NBER WORKING PAPERS SERIES

UNIVARIATE VS. MULTIVARIATE FORECASTS OF GNP GROWTH AND STOCK RETURNS: EVIDENCE AND IMPLICATIONS FOR THE PERSISTENCE OF SHOCKS, DETRENDING METHODS, AND TESTS OF THE PERMANENT INCOME HYPOTHESIS

John H. Cochrane

Working Paper No. 3427

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 1990

I thank Gene Fama for helpful comments. Parts of this research were completed while I was a Visiting Scholar the Hoover Institution, and this research was partially supported by NSF grant SES-88-09912. This paper is part of NBER's research program in Economic Fluctuations. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper #3427
September 1990

UNIVARIATE VS. MULTIVARIATE FORECASTS OF GNP GROWTH AND STOCK RETURNS: EVIDENCE AND IMPLICATIONS FOR THE PERSISTENCE OF SHOCKS, DETRENDING METHODS, AND TESTS OF THE PERMANENT INCOME HYPOTHESIS

ABSTRACT

Lagged GNP growth rates are poor forecasts of future GNP growth rates in postwar US data, leading to the impression that GNP is nearly a random walk. However, other variables, and especially the lagged consumption/GNP ratio, do forecast long-horizon GNP growth, and show that GNP has temporary components. Labor income and stock prices (using the dividend/price ratio) display the same behavior. This paper documents these facts and examines their implications for the persistence of shocks to GNP and time-variation in expected stock returns. I find that GNP has an almost entirely transitory response to a GNP shock that holds consumption constant. This is intuitive: if consumption does not change, permanent income did not change, so the change in GNP should be transitory. Similarly, a stock price shock that holds dividends constant suggests a discount rate change, and prices display a large transitory movement in response to this shock. The paper also examines implications of transitory variations in GNP and labor income for methods of extracting stochastic trends or "cyclically adjusting" GNP, and for explaining "excess smoothness" violations of the permanent income hypothesis.

John H. Cochrane
Department of Economics
University of Chicago
1126 E. 59th St.
Chicago IL 60637

Univariate vs. Multivariate Forecasts of GNP Growth and Stock Returns: Evidence and Implications for the Persistence of Shocks, Detrending Methods, and Tests of the Permanent Income Hypothesis.

1. Introduction

Lagged GNP growth rates are poor forecasters of future GNP growth rates in postwar US data and data from many other countries. A high GNP growth rate today does not seem to signal lower growth rates in the future that would bring the level of GNP back towards a trend. This observation has led many researchers to the conclusion that GNP is roughly a random walk. However, other variables are much better forecasters of future GNP growth, and this multivariate evidence shows that GNP does in fact contain strong mean-reverting or temporary components.

The consumption/GNP ratio is the most important variable for long-run GNP forecasts, since it is stable over long periods of time, since it is highly autocorrelated, and since consumption is nearly a random walk. GNP declines more than consumption in a recession, so the consumption/GNP ratio rises. Viewing such a rise in the consumption/GNP ratio, one can forecast that GNP must eventually rise again to reestablish the historical ratio. Thus a change in the consumption/GNP ratio can be used to forecast long term movements in GNP growth. Since consumption is nearly a random walk, GNP must do most of the adjusting: the ratio forecasts long term GNP growth rather than long term consumption growth. Its high autocorrelation means that the consumption/GNP ratio can pick up long horizon movements in GNP that a more choppy right hand variable might miss.

Roughly similar statements are true of labor income and stock prices. Lagged labor income growth is a poor forecaster of future labor income growth, and lagged returns are poor forecasters of future returns. The consumption/labor income ratio is a much better forecaster of labor income growth, and the dividend/price ratio is a much better forecaster of returns. Each implies much more temporary variation than is suggested by univariate regressions.

This paper documents this characterization of the data, extending the results of Cochrane and Sbordone (1988), King, Plosser, Stock and Watson (1989), and especially Fama (1990), and explores its implications for measurements of the persistence of shocks to GNP, for methods of detrending or "cyclically adjusting" GNP, for the excess smoothness of consumption found in tests of the permanent income hypothesis, and for time-variation in expected stock returns.

Section 2 studies the persistence of shocks to GNP. This issue has been the focus of a large body of empirical research (see Nelson and Plosser (1982), Campbell and Mankiw (1987) Clark (1987), Cochrane (1988), Cogley (1990) among many others). If GNP growth rates are not forecastable, then shocks are permanent, so there is no "business cycle" to study.

We will find two novelties in using multivariate information to study the persistence of shocks to GNP. First, with several variables there are several shocks, and these shocks contain more information than univariate shocks. A univariate "shock to GNP" is a movement in GNP growth not forecast by past GNP growth. But a multivariate shock to GNP is a movement in GNP growth not forecast by past GNP growth *and other variables*, and thus contains more information. Furthermore, GNP will generally respond to (multivariate) shocks to other variables as well as shocks to GNP, where we can only trace its response to its own shocks in a univariate regression.

I find that GNP behaves much like a random walk in its univariate representation, but that GNP displays transitory variation in response to multivariate shocks. GNP's response to a consumption shock is partly permanent but also partly temporary. More importantly, GNP's response to a GNP shock *holding consumption constant* is almost *entirely* transitory. This has a natural interpretation: If consumption does not change, permanent income must not have changed, so any such change in GNP must be entirely transitory. Thus, by isolating GNP shocks with no consumption change, the multivariate system is better able to document temporary components in GNP than regressions of GNP growth on lagged GNP growth.[1]

Second, even if we are only interested in the response of GNP to univariate shocks, an *estimate* of that response formed from a regression of GNP growth on its own past is different from an estimate formed by finding the univariate GNP process implied by a regression of GNP growth on other variables. This is a general proposition: a pth order vector autoregression implies different univariate processes than pth order univariate autoregressions (Zellner and Palm (1978)). But it is especially true when the lagged consumption/GNP ratio is a right hand variable. The VAR imposes that this ratio is stationary, so the permanent component of GNP cannot vary more than that of consumption. The VAR exploits this information to make an improved estimate of GNP's long-run response to univariate shocks, and this estimate shows more transitory variation than estimates based on univariate autoregressions.

Section 3 examines implications of multivariate GNP forecasts for detrending or "cyclically adjusting" GNP. The linear trends of the 60's broke down with the "productivity slowdown" of the 70's, and there has since been much interest in estimating *stochastic* trends for GNP. This issue (in part) originally motivated the literature on persistence and unit roots. A good stochastic trend should not respond to business cycles (if they exist), but should respond to long-term fluctuations in GNP in a sensible way, to allow definitions of cyclically adjusted budget deficits, or ratios of monetary aggregates to cyclically adjusted GNP.

---

[1] Cochrane and Sbordone (1988) show that the permanent income hypothesis implies that consumption and GNP are cointegrated and consumption is a random walk, so one can measure the variance of the permanent component of GNP by the variance of the permanent component of consumption, which is consumption itself. Fama (1990) explores the permanent income story in detail and uses it to interpret the response of GNP and investment to a consumption ("wealth") shock. He suggests the use of consumption to measure the permanent component of income, and documents the importance for forecasting GNP growth of the consumption/GNP ratio together with the observation that consumption is nearly a random walk.

As univariate and multivariate forecasts of GNP growth differ, univariate and multivariate estimates of stochastic trends differ. A stochastic trend based on a univariate GNP autoregression is essentially the same as GNP itself, since that autoregression has little power to forecast GNP growth. The multivariate estimate of a stochastic trend is very close to consumption multiplied by the mean GNP/consumption ratio. Thus it responds to long-term movements in GNP as consumption does, but moves little over business cycles. It can also be interpreted as an instance of the permanent income hypothesis: the multivariate trend in GNP is (approximately) permanent income, as revealed by consumption.

Section 4 examines implications for "excess smoothness" rejections of the permanent income hypothesis. The persistence of income shocks has had an immediate economic application in this area. If income really is a random walk, as the univariate evidence suggests, then consumption changes should equal income changes. Since the variance of consumption changes is a good deal less than that of income changes, the permanent income hypothesis has been rejected in favor of "excess smoothness" of consumption. (See Deaton (1987), Campbell and Deaton (1989), Campbell and Mankiw (1989). Flavin (1988), Hansen, Roberds and Sargent (1990) and Quah (1990) criticize this literature.) Section 4 verifies this finding based on univariate labor income regressions, but finds that the mean-reversion in labor income implied by *multivariate* estimates easily explains the smoothness of consumption growth.

