

February 2003

**A Direct Test of the Permanent Income Hypothesis
with an Application to the US States**

by

Joseph P. DeJuan
University of Waterloo
Department of Economics
Waterloo, Ontario N2L-3G1 Canada
email addresses: jdejuan@uwaterloo.ca

John J. Seater
North Carolina State University
Department of Economics
Raleigh, NC 27695 USA
email address: john_seater@ncsu.edu

Tony S. Wirjanto
University of Waterloo
Department of Economics
Waterloo, Ontario N2L-3G1 Canada
email address: twirjant@uwaterloo.ca

Abstract

This paper tests the prediction of the Permanent Income Hypothesis (PIH) that *news* about future income induce a revision in consumption equal to the revision in permanent income. We use time-series data from 48 contiguous US states to perform the test. The empirical results provide some support for the PIH across states.

JEL Classification: E21

Keywords: Permanent income, Consumption, US states

We thank Paul Evans and Mario Crucini for kindly providing us with the 1953-1990 state-level retail sales data. Without implicating we also thank an anonymous referee and Ken West (the Editor) for helpful comments and criticisms.

1. Introduction

Neoclassical consumption theory posits that consumers are forward-looking and base their consumption decisions not on current income but on the expected discounted value of lifetime resources which is known as the permanent income. A key implication of this rational expectations version of the permanent income hypothesis (hereafter, PIH) is that the size of the revision in consumption (β) due to an income innovation is equal to the size of the revision in permanent income (χ) due to the same income innovation. An indirect test of the equality of β and χ has been popular. It can be shown that if $\beta=\chi$, then the *variance* of the revisions in consumption must equal the *variance* of the revisions in permanent income. This latter relation has been tested extensively in the literature, notably by Deaton (1987), West (1988a), Campbell and Deaton (1989), and Gali (1991). Yet, a direct route to testing the equality of β and χ is available. One can first estimate β and χ for different sets of data and then use the resulting sets of estimates to test the equality of β and χ . This direct test has some advantages over the more popular indirect test. First, as part of the test procedure, it produces the magnitudes of β and χ , which are interesting economic quantities. Second and perhaps more important, the direct test is robust to measurement error, which can be severe in consumption data. The direct method uses consumption as a dependent variable. Noise in a dependent variable is just one more component of the estimation residual and does not bias the estimated coefficients of the independent variables. In contrast, the indirect test of equality of variances can be invalidated by measurement error, which artificially inflates one variable's variance relative to another and leads to a spurious rejection of the null hypothesis of equality.

The direct test seems to have originated with Bilson (1980), but has not been much used in the literature subsequently. Bilson (1980) performed the test on quarterly time series data from the U.S., U.K., and Germany and obtained support for the PIH. Flavin (1981) applied the test to aggregate quarterly U.S. data and rejected the PIH. Weissenberger (1986) subsequently argued that both Bilson (1980) and Flavin

(1981) imposed the restriction that the change in consumption is serially uncorrelated. Weissenberger (1986) fitted ARMA models to the data to obtain better-behaved consumption innovations. Using updated data for Germany and the U.K., Weissenberger (1986) rejected the PIH. More recently, Dawson, DeJuan, Seater and Stephenson (2001) conducted a cross-country study using the Penn World Table annual data and reported that the test supports the PIH for industrial countries but not for developing countries. They present evidence, however, that the rejections for the developing countries may have been an artifact of systematic differences in data quality between industrial and developing countries.

This paper builds on these previous studies by using time series data from 48 contiguous US states to perform a direct test of the $\beta=\chi$ implication of the PIH. US state level data are an underutilized source of information on consumer behavior and have a number of attractive features for testing consumption theories.¹ First, the quality of data across states is relatively uniform and has to meet the standard set by the U.S. *Bureau of Economic Analysis*. Consequently, test results based on these data are less prone to problems associated with systematic data quality variation that are inherent in many cross-country studies. Second, the US states have a national banking system, well-integrated capital and goods markets - see Asdrubali, Sorensen and Yosha (1996) and Crucini (1999) - a common form of government, and a common currency. Each state can be viewed as part of a highly integrated region, making it possible for states to adjust consumption in the face of income innovations. Test results based on state level data, therefore, provide a useful benchmark for subsequent studies applying the direct test to other data sets.

