer behavior can, in principle, not only reveal expectations of future labor income—as has been stressed in traditional implementations of the permanent income hypothesis—but also expectations of future returns on financial assets. While this observation has potentially important implications for both macroeconomics and finance, an obvious difficulty with implementing this approach empirically in order to assess theories concerning expected asset returns reflects the fact that human wealth (the discounted sum of expected future labor income) cannot be directly observed.

In a series of recent papers, Martin Lettau and Sydney Ludvigson (2001a, 2001b, 2002, 2004) have operationalized Campbell's framework by approximating the log of total wealth with a linear combination of the log of labor income and the log of observable tangible assets. Under the joint assumptions that the approximation is adequate and that consumption growth and the returns on human and asset wealth are stationary, this approach predicts the existence of a cointegrating relationship among the log-levels of consumption, assets, and labor income, in turn implying that the deviation of these variables from their common trend should forecast at least one of the growth rates of these series. Lettau and Ludvigson present empirical evidence against the hypothesis of no cointegration and in favor of the existence of a single cointegrating relationship. In addition, Lettau and Ludvigson (2001a) show that the residual associated with this relationship has predictive power for excess stock returns—a result that they suggest is consistent with the proposition that this residual does indeed summarize information about expected future asset returns—while Lettau and Ludvigson (2001b, 2002) use this residual as a direct proxy for expected asset returns in order to construct models of cross-sectional asset prices and capital investment.

In this paper, we re-examine two aspects of Lettau and Ludvigson's results. First, we consider whether the theoretical cointegrating relationship that is suggested by their framework actually exists. We conclude that, on statistical grounds, there is no reason to reject the hypothesis that cointegration is absent in postwar US data. The difference between Lettau and Ludvigson's assessment of the evidence and our own stems from our use of different empirical measures of real consumption, income, and assets. The measures of these variables that we use are defined so as to be consistent with the household budget constraint that underpins Campbell's theoretical framework. In contrast, we show that Lettau and Ludvigson's measure of real consumption—real outlays on nondurables and services excluding shoes and clothing—is not consistent with a budget constraint that includes their measures of real income and wealth, which were obtained by deflating nominal income and wealth by a price index for total consumption expenditures. Such a choice of variables—which appears to be informed by previous attempts to test theories of consumer behavior—is not appropriate in this context, where the underlying theoretical relationship does not depend on a specific theory of consumer behavior, but rather on an intertemporal budget constraint.

These results call into question whether Lettau and Ludvigson's log-linear approximation provides a sufficiently accurate characterization of the underlying aggregate budget constraint. Moreover, they suggest that it may not be correct to interpret their estimated linear combination of consumption, income, and assets as providing an appropriate empirical analogue to the aggregate consumption–wealth ratio.

[1] See Campbell and Mankiw (1989), Campbell (1993, 1996).

Absence of cointegration also implies that it is inappropriate to use the resulting proxy for the consumption–wealth ratio as a predictor of asset returns. Nevertheless, given the amount of attention that this aspect of Lettau and Ludvigson's analysis has garnered, it is of interest to consider whether these measurement issues matter in the context of returns predictability. We find that our empirical proxy for the consumption–wealth ratio remains a statistically significant *in-sample* predictor of the excess return on equities. However, when we use the framework advocated by Goyal and Welch (2004) to assess whether this proxy can be employed to construct useful *out-of-sample* forecasts for excess returns, we find that it cannot; in particular, when theoretically consistent measures of consumption, income, and assets are used, the resulting proxies for the consumption–wealth ratio yield out-of-sample forecasts for excess stock returns that underperform forecasts based merely on the prevailing historical mean.

The remainder of the paper is organized as follows. Section 2 covers some issues associated with the derivation of the log-linearized budget constraint, while Section 3 discusses how to test the cointegration hypothesis that is implied by this relationship. Section 4 presents results from cointegration tests that use either theoretically consistent measures of consumption, assets, and income or the measures employed by Lettau and Ludvigson, while Section 5 examines the issue of returns predictability using our theoretically preferable measure. Finally, Section 6 summarizes some implications of our findings.

2. Theory

2.1. The budget constraint

To illustrate our points clearly, it is useful to explicitly consider how a real budget constraint for consumption and total wealth can be derived. Begin by defining total nominal household wealth, $\widetilde{W}_t$, as the sum of the current-dollar value of household assets, $\widetilde{A}_t$, and the current-dollar value of human capital, $\widetilde{H}_t$. (Here and elsewhere, we will use tildes to denote nominal variables.) The evolution of nominal wealth is described by the following budget constraint:

$$\widetilde{W}_{t+1} = (1 + I_{w,t+1})\big(\widetilde{W}_t - \widetilde{C}_t\big), \tag{1}$$

where $I_{w,t}$ denotes the nominal rate of return on wealth. (Note that this equation differs from the usual nominal budget constraint in not featuring labor income; instead, this has been defined as a component of the return on $\widetilde{W}_t$.)

It is important to keep in mind that if any household expenditure is counted as adding to the nominal wealth measure $\widetilde{W}_t$, then it cannot also be considered "consumption" from the point of view of this budget constraint (in other words, it cannot be treated as a component of $\widetilde{C}_t$). Although somewhat obvious, this point is important in that it determines the consumption and asset measures that we should select when testing hypotheses derived directly from the budget constraint. For example, if one is using a measure of assets that includes the value of household durable goods, then expenditures on durables should not be included in the series on outlays used to measure $\widetilde{C}_t$. In contrast, if the measure of assets excludes consumer durables, then internal consistency requires that expenditures on these goods be included in $\widetilde{C}_t$.

To re-express the budget constraint in terms of real consumption, we need to divide both sides of the relation by the deflator for our consumption measure, P^C_{t+1}. This gives

$$\frac{\widetilde{W}_{t+1}}{P^C_{t+1}} = \frac{P^C_t(1 + I_{w,t+1})}{P^C_{t+1}} \cdot \frac{\widetilde{W}_t}{P^C_t} - \frac{P^C_t(1 + I_{w,t+1})}{P^C_{t+1}} \cdot C_t, \tag{2}$$

where C_t denotes real consumption. Defining real wealth, inflation, and the real rate of return on wealth by

$$W_t = \frac{\widetilde{W}_t}{P_t^C}, \qquad \frac{P_{t+1}^C}{P_t^C} = 1 + \pi_{t+1}, \qquad R_{w,t+1} = \frac{1 + I_{w,t+1}}{1 + \pi_{t+1}} - 1, \tag{3}$$

yields the following representation of the budget constraint in terms of real variables:

$$W_{t+1} = (1 + R_{w,t+1})(W_t - C_t). \tag{4}$$

This equation is the starting point for the analysis of the consumption–wealth relation. What should be noted about it at this stage is that real wealth has been defined using the same deflator that was used to construct real consumption.