Section 5 uses the same techniques to examine time-variation in expected stock returns, or equivalently, the existence of temporary components in stock prices. Fama and French (1988a) and Poterba and Summers (1988) found some evidence that lagged returns forecast future returns. Richardson (1990) and others have argued that the apparent forecast power of lagged returns is statistically insignificant. However Fama and French (1988b) find that other variables, and the dividend/price ratio in particular, are strong and statistically significant predictors of future returns. (See also Hodrick (1990) for a statistical investigation.) Similarly, Cochrane and Sbordone (1988) find that multivariate generalizations of Poterba and Summers' variance ratios that include dividends indicate much larger temporary components in prices.

I find that prices and dividends behave much like GNP and consumption. Returns have very little univariate predictability, so prices look like a univariate random walk. But prices display very different responses to the two shocks one can define in a bivariate system. The response of prices (and dividends) to a dividend shock is almost entirely permanent; the response of prices to to a price shock holding dividends constant is entirely transitory, while dividends show no response to this shock.

This also has a natural interpretation. A shock to dividends can come with no change in discount rates, and hence no change in expected returns. The shock to dividends has an entirely permanent effect on dividends. (One interpretation of this feature has managers setting dividends to "permanent earnings", inducing a random walk just like permanent income consumers.) Hence, the shock to dividends should and does have an entirely permanent effect on prices. On the other hand, a shock to prices with no contemporaneous change in dividends suggests a discount rate or risk premium change. This changes expected returns, and thus sets in motion expected changes in prices. Eventually, discount rates return to their mean and prices return to their customary multiple of dividends, which were unaffected by the discount rate shock. Thus, by isolating discount rate changes as shocks to prices with no change in dividends, the multivariate system is able to document time variation in expected returns that is missed by regressions of returns on lagged returns.

2. Measuring the persistence of GNP in postwar US data.

Table 1 presents a vector autoregression of log GNP and nondurable + services consumption growth on lagged log GNP and consumption growth and the lagged log consumption/GNP ratio. It also presents a univariate autoregression of GNP growth on lagged GNP growth. All the calculations that follow are based on these regression coefficients.[2]

The VAR includes two lags of each variable, while the univariate autoregression includes four lags. More lags do not change the qualitative results, but just add wiggles to the impulse-response and spectral density functions. Both regressions are generous by the usual specification tests, for example the last lag is statistically insignificant. I also experimented with a variety of extra right hand variables, including stock returns and term and default premia. These variables significantly forecast GNP growth, but they do not have much effect on the long-run impulse-response functions. Like more lags, extra variables basically just add wiggles to the impulse-responses at short horizons.

Several features of the regressions in table 1 are noteworthy. The multivariate GNP forecast is a good deal better than the univariate forecast: the lagged consumption/GNP ratio is highly significant in the VAR GNP forecasting equation, and the R^2 of the VAR GNP equation is higher than that of the (longer) univariate autoregression. Consumption growth is slightly predictable in the VAR, but with a much lower R^2 than GNP growth.

Fig. 1 presents impulse-response functions of the consumption-GNP VAR,

---

[2]The VAR *imposes* that GNP and consumption are cointegrated, or, equivalently that the log consumption/GNP ratio is stationary and does not contain a unit root or random walk component. Cochrane (1989) provides a critique of the methodology in which one conducts tests for cointegration and then imposes the results in subsequent analysis, which is why such tests are absent here. Plots of the consumption/GNP, consumption/labor income and dividend/price ratios show that they do not have trends, and suggest that the assumption that these ratios are stationary is not unreasonable. (The same is not true of all ratios, for example nondurables alone/GNP.) However, the whole point of these ratios is that they are slowly mean-reverting, and hence can forecast slowly mean-reverting behavior in GNP or stock prices. Hence, they are likely to spuriously fail to reject a unit root test if one does not allow for ample serial correlation.

All the left hand variables are one period growth rates. Hodrick (1990) argues that this gives better statistical performance than aggregated left hand variables that have been common in the mean reversion literature.

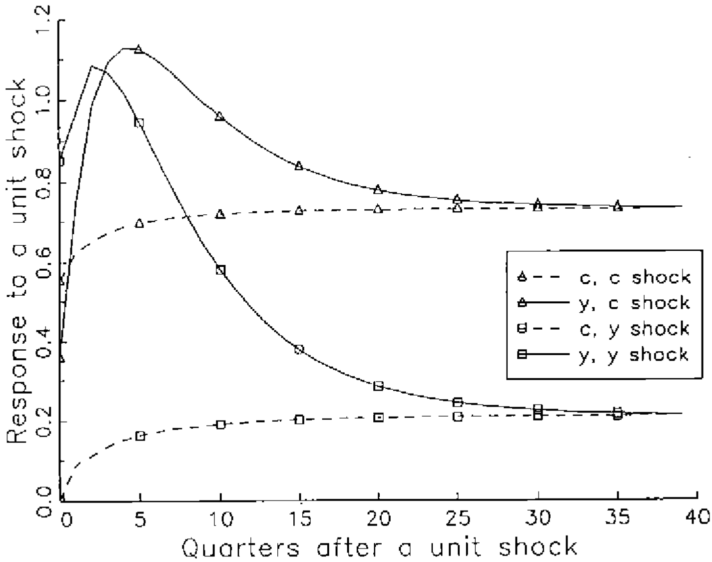

Fig. 1. Impulse-response functions for consumption-GNP VAR. Responses of consumption (c) and GNP (y) to unit consumption and GNP shocks. The VAR is a regression of consumption and GNP growth on the lagged consumption/GNP ratio and two lags of consumption and GNP growth (see Table 1). The shocks are orthogonalized to force to 0 the instantaneous response of consumption to a GNP shock.

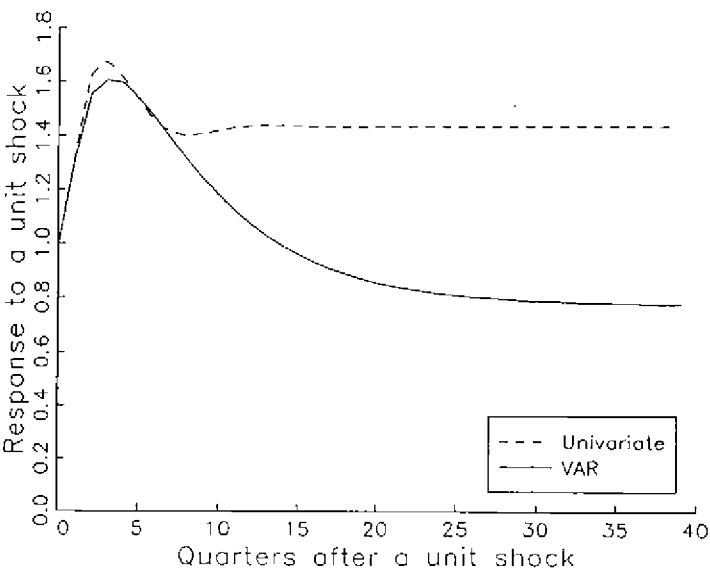

Fig. 2. Univariate GNP impulse-response functions. The "Univariate" response is estimated from a regression of GNP growth on 4 lags of GNP growth. The "VAR" response is estimated from a vector autoregression of GNP and consumption growth on two lags and the lagged consumption/GNP ratio.

calculated from the regressions in table 1. (The appendix details the calculations.) The VAR errors are orthogonalized with consumption first and then income. Equivalently, the instantaneous response of GNP to a consumption shock is forced to zero. This is also equivalent to including current GNP growth in the consumption growth regression, but not vice versa.[3]

Several features of fig. 1 are noteworthy. First, the eventual response of consumption is the same as that of GNP to each of the shocks. This results from the assumption that the consumption/GNP ratio is stationary, and hence can be included on the right hand side. If consumption and GNP ended up with different responses to a shock, the consumption/GNP ratio would not be stationary. Second, look at the responses of consumption and GNP to a consumption shock. Consumption is almost a random walk: its impulse response function is almost flat. GNP has an instantaneous response of about 2/3 the consumption change; this rises to about 1 1/2 times the consumption response after 4 quarters, and then declines to equal the consumption change by the time 40 quarters have passed. Fama (1990) interprets this as the response of investment to a wealth shock. Third, look at the response to a GNP shock. This is a shock to GNP that does not contemporaneously affect consumption. It has only a very small eventual impact on consumption, but GNP shows a strong mean-reverting response to this shock.[4] A shock to GNP that does not change consumption must not have changed permanent income, and thus must be

---

[3]Unfortunately, some identification assumption is always needed on VAR errors, since the system can always be equivalently reexpressed in terms of new errors that are a nonsingular linear combination of old errors. Blanchard and Quah (1989) explore an identification method in which the long run response to one shock is forced to zero. As it turns out, the two methods produce nearly the same result, as the income shock here has essentially no permanent component.

