

A Note on the Cointegration of Consumption, Income, and Wealth

Jeremy Rudd
Federal Reserve Board*

Karl Whelan
Central Bank of Ireland**

November 5, 2002

Abstract

Lettau and Ludvigson (2001) argue that a log-linearized approximation to an aggregate budget constraint predicts that log consumption, assets, and labor income will be cointegrated. They conclude that this cointegrating relationship is present in U.S. data, and that the estimated cointegrating residual forecasts future asset growth. This note examines whether the cointegrating relationship suggested by Lettau and Ludvigson's theoretical framework actually exists. We demonstrate that we cannot reject the hypothesis that cointegration is absent from the data once we employ measures of consumption, assets, and labor income that are jointly consistent with an underlying budget constraint. By contrast, Lettau and Ludvigson use a set of variables that do not belong together in an aggregate budget constraint, thereby testing a cointegrating relationship that is not implied by their theory.

*Corresponding author. Mailing address: Mail Stop 80, 20th and C Streets NW, Washington, DC 20551.
E-mail: jeremy.b.rudd@frb.gov.

**E-mail: karl.whelan@centralbank.ie. We are grateful to Martin Lettau for comments on an earlier draft of this paper. The views expressed are our own 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.

1 Introduction

In a recent paper, Martin Lettau and Sydney Ludvigson (2001) formulate an approximate log-linear aggregate budget constraint in which a linear combination of log consumption, assets, and labor income is related to expected future consumption growth and the expected future returns on assets and human capital. Under the assumption that consumption growth and the returns on human and asset wealth are stationary, Lettau and Ludvigson argue that log consumption, assets, and labor income will be cointegrated, and hence that the deviation of these variables from their common trend should forecast at least one of the growth rates of these series. They further argue that statistical evidence provides strong support for the existence of such a cointegrating relationship, and show that its associated residual has predictive power for the rate of growth of assets—a result that they suggest is consistent with the proposition that consumers adjust their spending in response to anticipated movements in future asset returns.

In this note, we re-examine whether the cointegrating relationship suggested by the log-linearized aggregate budget constraint actually exists. We conclude that, on statistical grounds, there is no reason to reject the hypothesis that cointegration is absent in postwar U.S. data. The difference between Lettau and Ludvigson’s results and our own stems from our using measures of real consumption, income, and assets that are consistent with an underlying budget constraint. 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 results concerning the forecastability of asset growth as providing evidence that consumption reflects rationally anticipated variations in future asset returns in a manner that is consistent with their theoretical framework.

The remainder of this note is organized as follows. Section 2 covers some issues asso-

ciated 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. Finally, section 5 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 (such as that used by Lettau and Ludvigson) 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:

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

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 .

Because consumer utility depends on quantities consumed, macroeconomists tend to re-express the budget constraint in terms of real consumption.¹ To do this, we need to

¹The following analysis is closely related to the discussion in Palumbo, Rudd, and Whelan (2002).

divide both sides by the deflator for our consumption measure, P_{t+1}^C . This gives

$$\frac{\widetilde{W}_{t+1}}{P_{t+1}^C} = \frac{P_t^C (1 + I_{w,t+1})}{P_{t+1}^C} \cdot \frac{\widetilde{W}_t}{P_t^C} - \frac{P_t^C (1 + I_{w,t+1})}{P_{t+1}^C} \cdot C_t, \quad (2)$$

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

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

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

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

This equation is the starting point for Lettau and Ludvigson's analysis. 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 equation (4) about a stationary steady-state to obtain

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

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 $\lim_{i \rightarrow \infty} \rho_w^{-i}(c_{t+i} - w_{t+i}) = 0$ allows equation (5) to be solved forward to yield

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

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 equation (6) conditional on time- t information therefore yields the following expression for the consumption-wealth ratio:

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

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

$$w_t \approx \omega a_t + (1 - \omega)h_t, \quad (8)$$

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 $r_{h,t}$

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

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

$$h_t = \mu + y_t + z_t, \quad (10)$$

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:

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

which forms the foundation for Lettau and Ludvigson's analysis. In particular, they argue that the right-hand side of equation (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.

3 Empirical Implementation

Consider now how we might test the hypothesis of cointegration among c_t , a_t , and y_t that is suggested by equation (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. There are two reasons for this. First, the income and asset measures 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. Second, even if the same deflator had been used to define each of these real variables, the measure of consumption is 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.

Is there a way to justify the joint use of these measures of consumption, assets, and income? Lettau and Ludvigson 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.”² 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 equation (11).