2.2. The log-linear approximation

Following Campbell and Mankiw (1989), Lettau and Ludvigson log-linearize Eq. (4) about a stationary steady state to obtain

$$\Delta w_{t+1} \approx r_{w,t+1} + \left(1 - \rho_w^{-1}\right)(c_t - w_t), \tag{5}$$

where r denotes the continuously compounded return $\ln(1 + R)$. (Here and elsewhere, log variables are denoted with lowercase letters and constants of linearization are ignored.) The parameter ρ_w is the steady state or average ratio of invested wealth $W - C$ to total wealth W; hence, $\rho_w < 1$. Using the identity

$$\Delta w_{t+1} = \Delta c_{t+1} + (c_t - w_t) - (c_{t+1} - w_{t+1}),$$

and imposing the condition that $\lim_{i\to\infty} \rho_w^{-i}(c_{t+i} - w_{t+i}) = 0$ allows Eq. (5) to be solved forward to yield

$$c_t - w_t \approx \sum_{i=1}^{\infty} \rho_w^i (r_{w,t+i} - \Delta c_{t+i}). \tag{6}$$

This equation holds ex post, but it should also hold if we replace actual future values with ex ante rational expectations. Taking the mathematical expectation of Eq. (6) conditional on time-t information therefore yields the following expression for the consumption–wealth ratio:

$$c_t - w_t \approx E_t \sum_{i=1}^{\infty} \rho_w^i (r_{w,t+i} - \Delta c_{t+i}). \tag{7}$$

Because aggregate wealth W_t is unobservable, Lettau and Ludvigson employ the following relations in order to further modify Eq. (7). First, they approximate the log of aggregate wealth as

$$w_t \approx \omega a_t + (1 - \omega) h_t, \tag{8}$$

where ω is the average share of asset holdings A in total wealth W. Second, the log return on aggregate wealth, $r_{w,t}$, is approximated by a weighted sum of the return on assets $r_{a,t}$ and the return on human capital $r_{h,t}$:

$$r_{w,t} \approx \omega r_{a,t} + (1 - \omega) r_{h,t}. \tag{9}$$

Finally, the nonstationary component of human capital is assumed to be captured by aggregate labor income Y_t, such that

$$h_t = \mu + y_t + z_t, \tag{10}$$

where μ is a parameter and z_t denotes a stationary zero-mean variable. (As with the definition of total real wealth, for this equation to be consistent with the underlying budget constraint, real labor income needs to be defined as nominal labor income divided by the same deflator that is used to construct real consumption.)

Putting these pieces together yields the following expression:

$$c_t - \omega a_t - (1-\omega)y_t \approx E_t \sum_{i=1}^{\infty} \rho_w^i \left[\omega r_{a,t+i} + (1-\omega) r_{h,t+i} - \Delta c_{t+i}\right] + (1-\omega) z_t, \tag{11}$$

which forms the foundation for Lettau and Ludvigson's analysis. In particular, they argue that the right-hand side of Eq. (11) is comprised of stationary variables; hence, the left-hand side of the equation should be stationary as well. It is this observation that serves as the theoretical basis for their hypothesis of a cointegrating relationship among log consumption, assets, and labor income. Moreover, if this equation is correct, it implies that if there are predictable and anticipated fluctuations in the rate of return on assets, $r_{a,t+i}$, then deviations of c_t, a_t, and y_t from their common trend should help to forecast these fluctuations.

It is useful at this point to briefly summarize the relations that underlie the derivation of an observable log-linear approximation to the aggregate budget constraint, and to identify what might affect the accuracy of the approximation. Our point of departure was the Campbell–Mankiw log-linear approximation to the budget constraint (Eq. (7)), the accuracy of which depends on the stability over time of the ratio of consumption to (unobserved) total wealth.[2] In addition, the derivation employs approximations to total wealth and the rate of return on total wealth (Eqs. (8) and (9)), whose accuracy in turn depends on the stability of the share of assets in total wealth as well as on an approximation for the level of human capital that requires the ratio of labor income to human capital to be stationary (Eq. (10)). Finally, substituting these various approximations into Eq. (7) yields an expression (Eq. (11)) that implies a cointegrating relationship among c_t, a_t, and y_t so long as $r_{a,t+i}$, $r_{h,t+i}$, Δc_{t+i}, and z_t are themselves stationary.

3. Empirical implementation

We now consider how we can construct empirical measures of c_t, a_t, and y_t that would permit us to explicitly test the hypothesis of cointegration among these variables that is suggested by Eq. (11).

Lettau and Ludvigson's empirical tests of this hypothesis defined consumption as real consumption of nondurables and services excluding shoes and clothing, with real assets and real labor income defined by dividing their nominal counterparts by the deflator for total consumption expenditures. Nominal assets were defined as total household net worth from the Flow of Funds accounts; this measure includes the value of the stock of consumer durables.

An immediate conclusion that follows from the preceding analysis is that this cointegration hypothesis cannot be derived directly from the aggregate budget constraint. This is because the

[2] Campbell (1993) discusses this approximation in detail, and notes that it will fail to be accurate when agents' intertemporal elasticity of substitution is high.

income and asset measures that were used were not defined by deflating their nominal counterparts by the price index for the measure of consumption that was employed: Real assets and income were defined relative to the deflator for total consumption expenditures, not the deflator for nondurables and services excluding shoes and clothing.[3]

Is there a way to justify the joint use of these measures of consumption, assets, and income? Lettau and Ludvigson (2001a) note that in using consumption of nondurables and services excluding shoes and clothing, they are "following in a tradition" set by previous studies such as Blinder and Deaton (1985). These studies employed this measure because the theories of consumer behavior that they sought to test applied to the *flow* of consumption enjoyed by consumers; expenditures on durable goods, by contrast, "are not a part of this flow because they represent replacements and additions to a stock, rather than a service flow from the existing stock."[4] This argument correctly characterizes the rationale for using this consumption series when testing behavioral relationships derived from a utility-maximization problem. However, this issue is not relevant in the context we are considering here: No theory of consumer behavior—for example, in the form of a consumption Euler equation—needed to be invoked in order to derive Eq. (11).