¹Other recent studies on consumption using US state-level data are Beaudry and van Wincoop (1996) and Ostergaard, Sorensen and Yosha (2002). Beaudry and van Wincoop (1996) examine the response of consumption growth to changes in expected interest rates using a panel of state-level consumption and income data. They find the estimates of the intertemporal elasticity of substitution in consumption to be statistically and economically significant, averaging at about unity. This result is in contrast to the less than 0.1 estimate obtained using US aggregate time-series data (see, e.g., Hall, 1988). The study by Ostergaard et al. (2002) is closely related to ours, in that they examine whether state-level consumption is excessively sensitive to lagged state-level income — a well-documented empirical regularity in US aggregate data. They find that when aggregate US income (or consumption) is controlled for, state-level consumption exhibits much less sensitivity to lagged income than when it is not controlled for.

Our test procedure can be summarized as follows. For each state, we jointly estimate (a) the univariate income generating process and (b) the equation describing the relation between consumption revision and income innovation. We then use the estimated parameters of the income process to quantify the magnitude of the revision in permanent income associated with an income innovation. Overall, we find β to be insignificantly different from χ across the US states, as the PIH predicts. The PIH may thus be a reasonable model for describing the response of consumption to income innovations at the state level.

The paper is organized as follows. Section 2 reviews the basic framework of the PIH. Section 3 describes the data. Section 4 presents the empirical results. Section 5 concludes the paper.

2. Permanent Income Innovations and Consumption Revisions

The testable implications of the rational expectations version of the permanent income model can be seen in the following standard model:

$$C_{it} = Y_{it}^p \quad (1)$$

$$Y_{it}^p = rW_{it} + \frac{r}{(1+r)} \sum_{j=0}^{\infty} \frac{1}{(1+r)^j} E_t Y_{i,t+j} \quad (2)$$

$$W_{it} = (1+r)W_{i,t-1} + Y_{i,t-1} - C_{i,t-1} \quad (3)$$

where C is consumption, Y^p is permanent income, W is nonhuman wealth, Y is labor income, r is real interest rate, E_t is the expectation operator conditional on all the information available to the consumer in period t , and all variables are subscripted by a state index i and a time index t . To avoid unnecessary complications in this motivational discussion, we assume that transitory consumption, which might arise from taste shocks, is zero (that is, equation (1) contains no error), the real interest rate is constant and equal to the rate of time preference, and income Y is exogenous with respect to the consumer's consumption decision. Under these

assumptions, permanent income changes only in response to new information about the future path of labor income:

$$\Delta Y_{it}^p = Y_{it}^p - Y_{it-1}^p = \theta_{it} \quad (4)$$

where

$$\theta_{it} = \frac{r}{(1+r)} \sum_{j=0}^{\infty} \frac{1}{(1+r)^j} (E_t - E_{t-1}) Y_{it+j} \quad (5)$$

Equation (5) shows that θ_{it} is equal to the annuity value of revisions in the expected income stream and is independent of all previously known information.

Substituting equation (2) into (1), combining the result with the wealth constraint in equation (3), and solving for C_{it} in terms of $C_{i,t-1}$ yields:

$$\Delta C_{it} = \Delta Y_{it}^p = \theta_{it} \quad (6)$$

where $\Delta C_{it} = C_{it} - C_{i,t-1}$. Equation (6) implies that an income innovation causes the same size revision in consumption as in permanent income; in other words, the change in consumption will be exactly the amount of the change in permanent income warranted by *news* about future income. We can test this implication if we can express the innovation in permanent income as a function of the innovation of the observables.

We make the usual assumption that income follows a linear stochastic process, for which there is a well-developed theory of estimation, inference, and prediction. More formally, suppose that ΔY_{it} follows an univariate ARMA process:

$$(1 - a_1 L - a_2 L^2 - \dots) \Delta Y_{it} = (1 + b_1 L + b_2 L^2 + \dots) \epsilon_{it} \quad (7)$$

where $\Delta Y_{it} = Y_{it} - Y_{i,t-1}$, L is the lag operator, a_{ij} are autoregressive parameters, b_{ij} are moving-average parameters, and ϵ_{it} is income innovation.² Given (7), the formula for the *revision* in permanent income, θ_{it} , is (see, e.g., Flavin, 1981):

$$\theta_{it} = \frac{1 + \sum_{j=1}^{\infty} \frac{b_j}{(1+r)^j}}{1 - \sum_{j=1}^{\infty} \frac{a_j}{(1+r)^j}} \cdot \epsilon_{it} = \chi_1(r, b_j, a_j) \cdot \epsilon_{it} \quad (8)$$

²Note that we are assuming here that income is difference-stationary. We justify this assumption later with formal statistical tests.