[4]The slightly positive long run response to the GNP shock is not robust to changes in specification. Slight changes in variables or sample period yield smaller or even negative responses to this shock. For example, if one uses private GNP (GNP - government purchases of good and services, CITIBASE series GGE82), the long run response to the GNP shock is -.5! In annual data, the positive serial correlation of consumption disappears, so both consumption responses are basically flat, and its response to a GNP shock essentially 0 at all horizons.

temporary.

Fig. 2 presents *univariate* impulse-response functions for GNP.[5] These functions are estimated from the univariate autoregression of table 1 and from the VAR. The impulse-response function estimated from the univariate autoregression displays a good deal of persistence: in response to a unit shock, GNP climbs to about 1.6 after a year, and then declines only to about 1.4. The univariate impulse-response estimated from the VAR has quite similar short run dynamics, but it displays much more mean reversion at long horizons, ending up at about half its peak value.[6]

Compare either of the univariate impulse-responses in fig. 2 to the multivariate impulse-responses in fig. 1. Clearly, whether shocks to GNP are persistent or not depends crucially on the information set one uses to forecast GNP growth. If one observes consumption as well as GNP, one can forecast that a shock to GNP which does not contemporaneously move consumption will almost entirely disappear. If one is restricted to only observing GNP itself, shocks induce a much larger persistent component.

Fig. 3 shows how the univariate GNP dynamics estimated from the VAR differ from those implied by the univariate autoregression in the frequency domain. The long-horizon behavior of GNP is reflected in the spectral density at frequencies near zero, or long periods. (The appendix details the connection between spectral densities at zero and impulse-response functions as univariate measures of persistence.) The VAR uses the information that the spectral densities of consumption and GNP growth must be equal at

---

[5] The vertical scales of the univariate (fig. 2) and multivariate (fig. 1) impulse-responses are not comparable. Loosely, fig. 1 presents the responses to one *standard deviation* shocks (unit shocks of the orthogonalized representation); fig. 2 presents the responses to one *percent* shocks. This follows the conventions in the literature.

[6] The value of the long horizon response is sensitive to variables and samples. For private GNP, it is less than .5. The basic pattern in which the VAR indicates much more mean reversion than the univariate autoregression is not sensitive.

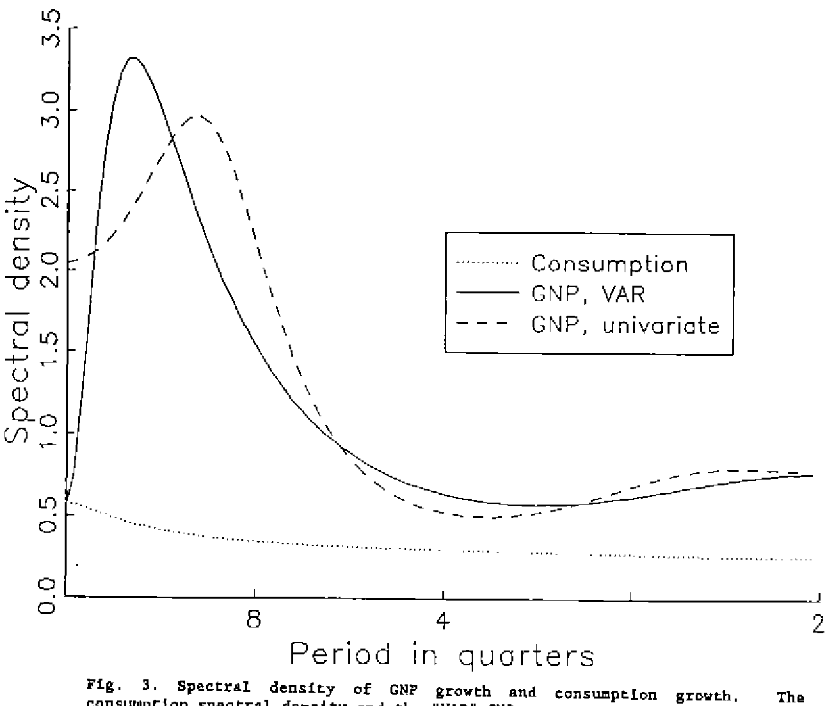

Fig. 3. Spectral density of GNP growth and consumption growth. The consumption spectral density and the "VAR" GNP spectral density are estimated from a vector autoregression of consumption and GNP growth on two lags and the lagged consumption/GNP ratio. The "univariate" GNP spectral density is estimated from a regression of GNP growth on four lags.

frequency zero, so that the consumption/GNP ratio is stationary. This results in the dip in the VAR estimated spectral density of GNP growth near 0, which the univariate autoregression does not pick up.[7] The univariate autoregression and VAR imply similar spectral densities at other frequencies.

3. "Detrending" or "cyclically adjusting" GNP

Given a single time series, there are many ways to break it into "trend" and "cyclical" components. Different GNP forecasts will affect most such decompositions. I will examine one attractive decomposition, due to Beveridge and Nelson (1981). This decomposition defines the stochastic trend in GNP as the level GNP will be after all transitory dynamics work themselves out. Equivalently, the trend in GNP is GNP plus all expected future changes in GNP.

Precisely, Beveridge and Nelson decompose GNP y_t into a stochastic trend z_t and a cyclical component s_t

$$y_t = z_t + s_t.$$

The trend is defined as

$$z_t = \lim_{k\to\infty} E_t(y_{t+k} - kE\Delta y).$$

or, equivalently,

$$z_t = y_t + E_t(\Delta y_{t+1} - E\Delta y) + E_t(\Delta y_{t+2} - E\Delta y) + \dots \quad (3.1)$$

where $E\Delta y$ is the unconditional mean growth rate of y_t. If GNP is expected to grow a lot in the future, GNP is below trend; if it is expected to decline, GNP is above trend. (The appendix relates this trend to the impulse response function and spectral density of GNP growth at frequency 0, and gives formulas for implementing (3.1) from a VAR.)

---

[7] Watson (1990) found a similar dip in the spectral density of output estimated from a cointegrated VAR and also the spectral density implied by the King, Plosser, Stock and Watson (1989) model. The contrast between Watson's results and univariate estimates of spectral densities implied by the persistence literature inspired this section of this paper.

Fig. 4 presents log GNP and the Beveridge-Nelson stochastic trend, estimated from the consumption-GNP VAR of table 1. (Precisely, the conditional expectations E_t in (3.1) are formed from GNP growth, consumption growth and the consumption/GNP ratio, using the VAR regression of table 1.) This stochastic trend responds to long-run movements in GNP growth during the 70's and 80's, yet shows the traditional NBER business cycles as transitory variations about that trend.

If consumption were a pure random walk, the Beveridge-Nelson trend would be exactly consumption plus the mean log GNP/consumption ratio. Fig. 4 also plots this quantity, and shows that it is almost the same as the trend. This provides a nice interpretation: consumption should be proportional to permanent income, which is a natural measure of the trend in actual income. By using consumption, we are in essence using consumer's forecasts of income growth to measure the trend. The Beveridge-Nelson trend differs slightly from consumption plus the mean ratio, as consumption data departs slightly from the predictions of the permanent income hypothesis. Consumption growth is slightly predictable, and the trend calculations exploit this predictability to define the trend as the level consumption will attain in the future (plus the mean ratio) rather than its value today.

Fig. 5 contrasts univariate and multivariate Beveridge-Nelson trends. The "VAR trend" is the multivariate trend from fig. 4. For the "VAR estimated" univariate trend, E_t is formed from current and lagged GNP growth, using the univariate impulse-response function implied by the VAR. In the other univariate trend, E_t is formed from GNP growth, using the parameters of the univariate autoregression. Since the VAR estimated univariate impulse-response displays more mean-reversion than the directly estimated univariate impulse-response, the two trends are different. Even if one insists on using past GNP only to form a detrended GNP, the univariate impulse-response estimated from the VAR shows the recessions of the 80's to be transitory, while the directly estimated univariate impulse-response views these movements in GNP as permanent.

