ALBERT JAEGER

Cross-Country Evidence on the Predictability of Consumption Growth

53

ALBERT JAEGER

Cross-Country Evidence on the Predictability of Consumption Growth

Juni 1992

WIFO Working Papers, 1992, (53)

Cross-Country Evidence on the Predictability of Consumption Growth*

Albert Jaeger

June 1992

Austrian Institute of Economic Research
P.O. Box 91
A-1103 Vienna, Austria

* I am grateful to Ken West and Michael Wueger for helpful comments on an earlier version of this paper. All errors are mine.

Abstract

Recent research on Hall's (1978) random-walk interpretation of the permanent income hypothesis has reported that predictability of consumption growth is statistically significant but restricted to a small portion of the total variation in consumption growth. This paper uses band spectrum regressions to study the predictability of consumption growth at different frequency bands. The main findings are that the predictive power of consumption growth forecasts is concentrated at the business-cycle frequencies and that predictability reflects tight co-movements of realized consumption growth and predicted income growth at these data frequencies. The findings suggest that time series studies based on data with substantial high-frequency fluctuations in consumption growth may severely understate the case against the random-walk interpretation of the permanent income hypothesis.

Keywords: Permanent Income Hypothesis, Macroeconomic Forecasts, Band Spectrum Regressions.

JEL Classification: 130, 210, 920.

Hall (1978) applied the logic of the efficient market hypothesis to derive a powerful implication of the permanent income hypothesis: To a first approximation, changes in consumption growth should be unpredictable. The random walk hypothesis of private consumption behaviour (RWH) has been tested extensively and the evidence is generally not supportive (see, e.g., Flavin, 1981; Muellbauer, 1983; Bean, 1986; Campbell and Mankiw, 1989). A statistically significant but typically small portion of the variation in consumption growth rates appears to be predictable when consumption growth is regressed on lagged variables including lagged income and consumption. While the statistical rejection of the RWH is widely accepted as robust, the economic interpretation of the empirical evidence remains controversial. One group of economists has interpreted the evidence as showing that deviations from the RWH are not just statistically but also economically substantive. Campbell and Mankiw (1989), for example, conclude that the world is roughly split between consumers following the RWH and consumers following the rule-of-thumb of consuming their current income. A second group of economists, however, has argued that the portion of predictable variation in consumption growth is negligible given the strong assumptions needed to derive the RWH. Cochrane (1991), for example, points out that the amount of predictable variation in quarterly consumption growth for the six OECD countries reported by Campbell and Mankiw (1991) is small - as small as 1 percent in the case of the U.S. data - and suggests that minor changes in model specification may account for the reported deviations from the RWH.

The purpose of this paper is to show that the typically small amount of predictable variation in consumption growth reported by time series studies may understate the case against the RWH for two reasons: First, the studies do not distinguish between predictability of high-frequency and low-frequency variations in consumption growth. Time series data may exhibit negligible overall predictability although predictability at the lower data frequencies is substantial. Second, the power of the usual test strategy of regressing consumption growth on a variety of lagged variables and testing for the statistical significance of the coefficients may be low. In this paper I employ semi-annual forecasting data for the seven large (G-7) industrialized countries, published by the Organization of Economic Co-operation and Development (OECD) over the time period 1968-1990. Because forecasters should have a strong interest in revealing any information that helps predict consumption growth, forecasting data may embody a significant amount of information in a single variable. Moreover, this approach to testing the RWH guards against the danger of finding spurious predictability through data mining. Using band spectrum regressions, I find that predictability of consumption growth is largely concentrated at the business-cycle frequencies of the data, defined as cycles in the data taking more than three years to complete. In contrast, high-frequency variations of consumption growth in all G-7 countries appear to be unpredictable. In addition, realized consumption growth appears to move closely with predicted income growth at the business-cycle frequencies but not at the higher frequencies.

The findings indicate that small statistical deviations from the RWH in time series data with short sampling interval are consistent with substantive predictable variations in consumption growth at the business-cycle frequencies. Similar findings are reported by Jaeger (1992) for quarterly U.S. consumption forecasts over a similar time period. In a cross-country study that provides evidence on co-movements of consumption and income growth at the zero frequency, Carrol and Summers (1991) also conclude that the evidence against the RWH may be severely understated by conventional time series studies of consumption growth.

I. Test Methodology

I take the strict RWH of private consumption behaviour to imply:

$$\Delta c_t = \epsilon_t, \qquad (1)$$

where Δc_t is the growth of real consumption adjusted for a constant mean, and ϵ_t is a regression disturbance orthogonal to all information available at time t-1.[1] Hall (1978) pointed out that forecasts of the mean-adjusted consumption growth rate based on information available at t-1, denoted by $\Delta c_{t,f}$, should be uncorrelated with actual consumption growth. Thus, in a bivariate regression of actual consumption growth on predicted consumption growth (both mean-adjusted):

$$\Delta c_t = \beta \Delta c_{t,f} + u_t. \qquad (2)$$

The coefficient estimate of β should be statistically insignificant and the R^2 of the regression should be zero.

Regression (2) assumes that the relationship between actual and predicted consumption growth is invariant across different frequencies of the data. Thus, low-frequency variation in consumption growth, typically associated with business-cycle fluctuations, are assumed to have the same relationship to predicted consumption growth as high-frequency variations. The band spectrum regression technique proposed by Engle (1974) allows to relax this assumption and to study predictability across different frequency bands. The regression results reported below are based on an estimation approach suggested by Harvey (1978). Let T denote the number of observations. The time series data are transformed to the frequency domain by using a real finite Fourier transform matrix (Z) with dimension TxT, whose elements are defined as follows:

$$z_{ts} = T^{-1/2} \qquad t=1; \qquad s=1,...,T;$$
$$z_{ts} = (2/T)^{1/2}\cos[\pi t(s)/T] \qquad t = 2,4,6,...,T-1; \qquad s=1,...,T;$$
$$z_{ts} = (2/T)^{1/2}\sin[\pi t(s)/T] \qquad t = 3,5,7,...,T; \qquad s=1,...,T.$$

Equation (2) is written in the frequency domain as:

$$\Delta c^* = \beta \Delta c^*_f + u^*, \qquad (3)$$

where $\Delta c^* = Z\Delta c_t$, $\Delta c^*_f = Z\Delta c_{t,f}$, and $u^* = Zu_t$.[2] These transformations create real valued variables with T entries which are indexed not by time, but by frequency (w). Because the sampling interval of the data used in the regressions is semi-annual, the highest frequency ($w=\pi$) corresponds to cycles of one-year-length. I define the high-frequency band of the data as comprising all cycles taking less than three years to complete ($\pi/3<w\leq\pi$). The low-frequency band comprises the cycles taking more than three years to complete and will therefore cover the remaining frequency range ($0\leq w\leq\pi/3$). Thus, the low-frequency band roughly covers the frequency range usually associated with business-cycle fluctuations.[3] Band spectrum regressions for the low-frequency range are equivalent to running ordinary least squares regressions using the first 1/3 entries of the transformed data. The regressions for the high-frequency range are based on the remaining 2/3 entries of the transformed data series. If regression (3) is run with all frequencies included, it is equivalent to regression (2) in the time domain.