One potential justification for Lettau and Ludvigson's empirical approach is that their consumption variable serves as a *proxy* for another consumption variable that does belong in the same budget constraint as their measures of income and wealth; such an approach may then allow the cointegrating hypothesis to be derived *indirectly* from the aggregate budget constraint. Indeed, Lettau and Ludvigson observe that the ratio of the log of total real consumption expenditures to the log of real nondurables and services consumption "appears to have exhibited little secular movement" over their sample period, which, if true, would permit the use of real nondurables and services consumption in place of real total consumption expenditures in testing for cointegration.[5]

However, Fig. 1 shows that this statement is not correct in that this ratio has exhibited a distinct upward trend over the postwar period; more formally, a regression of this ratio on calendar time and two own lags (to correct for residual autocorrelation) yields a t-statistic of 2.53 for the trend term's coefficient. Of course, rather than invoking the existence of a stable ratio, one could justify the cointegration hypothesis being considered here if there were a cointegrating relationship between log real total consumption expenditures and log real nondurables and services consumption. Specifically, a cointegrating relation of the form

$$c_t = \delta + \beta c_t^{nd} + v_t, \tag{12}$$

where c_t^{nd} denotes (log) real consumption of nondurables and services, would permit us to substitute for c_t in the left-hand portion of Eq. (11). This yields an alternative cointegrating relation of the form

$$c_t^{nd} - \beta^{-1}\omega a_t - \beta^{-1}(1-\omega)y_t \sim I(0), \tag{13}$$

which is the version tested by Lettau and Ludvigson.

In practice, however, we are unable to reject the hypothesis of no cointegration between c_t and c_t^{nd}. Table 1 gives the estimated t-statistics from applying an augmented Dickey–Fuller test

[3] Even if the same deflator had been used to define each of these real variables, the measure of consumption would still be inconsistent with the measure of assets because shoes and clothing are excluded from C_t even though the value of households' stocks of shoes and clothing is not itself included in the Flow of Funds measure of net worth.

[4] Lettau and Ludvigson (2001a, p. 822).

[5] The statement regarding the empirical properties of this ratio appears in Lettau and Ludvigson (2001a, p. 822, fn. 6).

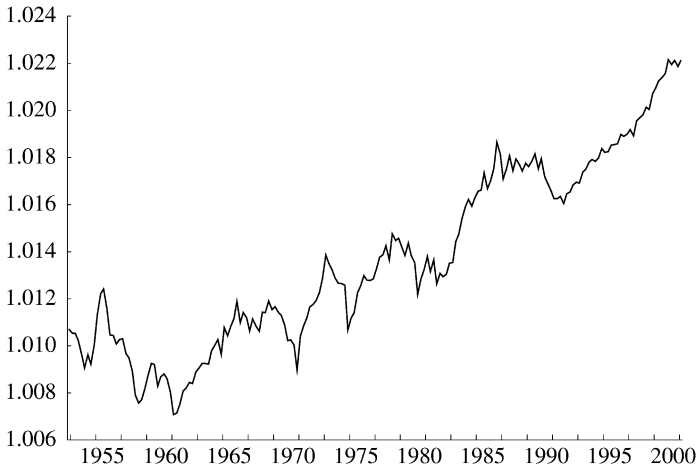


Fig. 1. Ratio of log real total consumption to log real nondurables and services consumption.

to the fitted residuals $\hat{v}_t$ from Eq. (12). The columns of the table are numbered from one to four; this corresponds to the number of lags of $\Delta\hat{v}_t$ that are used in the test regressions. The table reports results for two sample periods; the shorter sample (which runs from 1952:Q4 to 1998:Q3) corresponds to the dates used in Lettau and Ludvigson (2001a), while the second period extends the sample to the start of 2001. The five and 10 percent critical values for the test statistics are given as memo items in the table; they equal -3.43 and -3.13, respectively.[6] For either sample, it is evident that we are unable to reject the null hypothesis that cointegration is absent between total consumption expenditures and consumption of nondurables and services.[7]

An alternative way to justify Lettau and Ludvigson's approach is to assume an underlying budget constraint in which consumption is measured as the total flow of consumption inclusive of the service flow obtained from consumer durables. Although this service flow does not constitute a direct financial drain on asset accumulation, one can formulate a consistent budget constraint with C_t defined as this total flow measure of consumption so long as the rate of return $R_{w,t}$ is understood to include the implicit return from owning durables.[8] If consumption of nondurables and services were approximately proportional to this flow consumption measure, then its use could perhaps be justified as allowing us to approximate this alternative concept of the budget constraint.

There are two problems with this line of reasoning, however. The first is that there is little justification for assuming that consumption of nondurables and services is proportional to the (unobserved) series on total real flow consumption. As noted above, total real consumption expenditures have consistently grown faster than real consumption of nondurables and services. Moreover, standard proxies for the real service flow associated with consumer durables—such

[6] See Hamilton (1994, Table B.6) (note that these critical values are applicable for the case where the regressors in the cointegration model exhibit drift).

[7] The "general-to-specific" method of Campbell and Perron (1991) suggests that all lags of $\Delta\hat{v}_t$ should be excluded from the test equation. This yields t-statistics of -2.83 for the shorter period and -2.48 for the longer sample, which also fail to reject the null of no cointegration.

[8] In terms of accounting logic, the budget equality is maintained by adding an equal and offsetting adjustment on the income side (viz., the imputed rental income from durables ownership).

Table 1
Residual-based tests for cointegration of consumption measures

	Lag length			
	1	2	3	4
I. Sample: 1952:Q4 to 1998:Q3	−2.64	−2.86	−2.63	−2.36
II. Sample: 1952:Q4 to 2001:Q1	−2.27	−2.50	−2.26	−1.99
Memo:				
5 percent critical value	−3.43			
10 percent critical value	−3.13			

Notes: Figures are t-statistics for $\hat{\alpha}$ in regressions of the form $\Delta\hat{v}_t = \alpha\hat{v}_{t-1} + A(L)\Delta\hat{v}_{t-1}$, where $\hat{v}_t$ denotes the residual from a regression of log total consumption expenditures on log nondurables and services consumption. "Lag length" gives the number of lags of $\Delta\hat{v}_t$ used in the test regression. Critical values assume trending regressors.

as those used by Jorgenson and Stiroh (2000)—have grown even faster than real durable expenditures over this period. It is therefore likely that total real consumption inclusive of the service flow from consumer durables has grown even faster than total real consumption expenditure over this period, which would suggest an even steeper trend in the (unobserved) ratio of total real flow consumption to real nondurables and services consumption.