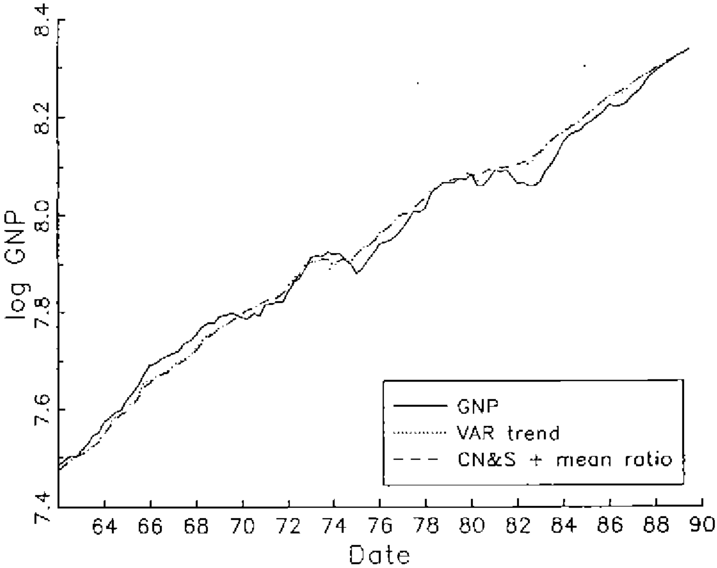

Fig. 4. GNP, Beveridge-Nelson trend, and Consumption. The Beveridge-Nelson stochastic trend is calculated from the consumption-GNP VAR as GNP plus all expected future changes to GNP. "CN&S + mean ratio" gives log nondurable + services consumption plus the mean log GNP/consumption ratio.

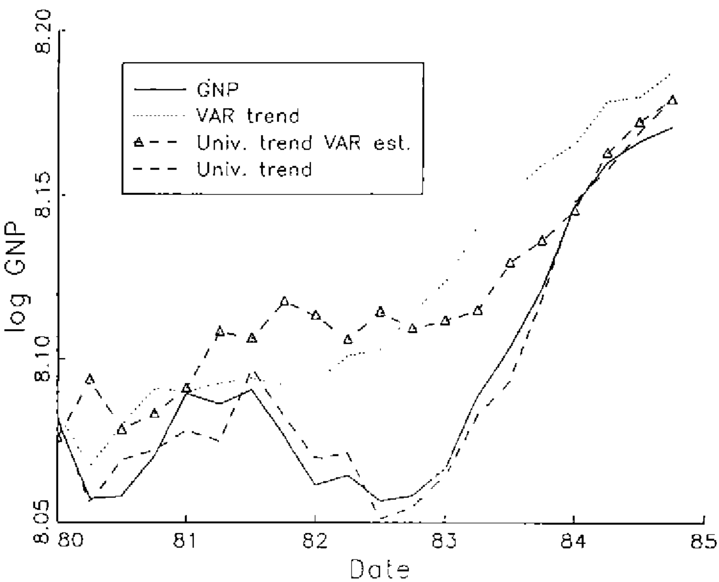

Fig. 5. GNP and trends. The "VAR" trend is calculated from the consumption-GNP VAR as in fig. 4. The "Univ." trend is calculated from a univariate autoregression of GNP growth on four lags. The "Univ. trend VAR est." is a trend based on the univariate representation of GNP implied by the VAR.

4. "Excess smoothness" and the permanent income hypothesis

If labor income e_t follows a process

$$\Delta e_t = \sum_{j=0}^{\infty} \rho_j w_{t-j} = \rho(L) w_t$$

where w_t is a possibly multidimensional white noise process generating all information observed by consumers, then the permanent income model predicts that the change in consumption should equal the change in the present value of future labor income,

$$\Delta c_t = (1-\lambda) \sum_{j=0}^{\infty} \lambda^j \left[E\left(e_{t+j} \mid w_t, w_{t-1}, \ldots\right) - E\left(e_{t+j} \mid w_{t-1}, w_{t-2}, \ldots\right)\right].$$

where $\lambda = \frac{1}{1+r}$ and r is the real interest rate. Sargent (1987) shows that the right hand side is equal to

$$\Delta c_t = \rho(\lambda) w_t \tag{4.1}$$

where $\rho(\lambda) = \rho_0 + \rho_1\lambda + \rho_2\lambda^2 + \ldots$.

If labor income follows a random walk, $\Delta e_t = w_t$, then consumption growth should equal income growth, $\Delta c_t = \Delta e_t = w_t$. A weaker implication is that the variance of consumption growth should equal the variance of income growth. As we have seen, income is nearly a random walk based on univariate autoregressions, but consumption varies a great deal less than income. This is the heart of "excess smoothness" of consumption. Conversely, if one assumes that income follows a stationary univariate process around a trend, as in Flavin (1981), one finds that consumption varies by more than it should under the PIH, which is the original finding of "excess sensitivity" (much simplified).

The observations on persistence of the last sections suggest a resolution of excess smoothness. Estimates of long run impulse-responses $\rho(1)$ in multivariate systems were much smaller than univariate estimates. Since $\lambda \cong 1$, this suggests that multivariate estimates of $\rho(\lambda)$ might also be lower than univariate estimates, so the predicted variance of consumption growth might be smaller.

To address this question, I repeated the VAR and univariate estimation of Table 1, using labor income in place of GNP. For labor income I used personal disposable income less dividend, interest and rent income (CITIBASE series GYD - GPRENJ - GPDIV - GPINT), converted to 1982 dollars using the personal disposable income deflator (GYD82/GYD). Table 2 presents the results. As before, the lagged consumption/income ratio is significant in the income regression but not in the consumption regression. Labor income is closer to a univariate random walk than GNP: the univariate regression of income growth on two lags has an R^2 of 0.000 with a p-value of .99.

Figures 6, 7, and 8 present the impulse-response functions and spectral densities, analogous to figures 1, 2, and 3 for GNP. These results present an even more dramatic case than GNP. The bivariate impulse-response function (fig. 6) shows that income and consumption have almost the same, and completely permanent, response to a consumption shock. The response of income to an income shock (with no contemporaneous consumption change) is almost entirely transitory. The univariate impulse-response estimated from a univariate autoregression (fig. 7) is almost completely flat, as is the spectral density of income growth estimated from the univariate autoregression (fig. 8). Based on univariate information, labor income is almost a perfect random walk. But the univariate impulse response (fig. 7) and spectral densities of income (fig. 8) estimated from the VAR show a substantial mean reversion at long lags. This occurs because, although both income and consumption look like random walks in univariate autoregressions, the variance of income growth is about three times that of consumption growth. The long run movements in the two series must be equal, or the consumption/income ratio would not be stable over time.

To examine the excess smoothness puzzle, I estimated the variance of revisions in permanent income, using the VAR forecasts of income, univariate forecasts of income, and univariate forecasts with the univariate process implied by the VAR. The VAR for consumption and income can be rewritten in orthogonalized moving average representation (see the appendix)

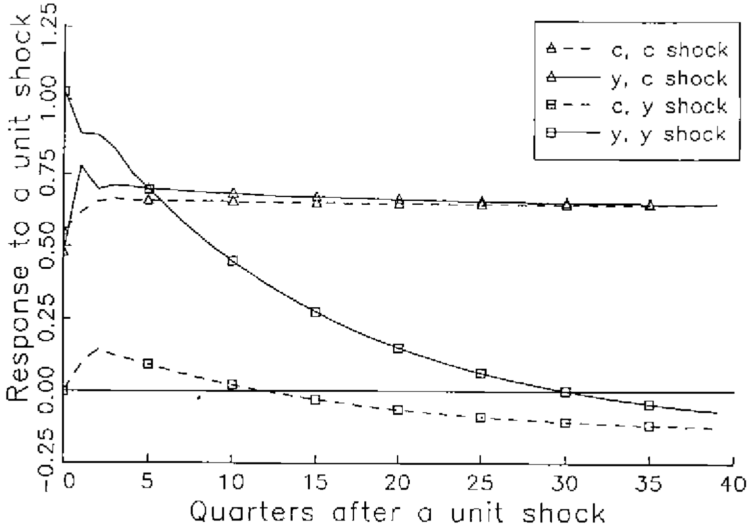

Fig. 6. Impulse-response functions for consumption-labor income VAR. Responses of consumption (c) and labor income (y) to unit shocks. The VAR includes the lagged consumption/income ratio and two lags of consumption and income growth. The shocks are orthogonalized to force to 0 the instantaneous response of consumption to an income shock.