II. The Data

The empirical analysis is based on forecasting data for consumption and income growth published by the OECD for the G-7 countries: the U.S., Japan, Germany, France, Italy, the U.K., and Canada. The forecasts of real consumption and income growth are taken from various issues of the semi-annual OECD publication OECD Economic Outlook. Real consumption includes consumption of nondurables and services as well as expenditures on durables. Real income is approximated by real GNP or GDP. The semi-annual forecasts are mid-of-year or end-of-year projections of seasonally adjusted real consumption and real GNP or GDP growth for the next half-year. The OECD did not issue a forecast for the U.K. in the December 1972 issue of the Economic Outlook. The observation for the first half year of 1973 was therefore omitted in the analysis of the U.K. data. The forecasts for all countries except France start in 1968.II (45 observations). The French data start in 1969.II. The realized series are taken from the OECD Economic Outlook database. The semi-annual series are seasonally adjusted. Following the practice of the OECD Economic Outlook, semi-annual growth rates are expressed at annual rates, i.e. if X_t denotes the level of the series in half year t, then the semi-annual growth rate at an annual rate is given by $100[(X_t/X_{t-1})^2-1]$.

Two data issues arise. First, consumption is measured as total consumption including expenditures on durables. Second, the semi-annual sampling interval may not be consistent with the decision interval of consumers leading to temporal aggregation bias. The first data issue is relevant because the random-walk interpretation of the permanent income hypothesis is consistent with predictability of variations in the growth rate of consumer durables. In the simplest extension of the permanent income hypothesis, Mankiw (1982) shows that expenditures on durables should follow a first-order moving average process:

$$\Delta c_t = \epsilon_t - (1-\delta)\epsilon_{t-1}, \tag{4}$$

where δ is the exponential depreciation rate on the stock of durables. Thus, changes in the growth rates of expenditures should be predictable. Unfortunately, consumption growth forecasts for all G-7 countries are only available for total consumption growth. But if the inclusion of durables is indeed the cause for the rejection of the RWH, equation (4) implies that for reasonable values of δ, predictability of consumption growth should be largely concentrated in the high-frequency range of the data. This conclusion follows because for a first-order moving average process with a large negative moving average parameter, most of the power of the spectrum is concentrated in the high-frequency range. I note, however, that this conclusion need not hold for models of durables with adjustment costs or non-exponential depreciation which can give rise to more complex time series processes (see, e.g., Bernanke, 1985).[4]

The effects of time aggregation of consumption data on tests of the RWH are studied by Christiano, Eichenbaum and Marshall (1991). If consumption follows a random walk at a time interval small than the decision interval of consumers, a famous result by Working (1960) implies that the growth rate of time-averaged consumption follows:

$$\Delta c_t = \epsilon_t + \alpha\epsilon_{t-1}. \quad (5)$$

For example, if monthly consumption is a random walk and Δc_t is the growth rate of semi-annual consumption, α is about 0.25. Thus, if time aggregation is indeed the source of predictability of consumption growth, the amount of predictability should be consistent with the following two considerations. First, if forecasters fully exploit the information contained in ϵ_{t-1}, the maximum achievable R^2 is $\alpha^2/(1+\alpha^2)$ or about 0.06 for α equal to 0.25. Second, the maximum R^2 of about 0.06 is likely to be an upper bound because ϵ_{t-1} is usually only incompletely known to forecasters.

III. Empirical Results

I consider first the predictability of consumption growth at different frequency bands using regressions (2) and (3). Analogous regression results are reported for real income growth approximated by real GDP or GNP growth for two reasons. First, if income growth is found to be unpredictable, the finding that consumption growth is unpredictable as well does not provide convincing evidence in favour of the RWH because a rival model which assumes that consumers simply set consumption equal to some fraction of income would also be consistent with the data (Mankiw and Shapiro, 1985).[5] Second, if the RWH is considered as a rough approximation that allows for a small amount of predictability, consumption growth should still be significantly less predictable than income growth.

Table 1 reports the regression results. The first and second column give the country and the series used in the regressions. The third, fourth and fifth columns report the regression results for equation (2) using all data frequencies. These three columns report the coefficient β with t-statistics in parentheses, the R^2 statistic and the Ljung-Box Q-statistic for general serial autocorrelation with marginal significance levels below the Q-statistic, respectively. The remaining two columns give the band spectrum regression coefficients β_{Low} and β_{High} with t-statistics in parentheses. The subscript "Low" indicates that the regression is restricted to the low frequency band which excludes all cycles in the data that take less than three years to complete. The subscript "High" indicates that the regression was run with data that excluded all cycles taking more than three years to complete.

The null hypothesis that variations in consumption growth are unpredictable across the whole frequency band is strongly rejected at the 1 percent significance level for all G-7 countries. A minimum of 16 percent and up to 41 percent of total variation in semi-annual consumption growth are explained by the forecasts. The size of the R^2's for all countries is much larger than the maximum upper bound of 0.06 which is implied by the view that temporal aggregation bias is the source of predictability of consumption growth. Results for income growth are roughly similar to the results for consumption growth. For four of the G-7 countries, namely Japan, Germany, Italy, and the U.K., the R^2's actually indicate that consumption growth is more predictable than income growth. Even for the U.S. and Canada, where almost 50 percent of the variation in income growth is predictable, the predictable portion of consumption growth is still roughly half of that of income growth. The Q-statistics for Italy and Canada suggest some serial correlation in the residuals for the consumption growth regressions. Also, given the sample period, heteroskedasticity of the residuals might be a problem. I re-estimated all regressions using the Newey-West procedure for correcting the variance-covariance matrix for heteroskedasticity and serial autocorrelation in the residuals up to lag 2. The outcome of the significance tests was unaffected by using the Newey-West corrected standard errors for β.