The second problem relates to deflation. According to the derivations presented earlier, even if the assumed proportionality between the log consumption measures did hold, the correct cointegrating relationship in this case would involve real income and assets defined relative to the (unobserved) price deflator for total flow consumption, *not* the price deflator for total consumption expenditures. Thus, this empirical approach does not correctly implement the log-linearized real budget constraint that underpins the analysis.

Taken together, these considerations imply that the particular cointegration hypothesis that Lettau and Ludvigson test empirically cannot be viewed as consistent with their theoretical framework. However, the preceding analysis also suggests a simple alternative methodology. Consider a budget constraint in which C_t is defined as total real consumption expenditures. Starting from such a constraint, one can follow the steps outlined above to generate a prediction of cointegration that involves the log of this series and the log of nominal income and assets defined relative to the deflator for total consumption outlays (although the asset measure in this case should not include the value of stocks of consumer durables). With the exception of this slight adjustment to the definition of assets, these are the same measures of a_t and y_t used by Lettau and Ludvigson.

We emphasize that it is the choice of deflation—rather than the difference in the specific definition of assets—that is most important here. Because the price of durable goods relative to the price of total consumption expenditures has declined rapidly over the postwar period, the relative price of nondurables and services has tended to rise over time (see Fig. 2).[9] In effect, therefore, a dollar's worth of real assets or income—when these measures are defined relative to the price index for *total* consumption outlay—is able to purchase a smaller and smaller amount

[9] Statistical tests indicate that there is strong evidence of a unit root in the relative price term, and no evidence of cointegration between the log price index for total consumption expenditures and the log price of nondurables and services consumption.

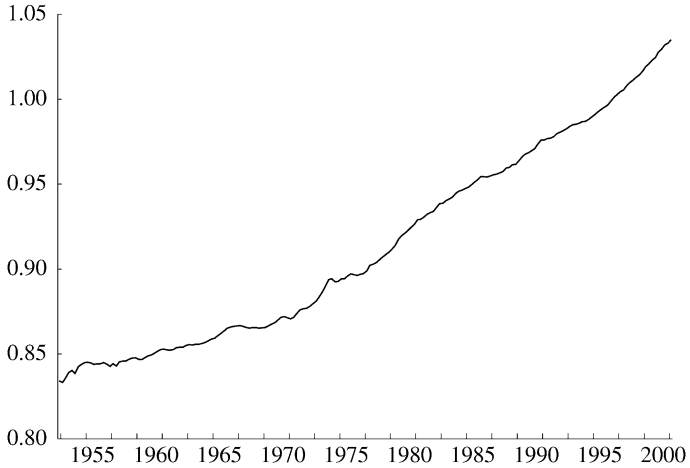


Fig. 2. Relative price of nondurables and services consumption.

of real nondurables and services consumption, and it is the presence of this wedge that leads to empirical inconsistencies in the consumer's budget constraint.[10]

In closing this discussion, we would note that the alternative approach for implementing the log-linearized aggregate budget constraint that we have outlined here does not require us to make any of the almost certainly inaccurate assumptions about the relationships between observable and unobservable measures of consumption that are required under Lettau and Ludvigson's approach. Moreover, it is worth keeping in mind that there is no theoretical requirement for us to adopt a budget constraint that features the (unobservable) total service flow from consumption. In fact, the $r_{a,t}$ concept associated with our approach does not include the unobservable implicit rental rate for consumer durables, and so is closer to the kinds of financial asset returns that have been related to the deviations of c_t, a_t, and y_t from their hypothesized common trend.

4. Cointegration tests

Ultimately, our choice of consumption and asset measures is only of interest if it significantly affects the statistical properties of the resulting empirical proxy for the consumption–wealth ratio. In this section, we document a crucial way in which this choice matters; specifically, how the results of tests for the cointegration of c_t, a_t, and y_t depend on the particular empirical measures we employ.

4.1. Residual-based tests

Table 2 presents the values of the t-statistics that we obtain from applying augmented Dickey–Fuller tests to the fitted residuals $\hat{u}_t$ from regressions of log consumption on log assets and labor income. Results are presented both for our proposed measures of c_t and a_t (total real consumption expenditures, and nominal net worth excluding durables divided by the deflator for total con-

[10] That the implications of this fact can be important in other, related contexts is documented by Palumbo et al. (2002), who show that tests of the Permanent Income Hypothesis as well as empirical estimates of wealth effects can be significantly affected by inconsistent deflation of the variables in the budget constraint.

Table 2
Residual-based tests for cointegration of c_t, a_t, and y_t

	Lag length			
	1	2	3	4
I. Sample: 1952:Q4 to 1998:Q3				
A. Lettau–Ludvigson c_t and a_t measures	−4.10	−3.88	−3.61	−3.55
B. Our c_t and a_t measures	−2.90	−2.81	−2.75	−2.75
II. Sample: 1952:Q4 to 2001:Q1				
A. Lettau–Ludvigson c_t and a_t measures	−3.80	−3.54	−3.18	−3.11
B. Our c_t and a_t measures	−3.08	−2.93	−2.79	−2.80
Memo:				
5 percent critical value	−3.80			
10 percent critical value	−3.52			

Notes: Figures are t-statistics for $\hat{\alpha}$ in regressions of the form $\Delta\hat{u}_t = \alpha\hat{u}_{t-1} + A(L)\Delta\hat{u}_{t-1}$, where $\hat{u}_t$ denotes the residual from a regression of a log consumption measure on log labor income and a log wealth measure. "Lag length" gives the number of lags of $\Delta\hat{u}_t$ used in the test regression. Critical values assume trending regressors.

sumption expenditures) as well as the measures used by Lettau and Ludvigson (real consumption of nondurables and services excluding shoes and clothing, and total net worth divided by the deflator for total consumption expenditures). Both approaches use the same measure of y_t; namely, the log of nominal labor income divided by the total consumption expenditures deflator.[11]

As before, the columns of the table are numbered from one to four, corresponding to the number of lags of $\Delta\hat{u}_t$ that are used in the test regressions; in addition, the table reports results for the two sample periods considered in Table 1. The five and 10 percent critical values for the test statistics are given as memo items in the table; they equal −3.80 and −3.52, respectively.[12]

The results are broadly similar for each period; we therefore focus on panel II, which presents the test results from the longer sample. (Note that because the NIPA and Flow of Funds data have each undergone several rounds of revisions since Lettau and Ludvigson's data set was put together, the results reported here for the shorter sample will not exactly match those found in Lettau and Ludvigson, 2001a.) First consider the line II.A of the table, which uses Lettau and Ludvigson's measures of consumption and assets. Consistent with their findings, we are able to reject the null hypothesis of no cointegration at the five percent level when one lag of $\Delta\hat{u}_t$ is used in the test equation.[13] The picture changes markedly, however, when we test for cointegration using our preferred measures of consumption and assets. As the line II.B of the table indicates, we are unable to reject the null of no cointegration at conventional significance levels: The largest t-statistic (in absolute value) has a p-value that is greater than 20 percent.[14] Thus, when theoretically consistent measures of c_t, a_t, and y_t are employed, the results from these tests suggest that there is no reason to reject the hypothesis that there is no cointegrating relation among these variables.