Conditional on the univariate income process and the real interest rate, χ provides a direct estimate of the size of the revision in permanent income associated with the realization of an innovation in current income. In particular, the size of χ depends on the *persistence* of innovations to the income process. Strong persistence in this case means that a positive (negative) shock to ΔY tend to be followed, on average, by another positive (negative) shock such that the effects of the shock persists over time. To take a simple example, suppose that ΔY follows a first-order autoregressive process with an autoregressive parameter a_j . With this income process, $\chi=1/(1-a_j/(1+r))$. For a given r , a highly persistent income innovation implies that $a_j>0$ (i.e., ΔY is positively autocorrelated) and hence, $\chi>1$. If, on the other hand, $a_j=0$ then $\chi=1$.

A natural means of testing the PIH is to estimate the following system of equations:

$$\begin{aligned} (1 - a_1L - a_2L^2 - \dots)\Delta Y_{it} &= (1 + b_1L + b_2L^2 + \dots)\epsilon_{it} \\ \Delta C_{it} &= \beta_i \epsilon_{it} + \xi_{it} \end{aligned} \tag{9}$$

both unconstrained and constrained, with the constrained system imposing the restriction given by the following nonlinear equation:

$$\beta_i = \frac{1 + \sum_{j=1}^{\infty} \frac{b_j}{(1+r)^j}}{1 - \sum_{j=1}^{\infty} \frac{a_j}{(1+r)^j}} = \chi_i(r, b_p, a_j) \tag{10}$$

where β_i is the marginal propensity to consume out of income innovation, and ξ_{it} is the random error in the consumption change equation. The first equation in (9) specifies the time-series process for income, and the second equation describes the relation between consumption revision and income innovation.

The large-sample Likelihood Ratio (LR) test is used to test the null hypothesis that the constraint (10) is binding. The LR statistics is calculated as $LR=2[\log L^u - \log L^c]$, where $\log L^u$ and $\log L^c$ are the log likelihood of the unconstrained and constrained system, respectively. The LR statistic is distributed

asymptotically as a chi-squared random variable with 1 degree of freedom under the null hypothesis of $\beta=\chi$. To interpret the test results, we report the marginal significance level (MSL), popularly known as the p -value, of the likelihood ratio statistics for testing the null hypothesis that $\beta=\chi$ for each state. It is calculated as $MSL=Pr(L>LR)$, where L is a chi-squared random variable with 1 degree of freedom, and LR is the actual value of the test statistic; i.e., MSL is the probability of observing a value of the LR at least as large as we did, given that the null hypothesis that $\beta=\chi$ is correct. The decision rule is to reject the null hypothesis at the 100α percent level if the MSL is less than or equal to the significance level α . If the null hypothesis is not rejected, then we conclude that consumption responds to income innovation as predicted by the PIH; that is, the marginal propensity to consume out of an income innovation (β) is equal to the marginal propensity to revise permanent income in response to the same innovation (χ). Other theories of consumption, such as Keynes's absolute income hypothesis or Campbell and Mankiw's (1990) rule-of-thumb consumption, do not have this prediction. Consequently, failure to reject the null hypothesis is not merely consistent with the PIH but helps distinguish it from competing theories.

Before proceeding with the estimation of equation (9), we need to address several issues concerning the data and specification that can affect the estimates of β and χ and hence the test for $\beta=\chi$. First, C is pure consumption and can be approximated by expenditures on nondurable goods and services plus the imputed services of the stock of consumer durables. However, reliable estimates of the stock of consumer durables across states are not available so we use total expenditure on goods and services as a measure of consumption. In the presence of convex adjustment costs, durables are likely to adjust gradually to a change in permanent income. See, for example, Bernanke (1985). Including them in the measure of consumption, therefore, will bias the estimate of β downward. Second, the univariate income generating process assumes that consumers use only their past income history to predict future incomes. West (1988a), Quah (1990), Flavin (1993) among others, however, have noted that consumers may also use other information in predicting

their future incomes; so the estimated income innovation, ϵ , is an errors-in-variable measure of the true income innovation. The use of ϵ as a regressor in (9) will then yield a downward biased estimate of β . Third, the empirical measure of χ requires the imposition of an assumed constant real interest rate. If the chosen rate is too high (low), then the estimate of χ will be biased downward (upward).