The band spectrum regression results reported in the last two columns of table 1 strongly indicate that predictability of consumption growth is a low-frequency phenomenon. The null hypothesis of unpredictability of consumption growth is rejected at the low-frequency band for all G-7 countries. In contrast, at the higher frequencies, unpredictability of consumption growth is not rejected for any of the G-7 countries. At least according to the simplest extension of the permanent income hypothesis to expenditures on durables given in equation (4), the results for the higher frequencies do not point to the inclusion of durables as the primary cause of the rejection of the RWH. Interestingly, income growth variations at the high-frequency band appear to be predictable in several of the countries, with predictability being especially strong for the U.S. and Canada. I note that these are also the two countries where the R^2's for the unrestricted income-growth regressions are substantially higher than the R^2's for the unrestricted consumption-growth regressions.

The estimates from the band spectrum regressions are subject to an econometric caveat. The regressions require the right-hand side regressor to be orthogonal to the error term at all leads and lags. This strong assumption is unlikely to be met exactly if forecasters adapt their forecasts based on past forecast errors. In particular, the forecast error of the previous period may be correlated with the forecast for the current period. The simplest route to gauge the severity of this problem is to evaluate the statistical significance of the cross-correlations between forecasts and past forecast errors. Approximate standard errors indicate, however, that significant cross-correlations are quite rare.[6] This evidence suggests that non-orthognonality of forecasts and past forecast errors is unlikely to account for the reported frequency-domain results.

The regression evidence collected in table 1 raises the question why consumption growth exhibits strong predictability at the low-frequency band of the data. One popular explanation for rejections of the RWH holds that consumption growth is predictable because liquidity constraints tie consumption growth to movements in income growth. If liquidity constraints explain the ability of forecasters to predict future consumption growth, predicted income growth should move with realized consumption growth. To test for a possible relationship between realized consumption and predicted income growth, the following regression is considered in the frequency domain:

$$\Delta c^* = \mu \Delta y^*_f + u^*, \qquad (6)$$

where $\Delta y^*_f = Z\Delta y_{t,f}$, and $\Delta y_{t,f}$ is the OECD forecast of real income growth for period t issued at t-1. Regression (6) is closely related to a model proposed by Campbell and Mankiw (1989) for estimating the portion of income received by rule-of-thumb consumers. According to this model, a share μ of total income accrues to consumers who set consumption equal to income and the remaining share of income $(1-\mu)$ accrues to consumers who obey the RWH. Campbell and Mankiw (1989) use an instrumental variables approach to generate the "predicted income series" and they do not consider estimates for regression (6) at specific frequency bands.[7]

Table 2 contains the results for regression (6). The regression coefficients based on the total frequency band show that realized consumption growth is significantly related to predicted income growth in all G-7 countries. Moreover, the R^2-estimates are similar to the R^2's found for the regressions of realized consumption growth on predicted consumption growth. The band spectrum estimates show that the evidence against the RWH is again largely concentrated at the low-frequency band of the data. At the high-frequency band, significant co-movements between realized consumption growth and predicted income growth are only found for Italy. This deviating result should, however, be treated with caution given the highly significant Q-statistic for the Italian estimates across all frequencies.[8]

The empirical evidence presented in tables 2 and 3 suggests that predictability of consumption growth is largely concentrated at the business-cycle frequencies. If this finding is correct, time-aggregation of consumption data, i.e. changing the sampling interval of the data from a short to a longer span, should lead to a substantial increase in the predictability of consumption growth. The intuition is that increasing the sampling interval reduces the importance of high-frequency movements and correspondingly increases the importance of low-frequency movements as a portion of the total variation of consumption growth.

Two pieces of evidence support the conjecture that lengthening the sampling interval increases the predictability of consumption growth. First, consider figures 1a, 1b, and 1c. Figure 1a plots the quarterly growth rate of total U.S. real consumption over the time period 1968.Q2-1990.Q4. Additionally, figure 1a also gives a quarterly forecast series issued by Wharton Econometrics Forecasting Associates (WEFA) over the time period 1970.Q3-1988.Q1.[9] Figure 1b plots realized and predicted U.S. total consumption growth using the semi-annual data analyzed in this paper. Finally, the OECD also publishes end-of-year forecasts of annual consumption growth for the next year. Figure 1c shows realized and predicted annual growth rates over the time period 1968-1990. The three figures illustrate that as the sampling interval of the data increases, low-frequency movements become more important compared to the overall variation in the series. Moreover, comparing realized and predicted consumption growth in figures 1a, 1b, and 1c, the forecasts track the movements in consumption growth increasingly better as the sampling interval increases and low-frequency movements become more important relative to the total variation in the realized series. Second, table 3 compares adjusted R^2's from regressing realized on predicted consumption growth using semi-annual and annual data for the G-7 countries. The R^2's for the semi-annual data are taken from table 1. The results corroborate for all countries that predictability of consumption growth increases substantially as the sampling interval of the data becomes longer. In fact, the portion of predictable consumption growth variations goes up for some countries to as much as 50 percent of the total variation.

IV. Conclusions

According to Hall's (1978) interpretation of the permanent income hypothesis, consumption should be well approximated by a random walk with drift. The starting point of this paper were two stylized facts. First, the random-walk implication is statistically rejected by most empirical work. But second, only a small portion of the future variation in consumption growth appears to be predictable. The paper employed consumption and income growth forecasts for the G-7 countries to study the predictability of consumption growth in the frequency domain. The key empirical finding is that unpredictability of consumption growth is largely concentrated at the low-frequency band of the data. In particular, realized consumption growth moves closely with predicted income growth at the business-cycle frequencies. The results support the conclusion that the statistically small amount of predictability of consumption growth in data with short sampling interval is consistent with substantive predictability at the business-cycle frequencies. It appears that consumption smoothing mainly takes place at the higher frequencies of consumption growth data.

The findings are in line with recent cross-country evidence on consumption behaviour presented by Campbell and Mankiw (1989, 1991) and Carrol and Summers (1991). These studies argue that the co-movements of consumption and income growth are too tight to be consistent with the RWH. The empirical failure of the RWH is unfortunate because it leaves the field without a generally accepted framework for organizing the study of private consumption behaviour (but see Deaton, 1991). On the other hand, as noted early on by Shiller (1984), applied macroeconomists have consistently been skeptical about the usefulness of the RWH as a description of private consumption behaviour. The findings of this paper -- which are based on the output of applied macroeconomists -- appear to rationalize their skeptical attitude.