[11] All variables are expressed in per capita terms; see the appendix for a complete description of the data.

[12] See Phillips and Ouliaris (1990, Table IIc) (note that these critical values are applicable for the case where the regressors in the cointegration model exhibit drift).

[13] The general-to-specific procedure suggests that all lags of $\Delta\hat{u}_t$ should be excluded from the test equation. Doing so yields a t-statistic of −4.02, which slightly strengthens the evidence in favor of cointegration.

[14] In this case, the general-to-specific procedure calls for one lag of $\Delta\hat{u}_t$ in the test equation.

4.2. Johansen trace test

In addition to the residual-based tests, we also present two sets of likelihood-based test statistics derived by Søren Johansen (1988, 1991). Table 3 reports the Johansen "trace" statistic, which tests the null hypothesis that the VAR system in c_t, a_t, and y_t contains no cointegrating relationship against the alternative hypothesis that one or more cointegrating vectors are present in the system. In constructing this test, we assume that the data are trending and that a constant is present in the cointegrating vector. As before, we consider two sample periods (both of which have the same effective starting date, 1954:Q1) and report results for various lag lengths in the underlying VAR. The table also reports the test's five and 10 percent critical values, which we obtained using the software described in MacKinnon et al. (1999).[15]

The results reveal that even when Lettau and Ludvigson's measures of consumption and assets are used (the lines I.A and II.A), the evidence against the null hypothesis of no cointegration is weak: For both samples, none of the estimated models rejects the null hypothesis at the five percent level, and only one of the models rejects the null at the 10 percent level. It is worth noting that this roughly matches the pattern reported in Lettau and Ludvigson (2001a), where none of the trace statistics were significant at the ten percent level. However, the appendix to that paper erroneously dismisses these results as irrelevant. Specifically, the paper states that the alternative hypothesis being considered in this case is there are three cointegrating vectors; i.e., that all of the variables are stationary.[16] In fact, the alternative is that there are one or more cointegrating vectors. Thus, trace test results that fail to reject the absence of cointegration cannot simply be

Table 3
Johansen trace tests for cointegration of c_t, a_t, and y_t

	Lag length			
	1	2	3	4
I. Sample: 1954:Q1 to 1998:Q3				
A. Lettau–Ludvigson c_t and a_t measures	26.86	27.33	20.53	20.65
B. Our c_t and a_t measures	19.04	21.94	18.67	19.97
II. Sample: 1954:Q1 to 2001:Q1				
A. Lettau–Ludvigson c_t and a_t measures	28.60	25.67	19.81	18.62
B. Our c_t and a_t measures	22.44	22.15	19.07	18.98
Memo:				
5 percent critical value	29.80			
10 percent critical value	27.07			

Notes: The table reports tests of the null hypothesis of no cointegrating relationships against the alternative of one or more cointegrating vectors. "Lag length" gives the number of lags in the estimated VAR system. Critical values generated using the computer program described in MacKinnon et al. (1999); a test statistic greater than the specified critical value suggests rejection of the null of no cointegration.

[15] Critical values for the Johansen cointegration tests are typically estimated using numerical techniques. As a result, the specific critical values reported in various studies can vary slightly depending on the details of the numerical exercise. Note, however, that the MacKinnon, Haug, and Michelis critical values that we report in the table are very close to those obtained by other studies, such as Johansen and Juselius (1990) and Osterwald-Lenum (1992).

[16] See Lettau and Ludvigson (2001a, p. 847).

dismissed on the grounds that unit root tests suggest that consumption, income, and assets are all nonstationary, as is done in the paper. Rather, according to Johansen's (1995) recommended procedure of "sequential" trace tests, these statistics directly imply that we cannot reject the null of no cointegration.

Table 3 also shows that when our preferred consumption and asset measures are used, the evidence against the null of no cointegration becomes extremely weak: The values of the test statistics decline in every case (the line II.B), and now never imply rejection of the null (the largest trace statistic is not even significant at the 25 percent level).

4.3. Johansen L-max test

Finally, Table 4 reports the Johansen maximal eigenvalue (or "L-max") statistic, which tests the null hypothesis of no cointegrating relationship against the more precise alternative that exactly one cointegrating vector is present. When Lettau and Ludvigson's measures of consumption and assets are used, we find that we are able to reject the null hypothesis of no cointegration at the five percent level in the one- and two-lag systems (though not in the three- and four-lag systems). However, when our preferred measures are employed, we again find that the test statistics decline in every case, with none implying rejection of the null at the 10 percent level.

One aspect of Table 4 worth highlighting is that the 10 percent critical value (of 18.89) that we report for the L-max statistic differs substantially from the corresponding critical value of 13.39 that was reported in Appendix Table AII of Lettau and Ludvigson (2001a). If this latter value were the right one, it would imply that all of our reported L-max statistics reject the null at the 10 percent level, in contrast to our conclusion that none do so. However, the critical value of 13.39 is not correct. A 10 percent critical value for the L-max statistic of 13.39 is reported by the CATS cointegration module of the RATS econometric package; however, the procedure that CATS uses in order to calculate this figure contains a conceptual error. Specifically, CATS computes this "critical value" as the difference between the 10 percent critical values for the trace test statistics under the null hypotheses of zero and one cointegrating vectors. While the L-max

Table 4
Johansen L-max tests for cointegration of c_t, a_t, and y_t

	Lag length			
	1	2	3	4
I. Sample: 1954:Q1 to 1998:Q3				
A. Lettau–Ludvigson c_t and a_t measures	22.35	23.48	16.55	16.09
B. Our c_t and a_t measures	14.68	17.15	13.86	14.49
II. Sample: 1954:Q1 to 2001:Q1				
A. Lettau–Ludvigson c_t and a_t measures	23.98	21.84	16.05	14.12
B. Our c_t and a_t measures	17.61	17.40	14.29	13.79
Memo:				
5 percent critical value	21.13			
10 percent critical value	18.89			