3. Data

We use annual data from 48 US states for the period 1953-1998.³ We exclude the pre-1953 period to avoid problems associated with the Korean War era noted by Campbell (1987) and Campbell and Mankiw (1990). State disposable personal income and population data are from *State Personal Income 1948-1998*, published by the U.S. Department of Commerce, Bureau of Economic Analysis. We approximate state private consumption by state total retail sales. These data are available in the *Statistical Abstract of the United States: Retail Sales by Type of Store and State*.⁴ All the data series are divided by state population and national consumer price index (1986 dollars CPI) to give real per-capita magnitudes. Ideally, state consumption and disposable incomes should be deflated using state price indices. However, such price indices are not available. Real per-capita consumption and real per-capita disposable personal income are used as measures of consumption and income, respectively. All the data series are expressed in logarithmic form.

4. Empirical Results

We begin by examining the time series properties of log disposable income per capita for each state. To this end, we use a modified version of the augmented Dickey-Fuller (1979) test developed by Elliott,

³Alaska and Hawaii are not included in the sample because of incomplete data.

⁴The 1953-1990 retail sales data are from Paul Evans and Mario Crucini. We updated their data set up to 1998 using the same source, *Statistical Abstract of the United States: Retail Sales by Type of Store and State*.

Rothenberg and Stock (1996) and a modification of the Phillips and Perron (1988) Z_{α} test proposed by Stock (1991).⁵ The former test is referred to as the DF-GLS test and the latter as the MZ_{α} test. These tests allow us to test formally the null hypothesis of a unit root series against the alternative hypothesis of no unit root (or stationary) series. The results of the DF-GLS and MZ_{α} tests are summarized in rows 1 and 2 of Table 1. For both tests, the null hypothesis of a unit root cannot be rejected for any states at the 5 percent level of significance. We also tested the first difference of the log of income using the same test statistics. The null of a unit root in these tests can be rejected for all states, even at the 10 percent level (see rows 3 and 4 of Table 1). From these results, we conclude that the state income series are well-characterized as nonstationary or $I(1)$ processes. These results also provide support for a difference-stationary specification for the income series and consequently, we perform our empirical analysis using first-differenced data.

Use of differenced data distinguishes our study from most previous applications of the direct test of the PIH. In particular Flavin (1981) and Weissenberger (1986) assume trend-stationarity.⁶ Mankiw and Shapiro (1985) showed that the direct test can be severely biased by inappropriately imposing trend-stationarity on difference-stationary data. However it is important to point out that some of the tests will be asymptotically valid if it is possible to rewrite the consumption regression involving lagged consumption and lagged incomes, so that the coefficients of lagged incomes are on trend stationary, zero mean regressors. See Sims, Stock and Watson (1990), Stock and West (1988) and West (1988b).

To determine a suitable time series model for the income series, we initially estimate an AR(2) model for the difference of the log of income for each state. Table 2 reports the estimation results. The estimated coefficients of twice lagged difference of the log of income are very small, with a mean value of -0.070, and

⁵Monte Carlo evidence suggests that these modified unit-root tests have better finite-sample properties in terms of both power and size relative to their standard counterparts (see, e.g., Elliott et al. 1996; and Stock 1991). The critical values for the test are calculated from the response surface estimates in MacKinnon (1991).

⁶Dawson, DeJuan, Seater, and Stephenson (2001) use first-differenced data.

are insignificantly different from zero in many states. By contrast, the estimated coefficient of once lagged difference of the log of income has a mean value of 0.204, with the absolute value of the t -statistics averaging at 1.736 which is statistically significant at the 10% level. These results suggest that state income series can be reasonably modeled as a simple AR(1) process in differences. Also, it is noteworthy that these results are in accordance with the statistical properties of the aggregate U.S. income series obtained by Campbell and Deaton (1989).