Table 1

Predictability of consumption and income growth

Country	Series	ß	$\bar{R}^2$	Q	$ß_{Low}$	$ß_{High}$
U.S.	Consumption	0.59**	0.31	23.26	0.86**	0.06
		(4.60)		(0.18)	(4.45)	(0.39)
	GNP	0.86**	0.49	22.89	0.83**	1.01**
		(6.62)		(0.19)	(4.49)	(4.43)
Japan	Consumption	0.55**	0.16	20.27	0.64*	-0.04
		(3.02)		(0.37)	(2.65)	(-0.17)
	GNP	0.34	0.05	24.80	0.50	-0.44
		(1.85)		(0.13)	(1.67)	(-1.54)
Germany	Consumption	0.78**	0.30	15.01	1.08**	-0.30
		(4.44)		(0.66)	(4.79)	(-1.84)
	GNP	0.72*	0.12	15.04	0.76*	0.50
		(2.65)		(0.66)	(2.12)	(0.85)
France	Consumption	0.44**	0.16	21.28	0.48**	0.10
		(2.97)		(0.27)	(2.86)	(0.23)
	GDP	0.61**	0.29	18.68	0.62**	0.43
		(4.27)		(0.41)	(3.84)	(1.46)
Italy	Consumption	0.41**	0.20	37.83	0.42*	0.09
		(3.44)		(0.01)	(2.42)	(0.52)
	GDP	0.44*	0.12	23.27	0.39	0.96*
		(2.65)		(0.18)	(1.72)	(2.24)
U.K.	Consumption	0.83**	0.41	12.88	0.97**	0.34
		(5.60)		(0.80)	(5.58)	(0.95)
	GDP	0.77**	0.33	16.94	0.77**	0.79*
		(4.71)		(0.53)	(3.39)	(2.73)
Canada	Consumption	0.87**	0.19	29.60	1.11**	-0.26
		(3.37)		(0.04)	(3.73)	(-0.44)
	GNP	1.21**	0.46	20.40	1.20**	1.24**
		(6.16)		(0.31)	(5.56)	(3.18)

Notes: Numbers below parameter estimates are t-statistics. The subscripts "Low" and "High" indicate the frequency band used in the regression. The low frequency band comprises all cycles in the data taking more than three years to complete. The high frequency band comprises all cycles in the data taking less than three years to complete. Q is the Ljung-Box test statistic for general serial autocorrelation with marginal significance levels in parentheses below the Q-statistic. * and ** indicate statistical significance at the 5 and 1 percent level in a two-tailed t-test.

Table 2

Co-movements of realized consumption growth and predicted income growth at different frequency bands

Country	μ	$\bar{R}^2$	Q	μ_{Low}	μ_{High}
U.S.	0.51** (5.22)	0.37	20.53 (0.30)	0.61** (3.99)	0.19 (1.22)
Japan	0.42* (2.43)	0.10	21.44 (0.26)	0.54* (2.31)	-0.23 (-0.64)
Germany	0.59* (2.58)	0.11	27.77 (0.07)	0.89* (2.45)	-0.45 (-1.45)
France	0.41** (2.76)	0.14	22.09 (0.27)	0.53** (3.24)	-0.23 (-0.66)
Italy	0.41** (3.62)	0.22	48.04 (0.00)	0.40* (2.19)	0.44* (2.51)
U.K.	0.69** (3.39)	0.19	15.77 (0.61)	0.85** (2.80)	0.30 (0.90)
Canada	1.06** (5.56)	0.40	25.90 (0.10)	1.28** (7.33)	0.52 (1.38)

Notes: See table 1.

Table 3

The predictability of semi-annual and annual consumption growth

Country	Adjusted R^2's	
	Semi-annual data	Annual data
U.S.	0.31	0.47
Japan	0.16	0.33
Germany	0.30	0.52
France	0.16	0.32
Italy	0.20	0.31
U.K.	0.41	0.56
Canada	0.19	0.37

Notes: The table reports the adjusted R^2's from regressions of realized consumption growth on predicted consumption growth using semi-annual and annual growth rates.