Notes: The table reports tests of the null hypothesis of no cointegrating relationships against the alternative of one cointegrating vector. "Lag length" gives the number of lags in the estimated VAR system. Critical values generated using the computer program described in MacKinnon et al. (1999); a test statistic greater than the specified critical value suggests rejection of the null of no cointegration.

test statistic in this case does equal the difference between these two trace statistics, this does not imply that its 10 percent *critical value* can simply be calculated as the difference between the critical values for these two specifications of the test.[17]

5. Forecasts of excess returns

By itself, our failure to find robust evidence of cointegration among consumption, assets, and labor income is significant inasmuch as it casts doubt on our ability to construct a theoretically consistent empirical proxy for the aggregate consumption–wealth ratio. It is also directly problematic for applications of this theoretical framework to certain practical questions. For example, if there exists a single cointegrating vector among three $I(1)$ variables, then it is well known that there exists a so-called "common trends" representation in which each of the variables is a function of two common stochastic trends and a single transitory shock. Lettau and Ludvigson (2004) use such a representation to divide c_t, a_t, and y_t into their stochastic trend and cycle components, and to explore the implications of this representation for the estimation of wealth effects on consumption. Of course, if cointegration is not present among these variables, then such a representation cannot be constructed.

Another application of this framework that has received a large amount of attention is its use in constructing a predictor of the excess return on equities. As was noted in our discussion of Eq. (11), predictable and anticipated fluctuations in the expected rate of return on assets should be forecastable by estimated deviations of c_t, a_t, and y_t from their common trend. Lettau and Ludvigson (2001a) document that these deviations—which they label $\widehat{cay}_t$—carry significant predictive power for excess returns on stocks, especially at long horizons.

In this section, we consider whether this finding is robust to our choice of real consumption, asset, and income measures. We note at the outset that our inability to uncover robust evidence of cointegration for our preferred consumption–wealth proxy a priori renders it an inappropriate predictor for excess returns; as Ferson et al. (2003) have argued, to the extent that *expected* excess returns are persistent, using highly persistent regressors in forecasting models for excess returns will yield spurious results.[18] Moreover, this problem is likely exacerbated when longer-horizon excess returns are used, as these are by construction more persistent. Hence, finding convincing evidence of cointegration among the components of an empirical proxy for the consumption–wealth ratio is a precondition for employing it in a forecasting model for excess returns, and we should therefore be wary of results based on our theoretically consistent consumption–wealth proxy (for which the hypothesis of a unit root cannot be rejected).[19]

5.1. In-sample forecasts

Table 5 gives coefficients and adjusted R^2 measures from regressions of H-period excess returns $r^e_{t+1,H}$ on $\widehat{cay}_t$, where $r^e_{t+1,H}$ is defined as $r^e_{t+1} + \cdots + r^e_{t+H}$ (here, r^e_t denotes the one-

[17] More generally, the value of the 10 percent tail for the difference between *any* two random variables X and Y cannot simply be calculated as the difference between the ten percent tails of the distributions for X and Y.

[18] This point is echoed by Lettau and Ludvigson (2005, p. 22), who state that "…it makes no sense to forecast returns with a unit root variable."

[19] The differential persistence of the two $\widehat{cay}_t$ measures is also illustrated by their sample autocorrelations: Under Lettau and Ludvigson's definitions, the point estimates for the sample autocorrelations of $\widehat{cay}_t$ cross zero after twelve quarters; by contrast, it takes nearly four times as long for this to happen if theoretically consistent measures of consumption, wealth, and income are used.

Table 5
In-sample excess returns forecasts

	Forecast horizon (quarters)						
	1	4	8	12	16	20	24
A. Lettau–Ludvigson c_t and a_t measures							
Coefficient	1.97	5.74	8.72	11.49	11.82	14.84	17.72
	(0.47)	(1.67)	(2.30)	(2.70)	(3.03)	(3.89)	(4.83)
	[4.20]	[3.44]	[3.78]	[4.25]	[3.90]	[3.82]	[3.67]
$\bar{R}^2$	0.09	0.17	0.22	0.27	0.24	0.28	0.30
B. Our c_t and a_t measures							
Coefficient	1.25	3.91	6.39	8.32	9.11	10.81	12.74
	(0.32)	(1.08)	(1.50)	(1.85)	(2.29)	(2.84)	(3.35)
	[3.94]	[3.61]	[4.24]	[4.49]	[3.98]	[3.81]	[3.80]
$\bar{R}^2$	0.06	0.16	0.27	0.37	0.38	0.38	0.41

Notes: The table gives the coefficient and $\bar{R}^2$ from a *cay*-based forecasting regression for the excess return on equities fit over the full sample (1952:Q4 to 2001:Q1). "Forecast horizon" denotes the period over which the excess return is calculated; that is, a horizon of H periods implies that the dependent variable in the forecasting regression equals $r_{t+1} + \cdots + r_{t+H}$, where r_t denotes the excess return on equities at time t. Newey–West standard errors are given in parentheses; corresponding t-statistics are given in square brackets.

period excess return on equities at time t).[20] The regressions are fit over the full sample (1952:Q4 to 2001:Q1). Comparing the two panels of the table reveals that using theoretically consistent measures of consumption, wealth, and income to compute $\widehat{cay}_t$ (panel B) yields forecasting regressions that have slightly lower adjusted R^2 values over short horizons (one to four quarters) relative to regressions that use Lettau and Ludvigson's definition of $\widehat{cay}_t$ (panel A). By contrast, over longer horizons there is a large *improvement* in relative fit from using the theoretically consistent measures. This result is most likely spurious, and reflects the much higher persistence of this variant of $\widehat{cay}_t$.

5.2. Out-of-sample forecasts

In recent work, Goyal and Welch (2004) advocate assessing returns predictability in the context of *out-of-sample* forecasting models. By using data in a recursive manner, such exercises help to assess whether a given predictor would have provided useful excess returns forecasts from the perspective of a practitioner acting in real time.[21] We therefore now consider measures of forecast performance that are obtained from rolling, out-of-sample regressions of excess returns (again computed over various horizons) on a candidate predictor.

Our specific procedure, which closely follows Goyal and Welch's analysis, involves computing the root mean-square error (RMSE) from a regression of H-period excess returns $r^e_{t+1,H}$ on $\widehat{cay}_t$, where $r^e_{t+1,H}$ is defined as before. Importantly, these forecasting regressions are true out-

[20] The excess return series is the value-weighted CRSP market return less the Treasury bill rate, and is derived using the data set provided on Professor Kenneth French's website. Following Lettau and Ludvigson, we estimate $\widehat{cay}_t$ using a dynamic OLS specification with eight leads and lags.