The two-equation system in (9) - without intercept terms - is jointly estimated for each state by nonlinear least squares.⁷ Based on the results in Table 2, the income generating process is restricted to an ARIMA(1,1,0) process. We experimented with longer autoregressive lags and found that the results of this paper are robust to this restriction. We estimate χ based on 3 percent rate but similar results are obtained using the alternative values of 1 percent and 6 percent.

The estimation results of equation (9) are summarized in Table 3. The estimates of β vary between 0.088 and 1.622 with a mean of 1.084. Ninety percent of the states have a t -statistics for β in excess of 1.645, the 10 percent significance level for a two-tailed test. These results imply that consumers in each state revise their current consumption as a result of innovations to current income. The estimates of χ are significantly positive for all states, and vary between 0.708 and 1.423 with a mean of 1.049. These results imply that current income innovations contain information about future income that leads consumers in each state to revise their estimated permanent income.

We now turn to our main empirical question: is β equal to χ ? We can get a quick picture of the relation between β and χ from Figure 1, which plots β against χ . If $\beta = \chi$, the points should fall along the 45 degree line through the origin. The cloud of points is widely scattered about the 45 degree line; nevertheless,

⁷Including intercept terms in (9) does not alter the estimation results appreciably.

when a line is fitted to the scatter plot, the joint null hypothesis that the intercept is zero and the slope is one cannot be rejected at the 5% significance level (see notes in Figure 1).

The LR test provides a more rigorous analysis. Row 5 of Table 3 summarizes the MSLs of the LR statistics for testing the null hypothesis that $\beta=\chi$ for each state. As mentioned in Section 2, the MSL is the largest significance level at which we could carry out the test and still fail to reject the null hypothesis. The 10th percentile of MSL is 0.106 which means that 90 percent of the MSLs are higher than 0.106. This reveals that more than 90 percent of the states cannot reject the null hypothesis of $\beta=\chi$ at the 10 percent level of significance. Figure 2 shows the frequency distribution of the MSL. It is apparent that the distribution of the MSL is not heavily skewed to the left or right, but it appears to be uniformly distributed over the range (0,1). In fact, under the null hypothesis that $\beta=\chi$, the distribution of MSL should be Uniform (0,1), whereas under the alternative it should be skewed toward low values of MSL. To formally test whether the empirical distribution of the MSL deviates significantly from the theoretical Uniform (0,1), we utilize the Kolmogorov-Smirnov (KS) test which is a non-parametric empirical distribution test. The KS test statistic is 0.086 with a corresponding MSL of 0.868. Clearly, this test fails to reject the null hypothesis of Uniform (0,1) distribution at any conventional significance level.

For comparative purposes, we also apply West's (1988a) excess smoothness test to our data set. As West (1988a) explains, if the PIH is true, we should expect the variance of consumption's actual change to be less than or equal to the variance of consumption's change predicted by the model. The reason is that households presumably have at least as much information as the econometrician about their income process, so that the true *news* about income is less than that estimated by the econometrician. In terms of our model, the household's residual ϵ is generally closer to zero than the residual obtained by the econometrician. As a result, the household will update its permanent income and, therefore, its consumption by less than the econometrician will predict, leading to a smaller variance for observed changes in consumption than for

predicted changes. See West (1988a) for details. In our model, the predicted change in consumption is $\beta\epsilon$, so following West (1988a), we should expect to see $\text{var}(\Delta C) \leq \text{var}(\beta\epsilon)$. In fact, as Table 4 shows, we have the opposite result; that is, $\text{var}(\Delta C)$ is much greater than $\text{var}(\beta\epsilon)$, always by at least a factor of two and often by a factor of five or more. Thus the use of West's (1988a) test leads to a rejection of the PIH.