Figure 1a: Quarterly U.S. Consumption Growth

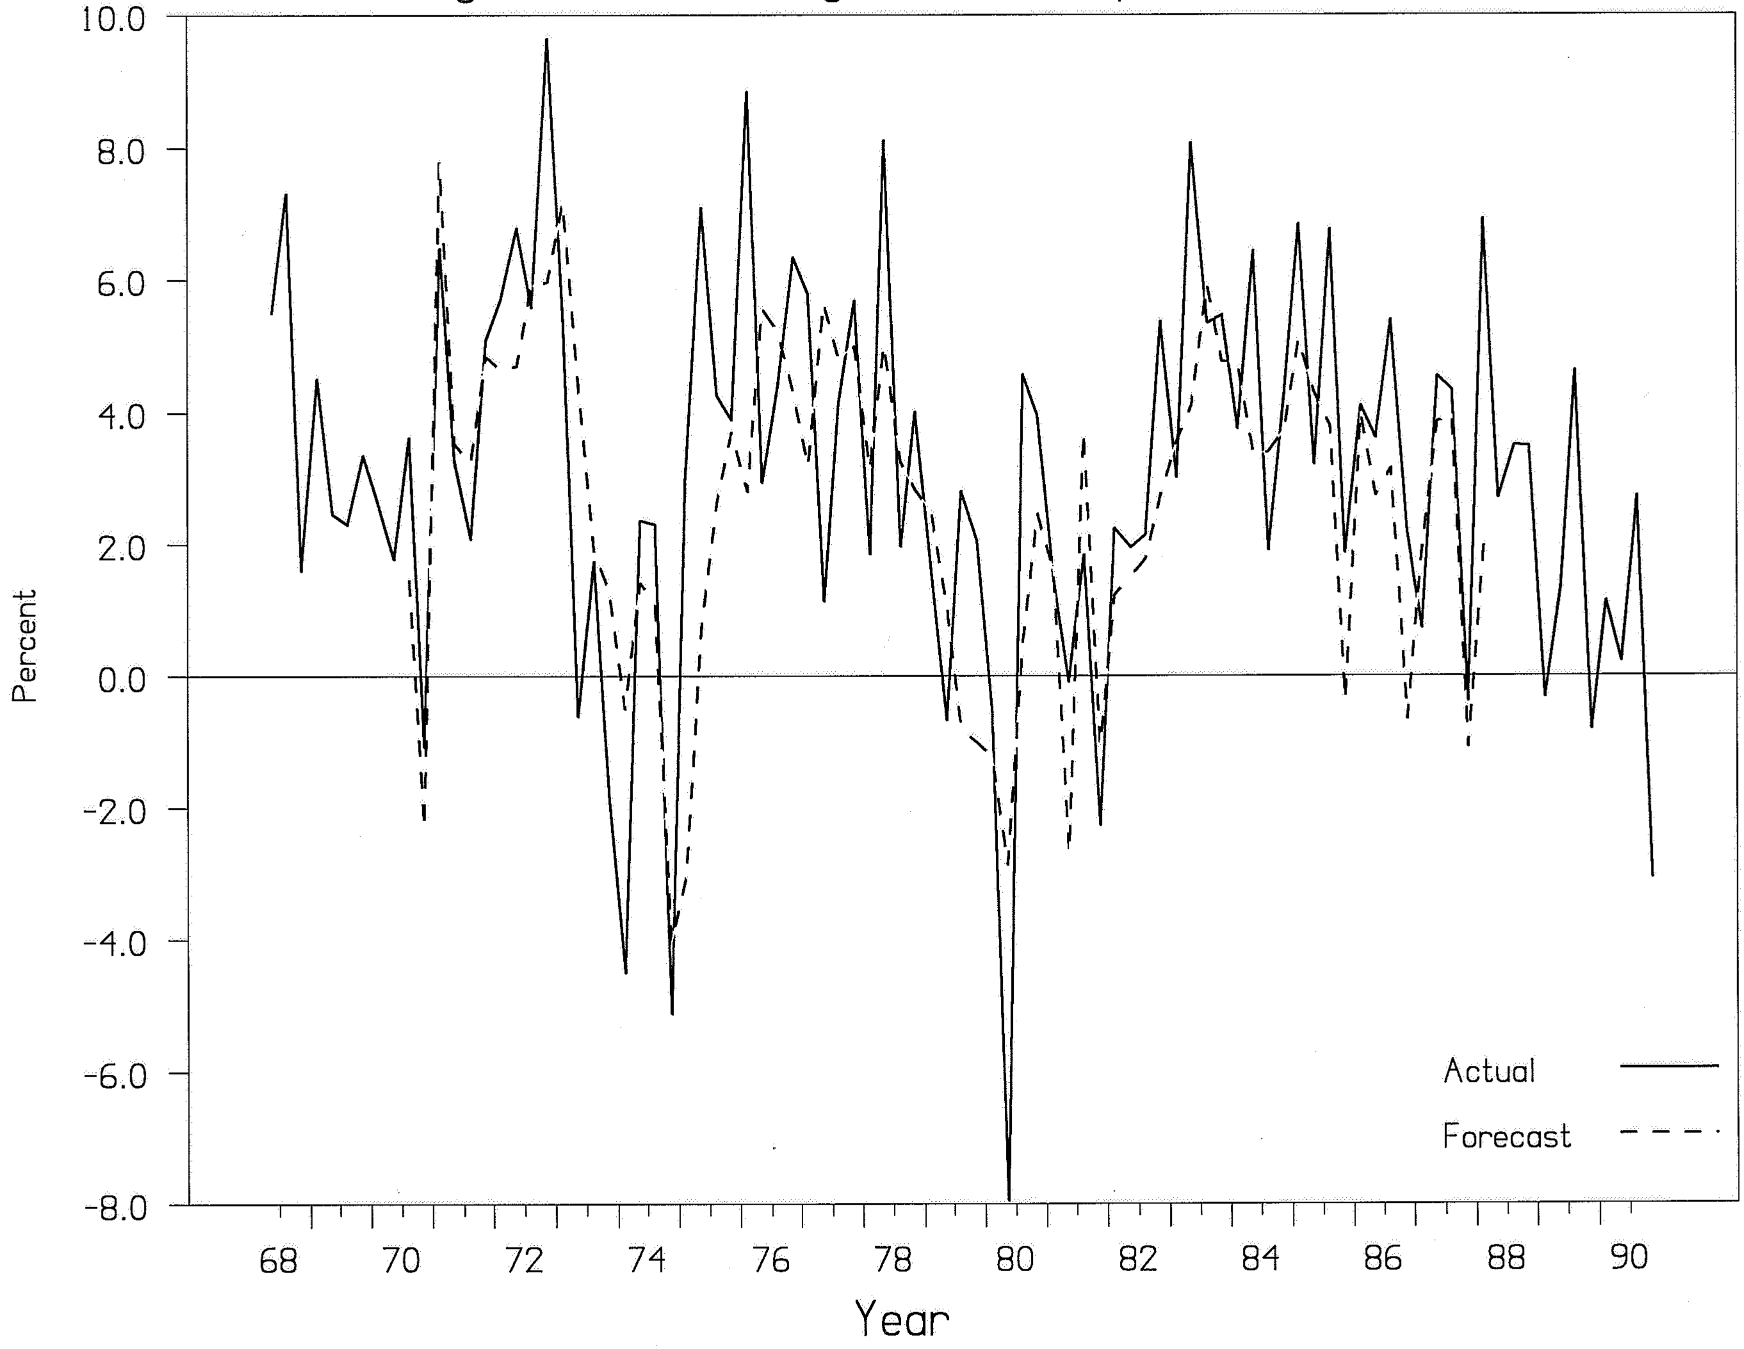

Figure 1b: Semi-annual U.S. Consumption Growth

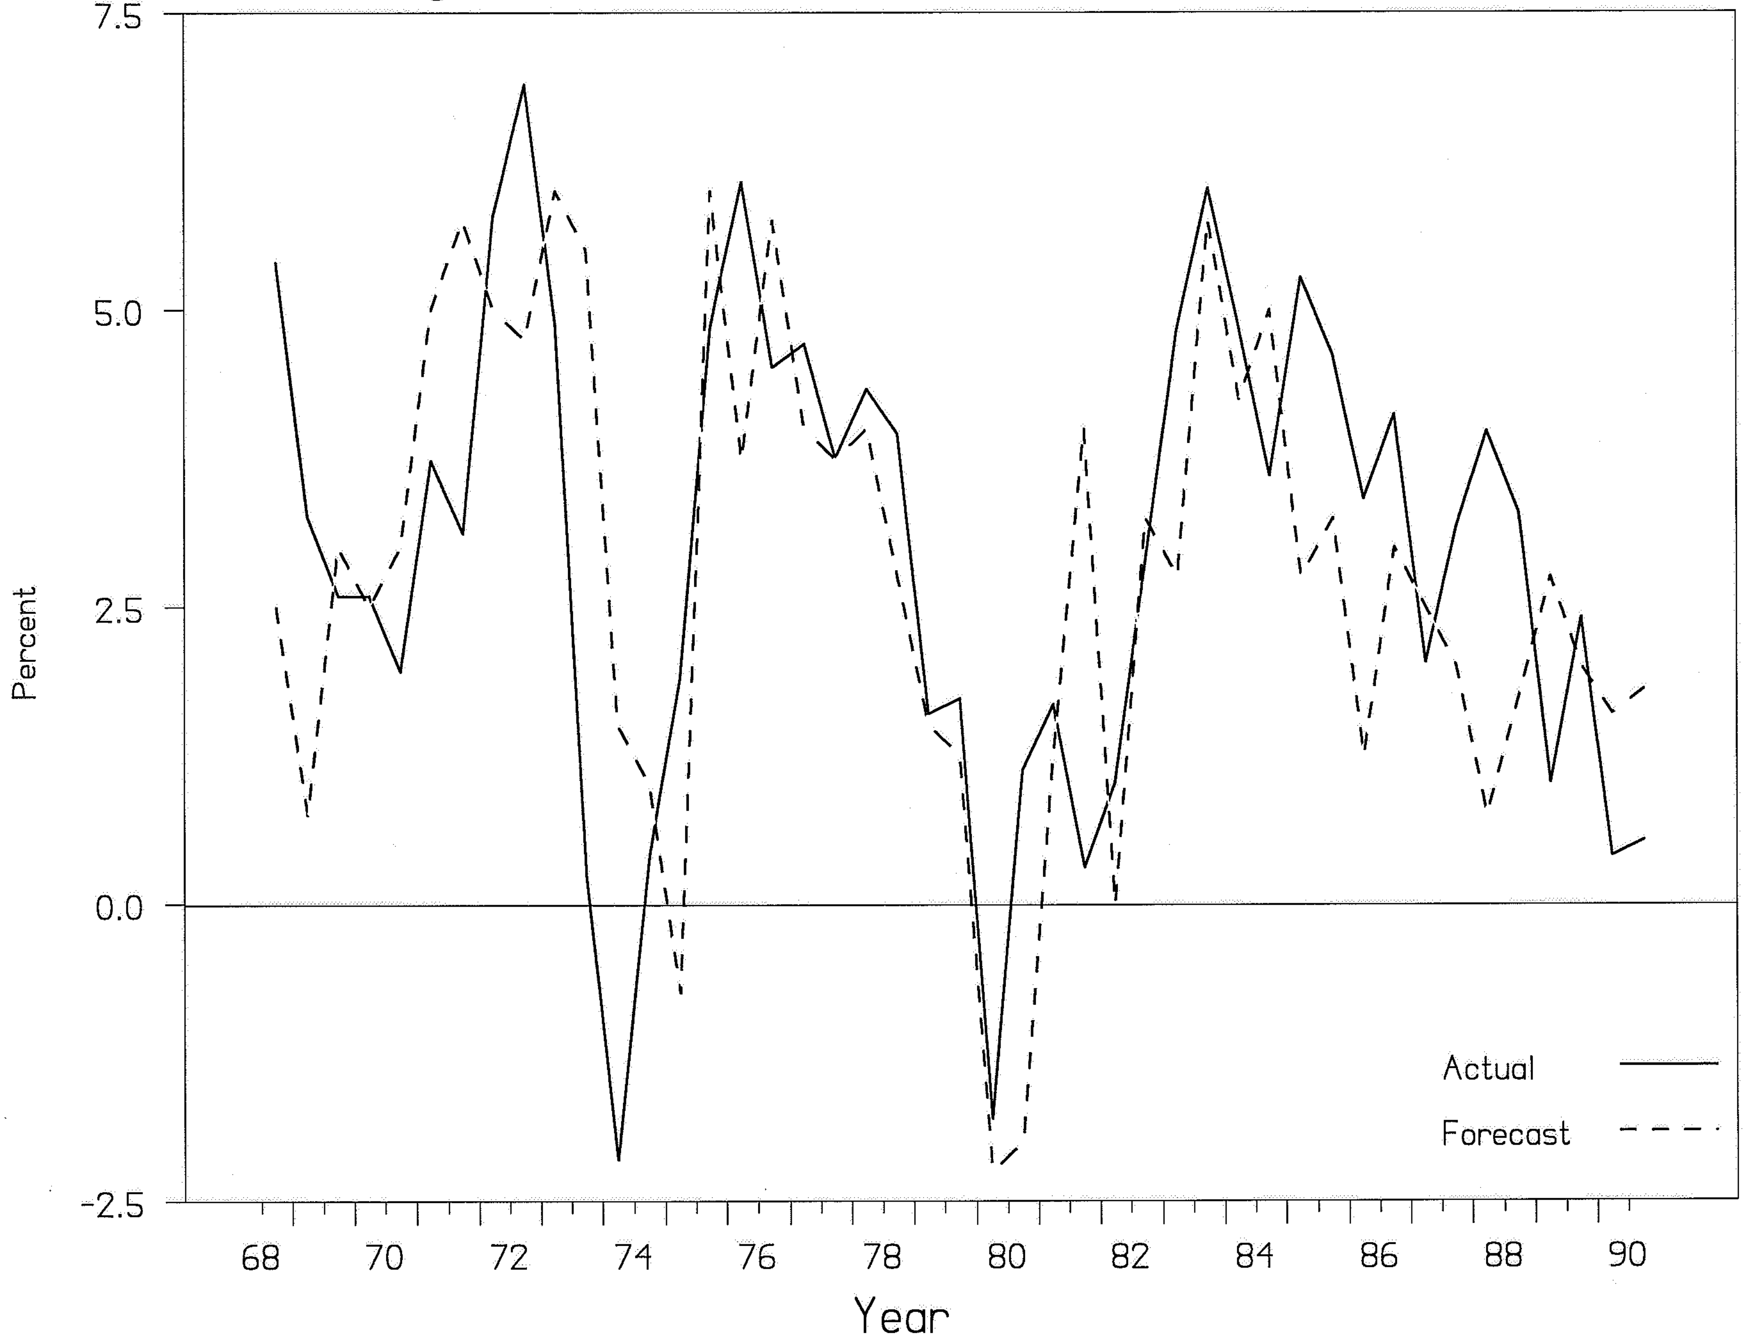

Figure 1c: Annual U.S. Consumption Growth

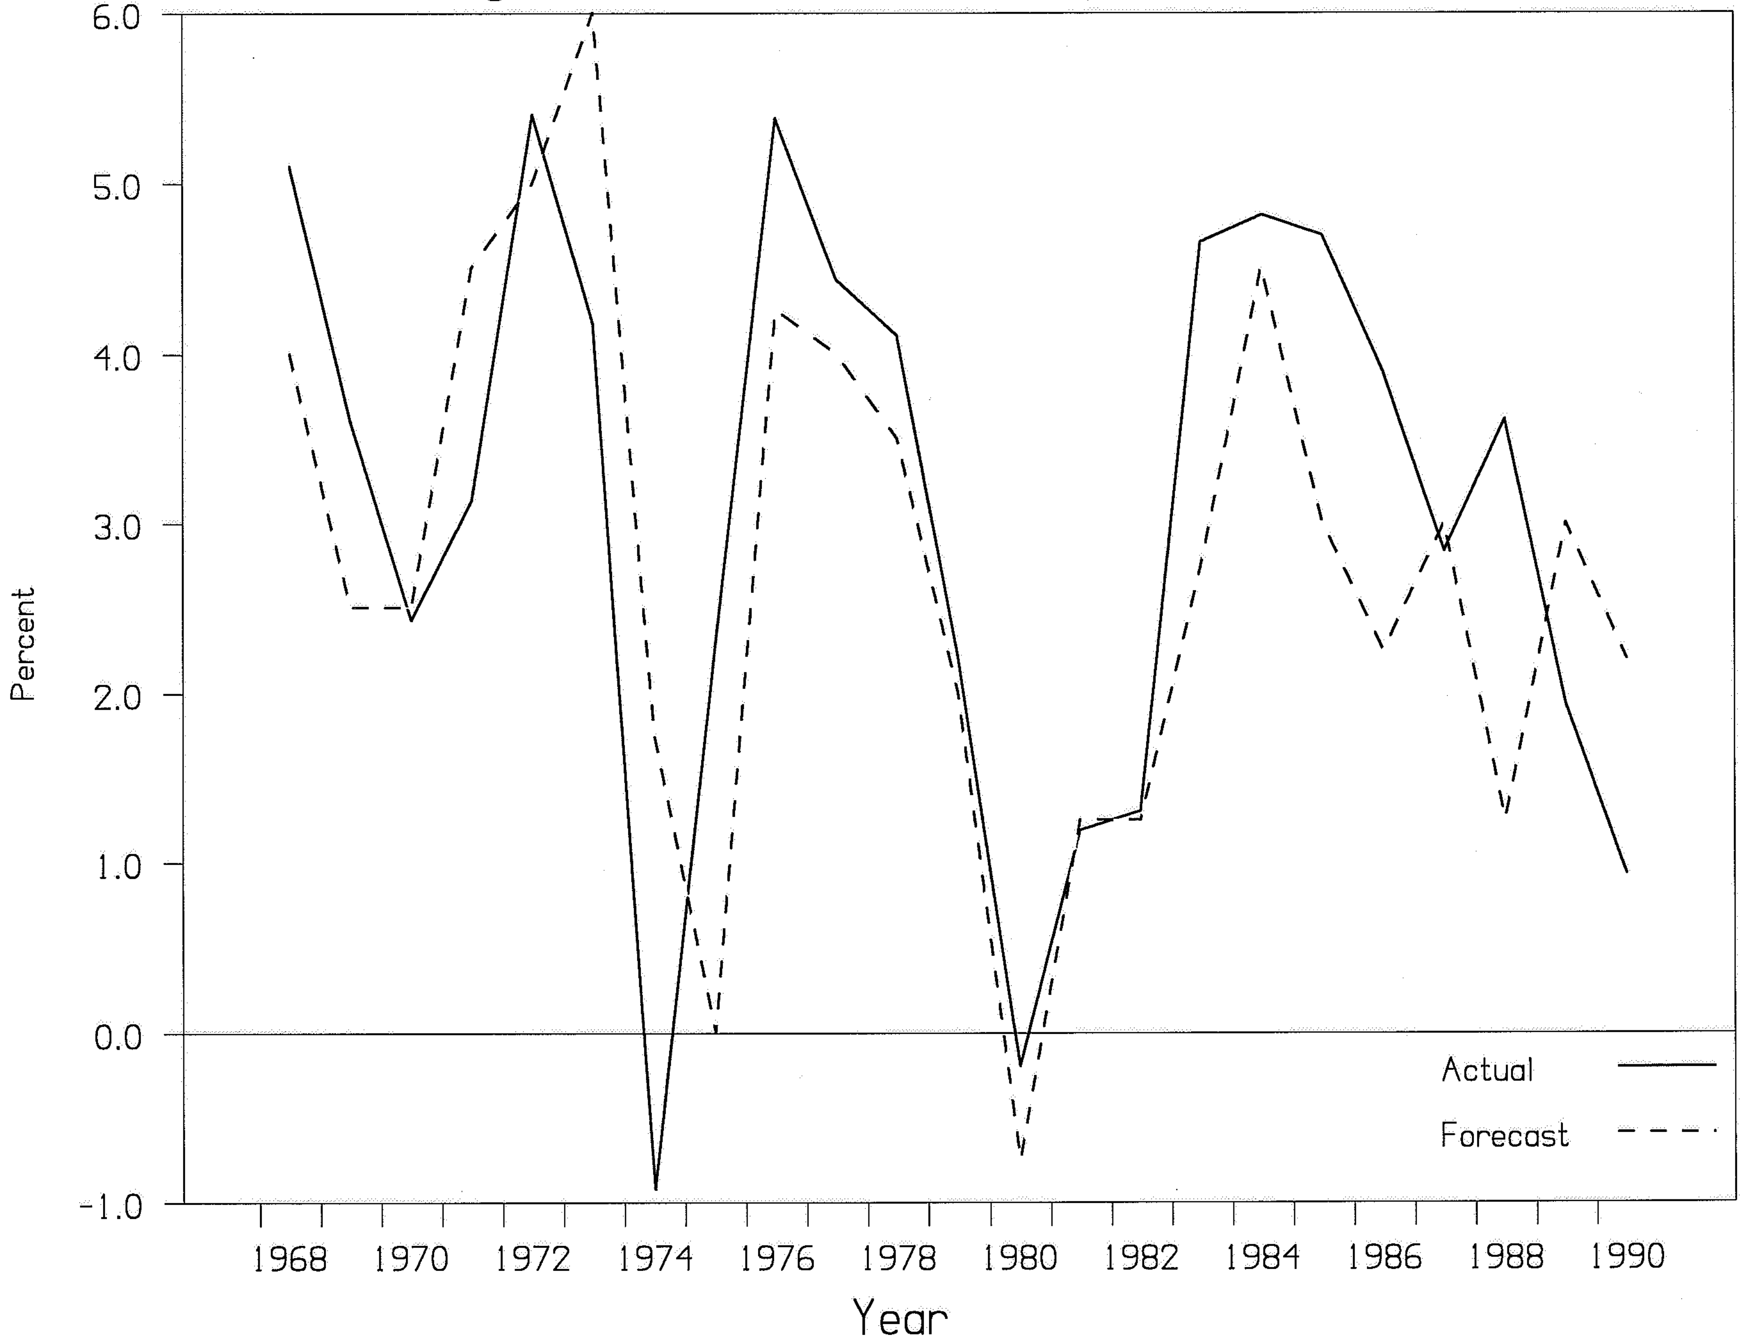

Notes

1 To be accurate, equation (1) says that consumption is a martingale.

2 The Fourier transformations are based on the assumption that the number of observations T is odd, as is the case in this paper. For T even, a slightly modified transformation matrix has to be used (see Harvey, 1978, p. 509).

3 The demarcation of the high- and low-frequency band is, to some extent, arbitrary. But the empirical results are unaffected if, for example, two or four years are used as alternative cut-off points between low- and high-frequency cycles.

4 For U.S. data, Mankiw (1982) and Startz (1989) report that U.S. durables data are well approximated by a random walk.

5 Hall's (1978) relatively favourable evidence on the RWH may be subject to this problem. Some researchers argue that U.S. quarterly income in the 1950s and 1960s behaved close to a random walk (see, e.g., Campbell and Mankiw, 1990).