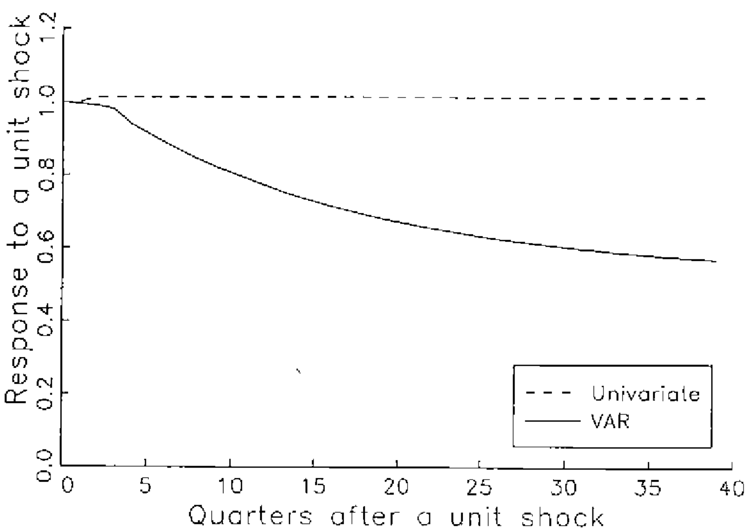

Fig. 7. Univariate labor income impulse-response functions. The "Univariate" response is estimated from a regression of income growth on 4 own lags. The "VAR" response is estimated from a vector autoregression of income and consumption growth on two lags and the lagged consumption/income ratio.

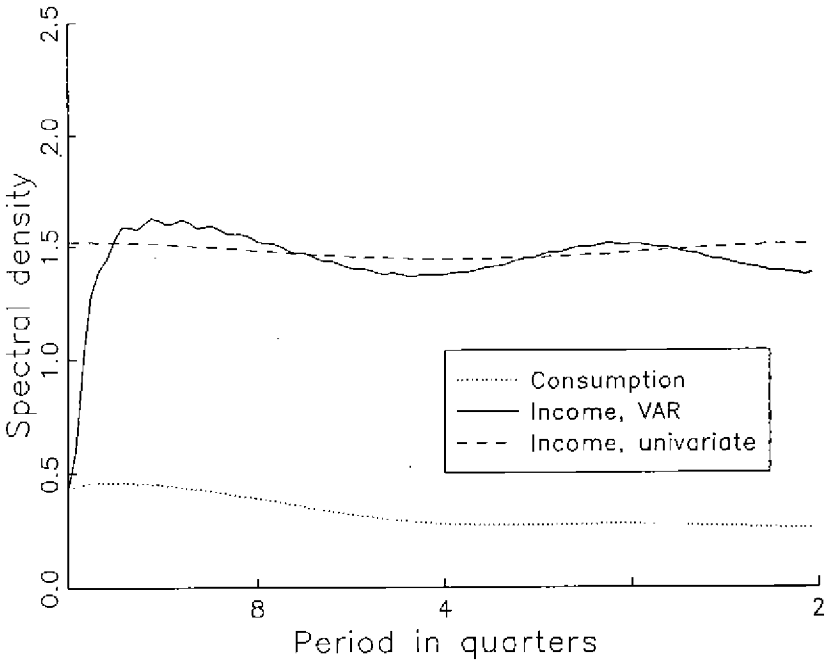

Fig. 8 Spectral density of labor income and consumption growth. The consumption spectral density and the "VAR" income spectral density are estimated from a vector autoregression of consumption and income growth on two lags and the lagged consumption/income ratio. The "univariate" income spectral density is estimated from a regression of income growth on four own lags.

$$\begin{bmatrix}\Delta e_t \\ \Delta c_t\end{bmatrix} = \mu + D(L)\nu_t \qquad ; E(\nu_t \nu_t') = I, \ E(\nu_t \nu_{t-j}') = 0$$

Denoting the elements of $\nu_t = [\nu_t^e \ \nu_t^c]'$, the permanent income hypothesis (4.1) implies[8]

$$\Delta c_t = D_{21}(\lambda)\nu_t^e + D_{22}(\lambda)\nu_t^c$$

and hence

$$\text{var}(\Delta c_t) = D_{21}(\lambda)^2 + D_{22}(\lambda)^2 \ .$$

When consumption growth is expressed in its univariate representation,

$$\Delta e_t = \mu + a(L)\epsilon_t \ , \quad E(\epsilon_t^2) = 1 \ ,$$

we predict

$$\text{var}(\Delta c_t) = a(\lambda)^2 .$$

Again, a(L) can be directly estimated from an income autoregression or inferred from the VAR.

Table 3 presents the results. Note that income growth has about twice the standard deviation of consumption growth. Thus, if income is a random walk, consumption is in fact excessively smooth. Line 2 of the table replicates this result: the income autoregression implies that consumption growth should have a standard deviation of 1.52%, compared to the actual value of 0.58%. However, when we estimate the parameters of the income autoregression with the VAR in line 3, the predicted standard deviation of consumption drops to 0.45%, less than the actual value. Similarly, when we forecast income from the VAR system itself, line 1, the predicted standard deviation of consumption is 0.65%. In either case, the evidence for excess

---

[8] One can make an even stronger prediction by noting that Δc_t is the same variable on the left and right hand side. According to the PIH, we should see $\Delta c_t = \nu_t^c$, so the stronger prediction is $D_{21}(\lambda) = 0$, $D_{22}(\lambda) = 1$ (See Hansen, Roberds and Sargent (1989).)

sensitivity vanishes when we use multivariate information to forecast labor income.[9]

These calculations are intended to illustrate the importance of multivariate rather than univariate forecasts of income, and the forecast power of the lagged consumption/income ratio in particular, using the rules of the game of recent permanent income studies. They are not intended as a resolution of all the many puzzles confronting empirical implementations of the permanent income hypothesis. In particular, I intentionally do not address the following issues:

1) *Non-testability*. Hansen, Roberds and Sargent (1990) show that the present value part of the permanent income model is not testable. (The Euler equation prediction that consumption growth should not be forecastable is, of course, testable.) If consumers really do only see past income in making consumption decisions, or any other set of variables observed by econometricians, then the model predicts an easily rejected stochastic singularity: equation (4.1) holds with no error. If consumers have access to variables not observed by the econometrician, and if there is a single nondurable consumption good, Hansen, Roberds and Sargent show that there is one testable restriction, namely that the present value of the change in forecasts of future income following a \$1 consumption change should be \$1. (This is a restriction on a regression of income on lagged consumption only.) Restrictions on other aspects of a VAR income forecast are not robust to the possibility that agents posses superior information. Furthermore, they show that if nondurables are only one component of consumption, even that one restriction is not testable.

---

[9] This explanation is related to Quah's (1990). Quah examined decompositions of a persistent univariate income process into components not observed by econometricians, but that agents could be imagined to observe, that would explain the relative variances of consumption and labor income. Here, consumers are assumed to only see the income process, and the puzzle is resolved by noting that VAR estimates of that process predict about the observed standard deviation of consumption.

2) *Non-cointegration of consumption and labor income.* The PIH model predicts that consumption and labor income should not be cointegrated, and hence that the consumption/labor income ratio contains a random walk. (consumption should be cointegrated with capital income). Yet the VAR imposes that labor income and consumption are cointegrated. The data suggest the same: the consumption/labor income ratio is stable over time, as the labor and capital shares of income are stable over time.

The prediction that consumption and labor income are not cointegrated comes from the linear technology adopted by the PIH. No matter how much capital consumers accumulate, this has no effect on their labor income. In growth models with nonlinear production functions, accumulating a large amount of capital raises wage rates and hence links the level of consumption (wealth) and labor income. This observation suggests that the instability of the consumption/labor income ratio is not a serious prediction of the PIH, which is meant as a local approximation. (However, human capital may be linearly accumulated, and the unskilled labor income/consumption ratio may not be stable over time. The non-cointegration prediction may apply better to this nonstandard interpretation of the variables.)

3) *Specification issues.* The model is in nonseasonally adjusted per-capita levels. I followed Campbell and Mankiw (1989) in applying it to seasonally adjusted logs. The specification above does not allow for time aggregation, nonseparable or nonquadratic preferences, durability in goods, and ignores growth. This "resolution" also ignores the fact that consumption growth is forecastable.

4) *Who cares?* The PIH is tested as a complete general equilibrium model. It is no longer a "consumption function"--a small part of a larger model. Since the state of the art in empirically oriented stochastic general equilibrium models has advanced beyond quadratic utility and linear technology, why bother testing the PIH? (Cochrane (1990) makes this point in detail.)