One possible 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 allow the cointegrating hypothesis to be derived *indirectly* from the aggregate budget con-

²Lettau and Ludvigson (2001), page 822.

straint. Indeed, Lettau and Ludvigson state that their procedure rests upon the assumption that their measure of log consumption is a constant multiple of the log of total consumption, by which is meant total flow consumption inclusive of the service flow obtained from stocks of consumer durable goods. Although the service flow from consumer durables 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.³

There are two problems with this argument, however. The first is that there appears to be no empirical justification for the assumption that the log of consumption of nondurables and services (excluding shoes and clothing) is proportional to the (unobserved) log of total real flow consumption. The only evidence provided by Lettau and Ludvigson is their observation that the ratio of 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.⁴ However, Figure 1 shows that this statement is incorrect: This ratio has exhibited a distinct upward trend over the postwar period.⁵ Moreover, because the stock of durable goods has historically tended to grow at a similar rate to expenditures on these items, it is very likely that the log of total real flow consumption has grown at least as fast as the log of total consumption expenditures.

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, Lettau and Ludvigson’s empirical approach does not correctly implement the log-linearized real budget constraint that underpins their analysis.

Taken together, these considerations imply that the particular cointegration hypothesis that Lettau and Ludvigson test empirically cannot be viewed as consistent with their

³In 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 owning durables).

⁴See footnote 6 on page 822 of Lettau and Ludvigson (2001).

⁵More relevant for our purposes, there is also no evidence that these series are cointegrated: The estimated *t*-statistic obtained from applying an augmented Dickey-Fuller test to the residuals from a regression of log real consumption outlays on log real nondurables and services consumption is -2.5 , compared with a 10 percent critical value of -3.13 .

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.

Importantly, this alternative approach 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. In addition, the $r_{a,t}$ concept associated with this approach does not include the unobservable implicit rental rate for consumer durables, and so is closer to the kinds of financial returns that Lettau and Ludvigson seek to forecast using 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 results from tests for the cointegration of c_t , a_t , and y_t . As we now demonstrate, it does.

Table 1 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 consumption 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.⁶

The columns of the table are numbered from one to four; this corresponds to the number of lags of $\Delta\hat{u}_t$ that are used in the test regressions. The table reports results for two sample

⁶All variables are expressed in per-capita terms; see the appendix for a complete description of the data.

periods; the shorter sample (which runs from 1952:Q4 to 1998:Q3) corresponds to the dates used in Lettau and Ludvigson’s original work, 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 -3.80 and -3.52 , respectively.⁷

The results are broadly similar for each period; we therefore focus on panel II, which presents the test results from the longer sample. First consider 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.⁸ The picture changes markedly, however, when we test for the cointegration using our preferred measures of consumption and assets. As 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.⁹ 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.

In addition to the residual-based tests, we also consider two likelihood-based test statistics derived by Søren Johansen (1988, 1991). Table 2 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.¹⁰ 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, Haug, and Michelis (1999).¹¹

⁷See 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).

⁸The “general-to-specific” method of Campbell and Perron (1991) 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.

⁹In this case, the general-to-specific procedure calls for one lag of $\Delta\hat{u}_t$ in the test equation.

¹⁰Note that the alternative here is *not*—as Lettau and Ludvigson (2001, p. 847) claim—that there are three cointegrating vectors. See Johansen (1995), chapter 12.

¹¹Critical 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

As before, the results are similar for either sample period; we therefore again focus on the results from the longer sample, which are given in panel II of the table. Even when Lettau and Ludvigson’s measures of consumption and assets are used (line II.A), the evidence against the null hypothesis of no cointegration is weak: None of the estimated models rejects the null hypothesis at the five percent level, and only one of the models (the one-lag system) rejects the null at the 10 percent level. Once our preferred consumption and asset measures are used, however, the evidence against the null of no cointegration becomes weaker still: The values of the test statistics decline in every case (line II.B), and now never imply rejection of the null (the largest trace statistic is not even significant at the 25 percent level).

Finally, Table 3 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 3 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 (2001). If this latter value were correct, 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 that Lettau and Ludvigson provide in their table appears to be an error.¹²

report in the table are very close to those obtained by other studies, such as Johansen and Juselius (1990) and Osterwald-Lenum (1992).

¹²A 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 *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. (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 .)

Hence, the results from the likelihood-based tests provide further indication that the evidence in favor of cointegration is substantially weakened once theoretically consistent measures of consumption and assets are employed.

5 Interpretation of Results

What should we conclude from our inability to 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?