[21] In addition, the out-of-sample performance of forecasting models for returns figures prominently in the exchange between Brennan and Xia (2005) and Lettau and Ludvigson (2005) regarding the usefulness of *cay*-based returns models.

of-sample exercises: They are constrained to only use data that are available through a specified period T, and the parameters used to construct $\widehat{cay}_t$ are estimated using data available at time T only. (The initial information set is defined to include data through T=1966:Q1—the full sample extends to 2001:Q1—and the sample begins in 1952:Q4.) As our benchmark, we compare the *cay*-based forecasts to those obtained from a naïve forecasting model in which the predicted value of the equity premium is assumed to equal the prevailing sample mean. Goyal and Welch emphasize that a number of well-known forecasting variables fail this seemingly weak test of returns forecastability (which is consistent with excess returns being i.i.d.). Our own results from this exercise are presented as follows: We compute the ratio of the RMSE from the *cay*-based forecasting regressions to the RMSE from the naïve specifications; if this ratio is less than one, then we conclude that the *cay*-based forecast is able to improve upon the mean-based forecast.

These ratios are reported in panel I of Table 6 for forecast horizons ranging from one quarter to six years. The line I.A of the table shows that when Lettau and Ludvigson's definitions of c_t, a_t, and y_t are used, the resulting empirical proxy for the consumption–wealth ratio yields out-of-sample excess-return forecasts that improve upon the naïve model's projections at horizons longer than two years. By contrast, when theoretically consistent measures of consumption, income, and wealth are employed (the line I.B), we are never able to improve upon the naïve forecast. Interestingly, the relative deterioration in forecast performance between the two models appears to increase with the forecast horizon; this is intuitive given that our choice of measurement affects a long-term property of the estimated consumption–wealth proxy (namely, whether cointegration is present or absent).

Recently, Campbell and Thompson (2004) have argued that the Goyal–Welch procedure does not provide a faithful representation of how real-world investors would employ forecasting models for equity returns. They point out that a practitioner would be unlikely to use a model whose coefficients ran counter to theory or a model that produced negative forecasts of excess returns (which would imply the existence of a negative equity premium), and so impose these a priori sign restrictions on their candidate models. Campbell and Thompson show that, once this is done, the forecasting performance of many variables that "fail" the Goyal–Welch test of returns predictability is improved considerably. We therefore repeat the out-of-sample prediction

Table 6
RMSE ratios from out-of-sample excess returns forecasts

	Forecast horizon (quarters)						
	1	4	8	12	16	20	24
I. Baseline models							
A. Lettau–Ludvigson c_t and a_t measures	1.02	1.07	1.01	0.92	0.90	0.89	0.85
B. Our c_t and a_t measures	1.03	1.07	1.14	1.24	1.22	1.20	1.21
II. With Campbell–Thompson restrictions							
A. Lettau–Ludvigson c_t and a_t measures	0.99	1.02	0.98	0.92	0.90	0.89	0.86
B. Our c_t and a_t measures	1.02	1.07	1.17	1.24	1.19	1.18	1.19

Notes: The table gives the ratio of the out-of-sample RMSE from a *cay*-based forecasting regression for the excess return on equities to the RMSE from a corresponding forecasting regression that uses the prevailing sample mean excess return. (A ratio greater than one therefore implies that the *cay*-based model does *worse* than the mean-based forecast, while a ratio less than one implies that it improves on the mean-based forecast.) "Forecast horizon" denotes the period over which the excess return is calculated; that is, a horizon of H periods implies that the dependent variable in the forecasting regression equals $r^e_{t+1} + \cdots + r^e_{t+H}$, where r^e_t denotes the excess return on equities at time t. The first forecast begins in 1966:Q1; the full sample period extends from 1952:Q4 to 2001:Q1. See text for additional details.

exercises using models in which the Campbell–Thompson sign and forecast restrictions are applied sequentially.[22] As can be seen from panel II of Table 6, however, imposing these restrictions leads to very little improvement in the forecasting performance of either model.

6. Interpreting the results

We have documented that one cannot reject the hypothesis that there is no cointegrating relationship among measures of log consumption, assets, and labor income that are mutually consistent with an underlying budget constraint. In addition, we have shown that deviations of these variables from an estimated common trend contain no out-of-sample predictive power for excess stock returns once theoretically consistent measures are employed.

A direct implication of these results is to weaken the theoretical and empirical case for the findings in Lettau and Ludvigson (2001a) that deviations of consumption, assets, and labor income from a common trend have forecasting power for excess stock returns. In response to Brennan and Xia's (2005) suggestion that this finding may represent a spurious relationship, Lettau and Ludvigson (2005) have argued that Eq. (11) provides a theoretical justification for their result, and that their evidence on cointegration supplies an empirical justification. In addition, Lettau and Ludvigson report robust evidence of an out-of-sample forecasting relation. However, our findings indicate that neither of these results—cointegration or out-of-sample excess-return predictability—are robust to the use of aggregate data that are compatible with the underlying theoretical framework.

It is worth asking why it is that we cannot reject the hypothesis that a cointegrating relationship among consumption, labor income, and assets is not present in US data. As Lettau and Ludvigson (2005) have pointed out, the budget constraint identity used to derive Eq. (11)—and, hence, the prediction of cointegration—"... must be a part of any sensible economic model," which suggests that our findings run counter to basic economic theory. However, we would argue that Eq. (11) does not necessarily imply that c_t, a_t, and y_t will be cointegrated. In particular, we can think of two possible explanations for why cointegration may be absent in practice.

The first possibility is that the expected return on human or asset wealth (or the growth rate of consumption) is not stationary. For example, this assumption (which was required in order to derive the prediction of a cointegrating relation among c_t, a_t, and y_t) could fail to hold if the economy undergoes periodic structural changes, such as shocks to trend productivity growth or demographic shifts. In this case, Eq. (11) may still be correct, but it does not follow that a stationary linear combination of c_t, a_t, and y_t will exist in all periods. If true, this possibility suggests that we will face a serious problem in implementing this framework empirically, given that attempts to identify ω based on a regression of c_t on a_t and y_t will suffer from the presence of $I(1)$ errors. Moreover, it seems likely that persistent shifts in expected returns on human or asset wealth would also lead to changes in ω (the average share of assets in total wealth). In practice, the existence of such breaks would make it very difficult to identify the relevant value for ω that holds over a given subperiod, because—as Lettau and Ludvigson (2005) have noted—samples smaller than the one used in their study will likely suffer from significant small-sample biases. (Note that the presence of structural breaks of this sort could also potentially explain why

[22] For Campbell and Thompson's set of indicators, the sequential restrictions yield the largest improvement in forecasting performance. We also considered what would happen if these restrictions were to be imposed separately; the results (not shown) are essentially identical to those reported here.

forecasting models for the excess return on equities manifest such different in- and out-of-sample performance when a theoretically consistent consumption–wealth proxy is used.)