How can the seemingly contradictory results of our direct test and West's (1988a) indirect test be reconciled? One possibility is measurement error. It is important to stress that the PIH is concerned with pure consumption and not consumption expenditure. However our data are total retail sales, which is a measure of expenditure and, therefore, an imperfect measure of pure consumption. It is well-known that measurement error in consumption can cause changes in measured consumption that do not reflect changes in permanent income; therefore it can increase $\text{var}(\Delta C)$ relative to $\text{var}(\beta\epsilon)$, tending to bias the excess smoothness test toward rejection of the PIH. In contrast, measurement error has no effect on the direct test of whether $\beta=\chi$. This is because the construction of the direct test treats consumption as a dependent variable in a regression, so that measurement error is merely one more component of the residual. As a result, no bias is introduced in the estimate of β or in the test of equality between β and χ . This property is an advantage of the direct test compared to the indirect excess smoothness test.⁸

All in all, then, the results across 48 US states indicate strong though not quite perfect support for the PIH model. It is important to note, however, that a key identifying assumption involved in the $\beta=\chi$ test is that state income can be modeled as a simple ARIMA(1,1,0) process. This specification may be limited to the extent that consumers can use information other than past incomes to predict their future incomes. Future research therefore can extend the test in at least two different ways. One extension is to include information

⁸Another possible explanation is that the direct test suffers from low power. Ideally, we should also perform a simulation exercise where we utilize data generating processes that do not quite satisfy the PIH to sort out which tests are informative and which are not in terms of the test power. We will pursue this line of research in the near future.

other than past income in the specification of the income process. Such an approach can lead to a more accurate estimate of the state income innovations and hence, to a more precise estimate of χ . Another interesting extension would be to decompose the income innovations into aggregate and state-specific components, and assess the response of consumption to each of these components. This decomposition may require specification of a more complicated income process, or the use of factor analysis models such as those undertaken in the study by Forni and Reichlin (2001).

5. Conclusion

Using time series data from 48 contiguous US states, this paper set out to implement a direct test of the implication of the Permanent Income Hypothesis (PIH) that the size of consumption revision (β) due to an income innovation is equal to the size of permanent income revision (χ) due to the same income innovation. For each state, we jointly estimated the univariate income generating process and the equation describing the relation between consumption revision and income innovation. We then obtained an estimate of the magnitude of the revision in permanent income associated with an income innovation using the estimated parameters of the stochastic process generating income.

Our estimates of β and χ across states were positive and statistically significant in most states. These results imply that innovations in current income contain new information about the expected future path of income that lead consumers in each state to revise both their consumption and permanent income. We then tested whether β is equal to χ for each state. The results indicated that the hypothesis of $\beta=\chi$ cannot be rejected in many of the 48 US states, suggesting that the PIH is a reasonable model for describing the response of state-level consumption to state-level income innovations.

This support for the PIH stands in contrast to the widely held view that credit constraints, myopia, or behavior toward risk are required to make sense of the aggregate consumption data. We have no concrete

explanation for this disagreement at this point, except to suggest that the favorable result obtained for the PIH at the state level need not extend to the aggregate level. One potential reason is expounded by Ostergaard, Sorenesen and Yosha (2002). They interpret their findings that state-level consumption exhibits much less sensitivity to lagged income when aggregate US income (or consumption) is controlled for as indicating that the degree of integration of states within the US is relatively high, while the US as a whole is not well-integrated with the outside world. This is a worthwhile future avenue of research to pursue.