One conclusion is that these results provide a reason to question the theoretical and empirical case for Lettau and Ludvigson’s key finding—that deviations of consumption, assets, and labor income from a common trend have forecasting power for stock returns. In response to Brennan and Xia’s (2002) suggestion that this finding represents a spurious relationship, Lettau and Ludvigson (2002) have argued that equation (11) provides a theoretical justification for their result, and that their evidence on cointegration supplies an empirical justification. However, our findings indicate that these cointegration results are not robust to the use of aggregate data that are compatible with the underlying theoretical framework. And, although we have found that the residuals from a regression involving the log of *total* consumption also forecast asset growth, this finding should be viewed with caution given that the results from our cointegration tests imply that the explanatory variable in this case appears to be nonstationary.

This leaves open the question of why we are unable to reject the hypothesis that a cointegrating relationship is not present in U.S. data. Lettau and Ludvigson (2002) have claimed that such a relationship “must be a part of any economic model where budget constraints are not routinely violated,” which suggests that our findings run counter to basic economic theory. However, we would argue that equation (11) does not in fact provide an airtight case for this claim. Specifically, 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, equation (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 (2002) have noted—samples smaller than the one used in their study will likely suffer from significant small-sample biases.

The second possibility is simply that the underlying relationship described in equation (11) may do a poor job of capturing reality. The derivation of this relationship relies on a host of approximations, starting with the Campbell-Mankiw log-linear approximation to the budget constraint (equation 7), the accuracy of which depends on the stability over time of the ratio of consumption to the unobserved total wealth series.¹³ In addition, the derivation relies on approximations to total wealth and the rate of return on total wealth (equations 8 and 9) whose accuracy 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 (equation 10). Thus, each of the approximations used to derive equation (11) rely on assumptions as to the stability over time 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 the interpretation of Lettau and Ludvigson’s results.

¹³Campbell (1993) discusses this approximation in detail, and notes that it will fail to be accurate when the intertemporal elasticity of substitution is high.

References

- [1] Blinder, Alan and Angus Deaton (1985). “The Time Series Consumption Function Revisited.” *Brookings Papers on Economic Activity*, 2, 465-511.
- [2] Brennan, M. J. and Y. Xia (2002). “*tay*’s as good as *cay*.” Unpublished manuscript (January 24th draft).
- [3] Campbell, John Y. (1993). “Intertemporal Asset Pricing without Consumption Data.” *American Economic Review*, 83, 487-512.
- [4] Campbell, John Y. and N. Gregory Mankiw (1989). “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence.” *NBER Macroeconomics Annual 1989*, 185-216.
- [5] Campbell, John Y. and Pierre Perron (1991). “Pitfalls and Opportunities: What Macroeconomists Should Know about Unit Roots.” *NBER Macroeconomics Annual 1991*, 141-201.
- [6] Johansen, Søren (1988). “Statistical Analysis of Cointegrating Vectors.” *Journal of Economic Dynamics and Control*, 12, 231-254.
- [7] Johansen, Søren (1991). “Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models.” *Econometrica*, 59, 1551-1580.
- [8] Johansen, Søren (1995). *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*, Oxford: Oxford University Press.
- [9] Johansen, Søren and Katarina Juselius (1990). “Maximum Likelihood Estimation and Inference on Cointegration—With Applications to the Demand for Money.” *Oxford Bulletin of Economics and Statistics*, 52, 169-210.
- [10] Lettau, Martin and Sydney Ludvigson (2001). “Consumption, Aggregate Wealth, and Expected Stock Returns.” *Journal of Finance*, 56, 815-849.
- [11] Lettau, Martin and Sydney Ludvigson (2002). “*tay*’s as good as *cay*: Reply.” Unpublished manuscript (May 8th draft).

- [12] MacKinnon, James G., Alfred A. Haug, and Leo Michelis (1999). "Numerical Distribution Functions of Likelihood Ratio Tests for Cointegration." *Journal of Applied Econometrics*, 14, 563-577.
- [13] Osterwald-Lenum, Michael (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.
- [14] Palumbo, Michael, Jeremy Rudd, and Karl Whelan (2002). "On the Relationships between Real Consumption, Income, and Wealth." Federal Reserve Board, Finance and Economics Discussion Series No. 2002-38.
- [15] Phillips, P. C. B. and S. Ouliaris (1990). "Asymptotic Properties of Residual Based Tests for Cointegration." *Econometrica*, 58, 165-193.

Appendix

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. 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 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.)

Table 1: Residual-Based Cointegration Tests

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

Note: 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.

Table 2: Johansen Trace Tests for Cointegration

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

Note: 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, Haug, and Michelis (1999); a test statistic greater than the specified critical value suggests rejection of the null of no cointegration.

Table 3: Johansen L-max Tests for Cointegration

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

Note: 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, Haug, and Michelis (1999); a test statistic greater than the specified critical value suggests rejection of the null of no cointegration.

**Figure 1:
Ratio of Log Real Total Consumption to
Log Real Nondurables and Services Consumption**