5. Mean reversion in stock returns

Table 4 presents a VAR of dividends and prices (cumulated returns), and a regression of returns on lagged returns. The data are from the CRSP value-weighted NYSE portfolio. They are annual to avoid the seasonal in dividends. More lags and other right hand variables (term premium, default premium, interest rate) add more wiggles to the short-run impulse-response functions, but again do not alter the pattern of the long-run impulse-response functions. The table and subsequent figures use the entire CRSP sample, from 1927 to 1988. Results using postwar data are quite similar.

The results in table 4 are similar to the previous results for GNP and consumption. The lagged dividend/price ratio significantly forecasts returns, but not dividend growth. Dividends look a lot like a random walk, as do returns when regressed only on lagged returns. (In postwar data, the dividend/price ratio forecasts both returns and dividend growth more strongly. The t statistics rise from 2.1 to 4.00, and 0.78 to 2.71 respectively. However, the impulse-response functions are quite similar.)

Fig. 9 and fig. 10 present multivariate and univariate impulse-response functions. Note the similarity of the multivariate impulse-response (fig. 9) to the consumption-GNP impulse-response in fig. 1. In response to a dividend shock, prices and dividends move immediately to their long run values. However, a price shock with no movement in dividends is completely transitory.

As I mentioned in the introduction, these results suggest a natural interpretation. A shock to dividends can occur with no effect on discount rates, and hence no effect on expected returns. Since the shock to dividends seems to have a permanent effect on dividends, it should and does have a permanent effect on prices. A pure shock to discount rates should affect prices and not earnings or dividends, as the price shock does in the VAR. A shock to discount rates changes expected returns, and thus induces future temporary variation in prices. As discount rates revert to their mean,

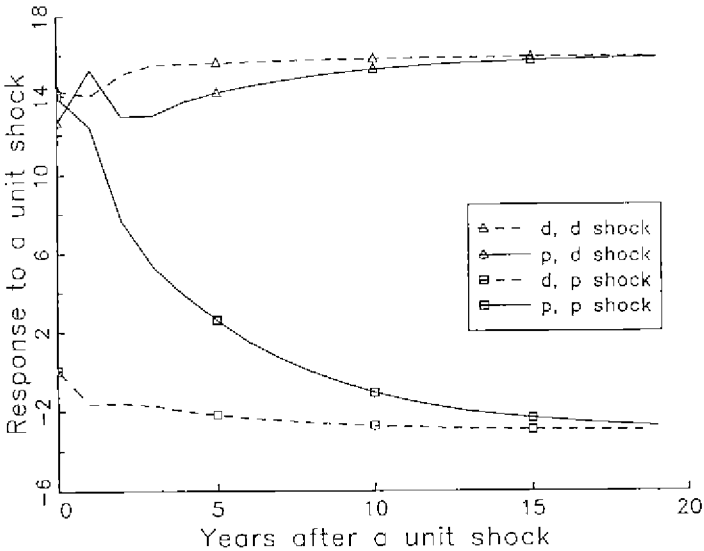

Fig. 9. Impulse-response function for dividend-price VAR. Responses of log dividends and log prices (cumulated returns) from the value weighted NYSE to one standard deviation dividend and price shocks. The VAR includes the lagged dividend/price ratio and 2 lags of dividend growth and returns. The shocks are orthogonalized to force to 0 the instantaneous response of dividends to a price shock.

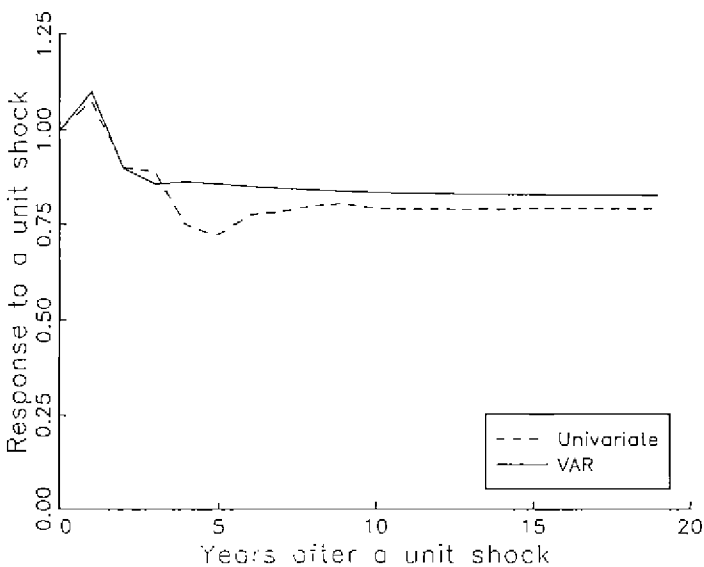

Fig. 10. Univariate stock price impulse-response function. The "Univariate" response is estimated from a regression of returns on 4 lagged returns. The "VAR" response is estimated from a vector autoregression of returns and dividend growth on 2 lags and the lagged dividend/price ratio.

prices revert to their normal multiple of (unchanged) dividends.

Fig. 10 presents univariate impulse-response functions for stock prices. One is directly estimated from the regression of returns on past returns, while the other is implied from the VAR. In contrast to the GNP regressions, these look the same. Thus, both the VAR and univariate estimate show little evidence for univariate mean reversion; evidence for mean reversion in prices (or predictability in returns) comes when you isolate a discount rate shock, as a price shock with no movement in dividends.

6. Concluding Remarks

Lagged growth rates of GNP, labor income and stock prices have little forecast power for future growth rates, and imply little mean-reversion in those variables. Yet multiple regressions using the consumption/GNP, consumption/labor income and dividend/price ratio imply much larger mean reversion in GNP, labor income and stock prices. Since these ratios are stable, consumption and dividends provide information about the "trend" to which GNP, labor income and stock prices must return.

In part, the ratios indicate mean reversion because they produce higher R^2 in forecasting regressions. But much more importantly, one can define multiple shocks with multiple forecasting variables. Thus, the response of GNP to a GNP shock that holds consumption constant is almost entirely transitory, as the permanent income hypothesis suggests; stock prices have a large transitory response to a shock to prices that holds dividends constant, as a change in discount rates suggests. On the other hand, shocks to consumption and dividends induce permanent changes in GNP, labor income and prices. Since univariate shocks are a combination of the two multivariate shocks, they mask the underlying mean reversion.

These observations help to document the existence of temporary components in GNP and stock prices, they help to define useful stochastic trends, and they indicate a resolution of the "excess smoothness" puzzle of consumption.

I conclude that if one is going to say that a variable is nearly a random walk, or that it displays some mean-reversion, it is crucially important to say what information set one has in mind. It is quite possible to find that a variable is nearly a random walk with respect to one information set (its own lags) but has large temporary variation with respect to another. Stable ratios (consumption/GNP, dividend/price) with near-random walks are some of the most important other variables to consider for the issue of long-run mean reversion.

Table 1.

Consumption and GNP Regressions

1. Vector autoregression

Left hand variable	const.	$y_{t-1}-c_{t-1}$	Δc_{t-1}	Δc_{t-2}	Δy_{t-1}	Δy_{t-2}	R^2	p-val
			OLS coefficients					
Δc_t	-0.43	-0.02	0.07	-0.02	0.09	-0.02	0.06	.007
Δy_t	5.19	0.08	0.52	0.16	0.22	0.14	0.27	.000
			t ratios					
Δc_t	-0.49	-1.23	0.90	-0.19	1.91	-0.40		
Δy_t	3.49	3.45	3.81	1.12	2.74	1.89		

(Right hand variable: const. through Δy_{t-2})

2. Univariate autoregression.

Left hand variable	const.	Δy_{t-1}	Δy_{t-2}	Δy_{t-3}	Δy_{t-4}	R^2	p-val
			OLS Coefficients				
Δy_t	0.56	0.33	0.19	-0.11	-0.11	0.18	0.000
			t ratios				
Δy_t	4.82	4.17	2.39	-1.37	-1.36		

(Right hand variable: const. through Δy_{t-4})

y_t is log real GNP. c_t is log of nondurable + services consumption. Δ denotes first differences, $\Delta y_t = y_t - y_{t-1}$. Data are quarterly, 1947:1-1989:3. "p-val" gives the probability value of an F-test for the joint significance of the right hand variables.

Table 2.