REFERENCES

- Asdrubali, P., Sorensen, B. and Yosha, O. (1996), "Channels of Interstate Risk Sharing: United States, 1963-1990," *Quarterly Journal of Economics*, 4, 955-1006.
- Beaudry, P. and Van Wincoop, E. (1996), "The Intertemporal Elasticity of Substitution: An Exploration Using a US Panel of State Data," *Economica*, 63, 495-515.
- Bernanke, B. (1985), "Adjustment Costs, Durables and Aggregate Consumption," *Journal of Monetary Economics*, 15, 41-68.
- Bilson, J. (1980), "The Rational Expectations Approach to the Consumption Function: A Multi-Country Study," *European Economic Review*, 13, 273-99.
- Campbell, J. (1987), "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis," *Econometrica*, 55, 1249-1273.
- Campbell, J. and Mankiw, N.G. (1990), "Permanent Income, Current Income and Consumption," *Journal of Business and Economic Statistics*, 8, 265-79.
- Campbell, J. and Deaton, A. (1989), "Why is Consumption So Smooth?," *Review of Economic Studies*, 56, 357-374.
- Crucini, M. (1999), "On International and National Dimensions of Risk Sharing," *The Review of Economics and Statistics*, 81, 73-84.
- Dawson, J., DeJuan, J., Seater, J. and Stephenson, E. F. (2001), "Economic Information versus Quality Variation in Cross-Country Data," *Canadian Journal of Economics*, 34, 988-1009.
- Deaton, A. (1987), "Life-Cycle Models of Consumption: Is the Evidence Consistent with the Theory?," in Truman F. Bewley (ed.), *Advances in Econometrics, Fifth World Congress, Vol. 2*, Cambridge University Press, 121-48.
- S) Q. (1992), *Understanding Consumption*, Oxford University Press.
- Dickey, D. A. and Fuller, A. (1979), "Distribution of the Estimator for Autoregressive Time Series with a Unit Root," *Journal of the American Statistical Association*, 74, 427-431.
- Elliott, G., Rothenberg, T.J. and Stock, J. H. (1996), "Efficient Tests for an Autoregressive Unit Root," *Econometrica*, 64, 813-836.
- Flavin, M. (1981), "The Adjustment of Consumption to Changing Expectations about Future Income," *Journal of Political Economy*, 89, 974-1009.
- S) Q. (1993), "The Excess Smoothness of Consumption: Identification and Interpretation," *Review of Economic Studies*, 60, 651-666.

- Forni, M. and Reichlin, L. (2001), "Federal Policies and Local Economics: Europe and the US," *European Economic Review*, 45, 109-134.
- Friedman, M. (1957) *A Theory of the Consumption Function*, Princeton, Princeton University Press.
- Gali, J. (1991), "Budget Constraints and Time Series Evidence on Consumption," *American Economic Review*, 81, 1238-1253.
- Hall, R. E. (1988), "Intertemporal Substitution in Consumption," *Journal of Political Economy*, 96, 339-357.
- Kormendi, R. and LaHaye, L. (1984), "Cross-regime Tests of the Permanent Income Hypothesis," University of Chicago mimeo.
- MacKinnon, J. G. (1991), "Critical Values for Cointegration Tests," in Robert F. Engle and Clive W.J. Granger (eds.), *Long-Run Economic Relationships: Readings in Cointegration*, Oxford University Press.
- Mankiw, N. G. and Shapiro, M. D. (1985), "Trends, Random Walks, and Tests of the Permanent Income Hypothesis," *Journal of Monetary Economics*, 16, 165-174.
- Newey, W. and West, K. (1987), "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703-8.
- Ostergaard, C., Sorensen, B. and Yosha, O. (2002), "Consumption and Aggregate Constraints: Evidence from US States and Canadian Provinces," *Journal of Political Economy*, forthcoming.
- Phillips, P. C. B. and Perron, P. (1988), "Testing for a Unit Root in Time Series Regressions," *Biometrika*, 75, 335-346.
- Pischke, J. S. (1995), "Individual Income, Incomplete Information, and Aggregate Consumption," *Econometrica*, 63, 805-40.
- Quah, D. (1990), "Permanent and Transitory Movements in Labor Income: An Explanation for Excess Smoothness in Consumption," *Journal of Political Economy*, 98, 449-75.
- Quah, D. (1990), "Permanent and Transitory Movements in Labor Income: An Explanation for Excess
- Sims, C. A., Stock, J. H. and Watson, M. W. (1990), "Inference in Linear Time Series Models with Some Unit Roots," *Econometrica*, 58, 113-114.
- Stock, J. H. (1991), "A Class of Tests for Integration and Cointegration," manuscript. Kennedy School of Government, Harvard University, Cambridge, MA.
- Stock, J. H. and West, K. D. (1988), "Integrated Regressors and Tests of the Permanent Income Hypothesis," *Journal of Monetary Economics*, 21, 85-95.

Weissenberger, E. (1986), "Consumption Innovations and Income Innovations: Evidence from United Kingdom and Germany," *The Review of Economics and Statistics*, 68, 1-8.

West, K. D. (1988a), "The Insensitivity of Consumption to News About Income," *Journal of Monetary Economics*, 21, 17-34.

West, K. D. (1988b), "Asymptotic Normality, When Regressors Have a Unit Root," *Econometrica*, 56, 1397-1418.