The second possibility is simply that the underlying relationship described in Eq. (11) may do a poor job of capturing reality. As noted above, the derivation of this relationship relies on a host of approximations, each of which in turn relies on assumptions as to the stability over time of a number of unobservable variables. Any one of these assumptions could be inaccurate enough to render this equation an unsatisfactory framework, which in turn could cause its predictions—such as the cointegration of c_t, a_t, and y_t—to be rejected in the data.

Of course, either explanation of our findings raises important concerns regarding our ability to empirically implement a log-linearized approximation to the aggregate consumption–wealth relation—and, by extension, our ability to formulate empirical models (such as models of expected asset returns) that are based on this relation.

Acknowledgments

We are grateful to Martin Lettau and an anonymous referee for comments on earlier drafts of this paper.

Appendix A. Data sources and definitions

All consumption, wealth, and income variables are expressed in per capita terms using the population measure described below. Real wealth and income are deflated with the price index for total personal consumption expenditure. All data are current as of January 2002 and, at the time of this writing, represent the latest vintage of data for which the National Income and Product Accounts and Flow of Funds Accounts are mutually consistent.

Consumption expenditures. Total personal consumption expenditure is taken from the National Income and Product Accounts (NIPAs). Consumption of nondurables and services excluding clothing is computed by combining NIPA personal consumption expenditures on nondurable goods with NIPA personal consumption expenditures on services, then removing NIPA personal consumption expenditures on clothing and shoes. All real measures are combined or subtracted using a Fisher chain-aggregation formula that replicates the procedure used by the Bureau of Economic Analysis in producing the NIPAs.

Wealth. Data on household net worth and the value of household stocks of durable goods are taken from the Flow of Funds Accounts of the Board of Governors of the Federal Reserve System, Table B.100.[23] Flow of Funds wealth measures are expressed on an end-of-period basis; we therefore associate the $t-1$ value of the data with period t wealth (that is, with A_t) in order to obtain a start-of-period measure.

Labor income. We define labor income as wage and salary disbursements (NIPA Table 2.1, line 2) *plus* transfers to persons (line 16) *plus* other labor income (line 9) *minus* personal contributions for social insurance (line 23) *minus* labor taxes. Labor taxes are defined by imputing a share of personal tax and nontax payments (line 24) to labor income, with the share calculated

[23] The Flow of Funds measure of net worth contains an estimate of the value of owner-occupied housing.

as the ratio of wage and salary disbursements to the sum of wage and salary disbursements, proprietors' income (line 10), and rental (line 13), dividend (line 14), and interest (line 15) income.

Population. Population from NIPA Table 8.7, line 16. (Note that this is the population measure used by the Bureau of Economic Analysis to compute official per capita income and consumption data.)

References

Blinder, A.S., Deaton, A., 1985. The time series consumption function revisited. Brookings Papers on Economic Activity 2, 465–511.

Brennan, M.J., Xia, Y., 2005. *tay*'s as good as *cay*. Finance Research Letters 2, 1–14.

Campbell, J.Y., 1993. Intertemporal asset pricing without consumption data. American Economic Review 83, 487–512.

Campbell, J.Y., 1996. Understanding risk and return. Journal of Political Economy 104, 298–345.

Campbell, J.Y., Mankiw, N.G., 1989. Consumption, income, and interest rates: Reinterpreting the time series evidence. In: NBER Macroeconomics Annual, pp. 185–216.

Campbell, J.Y., Perron, P., 1991. Pitfalls and opportunities: What macroeconomists should know about unit roots. In: NBER Macroeconomics Annual, pp. 141–201.

Campbell, J.Y., Thompson, S.B., 2004. Predicting the equity premium out of sample: Can anything beat the historical average? Mimeo. Harvard University, October 8th draft.

Ferson, W.E., Sarkissian, S., Simin, T.T., 2003. Spurious regressions in financial economics? Journal of Finance 58, 1393–1413.

Goyal, A., Welch, I., 2004. A comprehensive look at the empirical performance of equity premium prediction. Working paper 10483. NBER.

Hamilton, J.D., 1994. Time Series Analysis. Princeton Univ. Press, Princeton, NJ.

Johansen, S., 1988. Statistical analysis of cointegrating vectors. Journal of Economic Dynamics and Control 12, 231–254.

Johansen, S., 1991. Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. Econometrica 59, 1551–1580.

Johansen, S., 1995. Likelihood-Based Inference in Cointegrated Vector Autoregressive Models. Oxford Univ. Press, Oxford, UK.

Johansen, S., Juselius, K., 1990. Maximum likelihood estimation and inference on cointegration—With applications to the demand for money. Oxford Bulletin of Economics and Statistics 52, 169–210.

Jorgenson, D., Stiroh, K., 2000. Raising the speed limit: US economic growth in the information age. Brookings Papers on Economic Activity 1, 125–211.

Lettau, M., Ludvigson, S., 2001a. Consumption, aggregate wealth, and expected stock returns. Journal of Finance 56, 815–849.

Lettau, M., Ludvigson, S., 2001b. Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying. Journal of Political Economy 109, 1238–1287.

Lettau, M., Ludvigson, S., 2002. Time-varying risk premia and the cost of capital. Journal of Monetary Economics 49, 31–66.

Lettau, M., Ludvigson, S., 2004. Understanding trend and cycle in asset values: Re-evaluating the wealth effect on consumption. American Economic Review 94, 276–299.

Lettau, M., Ludvigson, S., 2005. *tay*'s as good as *cay*: Reply. Finance Research Letters 2, 15–22.

MacKinnon, J.G., Haug, A.A., Michelis, L., 1999. Numerical distribution functions of likelihood ratio tests for cointegration. Journal of Applied Econometrics 14, 563–577.

Osterwald-Lenum, M., 1992. A note with quantiles of the asymptotic distribution of the maximum likelihood cointegration rank test statistics. Oxford Bulletin of Economics and Statistics 54, 461–472.

Palumbo, M., Rudd, J., Whelan, K., 2002. On the relationships between real consumption, income, and wealth. Federal Reserve Board Finance and Economics Discussion Series, 2002-38.

Phillips, P.C.B., Ouliaris, S., 1990. Asymptotic properties of residual based tests for cointegration. Econometrica 58, 165–193.