Consumption and Labor Income Regressions

1. Vector autoregression

Left hand variable	Right hand variable							
	const.	$e_{t-1}-c_{t-1}$	Δc_{t-1}	Δc_{t-2}	Δe_{t-1}	Δe_{t-2}	R^2	p-val
	OLS coefficients							
Δc_t	0.702	0.010	0.031	-0.049	0.098	0.071	0.07	0.003
Δe_t	0.930	0.087	0.574	-0.162	-0.054	0.005	0.11	0.000
	t ratios							
Δc_t	5.88	0.57	0.37	-0.59	2.39	1.83		
Δe_t	3.78	2.48	3.31	-0.95	-0.64	0.06		

2. Univariate autoregression.

Left hand variable	Right hand variable				
	const.	Δe_{t-1}	Δe_{t-2}	R^2	p-val
	OLS Coefficients				
Δe_t	0.77	0.002	0.013	0.00	0.99
	t ratios				
Δe_t	6.07	0.03	0.17		

e_t is log labor income, where labor income = personal disposable income less dividend, interest and rent income (CITIBASE series GYD - GPRENJ - GPDIV - GPINT), converted to 1982 dollars using the personal disposable income deflator (GYD82/GYD). c_t is log of nondurable + services consumption, in percent units. The sample is 1947:1-1989:3. p-val gives the probability value of an F-test for the joint significance of the right hand variables.

Table 3.

Standard deviation of consumption growth and predictions from the permanent income model.

Std. dev. of consumption growth Δc	0.58 %
Std. dev. of labor income growth Δe	1.23 %

Estimation	Income forecast	Predicted s.d. Δc
1. VAR	VAR	0.65 %
2. Univariate	Univariate	1.52 %
3. VAR	Univariate	0.45 %

If labor income e_t has a moving average representation

$$\Delta e_t = a(L)w_t, \quad E(w_t w_t') = I,$$

then the PIH predicts the variance of consumption growth should be

$$\text{var}(\Delta c_t) = a(\lambda)a(\lambda)' \quad , \quad \lambda = 1/(1+r).$$

In line 1, a(L) is the income row of the moving average representation of a VAR in which Δe_t and Δc_t are regressed on two own lags and $c_{t-1}-e_{t-1}$, the lagged consumption/income ratio. In line 2, a(L) is the moving average representation of a univariate autoregression in which Δe_t is regressed on two own lags. In line 3, the parameters of the univariate autoregression are inferred from the VAR. The results are the same to two decimal places for interest rates r between 1% and 10% per year. Income is log personal disposable income less dividend, interest and rent income (CITIBASE series GYD - GPRENJ - GPDIV - GPINT), converted to 1982 dollars using the personal disposable income deflator (GYD82/GYD). Consumption is log consumption of nondurables and services (CITIBASE series GCN82 + GCS82). The sample is 1947:1-1989:3.

Table 4.

Dividend and Price Regressions

1. Vector autoregression

Left hand variable	const.	$p_{t-1}-d_{t-1}$	Δd_{t-1}	Δd_{t-2}	Δp_{t-1}	Δp_{t-2}	R^2	p-val
	Right hand variable							
	OLS coefficients							
Δd_t	20.01	0.038	0.046	0.062	-0.082	0.040	0.038	0.320
Δp_t	78.65	0.225	0.060	-0.086	0.114	-0.090	0.140	0.012
	t ratios							
Δd_t	0.78	0.47	0.25	0.34	-0.65	0.32		
Δp_t	2.34	2.11	0.25	-0.36	0.68	-0.55		

2. Univariate autoregression.

Left hand variable	const.	Δp_{t-1}	Δp_{t-2}	Δp_{t-3}	Δp_{t-4}	R^2	p-val
	Right hand variable						
	OLS Coefficients						
Δp_t	11.91	0.075	-0.179	0.015	-0.18	.061	.497
	t ratios						
Δp_t	3.46	0.57	-1.37	0.12	-1.37		

p_t is the log price (cumulated return) on the CRSP value- weighted NYSE portfolio. Thus, Δp is the log return. d_t is the corresponding log dividend (monthly dividends brought forward to the end of the year at the market return). Data are annual, 1927- 1988. "p-val" gives the probability value of an F-test for the joint significance of the right hand variables.

Appendix

1. Characterizing persistence

Start from a Wold moving average representation for first differences (log growth rates) Δy_t, which may be inferred from a regression of Δy_t on its past values.

$$\Delta y_t = \mu + \sum_{j=0}^{\infty} a_j \epsilon_{t-j}, \qquad \epsilon_t = \Delta y_t - \mathrm{Proj}(\Delta y_t | \Delta y_{t-1}, \Delta y_{t-2}, \ldots)$$

or, in lag operator notation,

$$\Delta y_t = \mu + a(L)\epsilon_t.$$

The a_j give the response of the growth rate at $t+j$ to a unit shock at t. Similarly, $\sum_{j=0}^{k} a_j$ gives the response of the level of y_{t+k} to a unit shock at t. The limiting value of the response of y_{t+k} to a unit impulse as k increases is $\sum_{j=0}^{\infty} a_j = a(1)$.

The series y_t may also be decomposed into a random walk and a purely stationary component. It turns out that no matter how one does this, the innovation variance of the random walk component is the same, and equal to the spectral density at frequency 0 of Δy_t. Both quantities are related to the univariate impulse response function by

$$\mathrm{var}(\Delta \text{random walk component}) = S_{\Delta y}(0) = a(1)^2 \sigma_\epsilon^2 .$$

In particular, the Beveridge-Nelson trend defined below is a random walk with this innovation variance. (See Cochrane (1988) for a derivation.)

Thus one may equivalently characterize the persistence of univariate shocks to y_t by the behavior of the univariate impulse-response function at high lags, the innovation variance of a random walk component, or the spectral density of growth rates at frequency 0. If y_t follows a random walk, the impulse-response function is one at all horizons and the spectral density at zero is equal to its value at other frequencies. Processes more persistent that a random walk feature impulse-response functions that rise past one, and a spectral density at zero larger than elsewhere. Mean-reverting processes have impulse response functions that fall, as y_t reverts following a shock, and spectral densities at zero lower than elsewhere. In the limit that the *level* y_t is stationary, the impulse-response function falls all the way back to zero, and the spectral density of Δy_t is zero at frequency zero.

The above generalizes naturally to multivariate systems. For example, when one variable Δc_t is added, we write the joint Wold representation as

$$\Delta x_t = \mu + C(L)\xi_t$$

where

$$\Delta x_t = \begin{bmatrix}\Delta y_t \\ \Delta c_t\end{bmatrix}, \quad \xi_t = \begin{bmatrix}\delta^y_t \\ \delta^c_t\end{bmatrix} = \Delta x_t - \mathrm{Proj}(\Delta x_t | \Delta x_{t-1}, \Delta x_{t-2}, \ldots)$$

and A(L) is a matrix of lag polynomials.

Now $\sum_{j=0}^{k} C_j$ gives the response of y_{t+k} (and c_{t+k}) to unit shocks at t, and $\sum_{j=0}^{\infty} C_j = C(1)$ measures the limiting value of this impulse-response function. Including $y_t - c_t$ on the right hand side of the VAR imposes that C(1) is singular, i.e. that the limiting responses of y_t and c_t to each shock is the same.

2. VAR estimation and transformations.

A. VAR

I started by estimating a cointegrated VAR in error-correction form,

$$\Delta y_t = \beta^y_0 + \beta^{yy}_1 \Delta y_{t-1} + \ \ldots + \beta^{yc}_1 \Delta c_{t-1} + \ldots + \beta^y (y_{t-1} - c_{t-1}) + \delta^y_t$$

$$\Delta c_t = \beta^c_0 + \beta^{cy}_1 \Delta y_{t-1} + \ \ldots + \beta^{cc}_1 \Delta c_{t-1} + \ldots + \beta^c (y_{t-1} - c_{t-1}) + \delta^c_t$$

Table 1 presents estimates of this VAR for consumption and GNP. In vector notation, the VAR may be written

$$A(L)\Delta x_t = \beta_0 + \gamma\alpha' x_{t-1} + \xi_t \qquad E(\xi_t \xi_t') = \Sigma \qquad (A.1)$$

where

$$\Delta x_t = \begin{bmatrix} \Delta y_t \\ \Delta c_t \end{bmatrix}, \quad \xi_t = \begin{bmatrix} \delta_t^y \\ \delta_t^c \end{bmatrix}, \quad \gamma = \begin{bmatrix} \beta^y \\ \beta^c \end{bmatrix}, \quad \alpha = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

and A(L) is a matrix of lag polynomials.