6 The estimated cross-correlations between the current forecast and the once-lagged forecast error are 0.05 (US), -0.19 (JAP), 0.27 (GER), -0.22 (FRA), -0.38 (ITA), 0.14 (UK), 0.10 (CAN). The cross-correlations between the current forecast and the twice-lagged forecast error are 0.15 (US), -0.06 (JAP), 0.22 (GER), -0.05 (FRA), -0.28

(ITA), 0.20 (UK), and 0.20 (CAN). If forecast errors are serially uncorrelated, the critical values at the 5 percent level are approximately ±0.30.

7 I also obtained regression results based on the Campbell-Mankiw approach for constructing $\Delta y_{t,f}$. The qualitative results reported below remained unaffected.

8 The coefficient estimates reported for μ_{Low} and μ_{High} differ substantially for some countries. Indeed, an F-test of the equality $\mu_{Low}=\mu_{High}$ rejects the null at the 1 percent level for Japan, Germany, and France. Furthermore, μ_{Low} is generally larger than the estimates for μ. This evidence indicates that unrestricted estimates of μ may provide downward biased estimates of the portion of income that accrues to rule-of-thumb consumers as their income fluctuates over the business-cycle.

9 I thank Stephen McNees for kindly supplying me with the quarterly forecasting data.

REFERENCES

Bean, C.R. (1986). "The estimation of `surprise' models and `surprise' consumption functions." Review of Economic Studies, 53, 497-516.