**Table 1 - Summary Statistics of the Unit Root Tests for
US State Per-Capita Disposable Income Processes**

	Statistics	Mean	Std. Dev.	Min	P10	Median	P90	Max
<i>Y</i>	DF-GLS	-1.644	0.507	-0.372	-1.067	-1.663	-2.492	-2.696
<i>Y</i>	MZ_{α}	-0.531	0.006	-0.542	-0.540	-0.531	-0.522	-0.517
ΔY	DF-GLS	-5.012	-1.027	-3.848	-4.009	-4.664	-7.083	-7.749
ΔY	MZ_{α}	-11.843	4.418	-24.136	-18.693	-11.077	-7.068	-5.405

Notes: *Y* is the log of real per-capita disposable income, Δ is the first-difference operator, P10 is the tenth percentile and P90 is the ninety percentile. DF-GLS test is the modified version of the Augmented Dickey-Fuller test developed by Elliott, Rothenberg and Stock (1996), with a data dependent lag length chosen according to the Akaike Information Criterion. MZ_{α} test is a modification of the Phillips and Perron (1988) Z_{α} test proposed by Stock (1991). For the MZ_{α} test, we estimate the spectral density using an AR(1) spectral estimator. The critical values for the DF-GLS and MZ_{α} tests are obtained using the response-surface method advocated by MacKinnon (1991).

TABLE 2 - Summary Statistics of the US State Level Disposable Income Processes

Model: $\Delta Y_{it} = \alpha_{i0} + \alpha_{i1}\Delta Y_{it-1} + \alpha_{i2}\Delta Y_{it-2} + \epsilon_{it}$

	Mean	Std. Dev.	Min	P10	Median	P90	Max
$\alpha_{i,1}$	0.204	0.208	-0.345	-0.183	0.238	0.442	0.473
$\alpha_{i,2}$	-0.070	0.093	-0.253	-0.177	-0.083	0.078	0.146
$t(\alpha_{i,1})$	1.736	0.860	0.060	0.680	1.670	2.990	3.190
$t(\alpha_{i,2})$	0.676	0.403	0.010	0.150	0.630	1.190	1.680

Notes: P10 is the tenth percentile, P90 is the ninety percentile, $\alpha_{i,1}$ is the AR(1) coefficient, $\alpha_{i,2}$ is the AR(2) coefficient, $t(\alpha_{i,1})$ is the absolute value of the t-statistic of the AR(1) coefficient, and $t(\alpha_{i,2})$ is the absolute value of the t-statistic of the AR(2) coefficient.

TABLE 3 - Summary Statistics for the Two-Equation System (9)

	Mean	Std. Dev.	Min	P10	Median	P90	Max
β	1.084	0.411	0.088	0.486	1.228	1.499	1.622
χ	1.049	0.157	0.708	0.833	1.046	1.241	1.423
$t(\beta)$	4.427	2.070	0.632	1.819	4.484	7.290	8.302
$t(\chi)$	8.080	1.357	5.974	6.633	7.744	9.843	11.659
$MSL(H_0: \beta=\chi)$	0.492	0.307	0.000	0.106	0.474	0.922	0.999


Notes: P10 is the tenth percentile, P90 is the ninety percentile, $t(\cdot)$ is the t-statistic, and MSL is the marginal significance level for the likelihood ratio statistic under the null hypothesis of $\beta=\chi$. Sample period: 1953-1998.

TABLE 4 - Summary Statistics of Variances of ΔC and $\beta\epsilon$

	Mean	Std. Dev.	Min	P10	Median	P90	Max
var (ΔC)	0.0028	0.0012	0.0012	0.0017	0.0027	0.0040	0.0084
var ($\beta\epsilon$)	0.0007	0.0002	0.0002	0.0005	0.0007	0.0010	0.0014

Notes: P10 is the tenth percentile, P90 is the ninety percentile. Sample period: 1953-1998.

Figure 1: Scatter Plot of β and χ



Notes: Using the data on β and χ shown above, we obtain the fitted values:

$$\hat{\beta}_1 = -0.41 + 1.42 \chi_1$$

std err: (0.34) (0.32)

Number of observations=48; $R^2=0.914$; $F_{(H_0: \text{intercept}=0 \text{ and slope coefficient}=1)}=1.14$ (MSL=0.32)

Figure 2: Histogram of MSL