B. Impulse-Response function

First, I orthogonalized the error terms by Choleski decomposing the variance covariance matrix of the innovations, in the order consumption, income, other variables. This is equivalent to including *current* consumption growth in the income regression, so all contemporaneous correlation between the consumption and income errors in the Wold representation is assigned to the consumption shock. Precisely, I found a triangular matrix R such that

$$RR' = E(\xi\ \xi') = \Sigma.$$

Then we can define new errors

$$\nu_t = R^{-1}\xi_t \qquad E(\nu_t\nu_t') = R^{-1}RR'R^{-1'} = I \quad .$$

Rewriting the VAR (A.1) in terms of these new errors,

$$A(L)\Delta x_t = \beta_0 + \gamma\alpha' x_{t-1} + R\nu_t \qquad E(\nu_t\nu_t') = I$$

To find the implied Wold moving average or VAR impulse-response function, I simulated the response of the VAR to the ν_t shocks without the constants and starting from initial conditions

$$y_{t-j} = 0, \ c_{t-j} = 0, \ \Delta y_{t-j} = 0, \ \Delta c_{t-j} = 0$$

This leads to the representation

$$\Delta x_t = \mu + D(L)\nu_t. \tag{A.2}$$

The elements of D(L) are the impulse-response functions plotted in Fig.s 1, 6 and 9.

C. Spectral density

I constructed spectral densities of GNP and consumption growth from (A.2) by

$$S_{\Delta x}(e^{-i\omega}) = D(e^{-i\omega})D(e^{i\omega})'. \tag{A.3}$$

D. Implied univariate impulse response

To find the univariate impulse response function for GNP implied by the VAR, I factored the spectral density of Δy_t. Precisely, the first row of (A.3) is

$$S_{\Delta y}(z) = D_{11}(z)D_{11}(z^{-1}) + D_{12}(z)D_{12}(z^{-1})$$

To find the Wold representation of Δy_t and hence its univariate impulse response, one must find a polynomial $a(z)$ whose roots are all on or outside the unit circle, and such that

$$a(z)a(z^{-1}) = D_{11}(z)D_{11}(z^{-1}) + D_{12}(z)D_{12}(z^{-1}).$$

To do this, I found the roots of the right hand side numerically, selected the roots outside the unit circle, and then constructed the polynomial $a(z)$ with those roots. This polynomial $a(L)$ is the univariate impulse-response function.

E. Beveridge-Nelson trend

Writing a one-lag VAR (without means) in companion form, we obtain

$$\begin{bmatrix} \Delta c_t \\ \Delta y_t \\ \Delta c_{t-1} \\ \Delta y_{t-1} \\ y_t - c_t \end{bmatrix} = \begin{bmatrix} ----\ \beta^c\ ------ \\ ----\ \beta^y\ ------ \\ 0\ 1\ \ ---\ 0\ -- \\ 0\ 0\ 1\ --\ 0\ -- \\ \beta_y - \beta_c + [00001] \end{bmatrix} \begin{bmatrix} \Delta c_{t-1} \\ \Delta y_{t-1} \\ \Delta c_{t-2} \\ \Delta y_{t-2} \\ y_{t-1} - c_{t-1} \end{bmatrix} + \begin{bmatrix} \delta_t^c \\ \delta_t^y \\ 0 \\ 0 \\ \delta_t^y - \delta_t^c \end{bmatrix}$$

$$x_t = B \quad x_{t-1} + \epsilon_t.$$

Where the β's are OLS regression coefficients. Then we can calculate the trend from (3.1) as

$$z_t = y_t + [\ 0\ 1\ 0\ 0\ ] \sum_{j=1}^{\infty} B^j\ x_t = [\ 0\ 1\ 0\ 0\ ]\ B\ (I-B)^{-1}\ x_t.$$

I followed this procedure, generalized to the number of lags in each VAR.

References

Beveridge, Stephen, and Nelson, Charles R. (1981), "A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the 'Business Cycle'" *Journal of Monetary Economics* 7 (March), 151-174.

Blanchard, Olivier Jean and Danny Quah (1989), "The Dynamic Effects of Aggregate Supply and Demand Disturbances", *American Economic Review* 79 (September), 655-673.

Campbell, John Y. and Angus Deaton (1989), "Why is Consumption So Smooth?" *Review of Economic Studies* 56 (July), 357-373.

Campbell, John Y. and N. Gregory Mankiw (1987), "Are Output Fluctuations Transitory?" *Quarterly Journal of Economics* 102 (November), 857-880.

Campbell, John Y. and N. Gregory Mankiw (1989), "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence", *NBER Macroeconomics Annual* 1989, Edited by Olivier J. Blanchard and Stanley Fisher, 185-216.

Clark, Peter K. (1987), "The Cyclical Component in U.S. Economic Activity", *Quarterly Journal of Economics* 102 (November), 797-814.

Cochrane, John H. (1988), "How big is the Random Walk in GNP?" *Journal of Political Economy* 96 (October), 893-920.

Cochrane, John H. (1989), "A Critique of the Application of Unit Roots Tests", *Journal of Economic Dynamics and Control*, Forthcoming.

Cochrane, John H. (1990), "Why Test the Permanent Income Hypothesis? Comments on 'The Response of Consumption to Income: a Cross Country Investigation' by John Y. Campbell and N. Gregory Mankiw." (Presented at the International Seminar on Macroeconomics, Mannheim, Germany, June 1990.) Manuscript, University of Chicago.

Cochrane, John H. and Argia M. Sbordone, (1988) "Multivariate Estimates of the Permanent Components of GNP and Stock Prices" *Journal of Economic Dynamics and Control* 12 (June/September), 255-296.

Cogley, Timothy (1990), "International Evidence on the Size of the Random Walk in Output", *Journal of Political Economy* 98 (June), 501-518.

Deaton, Angus S. (1987). "Life-Cycle Models of Consumption: Is the Evidence Consistent With the Theory?" In *Advances in Econometrics: Fifth World Congress*, Vol. 2, edited by Truman F. Bewely. New York: Cambridge University Press.

Fama, Eugene F. (1990), "Transitory Variation in Investment and Output", Manuscript, University of Chicago.

Fama, Eugene F. and Kenneth R. French (1988a), "Permanent and Temporary

Components of Stock Prices", *Journal of Political Economy* 96, 246-273.

Fama, Eugene F. and Kenneth R. French, (1988b), "Dividend Yields and Expected Stock Returns", *Journal of Financial Economics* 22, 3-25.

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

Flavin, Marjorie A. (1988) "The Excess Smoothness of Consumption: Identification and Interpretation", NBER Working Paper 2807 (December)

Hansen, Lars Peter, William Roberds, and Thomas J. Sargent (1990), "Time Series Implications of Present Value Budget Balance and of Martingale Models of Consumption and Taxes", Manuscript, University of Chicago.

Hodrick, Robert J. (1990), "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement", Manuscript, Northwestern University.

King, Robert G., Charles I. Plosser, James H. Stock and Mark W. Watson (1989), "Stochastic Trends and Economic Fluctuations", Manuscript.

Nelson, Charles R. and Charles I. Plosser (1982), "Trends and Random Walks in Macroeconomic Time Series: Some Evidence and Implications", *Journal of Monetary Economics* 10 (September), 139-162.

Quah, Danny (1990) "Permanent and Transitory Movements in Labor Income: An Explanation for 'Excess Smoothness' in Consumption" *Journal of Political Economy* 98:3 (June), 449-475

Poterba, James M., and Lawrence H. Summers (1988) "Mean Reversion in Stock Prices: Evidence and Implications " *Journal of Financial Economics* 22, 26-59.

Richardson, Matthew (1989), "Temporary Components of Stock Prices: a Skeptic's View", Manuscript, Stanford University.

Sargent, Thomas J. (1987), *Macroeconomic Theory*, 2nd Ed. New York: Academic Press.

Watson, Mark W. (1986), "Univariate Detrending Methods with Stochastic Trends," *Journal of Monetary Economics* 18 (July) 49-75.

Watson, Mark W. (1990), "Measures of Fit for Calibrated Models", Working paper, Northwestern University.

West, Kenneth D. (1988), "The Insensitivity of Consumption to News about Income" *Journal of Monetary Economics* 18 (July), 49-75.

Zellner, Arnold, and Franz Palm (1978) "Time-Series Analysis and Simultaneous Equations Models", *Journal of Econometrics* 2, 17-54.