Bernanke, B.S. (1985). "Adjustment costs, durables, and aggregate consumption." Journal of Monetary Economics, 15, 41-68.

Campbell, J.Y. and Mankiw, N.G. (1989). "Consumption, income, and interest rates: Reinterpreting the time series evidence." in: Olivier J. Blanchard, and Stanley Fischer (eds.), Macroeconomics Annual 1989 (Boston: MIT Press), 185-216.

Campbell, J.Y. and Mankiw, N.G. (1990). "Permanent income, current income, and consumption." Journal of Business & Economic Statistics, 8, 265-279.

Campbell, J.Y. and Mankiw, N.G. (1990). "The response of consumption to income: A cross-country investigation." European Economic Review, 35, 723-767.

Carrol, C.D. and Summers, L.H. (1991). "Consumption growth parallels income growth: Some new evidence." in: B. Douglas Bernheim and John B. Shoven (eds.), National Saving and Economic Performance (Chicago: University of Chicago Press), 305-343.

Christiano, L.J., Eichenbaum, M. and Marshall, D. (1991). "The permanent income hypothesis revisited." Econometrica, 59, 397-423.

Cochrane, J.H. (1991). "Comments on: 'The response of consumption to income: A cross-country investigation.' by J.Y. Campbell and N.G. Mankiw. Why test the permanent income hypothesis?" European Economic Review, 35, 757-764.

Deaton, A. (1991). "Saving and liquidity constraints." Econometrica, 59, 1221-1248.

Engle, R.F. (1974). "Band spectrum regression." International Economic Review, 15, 1-11.

Flavin, M.A. (1981). "Excess sensitivity of consumption to changing expectations about future income." Journal of Political Economy, 89, 974-1009.

Hall, R.E. (1978). "Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence." Journal of Political Economy, 86, 971-987.

Harvey, A.C. (1978). "Linear regression in the frequency domain." International Economic Review, 19, 507-512.

Jaeger, A., (1992). "Does consumption take a random walk: Some evidence based on U.S. forecasting data." Review of Economics and Statistics (forthcoming).

Mankiw, N.G. (1982). "Hall's consumption hypothesis and durable goods." Journal of Monetary Economics, 10, 417-426.

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.

Muellbauer, J. (1983). "Surprises in the consumption function." Economic Journal, 93, 34-40.

Shiller, R.J. (1984). "A review essay on `Rational Expectations and Econometric Practice' edited by R.E. Lucas and T.J. Sargent." Journal of Money, Credit, and Banking, 16, 120-123.

Working, H. (1960). "Note on the Correlation of First Differences of Averages in a Random Chain." Econometrica, 28, 916-918.